# Machine-Learning-Based Design of Quantum Systems for Extreme Sensing

by

Liang-Ying Chih

B.S., Physics, National Taiwan University, 2017

A thesis submitted to the Faculty of the Graduate School of the University of Colorado in partial fulfillment of the requirements for the degree of Doctor of Philosophy Department of Physics 2022

Committee Members:

Murray J. Holland, Chair Dana Z. Anderson Graeme Smith

Scott E. Parker

Joshua Combes

### Chih, Liang-Ying (Ph.D., Physics)

Machine-Learning-Based Design of Quantum Systems for Extreme Sensing

Thesis directed by Prof. Murray J. Holland

The constant improvement in metrological devices has advanced the frontiers of science. To push the limits of precision, the next generation of experiments will be comprised of highly engineered quantum systems. The intrinsic complexity of quantum mechanics often makes engineering quantum systems to perform a specified task a challenging design problem. We explore a new paradigm for the design and control of devices for performing quantum metrology tasks at levels that surpass what is achievable by conventional methods. In this thesis, we present three levels of design philosophy. The first is the periodic modulation of a control parameter that drives the system at a resonant frequency, leading to the amplification of a specific signal. We use the generation of momentum-squeezed particle-pairs from a driven Bose-Einstein condensate as an example. In the second level, we take advantage of a machine learning technique, reinforcement learning, to optimize the control of quantum states and operations. This approach is demonstrated through the design of the matterwave-optic components used for interferometry in a shaken optical lattice. Going one step further to the third level, we apply reinforcement learning in an end-to-end manner to design the whole protocol for a metrology purpose. This idea is illustrated through the design of a device that measures rotational signals, i.e., a gyroscope, in a two-dimensional shaken optical lattice. The design is not constrained to the conventional configuration, and leads to a remarkable gain in sensitivity. These design methodologies open new avenues to discovery with the potential to significantly advance the state-of-the-art quantum sensors.

## Acknowledgements

I am lucky to have spent five wonderful years in JILA, where I was provided with a comfortable research environment and met many awesome people. First of all, I am truly grateful to my advisor, Murray Holland, for kindly supporting me throughout my PhD and offering unlimited freedom for me to explore different topics. His endless curiosity and creativity has had great influence on me as a thinker. I am fortunate to have him as a mentor and a friend.

I would like to thank everyone in the Holland group for being fun to work with. I have always enjoyed our conversations, whether they are about physics or not. In particular, I have learned a lot from stimulating discussions with Haonan Liu, Athreya Shankar, John Bartolotta, Simon Jäger, Jarrod Reilly, and John Wilson. It has also been a pleasure working with Livia Gutierre, Gage Harmon, and Arnon Goldberg.

I would like to thank everyone on the gyro team, including Josh Combes, Penina Axelrad, Marco Nicotra, Alex Zouyula, Catie DeLesma, Kendall Mehling, Margaret Raybek, and Jieqiu Shao for offering a wide range of expertise. Special thanks go to our team leader, Dana Anderson, for leading interesting and fruitful discussions in the weekly gyro meeting.

I would like to thank friends from Boulder, including Yu-Ju Lee, Xin Xie, Kia Boon Ng, Yu-Ping Lin, Yang-Zhi Chou, Tianhan Lu, Chih-Ting Hsu, and many more. Not only have they offered a lot of help, but their company has also made my PhD life more colorful.

I owe a lot of what I have accomplished today to my high school physics teacher, Yi-An Chen. Without his inspiration and encouragement, I would not have embarked on a journey in physics. Lastly but not least, I would like to thank my parents for their support in my life.

## Contents

## Chapter

1	Intro	oduction	1
<b>2</b>	Qua	ntum metrology	<b>5</b>
	2.1	Bayes theorem	5
	2.2	Ramsey interferometry	6
	2.3	Fisher information	8
	2.4	Heisenberg limit	9
	2.5	Squeezed states	10
	2.6	Mach-Zehnder interferometer	11
	2.7	Matterwave interferometer	13
		2.7.1 Accelerometer	13
		2.7.2 Bragg interferometer	14
	2.8	Gyroscope and Sagnac effect	16
3	Mac	hine learning	17
	3.1	Three branches of machine learning	18
	3.2	Reinforcement learning framework	19
	3.3	Tabular Q-learning	23
	3.4	Deep learning basics	26
		3.4.1 Neural network	26

		3.4.2 Loss function	27
		3.4.3 Gradient descent and backpropagation	27
	3.5	Double deep Q-learning	28
		3.5.1 Improving the learning process	30
	3.6	Variants of reinforcement learning	31
	3.7	Summary	32
4	Driv	ing quantum-correlated atom pairs in Bose-Einstein condensates	34
	4.1	Introduction	34
	4.2	Formalism: General many-body field theory	37
	4.3	Renormalization of the Scattering Potential	40
	4.4	Self-consistent ground state solution	42
	4.5	Dynamics of the time-dependent HFB system	48
	4.6	Dynamically Generating Squeezed Quasiparticle States	52
	4.7	Interferometry with squeezed quasiparticles	56
	4.8	Quasi-2D system	61
	4.9	Conclusion	67
5	Rein	forcement-learning-based matter-wave interferometer in a shaken optical lattice	71
	5.1	Introduction	71
	5.2	Physical model	73
	5.3	Model-free reinforcement learning for design	74
	5.4	Component design	80
		5.4.1 Beam splitter	81
		5.4.2 Details of reinforcement learning for target state	82
		5.4.3 Mirror	84
		5.4.4 Details of reinforcement learning for target unitary	85
		5.4.5 Matter-wave interferometer	87

v

	5.5	Statistical analysis	87
	5.6	Conclusion	90
6	Prac	ctical considerations for the shaken lattice accelerometer	95
	6.1	From natural units to experimental units	95
	6.2	Noise sources in the experiment	96
		6.2.1 Laser linewidth	96
		6.2.2 Laser intensity fluctuation	99
		6.2.3 Quantum efficiency	102
		6.2.4 Dark counts	102
	6.3	Dynamic range	102
	6.4	The ultimate limit for sensitivity	106
	6.5	Outlook	106
7	How	to train your gyro: Reinforcement learning for rotation sensing with a shaken optical	
	Latt	ice	108
	7 1	Introduction	108
	7.1	Physical model	100
	7.2	Two path Sagnag matter wave interferometer	111
	7.3	Ffration 1D madel	111
	7.4		114
	7.5	End-to-end design with reinforcement learning	118
	7.6	Results for end-to-end design	119
	7.7	Conclusion and outlook	124
8	Sum	amary and outlook	125

## Bibliography

# Appendix

$\mathbf{A}$	Numerical methods for the Bose-Einstein condensate	134
	A.1 Imaginary-time propagation	134
	A.2 Spatial derivative in 1D	135
	A.3 Spatial derivative in 2D	135
в	Simulation techniques for the shaken lattice	137
	B.1 Bloch's theorem	137
	B.2 Momentum basis	137
	B.3 Accelerating reference frame	138

## Tables

## Table

5.1	Choice of hyperparameters for the splitting task. The exploration probability $\epsilon$
	decays linearly for each episode at the rate represented by $\epsilon$ decay. The optimizer
	we use, 'Adam' (which stands for adaptive momentum estimation), is a method for
	efficient stochastic gradient descent [42]
5.2	Choice of hyperparameters for the reflection task

## **Figures**

### Figure

- 2.1 (a) An illustration the Bloch vector (red arrow) for the state  $\cos \theta |\uparrow\rangle + e^{i\phi} \sin \theta |\downarrow\rangle$  on the Bloch sphere. (b) An illustration of the Ramsey sequence on the Bloch sphere. From left to right, the state is initialized in the  $|\downarrow\rangle$  state, then a  $\pi/2$  is applied that rotates the Bloch vector onto the equator. The signal couples to the spin as a precession about the z-axis by an angle  $\phi$ , and finally the angle is converted to a population measurement. (Figure credit: Athreya Shankar.) . . . . . . . . . . . . . . .
- 2.2 A sequence that generates the N00N state in a 4-spin system, represented as the Husimi Q-functions. (a) A coherent spin state pointing in the y-direction. (b) Starting with the coherent spin state shown in (a) and applying the one-axis twisting Hamiltonian for a duration τ such that χτ = π/2, we get the maximally-squeezed state. (c) Rotating the spin squeezed state in (b) along the x-axis for angle π/2, we get the N00N state. The important characteristic of the N00N state is that there are N fringes lying on the equator, and therefore one can discriminate a rotation angle that is N times smaller than that for a coherent spin state. This advantage comes at the cost of the reduced dynamic range due to the repetition in the distribution.
  12
  2.3 (a) The optical Mach-Zehnder interferometer. (b) The Bragg interferometer, which

- 3.1 The fundamental feedback loop in reinforcement learning. The agent selects an action based on the current state, and sends the action to the environment, where the action is executed. The environment then outputs the updated state, along with a reward that indicates how ideal the environmental configuration is after the update. 20
- 3.2 (a) A visualization of the Q-table for the game of tic-tac-toe. (b) The rates of winning, losing, and tying for the agent that plays first as a function of the total training episode.
  25

- Densities of the condensate and the non-condensate at t = 0.04 and  $0.08 \ (mR_0^2/\hbar)$ . 4.9The dark blue area is the range of simulation, and the white line indicates the trapping potential. The vellow arrows represent the density currents, with their lengths proportional to the magnitude. The system starts with the ground state with  $a_{
  m dc} = 1.99 \times 10^{-5} \ l_{
  m \perp}$  and then the scattering length is modulated with amplitude  $a_{\rm ac} = 1.99 \times 10^{-4} \ l_{\perp}$  for all time. The smoothing parameter of the finite circular well was set to  $\zeta = 0.2$ . At  $t = 0.04 \ (mR_0^2/\hbar)$ , the atoms in the condensate are pushed towards the edge of the trap, but because they are of low energy, they cannot escape the trap. The number of high momentum atoms is still relatively small, and therefore the density current is also small. The density currents at  $r = 0.4 R_0$  point outwards while those at  $r = 0.8 R_0$  point inwards. Outside the trap, there are few atoms, and the currents are essentially zero. At  $t = 0.08 \ (mR_0^2/\hbar)$ , a large fraction of noncondensate atoms with high energy escape the trap, and the density currents near the edge of the trap and outside the trap point outwards with large magnitudes. For clarity, the density currents at  $t = 0.04 \ (mR_0^2/\hbar)$  are scaled up 3 times compared to
- 4.10 The size of the non-condensate as a function of time. The size is defined as the radius R that encircles 96.2% of the non-condensate atoms. The system starts with initial chemical potential  $\mu = 26 \ (\hbar^2/mR_0^2)$  and is driven with modulation frequency  $\omega = 1000 \ (\hbar/mR_0^2)$ . The speed of the propagation v is given by the slope of the fitted line (red),  $v = 30.8 \ (\hbar/mR_0)$ . This value is approximately equal to  $\sqrt{(\hbar\omega 2\mu)/m}$ , the speed of a particle whose kinetic energy is half the photon energy minus the chemical potential. We have excluded the data points at times when the motion is not ballistic, including at early times where few non-condensate atoms escape the trap and the interaction between the condensate and the non-condensate is significant. 69

- 5.1 Interferometers are composed of (i) a beam splitter, (ii) mirrors, and (iii) a recombiner. Examples shown are; (a) an optical Mach-Zehnder interferometer (using half-silvered or conventional mirrors), (b) a Bragg interferometer (using three short light pulses of varying pulse area,  $\pi/2$  or  $\pi$ ), and (c) a shaken lattice interferometer. In (c), the shaken lattice mimics the interferometer components by splitting, reflecting, and recombining the atoms, through design of a specific shaking function for each case.
- 5.2Framework of reinforcement learning. The agent chooses an action, the environment responds to the action, and gives the next state and a reward as feedback. In our design task, the action is the translation of the optical lattice, the environment evolves according to the Schrödinger equation, the observation is the momentum population distribution, and the reward is a function of the quantum fidelity. Illustrated on the left is an example neural network as the decision-making agent. The neural network takes the state represented on its input layer, passes it through a hidden layer, and 765.3An example topology of the network that we used in the deep Q-algorithm. . . . . 795.4Shaking function for splitting. The learned lattice phase,  $\phi(t)$ , (top) is allowed to take on one of a discrete set of five possible values that span the range shown. The momentum probability distribution (bottom) is initialized to the ground state of the lattice and at the terminal time well approximates the desired superposition for a 91Shaking function for reflection. The learned lattice phase,  $\phi(t)$ , (top) is sinusoidal 5.5and the amplitude is allowed to take on any one of a discrete set of five values at each half-cycle. The momentum distribution (bottom) is prepared in  $-4\hbar k_L$  and 92

- 5.6 Time evolution of the matter-wave density throughout the entire interferometry sequence. The white dotted lines separate the plot into five regions. In region I, we apply the splitting protocol. In region II, we allow the matter-wave to propagate freely in the lattice. The appearance of two wave packets traveling in opposite directions shows that the beam splitter operates as expected. In region III, we apply the mirror shaking function. The matter-wave undergoes free propagation in region IV again, and the two wave packets switch directions, demonstrating the functionality of the mirror. Last, we apply the recombining protocol in region V. Apart from the main closed diamond-shaped paths, we observe that there are auxiliary paths that are fainter but still clearly evident. They arise due to the imperfection of the components, and also due to the side peaks arising from the third-excited Bloch state. 93
- 5.7 (a) Posterior probability density of the acceleration for the first 100 atoms. The true acceleration that we aim to reveal by the measurements is  $-3 \times 10^{-3} \omega_r v_r$  (red line), and the measurements are sampled from the momentum distribution at the end of the interferometry sequence, as shown in the inset. (b) Standard deviation of the acceleration estimated using Bayes theorem for up to  $10^4$  atoms. We show the results from both the shaken lattice interferometer (SLI) and the Bragg interferometer and conclude that the shaken lattice interferometer has a higher sensitivity. The standard deviations  $\sigma_a$  are roughly inversely proportional to the square-root of the number of measurements N, for  $N > 10^2$ . The the black dashed lines are the Cramér-Rao (CR) lower bounds and scale exactly as  $1/\sqrt{N}$ .

6.2	(a) Reconstruction of the true acceleration in the presence of intensity noise $\xi$ =
	$2 \times 10^{-4} \omega_r$ . (b) Standard deviation of the acceleration. (c) Mean squared error of
	the acceleration
6.3	(a) Reconstruction of the acceleration from a system with laser intensity noise as
	shown in (c). (b) The standard deviation of the acceleration as calculated from the
	reconstruction. (c) Time series of the noise. The noise is defined as the ratio between
	the intensity fluctuation and the mean intensity
6.4	(a) Reconstruction of acceleration with $\mathcal{E} = 0.9$ . (b) Standard deviation. We assume
	that there are originally 10000 atoms, but since $\mathcal{E} \neq 1$ , we only get 9035 counts on
	the detector. The inset shows a zoom-in of $N = 8000$ to $N = 9035$ 103
6.5	Error and standard deviation in the presence of dark counts. The absolute value of
	the error grows larger as the noise level increases, whereas the standard deviation
	remains almost the same
6.6	Momentum distribution for a wide range of acceleration. The free-propagation time
	is set to 10 $\omega_r^{-1}$
6.7	(a) A plot of the estimated acceleration versus the true acceleration. Each accel-
	eration is reconstructed from 10000 samples. The linear relation verifies that the
	estimations are unbiased. (b) A scatter plot of the error. It can be seen that the
	error is uniformly distributed across the whole range of true acceleration. (c) A
	histogram of the error, where the width of each bin is $4 \times 10^{-6} (\omega_r v_r)$
6.8	(a) Final momentum distribution for an interferometer where the free propagation
	time is 1000 $\omega_r^{-1}$ . (b) A zoom-in of (a) near $a = 0$ . (c) The error and the standard
	deviation of the estimation of the acceleration from plot (b) with 10000 samples. (d)
	Histogram of the estimation error

- 7.1 A two-dimensional optical lattice potential in a rotating frame, where atoms are trapped and their dynamics can be controlled by the shaking of the lattice. The blue arrows indicate possible trajectories associated with the splitting of the matterwaves. 110
- 7.3 The Fisher information for the machine-learned shaken-lattice gyroscope normalized with respect to the Fisher information that results from the gyroscope in which the matterwaves follow the ideal paths of the conventional two-path inteferometer. . . . 115

- 7.6 Phase and the spatial wavefunction as a function of time. The shaking protocol is generated from the reinforcement learning agent with  $\dot{\phi}$  as the actions. There are 125 steps in this protocol. The wavefunction is initialized in the ground state with a momentum width  $\sigma_p = 0.1$ , and evolve with the shaking protocol shown above. 121
- 7.7 Population distribution in the momentum basis and the resulting classical Fisher information for the machine-learned shaken-lattice gyroscope as a function of the rotation rate  $\Omega$ . The value of the Fisher information is normalized with the Fisher information of a conventional two-path interferometer with splitting  $\Delta p = 8\hbar k_L$ . . . 122

## Chapter 1

#### Introduction

The constant improvement of ultra-stable lasers with the linewidth now reaching below the 10 mHz level [67], and the ability to cool atoms below the  $\mu$ K regime, have opened up the door into the quantum engineering realm. The extreme level of control over ultracold quantum systems has lead to exciting applications such as quantum simulation, quantum computing, and quantum sensing. In particular, quantum sensing is currently achieving astounding levels of precision and accuracy, which provides a pathway to challenge the boundaries of physics. For example, we now have optical clocks that are capable of discriminating the gravitational redshift on the Earth's surface at the scale of millimeters [12], and atom interferometers that can support a coherent superposition of two clouds of atoms even when they are half a meter apart, which allows for the detection of the gravitational tidal force [47].

One of the key ingredients in cold-atom-based quantum sensors is the Bose-Einstein condensate (BEC). A BEC is the state of matter formed when a cloud of Bosonic particles is cooled to near absolute zero temperature, such that their de Broglie wavelength is much larger than the interparticle distance. In this case, their wave functions are mostly overlapped, and can be collectively described by a single macroscopic wave function. Since the matter wave of the BEC is highly coherent, it is an ideal medium to use for interferometric measurements, in analogy with how highly coherent lasers are used to carry out precise interferometry.

An example of using BEC to build an interferometer is the experiment by Weidner et al. [100]. In their work, they loaded a BEC into an optical lattice formed by two counter-propagating lasers, and sought to construct an accelerometer by 'shaking' the lattice with highly engineered patterns. Through a variety of optimal control methods, they found the shaking functions that gave rise to the splitting, symmetric translation, reflection, and recombination of the matter waves, and combined these components in the appropriate order for interferometry. The acceleration signals were then extracted via curve fitting from the interference patterns of the momentum distribution observed in the time-of-flight images.

When the interactions between atoms in a BEC come into play, even richer and more fascinating dynamics can occur. For instance, in the experiment by Clark et al. [18], they periodically modulated the interaction strength between atoms in a BEC, and observed firework patterns formed by the ejected atoms. This experiment is nicknamed the 'Bose firework'. Intrigued by the experimental observations, we set out to build a model to explain the phenomenon. From our analysis, we found that the high-momentum atoms generated from the BEC using this relatively straightforward experimental technique are actually entangled in momentum and form metrologically useful quantum states.

The examples above showcase how high-quality and diverse our toolbox is for the control of quantum systems. Nevertheless, designing quantum systems for a specific application is still a hard problem, due to the multiple interference paths or the exponential scaling of states. For all but the simplest quantum systems, it is in general difficult to figure out what is the best strategy to combine the tools we have to fully leverage the 'quantumness'.

But what if we ask the computer to do the engineering for us instead? Machine learning has celebrated spectacular success over the last decade, due to the exponential increase of computational power and the recent developments of sophisticated algorithms. Specifically, a branch of machine learning called reinforcement learning has been shown to be useful when a brute-force search of strategies is not possible. It has also demonstrated the ability to learn beyond human intuition and imagination. The design of quantum sensors fits into this paradigm, and therefore we believe the marriage of quantum sensing and machine learning can become powerful. Moreover, since the reinforcement learning approach does not rely on knowledge of the underlying structure of the system that we want to control, we think it is likely that this will open a new paradigm to quantum design, enabling scaling up of the complexity, and allowing the control of the real-world experimental platforms at JILA and elsewhere.

In this thesis, we present three example designs of quantum systems based on ultracold atoms for the purpose of high-precision sensing. These examples correspond to three levels of design philosophy, starting from the most intuitive one to one that often generates solutions that cannot easily be anticipated.

First, we investigate a control protocol that generates quasimomentum pairs from a Bose-Einstein condensate by the periodic modulation of the scattering length, as demonstrated in the Bose firework experiment. This protocol allows the selective amplification of quantum fluctuations with a predetermined momentum, where the momentum value can be spectroscopically tuned. Although this control protocol corresponds to the most basic design philosophy that does not involve any learning, it is still interesting in that the atoms generated form momentum-squeezed matter waves. These squeezed matter waves have reduced noise in one quadrature and could potentially be injected into an atom interferometer that aims to achieve a noise level below the standard quantum limit. The simplicity of this design philosophy also allows us to fully understand the underlying mechanism and explain the results.

Next, we take advantage of reinforcement learning to design the matterwave-optic components in a one dimensional lattice. Just like in the original shaken lattice experiment of Weidner, we aim to find the shaking functions that optimize the quality of the components. With reinforcement learning, we find the control protocol that leads to a larger momentum splitting of the matter wave. The improvement is partly due to the difference in the exploration space for the shaking functions. Indeed, our shaking functions bear more dramatic variations than those obtained from optimal control algorithms. This example demonstrates the potential effectiveness of reinforcement learning for the control of well-defined quantum states and operations.

Moving forward to the third level, we consider applying reinforcement learning in an endto-end manner to design the entire control protocol in a two-dimensional lattice with the goal of achieving the optimal precision for a gyroscopic device. This is the stage where the constraints set by what physicists have learned from the conventional configurations will be broken, and the fundamental limits to the achievable precision will be challenged.

This thesis is structured as follows. In the next two chapters, we will introduce some background knowledge that will help understanding the material that follows. Chapter 2 will focus on the basic theory of quantum metrology, which explains why quantum sensors can provide advantage in sensitivity over classical ones. We will also introduce a couple of text-book standard techniques. and they will serve as the baseline for comparison with the more complex techniques developed later. As mentioned previously, we have considered the use of machine learning for the design of quantum metrological devices. Seeing that machine learning is still a somewhat foreign field for many physicists, it will be necessary to provide a comprehensible guide to it. This will be carried out in Chapter 3, where we give an overview of machine learning before we dive into reinforcement learning, lay out the algorithms that we employ, and go through a few pedagogical examples. After the necessary background is introduced, we present our original research. In Chapter 4, we describe our theoretical investigation of the Bose firework. In Chapter 5, we present the design framework and the results for the shaken lattice accelerometer. Then in Chapter 6, we extend our study from the previous chapter to real-world applications by taking into consideration the geometry of the experimental apparatus and the noise sources that may affect the performance of the shaken lattice accelerometer. In Chapter 7, we present the gyroscope designed by the end-to-end reinforcement learning approach. Finally, we conclude the thesis in Chapter 8 and provide perspectives for future work.

## Chapter 2

### Quantum metrology

Quantum metrology is the study of making measurements with quantum systems to achieve a precision that cannot be reached by their classical counterparts. The central limit theorem in classical statistics tells us that the uncertainty of a parameter scales inversely with the square root of the number of independent measurements taken. By taking advantage of quantum entanglement, it is, however, possible to improve upon this classical scaling [29].

In this chapter, we will review the theoretical aspect of fundamental ideas in quantum metrology. Understanding statistics is the first step in extracting useful information from measurement data, and this is what we will start with. Next, we will introduce a few protocols commonly used in quantum metrology, including Ramsey interferometry and the 'squeezing' of quantum states. Finally, we will present several experimental techniques that put the abstract protocols mentioned previously into practice. The ideas introduced in this chapter layout the underlying principles of quantum metrology on which more complicated extensions or variants can be built.

## 2.1 Bayes theorem

The Bayesian philosophy is founded on the framework in which there is a fixed probability model for the data, or interchangeably, measurement outcomes, and from the data one can infer the underlying parameters of that model. Bayesian parameter estimation usually works when we know exactly what the forward distribution is, that is, we have a model that predicts the probability distribution of the measurement outcomes for each possible value that the parameter can take. Assume that one has a prior belief of the parameter  $\phi$ ,  $P(\phi)$ , where P(...) refers to the probability, and the model describes the probability of getting the data m conditioned on the parameter, i.e.,  $P(m|\phi)$ . Then, the probability distribution for  $\phi$  based on the measured data is found from the conditional probability as

$$P(\phi|m) = \mathcal{N}^{-1} P(m|\phi) P(\phi) , \qquad (2.1)$$

where  $\mathcal{N} = \int P(m|\phi')P(\phi')d\phi'$  is a normalizing factor. If we have multiple measurement outcomes,  $\{m_1, \ldots, m_N\}$ , then we can apply Eq. (2.1) iteratively, using the updated probability as the prior to the next measurement, and get

$$P(\phi|m_1,\ldots,m_N) \propto P(m_N|\phi)\ldots P(m_1|\phi)P(\phi).$$
(2.2)

The proportionality is resolved by having the posterior probability be normalized to unity.

Typically, the probability distribution becomes narrower as the number of measurements, N, increases. One finds, in practice, that if each measurement is independent, then the standard deviation typically converges as  $1/\sqrt{N}$  due to the central limit theorem.

### 2.2 Ramsey interferometry

Ramsey interferometry a measurement strategy very often applied in atomic physics [45]. Basically, it is made up of a sequence of operations that goes as follows. Initially, the atom is prepared in  $|\downarrow\rangle$  state. A  $\pi/2$  pulse associated with a rotation about the *y*-axis is applied and rotates the state to the  $|+x\rangle$  state. Then, the atom is subject to a rotation around the *z*-axis by an amount that is associated with the signal that we want to measure. After that, another  $\pi/2$  pulse is applied, and we can infer the phase accumulated during the free evolution from a population measurement in the *z* basis. The Bloch sphere representation of the sequence is shown in Fig. (2.1).

From a general point of view, the Ramsey sequence is made up of three elements: state preparation, interaction with the field that bears the signal, formally represented as a rotation induced by a generator, and finally, a unitary transformation to an appropriate basis followed by



Figure 2.1: (a) An illustration the Bloch vector (red arrow) for the state  $\cos \theta |\uparrow\rangle + e^{i\phi} \sin \theta |\downarrow\rangle$  on the Bloch sphere. (b) An illustration of the Ramsey sequence on the Bloch sphere. From left to right, the state is initialized in the  $|\downarrow\rangle$  state, then a  $\pi/2$  is applied that rotates the Bloch vector onto the equator. The signal couples to the spin as a precession about the z-axis by an angle  $\phi$ , and finally the angle is converted to a population measurement. (Figure credit: Athreya Shankar.)

a projective measurement. A thriving theme in quantum metrology research is to tailor the initial probe state and the measurement basis to the signal of interest such that the measurements are optimal with respect to a certain set of metrics, which is a non-trivial task in a multi-partite or multi-level system.

## 2.3 Fisher information

Fisher information is a metric that characterizes how much information one can gain about the value of a parameter from a measurement sampled from a distribution. The classical Fisher information is defined as

$$I^{(c)}(\phi) = \mathbb{E}_{\phi} \left[ \left( \frac{\partial}{\partial \phi} \log P(m|\phi) \right)^2 \right]$$
$$= \sum_{m} \frac{1}{P(m|\phi)} \left[ \frac{\partial P(m|\phi)}{\partial \phi} \right]^2, \qquad (2.3)$$

where  $\mathbb{E}_{\phi}$  is the expectation value evaluated at the parameter  $\phi$ , m denotes all the possible measurement outcomes, and  $P(m|\phi)$  is the conditional probability of getting a specific outcome m.

The Fisher information may also be calculated from

$$I^{(c)}(\phi) = \sum_{m} \left[ -\frac{\partial^2 \log P(m|\phi)}{\partial \phi^2} \right] P(m|\phi), \qquad (2.4)$$

which can be derived from Eq. (2.3) by making use of the identity  $\sum_{m} \frac{\partial^2 P(m|\phi)}{\partial \phi^2} = 0$ . With a sequence of N measurements,  $\{m_1, m_2, \ldots, m_N\}$ , the summation should be over all the possible measurement trajectories, weighted by the corresponding probabilities that can be obtained from Bayes theorem if independence of the measurements is assumed. The resulting Fisher information is then

$$I_{N}^{(c)}(\phi) = \sum_{\{m_{N}\}} \cdots \sum_{\{m_{1}\}} \left[ -\frac{\partial^{2} \log\left(\prod_{i} P(m_{i}|\phi)\right)}{\partial \phi^{2}} \right] \prod_{i} P(m_{i}|\phi)$$

$$= \sum_{\{m_{N}\}} \cdots \sum_{\{m_{1}\}} \left[ -\sum_{i} \frac{\partial^{2} \log P(m_{i}|\phi)}{\partial \phi^{2}} \right] \prod_{i} P(m_{i}|\phi)$$

$$= \sum_{i} \sum_{\{m_{i}\}} \left[ -\frac{\partial^{2} \log P(m_{i}|\phi)}{\partial \phi^{2}} \right] P(m_{i}|\phi)$$

$$= NI_{1}^{(c)}(\phi), \qquad (2.5)$$

where we have used the identity  $\sum_{\{m_N\}} \dots \sum_{\{m_{i+1}\}} \sum_{\{m_{i-1}\}} \sum_{\{m_1\}} \prod_{j \neq i} P(m_j | \phi) = 1$  to get from the second line to the third line. What the relation proven above tells us is that the Fisher information for N independent measurements is N times larger than the Fisher information for a single measurement.

The inverse of the Fisher information gives the lower bound of the variance of an unbiased estimator for the parameter, i.e.,

$$\sigma_{\phi}^2 \ge 1/I^{(c)}(\phi) \,. \tag{2.6}$$

This lower bound is called the Cramér-Rao bound. With this relation, we can see Fisher information as a metric for sensitivity, since the sensitivity of a metrological device is characterized by how small the variance is for the estimated parameter. As can be deduced from Eq. (2.5), the standard deviation scales with the number of measurements as  $1\sqrt{N}$ . Using a Bayesian estimator, the standard deviation approaches the Cramér-Rao bound at a large number of measurements.

The quantum analogue of the Fisher information tells us what is the maximum achievable information that one can extract from a quantum state. Even with an optimal measurement basis, the classical Fisher information can only reach the upper bound set by the quantum Fisher information. With a pure state as the probe state, the quantum Fisher information for the rotation angle can be calculated from the variance of the generator  $\hat{G}$  with

$$I^{(Q)}(\phi) = 4 \left[ \left\langle \hat{G}^2 \right\rangle - \left\langle \hat{G} \right\rangle^2 \right].$$
(2.7)

## 2.4 Heisenberg limit

Before we dive into the Heisenberg limit, let us define the standard quantum limit. The standard quantum limit, despite its name, is what can be achieved with classical measurements. It says that the standard deviation of a phase distribution for independent measurements scales as  $1/\sqrt{N}$ . The Heisenberg limit, on the other hand, is the ultimate limit that can be achieved in a multi-partite system with quantum entanglement. The name is motivated by the Heisenberg uncertainty principle, from which one would expect a general relationship  $\Delta N \Delta \phi \sim 1$ . With correlated particles, the standard deviation may therefore approach 1/N, where N is the number of particles in the system. This uncertainty principle scaling can also be proven with the quantum Fisher information. To make it more concrete, the Heisenberg limit can be understood with a Ramsey interferometry example. With N correlated atoms, such as the 'N00N state',  $|N00N\rangle \equiv \left(|\uparrow\rangle^{\otimes N} + |\downarrow\rangle^{\otimes N}\right)/\sqrt{2}$ , the phase accumulated during free evolution is N times larger, i.e.,  $\exp\left(-i\hat{S}_z\phi\right)|N00N\rangle = \left(e^{-iN\phi/2}|\uparrow\rangle^{\otimes N} + e^{iN\phi/2}|\downarrow\rangle^{\otimes N}\right)/\sqrt{2}$ . As a result, the phase uncertainty can be reduced by a factor of N, but at the expense of dynamic range due to the aliasing effect that arises.

Of course, this is a qualitative description as there is no formal definition of the phase operator. In fact, due to the periodicity of phase, its standard deviation cannot exceed  $2\pi$ . For example, it is clear that a Fock state, which has zero variance in number, violates the so-called Heisenberg inequality of phase and number. Nevertheless, these ideas are useful in many situations.

#### 2.5 Squeezed states

Squeezed states have reduced variance in one observable at the cost of a larger variance in the conjugate observable. They become metrologically useful when the direction with reduced variance aligns with the observable one would like to measure. In quantum optics, squeezed states of light can be generated in a nonlinear crystal by a physical process referred to as parametric down conversion. The Hamiltonian that describes the process can be written as

$$\hat{H} = -i\hbar \frac{\chi}{2} \left( \hat{a}^2 - \hat{a}^{\dagger \, 2} \right) \tag{2.8}$$

where  $\hat{a}$  and  $\hat{a}^{\dagger}$  are the annihilation and creation operators for the photon. Deriving the dynamics from the Hamiltonian in the Heisenberg picture, we get  $\hat{a}(t) = \hat{a}(0) \cosh \chi t + \hat{a}^{\dagger}(0) \sinh \chi t$ , and therefore the average number is  $\langle 0| \hat{a}^{\dagger}(t) \hat{a}(t) | 0 \rangle = \sinh \chi t^2$ . We introduce the two quadratures as

$$\hat{X}_1 = \hat{a} + \hat{a}^{\dagger}, \quad \hat{X}_2 = -i(\hat{a} - \hat{a}^{\dagger}), \quad (2.9)$$

then we arrive at the result that the variance in  $\hat{X}_1$  exponentially increases while in  $\hat{X}_2$  the variance exponentially decreases. For an initial vacuum coherent state, the variances evolve according to

$$\left\langle \hat{X}_{1}^{2} \right\rangle = e^{2\chi t}, \quad \left\langle \hat{X}_{2}^{2} \right\rangle = e^{-2\chi t}.$$
 (2.10)

In SU(2) (atomic) systems, the analogy to the squeezed state of the optical field systems is the spin-squeezed state, which has reduced variance in a direction of the collective spin vector. Spin squeezing can be generated by the one-axis twisting Hamiltonian,

$$\hat{H} = \frac{\hbar\chi}{2}\hat{S}_z^2, \qquad (2.11)$$

where  $\hat{S}_z = \frac{1}{2} \sum_{i=1}^{N} \hat{\sigma}_z^{(i)}$  is the collective spin in the z-direction. The reduced variance in the collective spin can be visualized by the Husimi Q-function. The Husimi Q-function represents the overlap of the state we have with the coherent spin state,  $|\theta, \phi\rangle\rangle \equiv |\theta, \phi\rangle^{\otimes N}$ , where  $|\theta, \phi\rangle$  is the single-qubit spin state labelled by its spin direction on the Bloch's sphere. That is,

$$Q(\theta, \phi) = |\langle \langle \theta, \phi | \psi \rangle|^2$$
(2.12)

When the one-axis twisting Hamiltonian is applied to a coherent spin state pointing in the ydirection for a time duration  $\pi/2\chi$ , we get a maximally squeezed state. If we then rotate the spins to align with the z-axis, the resulting state is the N00N state, up to a constant phase difference. The resulting state has the maximum possible Fisher information with respect to a rotation along the z-axis. The Husimi Q-functions of the coherent spin state, the squeezed state, and the N00N state are shown in Fig. (2.2).

#### 2.6 Mach-Zehnder interferometer

The Mach-Zehnder configuration is commonly used for measuring the phase difference between two optical paths through interference. First, the light is coherently split into two paths with a beam splitter, then a mirror reflects each path, and finally the two paths rejoin and a another beam splitter recombines the light to generate the interference. This configuration can be seen as a



Figure 2.2: A sequence that generates the N00N state in a 4-spin system, represented as the Husimi Q-functions. (a) A coherent spin state pointing in the y-direction. (b) Starting with the coherent spin state shown in (a) and applying the one-axis twisting Hamiltonian for a duration  $\tau$  such that  $\chi \tau = \pi/2$ , we get the maximally-squeezed state. (c) Rotating the spin squeezed state in (b) along the x-axis for angle  $\pi/2$ , we get the N00N state. The important characteristic of the N00N state is that there are N fringes lying on the equator, and therefore one can discriminate a rotation angle that is N times smaller than that for a coherent spin state. This advantage comes at the cost of the reduced dynamic range due to the repetition in the distribution.



Figure 2.3: (a) The optical Mach-Zehnder interferometer. (b) The Bragg interferometer, which is the matter-wave analog of the Mach-Zehnder configuration.

generalization of the Ramsey sequence, and the use of the mirrors is simply to bring the two paths back together in real space so that they can interfere with each other. The configuration is shown in Fig. (2.3).

## 2.7 Matterwave interferometer

The matterwave interferometer takes advantage of the wavelike nature of atoms to carry out interferometry. The de Broglie wavelength of a particle is given by  $\lambda = h/p$ , where p is the momentum of the particle, and h is Planck's constant. One of the benefits of using matterwaves over electromagnetic waves is that the wavelength of the matterwave can be much shorter than that of optical light when the momentum is large, which increases the sensitivity in principle. Particles are subject to gravitational and inertial forces, and therefore matterwave interferometry is useful for gravitational or inertial field sensing.

## 2.7.1 Accelerometer

The phase accumulated in a Mach-Zehnder matterwave accelerometer can be calculated as follows. Without loss of generality, we assume that the matterwave is split into a superposition of two momentum eigenstates,  $|\psi(t=0)\rangle = (|-p_0\rangle + |+p_0\rangle)/\sqrt{2}$ . We assume that the acceleration on the atom is a constant +a (relative to the reference frame that is accelerating in the negative direction), and in the first free propagation period  $t = 0 \sim T$ , the state evolves to

$$\exp\left[-i\left(\frac{\hat{p}^{2}}{2m}-ma\hat{x}\right)T/\hbar\right]|\psi(t=0)\rangle$$

$$=\exp\left(-i\hat{p}^{2}T/2m\hbar\right)\exp\left(ima\hat{x}T/\hbar\right)\exp\left(-aT^{2}[\hat{p}^{2},\hat{x}]/4\hbar^{2}\right)\exp\left(ima^{2}T^{3}[\hat{x},[\hat{p}^{2},\hat{x}]]/6\hbar^{3}\right)|\psi(t=0)\rangle$$

$$=\exp\left(ima^{2}T^{3}/3\hbar\right)\left[\exp\left(iaT^{2}p_{0}/2\hbar\right)\exp\left(-i(p_{0}+maT)^{2}T/2m\hbar\right)|p_{0}+maT\rangle\right.$$

$$+\exp\left(-iaT^{2}p_{0}/2\hbar\right)\exp\left(-i(-p_{0}+maT)^{2}T/2m\hbar\right)|-p_{0}+maT\rangle\right]/\sqrt{2}.$$
(2.13)

The phase difference between the two momentum states in this period is  $\Delta \phi_{\rm I} = -p_0 a T^2 / \hbar$ . Afterwards, a reflection is applied, switching the signs of  $\pm p_0$ , i.e.,

$$|p_0 + maT\rangle \rightarrow |-p_0 + maT\rangle$$
,  $|-p_0 + maT\rangle \rightarrow |p_0 + maT\rangle$ . (2.14)

In the second free propagation period  $t = T \sim 2T$ , we get  $\Delta \phi_{\rm II} = 3p_0 a T^2/\hbar$  by following the same derivation in Eq. (2.13). As a result, the total phase difference accumulated during the interferometry is  $\Delta \phi_{\rm tot} = 2p_0 a T^2/\hbar$ . Note that the total phase difference, which is proportional to the area enclosed in the space-time diagram, is proportional to the momentum splitting,  $2p_0$  and quadratic in the free propagation time.

#### 2.7.2 Bragg interferometer

A crucial part of constructing a Mach-Zehnder matterwave interferometer is the capabilities to coherently split the wavefunction into an equal superposition of two momentum states and swap the two momenta while preserving the relative phase between the states. The simplest way to do this is through Bragg diffraction [56], where a standing-wave laser pulse is applied to the atom, and the atom absorbs a photon in one direction, gets excited to a virtual state, and then emits a photon in the opposite direction and returns to the ground state. This process gives the atom a momentum kick of  $2\hbar k_L$ , where  $k_L$  is the wavenumber of the laser. The picture described above works in the Bragg regime, where the laser intensity is not strong enough to induce higher-order diffraction, i.e., diffraction processes that are not exactly resonant. With the appropriate choice of reference frame, we can assume that the only relevant momenta are  $+\hbar k_L$  and  $-\hbar k_L$ , since they have the same kinetic energy and therefore their energies are degenerate. The Rabi frequency of the laser pulse,  $\Omega$ , should be smaller than the frequency corresponding to the energy cap between the  $|\hbar k_L\rangle$  and the  $|3\hbar k_L\rangle$  states, i.e.,  $\Omega \ll 4\hbar k_L^2/m$ , so that this transition and the ones above are off the energy shell. These assumptions simplify the full Hamiltonian to allow an effective two-level Hamiltonian,

$$\hat{H} = \frac{\hbar\Omega}{2} \left( \left| -\hbar k_L \right\rangle \left\langle +\hbar k_L \right| + \left| +\hbar k_L \right\rangle \left\langle -\hbar k_L \right| \right).$$
(2.15)

If we see the two-momentum model as a spin-1/2 system, then this Hamiltonian is essentially proportional to the Pauli matrix  $\hat{\sigma}_x$ . Applying the laser pulse for a duration  $\tau$  is equivalent to rotating the pseudo-spin around the *x*-axis for an angle  $\Omega\tau$  on the Bloch sphere. This is reminiscent of the first field in Ramsey interferometry as introduced earlier. It is then straightforward to see that to split the atoms from the initial  $|-\hbar k_L\rangle$  momentum state to an equal superposition of the two momentum states, one should apply the laser pulse with integrated intensity such that  $\int_0^{\tau} \Omega(t')dt' = \pi/2$ . This transfers the pseudo-spin from the north pole to the equator of the Bloch sphere. Similarly, recombination is achieved by applying another  $\pi/2$ -pulse. To reflect the momentum states, which is equivalent to swapping the  $|\pm\hbar k_L\rangle$  states, one should apply the laser pulse such that  $\int_0^{\tau} \Omega(t')dt' = \pi$ . An illustration of the pulse sequence is shown in Fig. (2.3).

At the beginning of the Bragg interferometry sequence, the momentum of the atom cloud may not correspond precisely to the desired  $-\hbar k_L$ . If this is the case, one should adjust the laser frequencies so that the standing wave that results from the interference is moving at the velocity that matches the center-of-mass velocity after splitting.

## 2.8 Gyroscope and Sagnac effect

Gyroscopes usually operate based on the Sagnac effect, which gives rise to a phase difference between two counter-propagating paths when the system is rotating. For matter waves in free space, the Hamiltonian is

$$\hat{H} = \frac{\hat{\boldsymbol{p}}^2}{2m} - \boldsymbol{\Omega} \cdot \hat{\boldsymbol{L}}, \qquad (2.16)$$

where  $\hat{L} = \hat{r} \times \hat{p}$ . Since the first term in the Hamiltonian commutes with the second term, we only need to consider the phase shift induced by the rotation term. If we adopt the semiclassical picture and make the substitution  $p = m\dot{r}$ , then the phase difference between the two paths that follow the same trajectory but in the opposite directions is

$$\Delta \phi = \frac{1}{\hbar} \int m \boldsymbol{\Omega} \cdot \boldsymbol{r}^{(+)} \times \dot{\boldsymbol{r}}^{(+)} dt - \frac{1}{\hbar} \int m \boldsymbol{\Omega} \cdot \boldsymbol{r}^{(-)} \times \dot{\boldsymbol{r}}^{(-)} dt$$
$$= \frac{2}{\hbar} \oint m \boldsymbol{\Omega} \cdot \boldsymbol{r} d\boldsymbol{r}$$
$$= \frac{4m \boldsymbol{\Omega} \cdot \boldsymbol{A}}{\hbar}.$$
(2.17)

In the equation above, (+) and (-) labels the counterclockwise and clockwise paths respectively, and the contour integral is along the counterclockwise path, which encloses an area A.

Similarly, the Sagnac effect also gives rise to a phase difference between two counter-propagating paths of optical light in a fiber-optic gyroscope. The phase difference can be deduced from Eq. (2.17) by seeing the factor  $m/\hbar$  as  $p/\hbar v$ , which is  $k_L/c$  for light. The result is  $\Delta \phi = \frac{4k_L \boldsymbol{\Omega} \cdot \boldsymbol{A}}{c}$ , where  $k_L$  is the laser wavenumber and c is the speed of light.

## Chapter 3

#### Machine learning

Machine learning is the field of study that aims to 'teach' machines to make predictions and decisions based on data and observations without explicitly programming them to do so. As a powerful tool, machine learning has recently gained more and more attention in the physical science communities. One specific branch of machine learning, reinforcement learning, has grown in importance over the last few years with the introduction of a computer program, AlphaZero, which demonstrated the capability to beat the best players in highly non-trivial games like qo and chess [85, 87, 86]. The main challenge of developing computer programs that play these games is that the number of possible board configurations grows exponentially, and it quickly becomes impossible to determine what moves are the best through an exhausted search. The design and control of physical systems is very much like a game in some sense. The set of metrics associated with achieving certain design goals is like the score one gets in a game, and the design with the highest score is the winner. The games in this context are also highly non-trivial in that there are usually multiple knobs that can be turned, and the number of combinations for the setup and parameters to be tuned in a time sequence is intractably large. In this chapter, we will briefly introduce different branches of machine learning and their applications, and then dive deeper into a specific branch, reinforcement learning. We will start with an the general framework and then describe in detail how a specific type of reinforcement learning, Q-learning, works. We will also comment on how reinforcement learning can benefit the design of physical systems. The purpose of this chapter is not to present new research in machine learning, but to cover the fundamental knowledge one needs to learn to apply machine learning to the development of new designs for quantum systems.

## 3.1 Three branches of machine learning

Supervised learning is a branch of machine learning that learns from labelled data. One class of problems that supervised learning is used for is regression. An example is learning from historical data of temperature from a certain region, and trying to predict the temperature in that region by weighing in factors like pressure maps, humidity profiles, and terrain. Another class of problems is classification, where the machine learns from pre-labelled data to sort new data into categories. An example is training the machine to recognize hand-written digits with labelled images.

Unsupervised learning refers to a wide variety of machine learning algorithms that learn from unlabelled data. The goal is to find patterns from the distribution of data. One class of such a problem is clustering. Imagine information from customers that includes age, income, etc. represented by data points in a high dimensional space. Algorithms like *k*-means clustering put similar customers into the same group to get insight as to what products they may like according to their common traits with other people. In this case, there are no existing samples labelled by what group they belong to. Instead, how they are grouped together depends on the overall distribution of the customer base.

Another class of problem that unsupervised learning deals with is dimension reduction. To make an example easy to visualize, data points in a three-dimensional space may only be spread out on a two-dimensional plane and have little variance in the direction perpendicular to the plane. One can therefore find the two vectors that span the plane by finding the axes with the highest variances. This method is called the *principle component analysis*. More complex methods include the *autoencoder*, which is a systematic approach to find the latent space representation in the input data.

Reinforcement learning is a unique branch of machine learning that learns from trials and errors instead of from pre-existing examples, whether labelled or un-labelled. This is the branch that we will focus on in this thesis, and will be elaborated on more in the next section.

## 3.2 Reinforcement learning framework

Typically, reinforcement learning involves a feedback loop between an *agent* that invokes *actions* based on the observed *state* of an environment, and an *environment* that provides *rewards* to the agent based on the resulting outcomes. This framework can be formalized as a Markov decision process, where the probability of transition to the next state is only dependent on the current state and the current action, and not on prior history. The agent is tasked with the goal to discover the sequence of actions for which it receives the highest possible long-term reward [90] (see Fig. 3.1). It does this by trial and error, iteratively improving its actions in such a way as to corral the environment toward a target configuration.

Evolution in this framework is referred to as a *trajectory*, where the state is initially prepared, and then a sequence of actions and corresponding state updates are performed. The trajectory steps continue until the chain comes to an end when a predetermined condition is reached, which might be a certain number of steps, or a terminating state of the environment. An important concept is the idea of an *optimum trajectory* starting from any state, which is a trajectory such that the long-term reward is the maximum possible.

Decisions are made according to a policy framework, which we label as  $\pi$ , that instructs the agent how to carry out specific actions based solely on the current environmental state. That is,  $\pi$  takes an environmental state s and maps it to an action a either deterministically or probabilistically.

The learning paradigm is associated with evolving towards an optimal  $\pi$  that maximizes a *return* according to the following algorithm. At each time step, t, the agent receives a reward,  $R_t$ , that is a pure function of the environment and measures the degree to which the environment is in the desired configuration. We have been using the term 'long-term reward', but we refer to is formally defined as the *return*. The return is the cumulative discounted reward, which can be


Figure 3.1: The fundamental feedback loop in reinforcement learning. The agent selects an action based on the current state, and sends the action to the environment, where the action is executed. The environment then outputs the updated state, along with a reward that indicates how ideal the environmental configuration is after the update.

expressed in terms of the immediate reward and the future rewards,

$$G_t = \sum_{k=0}^{\infty} \gamma^k R_{t+k} \,. \tag{3.1}$$

This expression accounts for the fact that actions that are taken at time t affect all later times, and  $\gamma \in [0, 1]$  is the discount factor that characterizes how much the future rewards are valued due to the diminishing influence the current action has as the time goes further into the future.

To evaluate the policies, we define a *value function*, which calculates the expected return starting from a state s, i.e.,

$$V^{\pi}(s) = \mathbb{E}_{\pi} \left[ \sum_{k=0}^{\infty} \gamma^k R_{t+k} \middle| S_t = s \right] , \qquad (3.2)$$

where the expectation value is over all possible trajectories denoted by the sequential map of states to actions  $\{S_{t+k} \xrightarrow{\pi} A_{t+k}\}$ , for  $k \ge 0$ . The optimal value function is the maximum of the value function over all possible policies, i.e.

$$V^*(s) = \max_{\pi} V^{\pi}(s) \,. \tag{3.3}$$

Since the policy is an instruction set as to how to select an action, we also need a way to evaluate the desirability of each action. This leads us to defining the *action-value function*, also known as the Q-function, where Q stands for quality. It calculates the expected return starting from a state s, taking the action a, and then following the policy  $\pi$ , i.e.,

$$Q^{\pi}(s,a) = \mathbb{E}_{\pi} \left[ \sum_{k=0}^{\infty} \gamma^k R_{t+k} \middle| S_t = s, A_t = a \right] .$$
(3.4)

A simple relation between the value function and the action-value function is

$$V^{\pi}(s) = \sum_{a \in \mathcal{A}} \pi(a|s) Q^{\pi}(s,a) \,. \tag{3.5}$$

The optimal action-value function is

$$Q^*(s,a) = \max_{\pi} Q^{\pi}(s), \qquad (3.6)$$

and the optimal policy  $\pi^*$  can be found by maximizing Q(s, a) over all actions, i.e.

$$\pi^*(s) = \operatorname*{argmax}_{a} Q(s, a),$$
  

$$Q^* = \operatorname*{max}_{\pi} Q^{\pi}(s, a).$$
(3.7)

By isolating the immediate reward from the future rewards in Eq.(3.2), we can derive the Bellman expectation equation for the value function,

$$V^{\pi}(s) = \mathbb{E}_{\pi} \left[ R_t + \sum_{k=1}^{\infty} \gamma^k R_{t+k} \middle| S_t = s \right]$$
  
$$= \mathbb{E}_{\pi} \left[ R_t + \gamma V^{\pi}(S_{t+1}) \middle| S_t = s \right]$$
  
$$= \sum_{a \in \mathcal{A}} \pi(a|s) \left[ r(s,a) + \gamma \sum_{s' \in \mathcal{S}} P(s'|s,a) V^{\pi}(s') \right],$$
  
(3.8)

where s' represents the state that follows after the action a is applied to the state s. Similarly, the Bellman expectation equation for the action-value function can be written as

$$Q^{\pi}(s,a) = r(s,a) + \gamma \mathbb{E}_{\pi}[Q^{\pi}(s',a')].$$
(3.9)

By recognizing that when the policy is optimum the expected return should be maximized, we arrive at

$$Q^{*}(s,a) = r(s,a) + \gamma \max_{\pi} \mathbb{E}_{\pi}[Q^{\pi}(s',a')]$$
  
=  $r(s,a) + \mathbb{E}_{s'}[Q^{\pi^{*}}(s',\pi^{*}(s'))]$   
=  $r(s,a) + \mathbb{E}_{s'}[\max_{a'} Q^{\pi^{*}}(s',a')].$  (3.10)

In the scenario where the transition is deterministic, the above equation boils down to the Bellman optimality equation [7]

$$Q^*(s,a) = R_t + \gamma \max_{a'} Q^*(s',a'), \qquad (3.11)$$

which we use as our fundamental design equation. The learning algorithm we will focus on is *Qlearning*, which relies on learning to approximate the actual optimal Q-function as a way to obtain the optimal policy.

# 3.3 Tabular Q-learning

As mentioned previously, learning a good representation of the Q-function is the key to Qlearning. One of the simplest ways one can think of is to store the Q-value for each state-action pair, Q(s, a), in a lookup table where the rows correspond to the states, s, and the columns correspond to the actions, a. The Q-values in the table are updated through the temporal-difference method,

$$Q(s,a) \leftarrow Q(s,a) + \alpha \left( r(s') + \gamma \max_{a'} Q(s',a') - Q(s,a) \right)$$
  
=  $(1-\alpha) Q(s,a) + \alpha \left( r(s') + \gamma \max_{a'} Q(s',a') \right),$  (3.12)

where Q(s, a) moves toward the target return,  $Y(s) \equiv r(s') + \gamma \max_{a'} Q(s', a')$  with learning rate  $\alpha$ . As introduced above, machine learning algorithms typically involve many *hyperparameters* that can be tuned, making the learning tasks non-trivial. The simplicity of tabular Q-learning allows us to look under the hood and thereby to gain access to how the exploration for the optimal solution works. It also has the advantage of being quick, which can become crucial if one were to implement the learning algorithm as a real-time control method in an actual experiment. Tabular Q-learning is suitable for small-sized problems, where an exhaustive list of all the possible states can be stored and processed on the computer with reasonable memory requirements.

Here we show how tabular Q-learning works with a simple example, where we train the computer program to learn to play the game *tic-tac-toe*. In the program, there are two reinforcementlearning agents that represent players on the two sides, and an environment representing the  $3 \times 3$ board configuration. The values in the Q-tables for both agents are all initialized to 0.5, and the board in the environment is initialized to be empty. The agents choose the action associated with the highest value among the values in the row corresponding to the current state in the Q-table. Every time agent 1 makes a move, placing 'O' on the board, the environment evaluates whether the board configuration is terminal. If not, agent 2 will also make a move, placing 'X' on the board, and the resulting board configuration will be sent to this agent 1 as the next state. If the board configuration is terminal, then a reward is given to the agent. The reward is 1 if the agent wins, 0 if the agent loses, and 0.5 if the two agents draw. Once a game is finished, each agent updates their own Q-table based on the sequence of states and actions stored in their memory. The updates start from the last state before the terminal state and the action chosen in the last step. The Q-value for this state-action value is updated to the reward that the agent received for the terminal state, i.e.,  $Q(s_{T-1}, a_{T-1}) \leftarrow r(s_T)$ . Then, going backward in the sequence, the Q-values associated with the state-action pair at each step is updated following

$$Q(s_t, a_t) \leftarrow (1 - \alpha)Q(s_t, a_t) + \alpha \gamma \max_{a'} Q(s_{t+1}, a').$$

$$(3.13)$$

This process is repeated for many iterations until the Q-values in the table converge. We show the winning rate of agent 1, the winning rate of agent 2, and the rate that the two agents draw in Fig. (3.2. Initially agent 1 has a higher winning rate because the game favors the player who plays first if both sides are playing randomly, and after the agents are both fully trained they draw in almost every game.

Of course, teaching the computer to play *tic-tac-toe* may not seem all that impressive, as humans can learn to play this game optimally with few trials. However, a point worth emphasizing is that the agent does not need to know what the environment actually is. To the agent, the actions are just integers that enumerate elements from a set of possible actions. They are only meaningful to the agent through the updated state that the environment feeds back. The complete conceptual separation between the agent and the environment implies that one can potentially solve completely different problems, which are encoded in the environment, with the same learning agent. This is the essence of *model-free learning*. For example, one can replace the environment with a simulation of a physical system, and have the same agent that learned tic-tac-toe learn to control the system to achieve the desired configuration. In fact, an important benefit of having the model-free structure for physicists is that the environment can also potentially be replaced by an actual experimental apparatus, which may not be perfectly described by a theoretical simulation or model, but the agent would still be able to learn from observations from the experiment nonetheless. This is an extremely powerful concept that is not typically shared by gradient-based optimal control approaches that





Figure 3.2: (a) A visualization of the Q-table for the game of tic-tac-toe. (b) The rates of winning, losing, and tying for the agent that plays first as a function of the total training episode.

require the model to be known.

Understanding the simple tabular Q-learning algorithm is a prerequisite to moving forward to more complicated agents, such as the double deep Q-learning algorithm that follows.

## 3.4 Deep learning basics

Deep learning utilizes the cascade of multiple consecutive layers of non-linear functions to approximate functions with complicated behaviors. In reinforcement learning, the *Q*-function is the function we would like to approximate, and deep learning could therefore be useful for this purpose. In this section, we introduce the key ideas of deep learning, including the neural network approach, the loss function, and the learning stage through gradient descent.

#### 3.4.1 Neural network

A typical neural network structure is made up of an input layer, several hidden layers, and an output layer. Each layer is connected to the next one by a linear map followed by a non-linear activation function. This network topology was motivated by the highly connected active neurons in the brain. The most naïve neural network is the simple feed-forward neural network, where each node in a layer is connected to all the nodes in the next layer. In general, the formula for forward propagation can be written as

$$z^{(l)} = W^{(l)}a^{(l-1)} + b^{(l)}$$
  

$$a^{(l)} = g^{(l)}(z^{(l)}), \qquad (3.14)$$

where l labels the layer. In the formula above, W is the weight matrix with dimensionality  $n \times m$ , b is a  $m \times 1$  bias vector, and g(x) represents the non-linear activation function. The weights and biases are the parameters to be trained in the neural network, and are often denoted by  $\theta$  collectively. The initial condition  $a^{(0)}$  is given by the sample input vector  $x_i$ , and  $a^{(L)}$ , where L is the total number of layers, is equivalent to the output vector. Two common activation functions include ReLU, ReLU $(x) = \max(x, 0)$ , and the sigmoid function,  $\sigma(x) = \frac{1}{1+e^{-x}}$ , although there are many

others used. The use of the non-linear activation function is essential, for if it is omitted, the whole function would be a linear transformation, and the multi-layer structure would be redundant. The neural network forms a highly parameterized *hypothesis* function, denoted by  $h_{\theta}(x)$ , which we use to attempt to approximate any actual function with high fidelity.

#### **3.4.2** Loss function

To evaluate the efficacy of the hypothesis, we define a loss function that describes how faithful the hypothesis is to the desired actual function. The smaller the loss, the better the hypothesis is. Of course, we do not know the actual function, for that is exactly why learning is carried out in the first place. Therefore, we use samples generated from the actual functions and evaluate the loss by averaging over these samples. One of the most common loss functions is the mean squared error,

$$\mathbb{L}(\boldsymbol{\theta}) = \frac{1}{N} \sum_{i} |h_{\boldsymbol{\theta}}(x_i) - y_i|^2 , \qquad (3.15)$$

where *i* labels the training sample index, *N* is the number of samples used to calculate the loss,  $x_i$ and  $y_i$  are the input vector and the target vector respectively for each sample.

Typically, the learning cycle in machine learning algorithms involves techniques to minimize a loss function by variation of the adjustable parameters  $\theta$  in order to arrive at the optimal map for the purpose at hand.

### 3.4.3 Gradient descent and backpropagation

To minimize the loss function, the most common numerical approach is gradient descent. With each step of gradient descent, the parameters in the neural network move in the direction of the locally steepest descent of the loss function. The update rule corresponding to this approach is

$$\theta \leftarrow \theta - \alpha \, \frac{\partial \, \mathbb{L}}{\partial \, \theta} \,, \tag{3.16}$$

where  $\alpha$  is the learning rate that determines the step size. If the learning rate is too small, it will take a long time to converge to the minimum loss. However, if the step size is too large, then the gradient descent process may not even converge at all. Since the neural network has multiple connected layers, the gradients need to be calculated using the chain rule of differentiation, and the calculation of the gradients starting from the output layer all the way to the input layer is called *backpropagation*.

## 3.5 Double deep Q-learning

The objective of deep Q-learning is to train the deep Q-network so that it represents the return defined in Eq. (3.1) well and therefore satisfies the Bellman optimality equation,

$$Q(s,a) = r(s') + \gamma \max_{a'} Q(s',a').$$
(3.17)

We train the deep Q-network using the double deep Q-learning algorithm [94]. In this algorithm, an extra network called the target network is included to estimate the future return. The reason to do it this way is that Eq. (3.17) involves a self-consistency element. If during the training stage, the Q-values on both sides of the equation are determined by the same Q-network, then updating the Q-network will simultaneously change both the Q-value Q(s, a) and the return Y(s') that the Q-value is supposed to converge to. This creates a feedback cycle that can make the learning process extremely unstable, and in fact the Q-values may never converge. In order to mitigate this problem, we employ a Q-network,  $Q_{\theta}$ , just as described, from which the optimal action is chosen by the agent through their policy. However, we then calculate the expected return for the optimization of the Q-network using a Q-value corresponding to this action but obtained from a second neural network, the target network,  $Q_{\theta'}$ . The target network is only updated more slowly through a weighted average. This avoids the unstable feedback since the Q-network and target network are weakly correlated at early stages of the learning process [61]. Furthermore, this method resolves issues with overoptimism in the standard deep Q-learning algorithm [94]. A sketch of the algorithm is presented in Algorithm 1.

The algorithm is described as follows. First of all, we initialize the weights and biases in the Q-network randomly, and initialize the target network such that its network parameters are the same as those of the Q-network. Then, we iterate over a large number of episodes. In each episode,

Algorithm 1: Double Deep Q-Learning

Initialize Q-network  $Q_{\theta}$  and target network  $Q_{\theta'} \leftarrow Q_{\theta}$ ; for trajectory=1:episodes do  $s \leftarrow s_0;$ while d=0 do if  $rand() > \epsilon$  then  $a \leftarrow \arg \max_a Q_{\boldsymbol{\theta}}(s, a);$ else  $a \leftarrow rand(\mathcal{A});$ end  $[s', r, d] \leftarrow environment(s, a);$ Store (s, a, s', r, d) in replay buffer; Sample from replay buffer:  $\{(s_i, a_i, s'_i, r_i, d_i) | i = 1, ..., B\};$ 
$$\begin{split} \mathbb{L} &= \frac{1}{B} \sum_{i} \left[ Q_{\boldsymbol{\theta}}(s_{i}, a_{i}) - Y_{i} \right]^{2}, Y_{i} = r_{i} + \gamma \, Q_{\boldsymbol{\theta}'}(s_{i}', \arg \max_{a'} Q_{\boldsymbol{\theta}}(s_{i}', a_{i}')); \\ \text{Update } Q_{\boldsymbol{\theta}}: \, \boldsymbol{\theta} \leftarrow \boldsymbol{\theta} - \alpha \nabla_{\boldsymbol{\theta}} \mathbb{L}; \\ \text{Update } Q_{\boldsymbol{\theta}'}: \, \boldsymbol{\theta}' \leftarrow \tau \boldsymbol{\theta}' + (1 - \tau) \boldsymbol{\theta}; \end{split}$$
 $s \leftarrow s';$ end end

a trajectory of consecutive states and actions is generated. A trajectory starts with an initial state  $s_0$  and ends when we reach a terminal state or a threshold number of steps.

In each step, we choose the action that maximizes the Q-value for the current state most of the time, but with probability  $\epsilon$  we choose a random action. This policy is known as the  $\epsilon$ -greedy policy. In practice,  $\epsilon$  decays from 1 to a number that asymptotically approaches 0 over time, since exploitation becomes more important than exploration as the learning process proceeds. The action is then fed into the environment, and the environment outputs the subsequent state s', a reward r, and a Boolean d that indicates whether the trajectory terminates at this step. A tuple (s, a, s', r, d), which represents the experience we get from taking this step, is pushed into a memory stack known as the replay buffer.

To minimize the error of  $Q_{\theta}(s, a)$  for representing the return, we calculate a loss function defined as the squared difference between  $Q_{\theta}(s, a)$  and  $Y(s') \equiv \sum_{k=0} \gamma^k R(s_{k+1})$ , where  $s_k$  represents the state we get after k steps from s. The future return in Y(s') (with k > 0) is evaluated by the output value of the target network for the next optimal action. The next optimal action is determined by the action that maximizes the output values of the Q-network, given the next state s' as input. Since we need the Q-values to do well for all possible pairs of (s, a), we draw B(the batch size) samples from the replay buffer, where each sample is labelled by an index i, and calculate their average loss.

The parameters in the Q-network are updated through gradient descent in the direction of decreasing average loss. The parameters in the target network are updated by taking the weighted average of the previous parameters from the target network and the updated parameters from the Q-network.

### 3.5.1 Improving the learning process

Although we have outlined the learning process in the previous section, there are a few useful tricks to improve the performance of the learning process that are not discussed much in literature.

The first trick is to iterate the process multiple times, each time adding more complexity

to it. For example, one can consider dividing the time into more steps. Another example is to add more actions to the action set. In tabular Q-learning, this simply means that we can grow the Q-table horizontally, adding more columns associated with the new actions to the table. In deep Q-learning, this involves adding more nodes on the output layer, and initializing the weights and biases associated with these extra nodes randomly while copying the parameters in the smaller model trained in the the previous iteration to the new bigger model. The search space for the optimal protocol is exponentially huge, and one might find learning with a large action set and long sequence difficult. The iterative approach serves as a systematic way to guide the learning process, and often leads to better convergence.

The second trick is inspired by *transfer learning*, where a model trained for one scenario is applied to another scenario with some retraining. As an example, one can imagine that a deep neural network agent is trained first using the computer simulated environment, and then made to interact with an actual experiment. It is anticipated that the learned protocol would not be optimal for the experiment since there are usually imperfections in a real-world scenario. However, the simulated environment should still be close to the experimental environment, and therefore, we could allow only the last few layers of the neural network to vary and be retrained while fixing the rest of the parameters as static. Since the neural network is close to being fully trained, the retraining process is expected to take many fewer episodes. Due to the long run time of the experiment compared to the computer simulation, implementing machine learning on an actual experiment has naturally been a more difficult task than for the numerical counterpart. This trick could potentially save the training time on the experiment, and hopefully make reinforcement learning effective in the real-world scenario.

### **3.6** Variants of reinforcement learning

We have focused on Q-learning for the most part of this chapter. There are in fact a wide variety of alternative learning algorithms for the agent, which we will introduce briefly in this section.

32

In the policy gradient algorithm, the agent makes decisions based on the policy network, which takes the state s as the input and outputs the probability to select each action a, i.e.,  $\pi(a|s)$ . The agent selects a random action at each step with weights based on this probability distribution. The policy is parametrized by a set of variables denoted by  $\theta$ , and the learning process entails maximizing an objective function with respect to  $\theta$ . A common choice of the objective function is defined as the expected value function for the initial state, i.e.,

$$J(\boldsymbol{\theta}) \equiv \mathbb{E}_{\pi_{\boldsymbol{\theta}}}\left[V(s_0)\right] \,. \tag{3.18}$$

The expectation value is evaluated over  $\{\tau\}$ , a set of trajectories sampled over  $\pi_{\theta}$ . Thus, the objective function can be written as

$$J(\boldsymbol{\theta}) = \int \pi_{\boldsymbol{\theta}}(\tau) R(\tau) d\tau$$
  

$$\approx \sum_{i \in \{\tau\}} \sum_{t} R_{t}^{(i)}.$$
(3.19)

To maximize the objective function, we calculate the gradient of it with respect to  $\theta$ ,

$$\nabla_{\boldsymbol{\theta}} J(\boldsymbol{\theta}) = \int \pi_{\boldsymbol{\theta}}(\tau) \nabla_{\boldsymbol{\theta}} \log \pi_{\boldsymbol{\theta}}(\tau) R(\tau) d\tau$$
$$\approx \sum_{i \in \{\tau\}} \sum_{t} \nabla_{\boldsymbol{\theta}} \log \pi_{\boldsymbol{\theta}}(a_t^{(i)} | s_t^{(i)}) \sum_{t} R_t^{(i)}, \qquad (3.20)$$

and update the parameters  $\theta$  by

$$\boldsymbol{\theta} \leftarrow \boldsymbol{\theta} + \alpha \nabla_{\boldsymbol{\theta}} J(\boldsymbol{\theta}) \,. \tag{3.21}$$

The actor-critic model for the agent can be seen as a hybrid of Q-learning and policy gradient. In this framework, there is an actor that selects an action based on its policy network, which outputs the probability distribution of the actions, and a critic that tells the actor how good the selected action is based on evaluation of the corresponding Q-value from its Q-network.

### 3.7 Summary

Machine learning has been demonstrated to be a powerful tool. The power comes at the cost of complexity; there are many moving parts in a machine learning algorithm that can be improved, which is also why machine learning is such an active research field. From a broader perspective, reinforcement learning is the branch of machine learning most relevant to quantum design problems for complex systems since it learns a pathway through a multi-parameter landscape by trial and error. In such a situation, a brute-force search would be prohibitive.

## Chapter 4

#### Driving quantum-correlated atom pairs in Bose-Einstein condensates

#### 4.1 Introduction

The ability to tune the two-body scattering length in a Bose-Einstein Condensate (BEC) by varying the magnitude of a magnetic field in the vicinity of a Feshbach resonance has been employed in a number of seminal experiments that aim to investigate controlled non-equilibrium quantum dynamics. One such example is the so-called 'Bosenova' experiment by Donley *et al.* [21], in which a BEC was subject to a sudden change of the scattering length from a small positive value to a large negative value. This resulted in a change of the sign of the mean-field interactions from repulsive, where the gas is mechanically stable, to attractive, where the compressibility may be negative and the gas is then unstable [79]. What was observed experimentally after this abrupt change in the scattering length was a collapse and subsequent explosion of the quantum gas in a manner that resembled an astrophysical supernova. Theoretical models were developed and illustrated that the emergence of a pairing field in the underlying quantum many-body system can explain the observed burst of non-condensate atoms [59].

More recently, Bose 'firework' experiments [18] have observed pairs of high momentum atoms emitted as jets from a condensate driven by a periodic modulation of the two-body *s*-wave scattering length. These experiments demonstrated a protocol for resonantly amplifying quantum fluctuations with well-controlled momenta when starting from a stationary BEC. The fact that the jets were observed to be correlated in emission direction motivates us to consider whether the many-body pairing field played an important role in the dynamics, in a similar manner to the Bosenova system previously studied. This aspect is related to other calculations that explore the second order coherence of the gas [4].

In the case of a dilute quantum gas, the Gross-Pitaevskii equation (GPE) provides an accurate description of the equilibrium and time-dependent behavior of the BEC. In this framework, the interacting condensate is completely described by a mean-field superfluid order parameter. The GPE framework has been extensively applied to model the behavior of BECs at zero temperature, and also to their coherent manipulation through externally applied potentials. However, when there is a significant portion of non-condensate atoms, the GPE will fail to provide an accurate description of the system. A small amount of non-condensate atoms is always present even at zero temperature in a dilute quantum gas arising from the beyond mean-field fluctuations that are due to the finite interaction strength. It is of course possible to generate substantial fractions of non-condensate atoms by driving a pure condensate in a variety of ways, and this is typically unavoidable when the currents that generate the magnetic confinement fields contain stochastic noise. Furthermore, non-condensate atoms are always present in systems with finite temperature since they embody the thermal excitations. In order to capture the essential dynamics associated with the non-condensed component, a theory that goes beyond the mean-field approximation is necessary.

A systematic extension to the simplest mean-field approach given by the GPE is the Hartree-Fock Bogoliubov (HFB) formalism that takes into account the interactions between three components. These components are the condensate, the non-condensate, and the pairing field of the fluctuations. In this formalism, the elementary excitations are described as Bogoliubov quasiparticles [10], and the ground state condensate is a vacuum of such quasiparticles. The quasiparticle creation operator is a linear combination of the atom creation operator (particles) and the atom annihilation operator (holes). The vacuum state due to interactions possesses a portion of non-condensate atoms referred to as quantum depletion. The 'firework' experiments that tune the scattering length by applying an appropriate external magnetic field potentially allow all of these components—the mean-field, non-condensate, and pairing field—to be controlled, manipulated, and engineered. In this chapter, we derive the solutions of the HFB theory as applied to well-controlled experimental geometries in order to determine the efficacy of this framework for providing a theoretical basis for the observations from the 'firework' experiments.

One approach for describing collective excitations of a condensate involves solving the Bogoliubov de Gennes equations [109], which is most accurate when the excitations are weak. When the excitations are not weak, and the non-condensate fraction can be significant, a more complete approach must be used such as the time-dependent self-consistent HFB equations [16], and that is the method we will focus on in this chapter. Note that one alternative approach that can incorporate the excitations is through the addition of a noise source to introduce fluctuations directly into the Gross-Pitaevskii equation [25]. However, this assumes by construction that the many-body state can be accurately described by a unique macroscopic wavefunction, and therefore a more complete theory is needed to describe two-body correlations.

We emphasize the importance of the pairing field in our analysis. Pairing gives rise to an anomalous density that allows us to investigate the coherence of the system and explore methods to probe phase-sensitive quantities. However, to incorporate the pairing field in our simulations requires a number of important considerations. Since our modeling assumes contact interactions, numerical studies have to account for the potentially divergent nature of the pairing field at both short and long length scales by appropriate renormalization of the scattering potential. We demonstrate how to renormalize the scattering potential when momentum is represented on a discrete grid. Furthermore, when solving for the initial condition of the system, instead of using an approximation that ignores the pairing field [64] in order to remedy issues associated with the gapless energy spectrum, we take an alternative approach in which the condensate, depletion and pairing field are accounted for and we solve for the HFB theory self-consistently.

We investigate the metrological use of the atom pairs generated from the periodic driving protocol. We analyze the Hamiltonian for the excitation fields and associate it with that of optical squeezing. The analogy implies that the variance in of the the quadratures of the excitation field can go beyond the standard quantum limit. Since at high momentum, the quasiparticle basis is close to the momentum basis, we point out that the atoms are ideal candidates for the momentum superposition state used in matterwave interferometry.

The chapter is outlined as follows. We present the model in Section 4.2 and provide details of the renormalization process in Section 4.3. At first we limit our discussion to the most straightforward case of quasi-1D systems. In Section 4.4, we outline the numerical procedures necessary to obtain a self-consistent ground state solution to the HFB theory for a weakly-interacting trapped quantum gas, and quantify the quantum depletion as well as the pairing field amplitude. In Sections 4.5 and 4.6, we use the time-dependent HFB theory to show that the modification of the interaction strength through modulation of the scattering length parametrically amplifies a certain quasiparticle mode and generates a matter-wave solution that is analogous to a squeezed state of light. In Section 4.7, we use these results to explore the possibility of future experiments that utilize interferometry to probe the pair correlation amplitude. We consider two methods that create a phase difference between the driving field and the pairing field, and consequently lead to the possibility for constructive and destructive interference in the matter-wave density. Finally, in Section 4.8, we extend the results to quasi-2D so that they can be compared with the experimental observations of angular correlations in the firework pattern, where the atoms were confined in a pancake-shaped confining potential well and were ballistically expanded.

## 4.2 Formalism: General many-body field theory

We begin from the many-body Hamiltonian that describes a weakly interacting Bose gas with pairwise contact interactions:

$$\mathcal{H} = \int d^3 x \, \hat{\psi}^{\dagger}(\boldsymbol{x}) \left( -\frac{\hbar^2}{2m} \nabla^2 + V_{\text{ext}}(\boldsymbol{x}) \right) \hat{\psi}(\boldsymbol{x}) \\ + \frac{V}{2} \int d^3 x \, \hat{\psi}^{\dagger}(\boldsymbol{x}) \hat{\psi}^{\dagger}(\boldsymbol{x}) \hat{\psi}(\boldsymbol{x}) \hat{\psi}(\boldsymbol{x}) \,, \qquad (4.1)$$

where *m* is the mass of the atom and  $V_{\text{ext}}$  is the external trapping potential. The field operators,  $\hat{\psi}(\boldsymbol{x})$  and  $\hat{\psi}^{\dagger}(\boldsymbol{x})$ , are bosonic operators that annihilate and create particles and obey commutation relations  $[\hat{\psi}(\boldsymbol{x}), \hat{\psi}^{\dagger}(\boldsymbol{x}')] = \delta(\boldsymbol{x} - \boldsymbol{x}')$ . The strength of the interaction potential, *V*, is related to the s-wave scattering length, a, by  $V = T\Gamma$ , where  $T = 4\pi\hbar^2 a/m$  is the three-dimensional T-matrix (here it is actually a simple scalar and not a matrix since we consider the regime in which there is no dependence of the scattering phase shift on energy) and  $\Gamma$  is the dimensionless renormalization factor that will be fully discussed in Section 4.3.

Since we intend to explore excitations from a BEC, we assume that the field operator is well described by a mean field amplitude describing the atom condensate,  $\phi_a(\boldsymbol{x})$ , and a fluctuating component, i.e.,

$$\hat{\psi}(\boldsymbol{x}) = \langle \hat{\psi}(\boldsymbol{x}) \rangle + \delta \hat{\psi}(\boldsymbol{x}) = \phi_a(\boldsymbol{x}) + \delta \hat{\psi}(\boldsymbol{x}) \,. \tag{4.2}$$

where  $\langle \delta \hat{\psi}(\boldsymbol{x}) \rangle = 0$ . The second-order terms—normal and anomalous densities—are defined respectively as,

$$G_N(\boldsymbol{x}, \boldsymbol{x}') = \langle \delta \hat{\psi}^{\dagger}(\boldsymbol{x}') \delta \hat{\psi}(\boldsymbol{x}) \rangle,$$
  

$$G_A(\boldsymbol{x}, \boldsymbol{x}') = \langle \delta \hat{\psi}(\boldsymbol{x}') \delta \hat{\psi}(\boldsymbol{x}) \rangle.$$
(4.3)

Both of these play an important role in the dynamics of the non-condensate component of the system we are interested in. In particular, the diagonal elements of the normal density,  $G_N(\boldsymbol{x}, \boldsymbol{x})$ , represent the physical non-condensate atom densities at position  $\boldsymbol{x}$  and are therefore positive semi-definite. The off-diagonal elements represent the matter-wave correlations of the non-condensate atoms that are characterized by quantities such as the de Broglie wavelength and effective temperature. The anomalous density,  $G_A(\boldsymbol{x}, \boldsymbol{x})$ , is the pairing field that characterizes the two-particle correlations in the system.

If we assume that the field fluctuations are Gaussian, one can drop the third-order cumulants, and expand the fourth-order quantities in terms of the second-order cumulants when deriving the evolution equations. In practice, this involves repeated application of Wick's theorem [73]. The

$$i\hbar \frac{\partial \phi_a(\boldsymbol{x})}{\partial t} = \left( -\frac{\hbar^2}{2m} \nabla^2 + V_{\text{ext}}(\boldsymbol{x}) \right) \phi_a(\boldsymbol{x}) + V \left[ |\phi_a(\boldsymbol{x})|^2 + 2G_N(\boldsymbol{x}, \boldsymbol{x}) \right] \phi_a(\boldsymbol{x}) + V G_A(\boldsymbol{x}, \boldsymbol{x}) \phi_a^*(\boldsymbol{x}) , \qquad (4.4)$$

for the normal density as;

$$i\hbar \frac{\partial}{\partial t} G_N(\boldsymbol{x}, \boldsymbol{x}') = \mathcal{H}'(\boldsymbol{x}) G_N(\boldsymbol{x}, \boldsymbol{x}') - \mathcal{H}'(\boldsymbol{x}') G_N(\boldsymbol{x}, \boldsymbol{x}') + \Delta(\boldsymbol{x}) G_A^*(\boldsymbol{x}, \boldsymbol{x}') - \Delta^*(\boldsymbol{x}') G_A(\boldsymbol{x}, \boldsymbol{x}')$$
(4.5)

and for the anomalous density as;

$$i\hbar \frac{\partial}{\partial t} G_A(\boldsymbol{x}, \boldsymbol{x}') = \mathcal{H}'(\boldsymbol{x}) G_A(\boldsymbol{x}, \boldsymbol{x}') + \mathcal{H}'(\boldsymbol{x}') G_A(\boldsymbol{x}, \boldsymbol{x}') + \Delta(\boldsymbol{x}) \left[ G_N^*(\boldsymbol{x}, \boldsymbol{x}') + \delta(\boldsymbol{x} - \boldsymbol{x}') \right] + \Delta^*(\boldsymbol{x}') G_N(\boldsymbol{x}, \boldsymbol{x}').$$
(4.6)

Here we have simplified the notation by introducing two energy functionals,

$$\mathcal{H}'(\boldsymbol{x}) = -\frac{\hbar^2}{2m} \nabla^2 + V_{\text{ext}}(\boldsymbol{x}) + 2V \left[ |\phi_a(\boldsymbol{x})|^2 + G_N(\boldsymbol{x}, \boldsymbol{x}) \right] ,$$
  
$$\Delta(\boldsymbol{x}) = V \left[ \phi_a(\boldsymbol{x})^2 + G_A(\boldsymbol{x}, \boldsymbol{x}) \right] , \qquad (4.7)$$

for the single-particle self-energy and the gap, respectively. Due to the fact that we neglect explicit three-particle and higher correlations, the validity of this approach is restricted to the dilute gas regime. Note that Eq. (4.4) can be simplified to the Gross-Pitaevskii equation if terms involving  $G_N(\boldsymbol{x}, \boldsymbol{x})$  and  $G_A(\boldsymbol{x}, \boldsymbol{x})$  are dropped. In this case, the time-independent energy eigenvalue represents the chemical potential,  $\mu$ , so that  $i\hbar\dot{\phi}_a = \mu\phi_a$ . The delta function in Eq. (4.6) arises from the bosonic commutation relation of the field operators and can therefore be interpreted as a quantum effect. A number of quantities are conserved in this evolution; in particular, the total atom number

$$\mathcal{N} = \int d\boldsymbol{x} \left( |\phi_a(\boldsymbol{x})|^2 + G_N(\boldsymbol{x}, \boldsymbol{x}) \right)$$
(4.8)

is invariant under time evolution governed by Eqs. (4.4)–(4.7).

In order to see how the anomalous density is related to the vacuum pair wavefunction for the interatomic separation of two atoms, we may neglect the mean-field density and the normal density in Eq. (4.6), and then the eigenvalue equation is simplified to

$$-\frac{\hbar^2}{m}\nabla^2 G_A(\boldsymbol{r}) + V\delta(\boldsymbol{r})G_A(\boldsymbol{r}) = 2\mu G_A(\boldsymbol{r}), \qquad (4.9)$$

where  $\mathbf{r} = \mathbf{x} - \mathbf{x}'$ . Eq. (4.9) can be identified as a one-dimensional Schrödinger equation of a fictitious particle of reduced mass m/2 scattering off a potential  $V\delta(\mathbf{r})$ . Then  $G_A(\mathbf{r})$  is interpreted as the resulting eigenstate wavefunction corresponding to the familiar two-particle scattering solution of the equation written in terms of the relative coordinate.

## 4.3 Renormalization of the Scattering Potential

The Dirac delta function in Eq. (4.9) implies that we are implicitly building a scattering model from a contact interaction. This is convenient as it simplifies the resulting field theory, but care must be taken to account for divergences that can arise at small and large scales. In general, this is remedied by renormalization of the potential strength. In order to carry out this renormalization procedure, we begin from the formal scattering theory [91], where we define the bare scattering potential operator,  $\hat{V}$ , which has units of energy, and thereby expand the *T*-matrix in an order-by-order series;

$$\hat{T} = \hat{V} + \hat{V}G_0\hat{V} + \hat{V}G_0\hat{V}G_0\hat{V} + \dots$$
  
=  $\hat{V} + \hat{V}G_0\hat{T}$ . (4.10)

Here  $G_0$  is the bare single particle propagator,

$$G_0 = \frac{1}{E - \hat{H}_0 + i\epsilon},$$
(4.11)

the scattering energy is E, the dispersion relation is  $\hat{H}_0 = \hat{p}^2/(2m)$  with  $\hat{p}$  the momentum, and we need to implicitly consider the limit  $\epsilon \to 0$ . The *T*-matrix elements are  $T = \langle \mathbf{k}' | \hat{T} | \mathbf{k} \rangle$ , where  $| \mathbf{k} \rangle$  is

the wavenumber basis state. For the low energy scattering limit, the *T*-matrix becomes independent of *E*, and does not depend on  $\mathbf{k}$  or  $\mathbf{k}'$ . In this case the *T*-matrix is well characterized by a constant scalar associated with the *s*-wave scattering length, as mentioned earlier, i.e.,  $T = 4\pi\hbar^2 a/m$ .

Further considerations have to be made when one or more dimensions are effectively frozen out due to imposing a strong confining potential in these dimensions. Without loss of generality, let us consider the strong confining potential to be a harmonic potential with oscillator length given by  $l_{\perp}$ . If one dimension is frozen out, an effective quasi-2D geometry is realized, and if two dimensions are frozen out, an effective quasi-1D system is generated. If we denote the number of free dimensions by  $n \in \{1, 2, 3\}$ , the appropriate *T*-matrix expression,  $T_n$ , for the reduced dimensional case can be related recursively by  $T_3 = T$  and  $T_{n-1} = T_n/(\sqrt{2\pi}l_{\perp})$  [8][75].

The process of renormalization connects the *T*-matrix,  $T_n$ , to the strength of the potential,  $V_n$ , by expanding Eq. (4.10) in the momentum basis, and this connection depends on the dimensionality of the system,

$$T_n = V_n + V_n \int_{K_-}^{K_+} \frac{d^n k}{(2\pi)^n} \frac{T_n}{E - \frac{\hbar^2 k^2}{2m}},$$
(4.12)

where the critical element here is the introduction of  $K_{-}$  and  $K_{+}$  as infrared and ultraviolet momentum cutoffs, respectively. The cutoffs have to be chosen from an appropriate asymptotic limit in order to accurately capture the dynamics of interest. The renormalization procedure can be represented by the introduction of a parameter,  $\Gamma_n$ , defined by solving Eq. (4.12) for  $V_n$ . This gives the solution,

$$V_n = \frac{T_n}{1 - \alpha_n T_n} \equiv T_n \Gamma_n , \qquad (4.13)$$

where

$$\alpha_n = -\int_{K_-}^{K_+} \frac{d^n k}{(2\pi)^n} \frac{1}{E - \frac{\hbar^2 k^2}{2m}} \,. \tag{4.14}$$

In order to illustrate the behavior of  $\alpha_n$ , we consider the solution to scattering equation, Eq. (4.9), where the stationary energy eigenvalue is  $E = 2\mu$  and the mass is replaced by the reduced mass of two particles,  $m \to m/2$ . Solving this system of equations has a character that depends on the dimensionality. In three dimensions, we may set  $K_- = 0$ , and perform the integral to give  $\alpha_3 = mK_+/(2\pi^2\hbar^2)$  for a particle scattering at low energy,  $E \to 0$  [46]. In 2D, the integral scales logarithmically and has both ultraviolet and infrared divergences. In 1D, there is an infrared divergence so  $K_-$  must be non-zero but we may set  $K_+$  to infinity.

We do not provide all the details here, since, in practice, these are formal considerations that do not actually affect our numerical simulations. Indeed there are actually no divergences introduced that require the introduction of momentum cutoffs to rectify when the momenta are restricted to values on discrete and finite grids. This is always the case in a numerical computer model that aims to describe a realistic experiment. In such a discrete representation of possible momenta, it is preferable to simply calculate a finite sum over a specific partition instead of evaluating the continuous integral analytically. This implies a numerical evaluation of  $\sum_{i=0}^{N-1} (\hbar^2 k_i^2/m - 2\mu)^{-1}$ , giving  $\alpha_n$ , and therefore determining  $V_n$  for a given *T*-matrix, which replaces *V* in the HFB equations, i.e., Eqs. (4.4)–(4.6). Here the subscripts *i* label individual discrete momenta, and thus  $\{k_i\}$ represents the momentum grid, with *N* is the total number of grid points.

We carry out this renormalization procedure for all the results that we present in this chapter. For each calculation, we verify that the numerical results are independent of the details of the momentum grid on which the field theory is represented.

#### 4.4 Self-consistent ground state solution

In order to find a self-consistent solution to prepare an initial condition for the subsequent time evolution, the first step will be to consider the non-condensate component to be absent, and to find a ground state representation of the condensate by solving the GPE. We then use this condensate field as input into the time-independent equations for the normal and anomalous densities, and diagonalize the resulting HFB self-energy matrix to find the quasiparticle basis. As we will see, this process exhibits a defect in the zero-energy subspace (i.e., the eigenvectors do not span the space). The intepretation is that the eigensolution is not stationary and cannot be used as an accurate description of the initial condition for subsequent time evolution. We therefore reintroduce the non-condensate terms that we have just found into the equations for the condensate, normal density, and anomalous density and solve again the system of equations, giving rise to an iterative method that generates an accurate self-consistent initial condition.

Our approach will be to begin by first fully describing the necessary procedure using the simple case of quasi-1D where the problem is most easily tractable. However, higher dimensions can be treated in a similar method to the manner we present (we will consider quasi-2D later in Section 4.8). The reduction to one-dimensional behavior requires the transverse confinement condition

$$\frac{a}{n_{1\mathrm{D}}l_{\perp}^2} \ll 1 \tag{4.15}$$

to be satisfied, where  $n_{1D}$  is the one-dimensional density [68][51], and as defined previously,  $l_{\perp}$  is the harmonic oscillator length in the two strongly confining directions, here assumed to be equal.

The first part of our numerical algorithm is to solve for the ground state of the GPE,

$$i\hbar \frac{\partial \phi_a(x)}{\partial t} = \left(-\frac{\hbar^2}{2m}\nabla^2 + V_{\text{ext}}(x)\right)\phi_a(x) + T_1|\phi_a(x)|^2\phi_a(x).$$
(4.16)

We use imaginary-time propagation to derive the lowest energy solution and its energy eigenvalue  $\mu$  representing the associated chemical potential. Then, the mean field solution,  $\phi_a(x)$ , can be used as a parameter to construct the self-energy matrix;

$$\Sigma = \begin{pmatrix} \Sigma_N & \Sigma_A \\ -\Sigma_A^* & -\Sigma_N^* \end{pmatrix}, \qquad (4.17)$$

where

$$\Sigma_N = -\frac{\hbar^2}{2m} \nabla^2 - \mu + 2T_1 |\phi_a(x)|^2 ,$$
  

$$\Sigma_A = T_1 \phi_a(x)^2 . \qquad (4.18)$$

The self-energy matrix has dimensionality  $2N \times 2N$  where N is the size of the single particle basis, as defined previously. This energy operator is most simply expressed in the position basis, where  $x \in (0, L]$ , since in that representation the potential terms including the mean-field appear as diagonal blocks. The eigenstates of  $\Sigma$  are the Bogoliubov quasiparticles. Since the matrix satisfies  $\sigma_z \Sigma = \Sigma^{\dagger} \sigma_z$ , where  $\sigma_z = \text{diag}(I_{N \times N}, -I_{N \times N})$ , the eigenenergies come in pairs of positive and negative values,  $\pm \epsilon_k$ , and the corresponding eigenstates are  $w_k = (u_k(x), v_k(x))^T$  and  $w_{-k} = (v_k^*(x), u_k^*(x))^T$ . The eigenstates are normalized by satisfying the constraint

$$\int_0^L dx \left( |u_k(x)|^2 - |v_k(x)|^2 \right) \longrightarrow \sum_{j=1}^N \frac{L}{N} \left( |u_k(jL/N)|^2 - |v_k(jL/N)|^2 \right) = 1, \quad (4.19)$$

Although this construction may appear standard and straightforward, there is a well-known and implicit subtlety when examining the solutions to this eigensystem. When investigating the zeroenergy eigensolutions, one finds a pair of eigenstates that are colinear (equal up to a multiplicative scalar) that have the form  $\mathcal{P} = (\phi_a(x)/\sqrt{2N}, -\phi_a^*(x)/\sqrt{2N})^T$ . This solution creates two significant issues. First, the colinear eigenstates cannot be normalized by Eq. (4.19). Second, they do not span the two-dimensional subspace of the Hilbert space corresponding to zero-energy.

The origin of this mathematical fact has an intuitive explanation. It arises from the approximations that lead to this self-energy matrix, that is, by fixing the condensate solution as an unchanging parameter, one builds an unphysical model that implicitly allows the unconstrained growth of a zero energy mode as a function of increasing time. Consequently there is no stationary solution. This has to be remedied, for example, through a self-consistent approach in which the condensate is treated as a variational parameter, in order to allow us to extend the formalism so that it may be applied to our system of interest.

We begin by determining the remaining eigenvector to fully span the zero-energy subspace by employing the Gram-Schmidt orthogonalization method to numerically calculate the remaining basis vector. In this way we determine an eigenvector solution  $\mathcal{Q} = (q(x), -q^*(x))^T$  such that  $\frac{L}{N}\mathcal{Q}^{\dagger}\sigma_z w_k = 0$  for all  $k \neq 0$ , and normalize it to  $\frac{L}{N}\mathcal{Q}^{\dagger}\sigma_z\mathcal{P} = i$  [97]. The addition of this vector to the eigenvectors of the self-energy completes the basis of the vector space. The reason that this is important is that it allows the field operator to be expanded as

$$\delta\hat{\psi}(x) = \sum_{k=1}^{N-1} (u_k(x)\hat{b}_k + v_k^*(x)\hat{b}_k^{\dagger}) - i\frac{\phi_a(x)}{\sqrt{2N}}\hat{\theta} + iq(x)\hat{L}$$
(4.20)

where  $\hat{b}_k$  and  $\hat{b}_k^{\dagger}$  are bosonic annihilation and creation operators for the quasiparticles. We have introduced  $\hat{\theta}$  and  $\hat{L}$  as a pair of canonically conjugate operators that fully describe the zero-energy mode and obey the canonical commutation relation  $[\hat{\theta}, \hat{L}] = i$ . It is convenient to identify two special combinations of  $\mathcal{P}$  and  $\mathcal{Q}$  in order to give a concise expression for the completeness relation. We define  $w_{\pm} = (\mp \mathcal{P} + i\mathcal{Q})/\sqrt{2}$ , along with the matrix

$$W = \left(w_{+}, w_{1}, ..., w_{N-1}, w_{-}, w_{-1}, ..., w_{-(N-1)}\right)$$
(4.21)

so that the following completeness relation is satisfied;

$$\frac{L}{N}W^{\dagger}\sigma_{z}W = \sigma_{z}. \qquad (4.22)$$

This allows the particle annihilation operator to be written as

$$\delta\hat{\psi}(x) = \sum_{k\in\mathcal{S}} (u_k(x)\hat{b}_k + v_k^*(x)\hat{b}_k^\dagger), \qquad (4.23)$$

where the sum is over the elements of the index set  $S = \{+, 1, 2, ..., N - 1\}$ , and  $\hat{b}_+$  is the annihilation operator for the zero-energy mode given by  $\hat{b}_+ = (i\hat{\theta} + \hat{L})/\sqrt{2}$ .

At this point, we have determined the quasiparticle basis, and can populate that basis with a given set of probabilities in order to generate particle distributions. In particular, we would like to derive the normal  $G_N(x, x')$  and anomalous  $G_A(x, x')$  densities that are essential elements of the HFB theory. To begin with we construct the Hermitian density matrix:

$$G = \begin{pmatrix} G_N(x, x') & G_A(x, x') \\ G_A^*(x, x') & \delta(x - x') + G_N^*(x, x') \end{pmatrix}$$
  
=  $W\Pi W^{\dagger}$ , (4.24)

where the population matrix  $\Pi$  has the form

$$\Pi = \begin{pmatrix} p & q \\ q^* & I + p \end{pmatrix}.$$
(4.25)

The diagonal elements of p are the populations of each quasiparticle  $\langle \hat{b}_k^{\dagger} \hat{b}_k \rangle$ , and the off-diagonal elements represent the correlations between different quasiparticles. In the ground state, p = 0 and q = 0. The identity, I, on the lower-right block is interpreted as a bosonic analog to the Dirac sea [33], in which the negative energy states are occupied by boson holes. When there is an excitation,

a pair of one particle and one hole is created, and therefore p appears in both the upper-left and the lower-right block, as shown in Eq. (4.25).

This formalism now allows an extremely concise representation of the full dynamical evolution encapsulated in Eqs. (4.5) and (4.6);

$$i\hbar\frac{\partial G}{\partial t} = \Sigma G - G\Sigma^{\dagger} \tag{4.26}$$

where G is defined according to Eq. (4.24). The consequence of completing the basis by establishing the missing eigenvector through Gram-Schmidt orthogonalization is now evident. If we begin with the bare  $\Sigma$ , as defined in Eq. (4.17) and initialize G to the ground state (meaning p = 0 and q = 0) of the corresponding eigenbasis, then when Eq. (4.26) is propagated from this initial condition, it is evident that the solution is not stationary. The number of non-condensate atoms is seen to grow as  $\sim t^2$ , as shown in Fig. 4.1. This implies that we have not in fact determined the correct ground state.

This problem arises because, using the language of quantum optics, we are effectively assuming that the condensate is a coherent field that may act as an infinite classical pump and can provide a reservoir source for introducing an infinite number of atom-pairs. Furthermore, it does not cost any energy to introduce a zero-energy quasiparticle within this framework. This is clearly unphysical for a number of reasons including the fact that, as can be seen in Eq. (4.4), the factor of two in front of the interaction between the condensate and the non-condensate atoms means that it actually costs energy to take away atoms from the condensate and move them into the noncondensate fraction, providing the interactions are repulsive (scattering length positive). There is some literature that suggests simply dropping the zero-modes entirely to remedy this problem, for example, Ref. [9]; however, this generally violates the fundamental commutation relations of the bosonic field operator and therefore the uncertainty principle, so we do not employ that approach here.

We instead employ an alternative solution by including the second-order terms to generalize the self-energy matrix. This means that we modify Eq. (4.18) to include the effects of the normal



Figure 4.1: Quantum depletion  $\left(\frac{1}{N}\int G_N(x,x)\,dx\right)$  as a function of time (the proportion of non-condensate atoms at zero-temperature) simulated with the gapless HFB theory (i.e., using Eqs.(4.16), (4.18), (4.21), (4.24), and (4.26)) gives a depletion proportion that initially scales as  $\sim t^2$ .

and anomalous densities, and then introduce the renormalization of the T-matrix to give

$$\Sigma_N = -\frac{\hbar^2}{2m} \nabla^2 - \mu + 2V_1[|\phi_a(x)|^2 + G_N(x, x)],$$
  

$$\Sigma_A = V_1[\phi_a(x)^2 + G_A(x, x)].$$
(4.27)

with both Eq. (4.17) and Eq. (4.26) unmodified. In order to be consistent, however, we must also generalize the GPE, Eq. (4.16), to

$$i\hbar \frac{\partial \phi_a(x)}{\partial t} = -\frac{\hbar^2}{2m} \nabla^2 \phi_a(x) - \mu \phi_a(x) + V_1[|\phi_a(x)|^2 + 2G_N(x,x)]\phi_a(x) + G_A(x,x)\phi_a^*(x,x).$$
(4.28)

Note that the ground state solution of the GPE is stationary, and thereby determines the value of the chemical potential that enters the renormalization (see Section 4.3). Since  $G_N(x, x)$  and  $G_A(x, x)$  are functionally dependent on the eigenstates themselves, the problem is nonlinear, and it is necessary to solve the generalized self-energy, Eq. (4.27), and the generalized GPE, Eq. (4.28), iteratively until the equations are self-consistent [31]. We point out that this iterative process will typically create a small gap in the energy spectrum of the system around zero energy, and the problem of the unphysical non-stationary eigensolution that is caused by the zero-energy subspace is no longer present. The resulting self-consistent solution is stationary under the evolution given by Eq. (4.26) and provides an accurate ground state initial condition for the subsequent timedependent simulations that follow.

### 4.5 Dynamics of the time-dependent HFB system

In the experiment by Clark *et al.* [18], an external sinusoidally oscillating magnetic field is applied, and therefore the scattering length is modulated in the form

$$a(t) = a_{\rm dc} + a_{\rm ac} \sin \omega t, \tag{4.29}$$

where  $a_{dc}$  is the initial scattering length, and  $a_{ac}$  is the amplitude of the oscillating component of the scattering length at angular frequency  $\omega$ . The dynamics of the system under this modulation



Figure 4.2: Solutions for a system with total atom number  $\mathcal{N} = 6 \times 10^5$  in a 1D infinite potential well of size L, with the scattering potential between atoms given by  $a = 10^{-4} l_{\perp}^2/L$ . (a) Ground state condensate density found from the gapped self-consistent generalized GPE theory (i.e., replacing Eq. (4.16) with Eq. (4.28) and Eq. (4.18) with Eq. (4.27)). The length scale over which the condensate density falls to zero at the edges of the box is known is the healing length. (b) Solution to the normal density  $G_N(x, x')$  in the ground state as found from the self-consistent HFB theory. (c) Solution to the absolute value of the anomalous density  $|G_A(x, x')|$  in the ground state as found from the self-consistent HFB theory.

is interesting to consider because the oscillating external field will inject energy into the system, and this will result in exciting atoms from the ground state into higher quasiparticle levels.

We begin our simulations by preparing the system in the self-consistent ground-state of the HFB theory for a small positive value of the scattering length using the procedure just described. An illustration of the resulting condensate, normal and anomalous densities are shown in figure 4.2. After preparing the system in the ground state, we solve Eq. (4.17) and Eq. (4.26) using the generalized Eqs. (4.27) and (4.28) with a sinusoidal modulation of the scattering potential, i.e.,

$$V \to V(t) = V_{\rm dc} + V_{\rm ac} \sin \omega t \,. \tag{4.30}$$

In order to interpret our results, we display the occupation probabilities via the projection of G(t) onto the initial quasiparticle basis found from the self-consistent HFB Hamiltonian at time t = 0. The procedure is as follows. Since the quasiparticle eigenbasis matrix, W, satisfies  $\frac{L}{N}W^{\dagger}\sigma_{z}W = \sigma_{z}$ , we may write

$$W^{-1} = \frac{L}{N} \sigma_z W^{\dagger} \sigma_z \,. \tag{4.31}$$

Then, according to Eq. (4.24),

$$\Pi = W_0^{-1} G W_0^{\dagger -1}$$

$$= \frac{L^2}{N^2} \sigma_z W_0^{\dagger} \sigma_z G \sigma_z W_0 \sigma_z$$

$$(4.32)$$

where  $W_0$  is the original self-consistent quasiparticle basis determined for the initial condition. The resulting population is shown in figure (4.3). The height of the peak in the off-diagonal block (i.e., q) is notable since the coherence saturates the upper bound of the Cauchy-Schwartz inequality,

$$\left[\delta(x-x') + G_N(x,x)\right] G_N(x',x') \ge |G_A(x,x')|^2, \tag{4.33}$$

which in turn can be interpreted as confirming that the process of exciting quasiparticles from the condensate is maximally coherent. The diagonal elements,  $p_k$ , can be measured by time of flight, since the quasiparticles transform into regular particles that can be detected during ballistic expansion. In other words, when the kinetic energy greatly exceedes the interaction energy, the



Figure 4.3: Snapshots of (a) absolute value of the matrix elements of the upper-left block of the population matrix, |p|, and (b) absolute value of the matrix elements of the upper-right block of the population matrix, |q|, at  $t = 0.1 (mL^2/\hbar)$ , starting from the initial condition shown in figure 4.2 and then continuously driven with amplitude  $a_{\rm ac} = 10^{-4} l_{\perp}^2/L$  and frequency  $\omega = 1000 (\hbar/mL^2)$ . This frequency resonates with the quasiparticles with energy  $\epsilon_{k^*} = 500 (\hbar^2/mL^2)$ , corresponding to the resonant wavenumbers shown for reference as white lines (at k-index  $(\pi k)^2 \approx 500$ ). At the resonant quasiparticle excitation a clear spike is evident. (c) Density profile of the condensate, revealing in general form the spatial dependence of the eigenmode function of the resonant quasiparticle excitation.

k's then effectively label the free momentum, i.e.,  $k\hbar\pi/L$ . As shown in figure (4.3), when the periodic drive is turned on continuously for many cycles, essentially only one quasiparticle mode is resonantly amplified. That is consistent with the narrow spectrum. This physical process can be interpreted as being due, as a consequence of the oscillating drive, to a photon with energy  $\hbar\omega$  being absorbed by a pair of atoms, with each of them getting half the energy,  $\epsilon_k = \hbar\omega/2$ . In addition, the phonon-like collective excitations that correspond to the observed wave-like patterns seen in the condensate density can be interpreted as the Faraday patterns that typically manifest in different kinds of parametrically driven fluids [88]. The pattern resembles the wavefunction density for the single quasiparticle mode on resonance. This simulation illustrates that by careful engineering of the drive, one can potentially prepare a variety of quantum states, selectively exciting atoms from the condensate field. We now show a few illustrative examples of interesting cases that employ this technique.

### 4.6 Dynamically Generating Squeezed Quasiparticle States

A squeezed state refers to a quantum state that has a reduced uncertainty in one degree of freedom ('squeezed') at the expense of increased uncertainty in a canonically conjugate variable [44, 104, 96, 105, 80, 81, 84]. Such states have been extensively studied in quantum optics and atomic physics due to their utility in quantum metrology for producing measurement precision that exceeds the limits derived from classical states. Here we will show how to use the resonant quasiparticle excitation in order to generate a squeezed matter-wave state, anticipating that this could potentially be applied to quantum matter-wave interferometry.

By driving the system resonantly, we are effectively producing resonant pairs with well defined energy, and this is reminiscent of nonlinear optical devices that down-convert pump photons into signal and idler pairs. Here we will demonstrate that this correspondence is robust and quantitative by demonstrating how one may calculate the squeezing parameter associated with the analogous quantity that is regularly computed in the quantum description of light.

In order to do this we assume a weak excitation limit, so that  $G_N(x, x)$  and  $G_A(x, x)$  are small

compared to  $|\phi_a(x)|^2$ . Furthermore, we consider the kinetic energy term in the time-dependent GPE to be small, and then we can find a general solution for the condensate that has the form

$$\phi_a(t) = \phi_0 e^{iA\cos(\omega t)} = \phi_0 \sum_{n=-\infty}^{n=\infty} J_n(A) e^{in\omega t}, \qquad (4.34)$$

where  $A = V_{ac} |\phi_a|^2 / \hbar \omega$  and  $J_n(...)$  is the Bessel function of the first kind. We will limit our discussion to the case of high modulation frequency, in which the photon energy associated with the drive,  $\hbar \omega$ , greatly exceeds the mean field shift associated with the drive amplitude,  $V_{ac} |\phi_a|^2$ , so that  $A \ll 1$ . In this case the n = 0 term completely dominates the series expansion and we can drop all other terms.

The initial stationary Hamiltonian for the fluctuations can be written as  $H_0 = \sum_k \epsilon_k \hat{b}_k^{\dagger} \hat{b}_k$ , where  $\epsilon_k$  is the energy of the k-th quasiparticle, and the transformation to a rotating frame involves making the replacement of the quasiparticle operators

$$\hat{b}_k \to \hat{b}_k e^{i\epsilon_k t/\hbar} \,. \tag{4.35}$$

The contact interaction term in the Hamiltonian can be derived from the interaction term of Eq. (4.1),

$$H_I = \frac{V_{ac}}{2} \sin \omega t \, \int dx \, \hat{\psi}^{\dagger}(x) \hat{\psi}^{\dagger}(x) \hat{\psi}(x) \hat{\psi}(x)$$

$$\tag{4.36}$$

From this point, we keep only the second-order terms in  $\delta \hat{\psi}$ , because these terms correspond to exponential growth and therefore dominate the solution. In order to simplify the problem further, we assume that the drive frequency  $\omega$  corresponds to the resonance condition  $\omega = 2\epsilon_k/\hbar$ , and introduce the rotating wave approximation, which allows us to keep only terms with  $e^{\pm i(\omega - 2\epsilon_k/\hbar)t}$ . By representing  $\delta \hat{\psi}$  in the quasiparticle basis, we obtain an effective interaction Hamiltonian

$$\begin{split} H_{I} &\approx \frac{V_{ac}}{2} \sin \omega t \int dx \left( 4|\phi_{a}|^{2} \delta \hat{\psi}^{\dagger} \delta \hat{\psi} + \phi_{a}^{2} \delta \hat{\psi}^{\dagger} \delta \hat{\psi}^{\dagger} + \phi_{a}^{*2} \delta \hat{\psi} \delta \hat{\psi} \right) \\ &\approx \frac{V_{ac}}{2} \sin \omega t \int dx \left[ 4|\phi_{0}|^{2} \\ &\times \sum_{k'} \left( u_{k'}^{*} \hat{b}_{k'}^{\dagger} e^{i\epsilon_{k'}t/\hbar} + v_{k'} \hat{b}_{k'} e^{-i\epsilon_{k'}t/\hbar} \right) \sum_{k''} \left( u_{k''} \hat{b}_{k''} e^{-i\epsilon_{k''}t/\hbar} + v_{k''} \hat{b}_{k''}^{\dagger} e^{i\epsilon_{k''}t/\hbar} \right) \\ &+ \phi_{0}^{2} \sum_{k'} \left( u_{k'}^{*} \hat{b}_{k'}^{\dagger} e^{i\epsilon_{k'}t/\hbar} + v_{k'} \hat{b}_{k'} e^{-i\epsilon_{k'}t/\hbar} \right) \sum_{k''} \left( u_{k''} \hat{b}_{k''}^{\dagger} e^{i\epsilon_{k''}t/\hbar} + v_{k''} \hat{b}_{k''} e^{-i\epsilon_{k''}t/\hbar} \right) \\ &+ \phi_{0}^{*2} \sum_{k'} \left( u_{k'} \hat{b}_{k'} e^{-i\epsilon_{k'}t/\hbar} + v_{k'}^{*} \hat{b}_{k'}^{\dagger} e^{i\epsilon_{k'}t/\hbar} \right) \sum_{k''} \left( u_{k''} \hat{b}_{k''} e^{-i\epsilon_{k''}t/\hbar} + v_{k''}^{*} \hat{b}_{k''}^{\dagger} e^{i\epsilon_{k''}t/\hbar} \right) \\ &+ \phi_{0}^{*2} \sum_{k'} \left( u_{k'} \hat{b}_{k'} e^{-i\epsilon_{k'}t/\hbar} + v_{k'}^{*} \hat{b}_{k'}^{\dagger} e^{i\epsilon_{k'}t/\hbar} \right) \sum_{k''} \left( u_{k''} \hat{b}_{k''} e^{-i\epsilon_{k''}t/\hbar} + v_{k''}^{*} \hat{b}_{k''}^{\dagger} e^{i\epsilon_{k''}t/\hbar} \right) \\ &= \frac{V_{ac}}{4i} \left[ \int dx \left( 4|\phi_{0}|^{2} v_{k} u_{k} + \phi_{0}^{2} v_{k}^{2} + \phi_{0}^{*2} v_{k}^{*2} \right) \hat{b}_{k}^{\dagger} \hat{b}_{k}^{\dagger} \right] \end{aligned}$$

$$(4.37)$$

This corresponds to the interaction Hamiltonian of a parametric amplifier in nonlinear quantum optics, namely  $H_I = -i\hbar_2^{\chi}(\hat{a}^2 - \hat{a}^{\dagger 2})$ , where  $\chi$  represents the second-order nonlinear susceptibility that corresponds to the squeezing rate. We refer to the resulting time-evolved state as a squeezed quasiparticle state since the analog is an archetypal system for creating squeezed states of light. This mapping allows us to extract the squeezing rate, i.e.,

$$\chi = \frac{V_{ac}}{2\hbar} \int dx \left( 4|\phi_a|^2 v_k u_k + \phi_a^2 v_k^2 + \phi_a^{*2} u_k^2 \right) \,, \tag{4.38}$$

and the squeezing parameter increases with time at this rate, i.e,  $\xi = \chi t$ . If we choose the phases of  $\phi_a$ ,  $u_k$  and  $v_k$  appropriately, then  $\chi$  is real. As expected from the known optical solutions, the population in the k-th quasiparticle mode grows proportional to  $\sinh^2(\chi t)$ . Figure 4.4 shows the population as a function of time at different modulation amplitudes. Since  $\sinh^2(\chi t) \rightarrow e^{2\chi t}/4$  at large t, one can extract the squeezing rate from the asymptotic slope of log  $p_{kk}$ . We confirm that the squeezing rate is proportional to the modulation amplitude, as indicated by Eq. (4.38).

Squeezed states are characterized by reduced variance in one quadrature at the expense of increased variance in the other quadrature perpendicular to it. We define the quadrature for the resonant quasiparticles as

$$X_{\theta} \equiv \hat{b}_k^{\dagger} e^{i\theta} + \hat{b}_k e^{-i\theta} , \qquad (4.39)$$



Figure 4.4: Population in the resonant quasiparticle mode (labelled k) driven with frequency  $\omega = 5500 \ (\hbar/mL^2)$  as a function of time for different modulation amplitudes. The amplitudes are  $V_{\rm ref}$  (yellow),  $2V_{\rm ref}$  (red), and  $4V_{\rm ref}$  (blue), where  $V_{\rm ref} = 1.25 \times 10^{-4} \ (\hbar^2/mL)$ . The slope of log  $p_{kk}$  at large time is equal to twice the squeezing rate, which is proportional to the modulation amplitude as shown in Eq. (4.38). The squeezing rates calculated from the slopes of the curves in the interval  $t = [0.3, 0.4] \ (mL^2/\hbar)$  are 4.3, 7.2, 14.1  $(\hbar/mL^2)$ , and the squeezing rates calculated from Eq. (4.38) are 3.6, 7.1, 14.3  $(\hbar/mL^2)$  respectively for amplitudes  $V_{\rm ref}$ ,  $2V_{\rm ref}$ ,  $4V_{\rm ref}$ .
where  $\theta$  is the angle of the orientation of the quadrature. Then the variance is

$$\langle (\Delta X_{\theta})^2 \rangle = \left\langle \left( \hat{b}_k^{\dagger} e^{i\theta} + \hat{b}_k e^{-i\theta} \right)^2 \right\rangle - \left\langle \hat{b}_k^{\dagger} e^{i\theta} + \hat{b}_k e^{-i\theta} \right\rangle^2$$

$$= 2p_{kk} + 1 + 2Re\{q_{kk} e^{-i2\theta}\}.$$

$$(4.40)$$

The variance as a function of  $\theta$  is shown in figure 4.5. We see that the variance at certain quadrature phase angles,  $\theta$ , of states produced by modulation of the scattering potential can fall below the standard quantum limit. The standard quantum limit is the level generated by the uncertainty principle under the assumption that the variance in all angles  $\theta$  is uniform.

Although direct measurement of the squeezing may not be as straightforward to implement as in its optics counterpart, it may be possible to observe directly the atom coincidence (since the particles are produced in pairs) on detectors placed in directions corresponding to opposite momenta, and thereby measure the second-order coherence. The direct analogue of phase sensitive photodetection (homodyne and heterodyne detection, for example) is generally more complicated to implement with atoms than light, but in the next section we propose a possible experiment that could be used to perform an analogue of such interference measurements on the squeezed quasiparticle distributions that are generated.

## 4.7 Interferometry with squeezed quasiparticles

In principle, the diagonal elements of the normal density are the quantities that can be directly probed with standard atomic density images, for example in dispersive, absorption, or fluorescence imaging techniques. On the other hand, the off-diagonal elements of the normal density and the anomalous density cannot be directly observed since they are phase dependent quantities and have complex values that require an interferometric method to determine the phases. We investigate the phase dependence of the quasiparticle production by analysing two distinct methodologies. One approach is a potential experiment that is capable of performing the phase measurement through the use of a protocol that is based on the Ramsey sequence widely used in atomic physics [45]. A second alternative approach is closely associated with a recent experiment by Hu *et al.* [36],



Figure 4.5: The variance of the quadrature as a function of the angle, at  $t = 0.05 \ (mL^2/\hbar)$  with modulation amplitude  $V_{ac} = 2 \times 10^{-3} \ (\hbar^2/mL)$ , evaluated using Eq. (4.40). The dotted line is the standard quantum limit, where the variance is equal to 1. For a certain range of angles, the variance falls below the standard quantum limit. More specifically, at  $\theta = 0.88\pi$  the variance has minimum, which means measurements of the quadrature along this direction will have the greatest precision.

who demonstrated that applying a phase shift to the oscillatory field after driving the system for a period of time will suppress the non-condensate atom number, and that a  $\pi$  phase shift results in the greatest suppression.

Our Ramsey protocol is as follows. First, we apply a non-zero  $V_{ac}$  for a period of time  $\tau$ to implement the first oscillatory field in the Ramsey sequence. We then set  $V_{ac}$  zero for a brief waiting period of time  $\Delta t$ . During this interval the anomalous density evolves freely at the resonance frequency,  $2\epsilon_k/\hbar$ , and because there is no external work done on the system, the number of noncondensate atoms remains essentially constant. Next,  $V_{ac}$  is set to the same nonzero value as earlier to implement the second oscillatory field, again for the same period of time  $\tau$ . This sequence is illustated in figure 4.6(a). From our simulation results, shown in figure 4.6(b), we observe that the number of non-condensate atoms oscillates as a function of the free evolution time,  $\Delta t$ . This is because a phase difference  $\theta = \omega \Delta t$  accumulates between the anomalous density and the driving field during the free evolution period. We account for this behavior by showing that the oscillations observed are a consequence of the driving field in the second zone either amplifying or attenuating the anomalous density depending on the accumulated relative phase.

For comparison, we now examine an abrupt phase change protocol based on the Hu *et al.* [36] experiment. We consider the effect of the phase shift by first modulating the interaction for a period of time  $\tau$ , then applying a phase shift  $\theta$  to the oscillating drive, and repeating again the interaction for a period of time  $\tau$ , as shown in figure 4.6(c). The result of the final non-condensate atom number as a function of the phase shift is shown in figure (4.6)(d). It is interesting to compare this protocol and the resulting fringe pattern to that found from the first method. The explanation is that the two methods both operate in a manner that is analogous to a Mach-Zehnder interferometer, where interference fringes are seen in the recombination of light propagating along two paths as the relative accumulated phase is varied. In the first protocol that we have presented, the phase is accumulated in the anomalous density, whereas in the second method, a direct phase shift is applied to the external field. We have observed that both methods result in an interference pattern with high visibility fringes that allow direct access for the observer to probe the phase behavior.

The two methods, the complete Ramsey sequence or the abrupt intermediate phase shift change, can be understood in a similar formalism. Both the phase shift change and the Ramsey wait-time effectively generate a phase shift in the direction of squeezing. This manifests as a change in the phase of the squeezing rate, i.e.,  $\chi \to \chi e^{i\theta}$ , and is associated with the resonant quasiparticle state evolving under the unitary operator,

$$U_{\theta}(t) = e^{-\frac{\chi}{2}e^{i\theta}(\hat{b}_{k}^{2} - \hat{b}_{k}^{\dagger 2})t}$$
(4.41)

during the subsequent time evolution period. In the Heisenberg picture, the time-evolved operator  $\hat{b}_k$  for the quasiparticle at index k at the end of the sequence is therefore given by

$$\hat{b}_{k}(2\tau) = U_{\theta}^{\dagger}(\tau)U_{0}^{\dagger}(\tau)\hat{b}_{k}(0)U_{0}(\tau)U_{\theta}(\tau)$$

$$\left\langle \hat{b}_{k}^{\dagger}(2\tau)\hat{b}_{k}(2\tau)\right\rangle = \cos^{2}(\chi\tau\sin\theta)\sinh^{2}(\chi\tau(1+\cos\theta))$$

$$+\sin^{2}(\chi\tau\sin\theta)\cosh^{2}(\chi\tau(1+\cos\theta)). \qquad (4.42)$$

Note that at the special point  $\theta = \pi$ ,  $U_{\pi}(t) = U_0(-t)$ , and the second period of modulation simply reverses the effect of the first period of modulation, so that the final population is zero. However, we can see that in the numerical simulation, the final number of non-condensate atoms at a phase shift of  $\pi$  is non-zero. This is because the analytic result is derived using the rotating-wave approximation, and in the full simulation, the populations of the off-resonance quasiparticles are not fully reversed due to the influence of the other terms that were dropped. In this case, the second period of the modulation may further increase their populations even at the special point,  $\theta = \pi$ , leading to the observed finite non-condensate population. As a consquence, the degree to which the excitations can be fully reversed can be interpreted as a measurement of the fidelity of the protocol for producing quasiparticle squeezing. fidelity of the preparation of the squeezed quasiparticle state.



Figure 4.6: The two methods for probing the phase of the quasiparticle squeezing. In both cases, the scattering length at first oscillates at frequency  $\omega = 5500 \ (\hbar/mL^2)$  for a period of time  $\tau$ . At  $t = \tau$ , in (a) the amplitude of the oscillation remains at zero for a time interval  $\Delta t$ , and then the scattering length again oscillates for another period of time  $\tau$ , and in (c) a phase shift  $\theta$  is applied to the oscillation . Panels (b) and (d) show the resulting non-condensate fraction at the final time as a function of  $\Delta t$  or phase shift  $\theta$ , respectively, for the two cases.

#### 4.8 Quasi-2D system

In order to make a more robust connection with the recent Bose firework experiment [18], we would like to generalize the formalism we have presented from a quasi-1D gas trapped in a box potential to a quasi-2D gas that is initially trapped by a circular potential with the third out-of-plane direction frozen. Although this geometry adds new degrees of freedom to our previous analysis, we may exploit the fact that the circular system possesses cylindrical symmetry, so that the wavefunction of the condensate can be solved effectively as a 1D problem in the radial coordinate. Note that the quasi-2D system differs from the quasi-1D system in a number of important ways. The momentum correlations will manifest as angular correlations that may be detected by looking for atom-atom coincidence on two detectors aligned in opposite directions. Furthermore, the divergence properties of the renormalization problem are qualitatively different in two dimensions, as discussed previously.

We begin by writing the fluctuations in the field operator in the quasiparticle basis using appropriate indices for two dimensions,

$$\delta\hat{\psi}(r,\theta) = \sum_{k,l} u_{k,l}(r)e^{il\theta}\hat{b}_{k,l} + v_{k,l}^*(r)e^{-il\theta}\hat{b}_{k,l}^{\dagger}, \qquad (4.43)$$

where k corresponds to the excitation in the radial coordinate and l represents the angular momentum quantum number. The angular momentum will modify the form of the kinetic energy for the 2D quasiparticles by including a new centrifugal term,  $\hbar^2 l^2 / 2mr^2$ , that arises physically from circulation about the trap center. Due to cylindrical symmetry, the normal and anomalous densities should be functions of only three real variables, two radii and a relative angle, which we denote by  $r_1$ ,  $r_2$ , and  $\phi \equiv \theta_2 - \theta_1$ , respectively. For both normal and anomalous densities, we write the functions in terms of their expansion in angular momentum,

$$G_{N,A}(r_1, r_2, \phi) = \sum_{l} G_{N,A}^{(l)}(r_1, r_2) e^{il\phi}.$$
(4.44)

The time evolution can then be solved by substituting this expansion into Eqs. (4.4)-(4.6) and using the appropriate form for the two dimensional kinetic energy.

The first case we consider is for the situation in which the circular trap potential well is infinite and has radius  $R_0$ ,

$$V_{\text{ext}}(r) = \begin{cases} 0 & r < R_0 \\ \infty & \text{otherwise} \end{cases}$$
(4.45)

We prepare the quantum gas in the ground state with a small repulsive scattering length a in order to stabilize the system mechanically. The repulsive interactions are characterized by the appropriate 2D *T*-matrix, as discussed in Section 4.3. Procedurally, we carry out a similar sequence of steps to those previously discussed for quasi-1D. First we solve the GPE using imaginary-time propagation, and use that mean-field solution as the first iteration for the solution of the HFB equations, ignoring the non-condensate terms in the HFB self-energy. As before, this solution is non-stationary and we must iterate between the GPE and HFB solutions in order to find a self-consistent solution whose resulting evolution gives rise to densities that do not depend on time. The resulting three components, the condensate, the normal density, and the anomalous density, are illustrated in figure (4.7). Note that the anomalous density diverges in general as the Hankel function of the first kind as a function of the relative distance  $|\mathbf{r}_1 - \mathbf{r}_2|$  close to the origin. This is an analytic result that can be derived by solving the scattering equation Eq. (4.9) in 2D [3]. This emphasizes an important point; the anomalous density cannot be accessed directly in experiment and does not form an observable.

Now that we have prepared an accurate initial state, we can then begin to examine its time evolution when subjected to a drive via a modulation of the scattering length. Similar to what we saw in quasi-1D, the modulation leads to excitation of quasiparticles with energies on resonance with the modulation frequency. Figure (4.8) shows the normal and anomalous density as a function of time and relative distance at the center of the trap. A principal feature of the radial density dependence is the appearance of phonon-like excitations with well defined wave-number. The noncondensate density increases monotonically with time, as is consistent with the squeezing picture discussed earlier. On the other hand the anomalous density oscillates in time tracking the external field.



Figure 4.7: Ground state solutions for a system with total atom number  $\mathcal{N} = 6 \times 10^5$  in an infinite circular box of radius  $R_0$ . The scattering length is set to  $a = 3.99 \times 10^{-5} l_{\perp}$ . (a) Condensate density as a function of radial position. (b) Quantum depletion density as a function of radial position. (c) Anomalous density with the center of mass position at the center of the trap i.e.,  $\mathbf{r}_1 + \mathbf{r}_2 = 0$ , as a function of the relative distance  $|\mathbf{r}_1 - \mathbf{r}_2|$ . The divergence that scales as the Hankel function of the first kind close to the origin is a result of 2D scattering theory.



Figure 4.8: Normal and anomalous densities in a system that is prepared in the ground state of the self-consistent HFB solution with a positive scattering potential  $a_{dc} = 3.99 \times 10^{-5} l_{\perp}$ , and then subjected to the modulating drive with angular frequency  $\omega = 1200 (\hbar/mR_0^2)$  and constant amplitude  $a_{ac} = 3.99 \times 10^{-5} l_{\perp}$ . (a) Non-condensate density as a function of radial potition  $r = |\mathbf{r}_1 + \mathbf{r}_2|/2$  (the origin is on the right) and time, i.e.,  $G_N(r, |\mathbf{r}_1 - \mathbf{r}_2| = 0, t)$ . (b) Magnitude of the anomalous density as functions of the relative coordinate  $|\mathbf{r}_1 - \mathbf{r}_2|$  (the origin is on the right) and time, i.e.,  $|G_A(r = 0, |\mathbf{r}_1 - \mathbf{r}_2|, t)|$ . Only a small time interval beginning at  $t = 0.6 (mR_0^2/\hbar)$  is shown, so that the oscillations are resolved.

In order to capture the dynamics of the high momentum atoms emitted outwards in the Bose fireworks experiment, we extend our simulations to a system with a finite trap potential that is higher than the initial chemical potential but lower than the kinetic energy of the excited atoms, using a smooth hypertangent functional of form,

$$V_{\text{ext}}(r) = \frac{V_{\text{well}}}{2} \left( 1 + \tanh\left(\frac{r - R_0}{\zeta R_0}\right) \right)$$
(4.46)

where  $V_{\text{well}}$  and  $\zeta$  are positive constants. The form of this external confining potential was chosen to reduce numerical artifacts. Figure (4.9) shows snapshots of the condensate and non-condensate densities at time t = 0.04 and 0.08  $(mR_0^2/\hbar)$  during the drive. At t = 0.04  $(mR_0^2/\hbar)$ , we observe that the condensate density is pushed towards the edge of the trap, and that some non-condensate atoms are generated. At t = 0.08  $(mR_0^2/\hbar)$ , we see phonon-like patterns in the condensate and non-condensate densities that appear as ripples or waves. We see qualitatively a residual excited condensate that represents a component that does not have sufficient energy to overcome the potential barrier, and a non-condensate density containing much more energetic atoms that is observed to propagate outwards and leave the finite trap region. These are due to the energetic quasiparticles created by the drive and form an experimentally observable quantity in ballistic expansion images.

Figure (4.9) also shows the density currents, which indicate the flow of atoms, including those in the condensate and the non-condensate. They are computed from

$$J(r) = -\frac{1}{r} \frac{d}{dt} \int_0^r \rho_{\text{tot}}(r') \cdot r' \, dr' \,, \qquad (4.47)$$

where  $\rho_{\text{tot}}$  is the total density. At  $t = 0.04 \ (mR_0^2/\hbar)$ , the number of high momentum atoms is still relatively small, and therefore the density current is also small. The density currents at  $r = 0.4 R_0$ point outwards while those at  $r = 0.8 R_0$  point inwards. Outside the trap there are few atoms, and the currents are essentially zero. At  $t = 0.08 \ (mR_0^2/\hbar)$ , the density currents near the edge of the trap and outside the trap point outwards with large magnitudes, which represents the emission of high momentum atom pairs. Unlike in the experiment, where jet-like patterns were observed, our simulation results show isotropic images. This is anticipated since the functionals we calculate for the condensate, normal, and anomalous densities, represent probability densities and not individual realizations (i.e., they are ensemble averages over many experimental realizations.) On the other hand, a given experiment is fundamentally different in that it represents a single trial that exhibits shot-to-shot noise associated with the projection that occurs in a single quantum measurement. In order to model this projection noise it would be necessary to simulate quantum trajectories [34], rather than solving for the density matrix evolution, and this may be done by adding white noise to the initial condensate wavefunction (see for example Ref. [25]).

In order to demonstrate a quantitative comparison of the energy of the generated quasiparticles with respect to their ballistic motion, we present a numerical 'time-of-flight' calculation. In this simulation, we measure the momentum of the atoms in the system by evaluating the speed at which the gas expands. We define the effective size of the gas as the radius encircling a large fixed fraction of the non-condensate (say more than 90%), so that at t = 0 we begin with size  $R_0$ . Using this metric, from our simulations, we observe that initially a large amount of non-condensate density is generated close to the center of the trap and most of the non-condensate atoms have not left the finite trapping region, so that the size of the gas appears to be shrinking. However, later in the evolution and after a significant fraction of non-condensate atoms escapes the trap, the expansion of the size of the gas becomes essentially ballistic (expansion size increasing linearly in time), and the speed of the expansion is approximately  $\sqrt{(\hbar \omega - 2\mu)/m}$ . The reason is as follows. From the HFB equations, we know that

$$\left(-\frac{\hbar^2}{2m}\nabla^2 - \mu + 2V|\phi_a|^2\right)u_k + V|\phi_a|^2v_k = \epsilon_k u_k.$$
(4.48)

Using the approximations  $V |\phi_a|^2 \approx \mu$  and  $v_k \approx 0$  for large k, we may write

$$-\frac{\hbar^2}{2m}\nabla^2 u_k \approx (\epsilon_k - \mu) \, u_k \,. \tag{4.49}$$

Furthermore, since the resonant quasiparticle energy is  $\hbar\omega/2$ , we find

$$\frac{1}{2}mv^2 \approx \left(\frac{\hbar\omega}{2} - \mu\right) \,. \tag{4.50}$$

The speed of the expansion is evaluated from the slope of the fitted curve in figure 4.10. The fit only includes data points from time  $t = 0.065 \sim 0.08 \ (mR_0^2/\hbar)$  so that we avoid the initial transients where interaction energy within the condensate and between the condensate and noncondensate is significant, and the motion is not ballistic. The speed we get from the slope agrees well with the analytical result derived above.

## 4.9 Conclusion

In this chapter, we have developed a description of a condensate and non-condensate system starting from the many-body field theory Hamiltonian and deriving the evolution equations for the condensate, normal density and anomalous density. Since we assumed contact interactions, the contact potential may lead to divergences in the field theory at small and large momenta. We took care of this issue by properly renormalizing the scattering potential.

We solved the quantum fluctuations in the initial stationary state in 1D using the selfconsistent HFB theory, which does not involve any free parameters. Then, we simulated the amplification of the quantum fluctuations with a well-defined energy using the time-dependent equations. The amplification has aspects similar to the generation of squeezed states of light, and we were able to verify that the variance of the quadrature phases can fall below the standard quantum limit. We proposed to observe phase sensitive quantities through two alternate approaches including Ramsey interferometry and discrete phase jumps. We showed how this is able to provide information on the characterization of quasiparticle squeezed states. Finally, we showed simulation results in 2D, and found that the excited non-condensate atoms eventually leave the trap and propagate outwards at a well-defined speed, consistent with the experimentally observed time-offlight results. Although we showed simulation results for only quasi-1D and quasi-2D systems, 3D systems would be interesting and can be analyzed systematically using similar approaches.



Figure 4.9: Densities of the condensate and the non-condensate at t = 0.04 and  $0.08 (mR_0^2/\hbar)$ . The dark blue area is the range of simulation, and the white line indicates the trapping potential. The yellow arrows represent the density currents, with their lengths proportional to the magnitude. The system starts with the ground state with  $a_{dc} = 1.99 \times 10^{-5} l_{\perp}$  and then the scattering length is modulated with amplitude  $a_{ac} = 1.99 \times 10^{-4} l_{\perp}$  for all time. The smoothing parameter of the finite circular well was set to  $\zeta = 0.2$ . At  $t = 0.04 (mR_0^2/\hbar)$ , the atoms in the condensate are pushed towards the edge of the trap, but because they are of low energy, they cannot escape the trap. The number of high momentum atoms is still relatively small, and therefore the density current is also small. The density currents at  $r = 0.4 R_0$  point outwards while those at  $r = 0.8 R_0$ point inwards. Outside the trap, there are few atoms, and the currents are essentially zero. At  $t = 0.08 (mR_0^2/\hbar)$ , a large fraction of non-condensate atoms with high energy escape the trap, and the density currents near the edge of the trap and outside the trap point outwards with large magnitudes. For clarity, the density currents at  $t = 0.04 (mR_0^2/\hbar)$  are scaled up 3 times compared to those at  $t = 0.08 (mR_0^2/\hbar)$ .



Figure 4.10: The size of the non-condensate as a function of time. The size is defined as the radius R that encircles 96.2% of the non-condensate atoms. The system starts with initial chemical potential  $\mu = 26 \ (\hbar^2/mR_0^2)$  and is driven with modulation frequency  $\omega = 1000 \ (\hbar/mR_0^2)$ . The speed of the propagation v is given by the slope of the fitted line (red),  $v = 30.8 \ (\hbar/mR_0)$ . This value is approximately equal to  $\sqrt{(\hbar\omega - 2\mu)/m}$ , the speed of a particle whose kinetic energy is half the photon energy minus the chemical potential. We have excluded the data points at times when the motion is not ballistic, including at early times where few non-condensate atoms escape the trap and the interaction between the condensate and the non-condensate is significant.

Note that there is one other important consideration; our model does not include the collisions terms in the kinetic theory that result in equilibration of the gas to its thermal state [53]. Neglecting collisions is a good approximation for a dilute gas at low temperature, but implicitly requires us to limit our discussion to the regime in which the time-scale between consecutive two-body collision events greatly exceeds the time-scale of the quantum dynamics that we investigate.

We have demonstrated a method to generate momentum squeezed states that may be useful for metrology applications. This motivates us to further consider engineering the scattering length as a function of time to generate two-mode squeezed states in quasimomentum that could be injected into matter-wave interferometry. The entanglement properties of such states would be interesting to investigate along with the metrological gain that arises from the quantum advantage.

The importance of pairing in this work also motivates us to consider a similar experiment on fermions, where the interactions could be modulated by variation of the scattering length in the BEC-BCS crossover regime. The motivation for this is simply that the pairing physics is closely connected with the previously observed fermionic condensation. These considerations will be the subject of future studies.

## Chapter 5

# Reinforcement-learning-based matter-wave interferometer in a shaken optical lattice

## 5.1 Introduction

Extensive efforts have been made to develop metrological devices based on interference and detection of either electromagnetic waves or matter waves. By exploiting the quantum aspects of superposition and entanglement, one can potentially achieve high sensitivity to phase shifts, and this has inspired a wide variety of applications, including detecting gravitational waves [1], measuring the fine-structure constant [70], testing the universality of free fall [24], and inertial sensing for GPS navigation [32, 76].

The usual direct-design approaches to engineering complex quantum systems that consist of many degrees of freedom or many particles are typically founded on experience with simpler analogs or intuition for underlying mechanisms. This naturally leads to a paradigm that is most accessible in terms of understanding, but incorporates human bias that may potentially generate non-optimal solutions. With this perspective in mind, we point out there are a few purely systematic methods that are often used as a way to develop unbiased strategies, including optimal control [72] and optimization algorithms such as the Nelder-Mead simplex [63] or simulated annealing [43]. Utilizing these methods can allow one to explore solutions with more complicated forms with the potential to reach closer-to-optimum control protocols.

Recently, it has been shown that quantum design may benefit tremendously from a branch of machine learning based on trial and error, known as reinforcement learning, that aims to employ



Figure 5.1: Interferometers are composed of (i) a beam splitter, (ii) mirrors, and (iii) a recombiner. Examples shown are; (a) an optical Mach-Zehnder interferometer (using half-silvered or conventional mirrors), (b) a Bragg interferometer (using three short light pulses of varying pulse area,  $\pi/2$  or  $\pi$ ), and (c) a shaken lattice interferometer. In (c), the shaken lattice mimics the interferometer components by splitting, reflecting, and recombining the atoms, through design of a specific shaking function for each case.

machines to find the optimal strategy for accomplishing a specific task. One of the reasons for its success is that the learning framework is decoupled from human intuition, and therefore may explore novel solutions that have been previously undiscovered. Moreover, in situations where problems are so complex or extensive that naïve brute-force algorithms are considered unfeasible, such as playing games like *go* and *chess* [85, 87, 86], sophisticated reinforcement-learning algorithms have been developed to enable the machines to perform at a level that exceeds human capability. There are many systems in quantum physics that fit naturally within this scope due to their underlying complexity. For example, reinforcement learning has been applied to control and study phase transitions of many-body quantum systems [14], to find strategies for quantum error correction [23, 62], to design quantum circuits [65], to prepare novel quantum states [13, 55, 58], to find protocols for quantum communication [95], and to improve quantum sensors [83].

In this chapter, we will investigate the potential for reinforcement learning to improve on existing matter-wave interferometery. In Sec. 5.2, we describe the physical model for the interferometric system. We outline our reinforcement learning approach in Sec. 5.3 and explain the key concepts crucial to understanding its application. In Sec. 5.4, we demonstrate how the reinforcement-learning framework may be used to design a beam splitter and a mirror, and present the protocols that are learned. Following this we show how the components can be cascaded together to form a complete interferometer. In Sec. 5.5, we evaluate the resulting performance of the interferometer through a Bayesian statistical analysis.

# 5.2 Physical model

Although the topology of interferometers can vary somewhat, the archetypal design can be thought of as the Mach-Zehnder interferometer [54, 107]. This device is sensitive to the differential phase accumulated between two alternate paths [see Fig. 5.1(a)]. Mach-Zehnder interferometers are composed from three essential components: A beam splitter that separates the wave coherently into two directions, mirrors that reflect both parts, and a recombiner where the waves are brought back together and constructively or destructively interfere. While in the optical case, components that split or reflect light beams are readily available (e.g., half-silvered or conventional mirrors), for matter-wave optics, the components have to be generated through the careful control of laser-atom interactions. The traditional way to do this is through Bragg diffraction [see Fig. 5.1(b)] [11, 40], in which a sequence of three short light pulses are used to separate, reflect and recombine the matter wave.

An alternative approach is possible for atoms moving in an optical lattice potential where the intensity pattern can be shifted backward and forward in time [102, 100, 101]. The elementary components, i.e., a beam splitter, reflectors, and a recombiner, in that case may be implemented by "shaking" the lattice in a tailored pattern. If we denote the canonical position and momentum operators of an atom by  $\hat{x}$  and  $\hat{p}$ , respectively, with m the atom mass, the optical lattice system can be described by the Hamiltonian

$$\hat{H}(t) = \frac{\hat{p}^2}{2m} - \frac{V_0}{2} \cos\left[2k_L\hat{x} + \phi(t)\right],$$
(5.1)

where  $k_L$  is the laser wavenumber, and  $\phi(t)$  is the time-dependent phase difference between two counter-propagating lasers. The lattice, with constant amplitude  $V_0$ , is shaken through the variation of  $\phi(t)$ , since that parameter determines the position of the nodes and the anti-nodes of the standing wave intensity.

Unlike a typical Mach-Zehnder interferometer, which is formally a two-port device and transforms quantum states according to simple SU(2) group rotations, the relevant eigenstates of the shaken lattice potential consist of many Bloch states that can be coupled by the time-dependence of the laser phase,  $\phi(t)$ . In some sense this situation represents a highly multi-path form of interferometry since the Bloch basis provides many accessible paths for the quantum wave function to explore. While this establishes a rich evolution and could provide metrological benefit, it makes it more difficult to design an intuitive control protocol. It is for this goal of obtaining a high performance solution among a complex landscape that we are led to consider reinforcement learning as a design approach.

## 5.3 Model-free reinforcement learning for design

In order to understand our design philosophy, we first need to address the important considerations to make when using machines to find solutions to design tasks. In this section we will present the principal ideas and concepts that will form the underlying foundations of our learningbased methodology. Since there are a variety of methods available, we begin with a discussion of the main structure that we will employ.

Reinforcement learning consists of a closed loop in which an *agent* invokes *actions* based on the observed *state* of an *environment*, and an environment provides *rewards* to the agent based on the observed outcome of its actions. The agent is tasked with the goal to discover the sequence of actions for which it receives the highest possible terminal reward [90] (see Fig. 5.2). It does this by trial and error, iteratively improving its actions in such a way as to corral the environment toward a target configuration.

We are primarily interested in a specific kind of reinforcement learning, known as modelfree learning [90], where the agent has no detailed insight as to the structure of the environment, and cannot know *a priori* what effect an action will have on the environment state. Here the agent learns and makes decisions even when the environmental model may not be fully known. This general setting can even be applied to the situation in which the environment is represented by an experimental apparatus that cannot be fully understood, and the reward is derived from experimental observation and feedback. Using this approach significantly expands the potential scope of our work, since it is precisely this aspect that distinguishes model-free reinforcement learning from classical optimal control methods. In classical optimal control, the optimization depends heavily on the mathematical form of the model and, for complex systems, can be difficult to implement in practice.

Typically, reinforcement learning problems are Markov decision processes, where the probability of transition to the next state is only dependent on the current state and the current action, and not on prior history. Evolution in this framework is referred to as a *trajectory*, where the state is initially prepared, and then a sequence of actions and corresponding state updates are performed. The trajectory steps continue until the chain comes to an end when a predetermined condition is reached, which might be a certain number of steps, or a terminating state of the environment. An important concept is the idea of an *optimum trajectory* starting from any state, which is a trajectory such that the terminal reward is the maximum possible. We use what is known as Q-learning to make the action decisions, as we now describe. Although this approach is not the only possible choice, it is appropriate when the effect of an action on the environment state is deterministic, which will be the case here.

Since there is potentially an enormous dimensionality of the state space and the action space, a brute-force search is not possible, and we employ a neural network to learn and to approximately optimize the crucial decision-making task. That is, we define the input nodes of a neural network to be a representation of an arbitrary vector in the state space. We ascribe to the output of the neural network a vector of quality factors, simplified to Q-values,  $Q(s, a) \in \mathbb{R}$ , where each element represents the desirability for a possible action, a, given an input state, s. Ideally the neural network should take *any* arbitrary state as input and output a distribution in which the maximum Q-value corresponds to the best next action to perform.

We train the neural network in the following manner. The agent will utilize a policy in order



Figure 5.2: Framework of reinforcement learning. The agent chooses an action, the environment responds to the action, and gives the next state and a reward as feedback. In our design task, the action is the translation of the optical lattice, the environment evolves according to the Schrödinger equation, the observation is the momentum population distribution, and the reward is a function of the quantum fidelity. Illustrated on the left is an example neural network as the decision-making agent. The neural network takes the state represented on its input layer, passes it through a hidden layer, and generates a vector of Q-values at its output layer.

to decide on the next action for a given state, for example, by taking the action associated with the maximum Q-value from the neural network output vector. The environment, in some initial state s, receives the action a from the agent, and then updates the state, symbolically denoted as

$$s \xrightarrow{a} s'$$
.

The environment then reports back to the agent the reward, r(s'). As a step in a sequence, this leads to the important concept of the discounted cumulative reward, known simply as the *return*. The return, Y(s'), is defined as the combination of a current reward, r(s'), and every future reward that would be generated in an optimum trajectory seeded by s'. Since any current action will have decreasing influence on future decision making as the steps become more distant, it is useful to deweight each consecutive reward by a discount factor,  $\gamma \in [0, 1]$ , i.e.,

$$Y(s') \equiv r(s') + \gamma r(s'') + \gamma^2 r(s''') + \dots$$
(5.2)

where

$$s' \xrightarrow{a'} s'' \xrightarrow{a''} s''' \xrightarrow{a'''} \dots$$

are steps along an optimum trajectory.

The learning process is now framed as the minimization of the squared difference between the return, Y(s'), and the network output, Q(s, a), where the minimization is accomplished by varying the internal weights and biases that constitute the neural network. Should the learning be perfected, the squared difference would be zero, implying that the Q-value and corresponding return would be identical. We emphasize that, while mathematically complete, this learning definition is a formal construction that is intractable as written, since if we knew the optimum trajectory there would be no need to carry out the learning at all. In fact, the only information that is accessible about the future component of Y(s') is the inference one can make from the neural network output itself, introducing a self-consistency element to the optimization problem. In other words, the series in Eq. (5.2) is approximated as

$$r(s'') + \gamma r(s''') + \ldots \approx \max_{a'} [Q(s', a')]$$
 (5.3)

so that

$$Y(s') = r(s') + \gamma \max_{a'} [Q(s', a')].$$
(5.4)

The final result given in Eq. (5.4) encapsulates all the principal concepts that establish a closed learning formulation that can be computationally implemented.

We employ several technical but important improvements that increase the convergence, such as to replace the neural network used for the calculation of the future return, as approximated in Eq. (5.3), with a second identical neural network, the target network. The target network is updated at each step from the Q-network by a weighted average leading to increased stability [61, 94]. Another useful trick is to employ an  $\epsilon$ -greedy policy in the training stage [90], which simply means that the action associated with the highest Q-network output value is not always chosen, but with  $\epsilon$ likelihood, a random action is selected. This allows the neural network to explore a more extensive state space and to avoid becoming trapped in a locally optimal solution. Finally, as is common to many forms of machine learning, a variety of batch sampling methods, in this case for a replay buffer of previously experienced transitions, (s, a, s'), can be used to assist in the efficiency and stability of the learning minimization stage [52]. For further details of the algorithmic structure we refer the reader to Chapter 3.

The Q-network we will use for includes an input layer, one hidden layer, and an output layer. Each layer is connected to the next one by a linear function followed by a ReLU (Rectified Linear Unit) activation function, which represents the function  $\max(x, 0)$  with the argument x being a real number. The relations between layers are explicitly written out as follows,

$$\boldsymbol{h} = \operatorname{ReLU}(\mathbf{W}_1 \boldsymbol{x} + \boldsymbol{b}_1) \tag{5.5}$$

$$\boldsymbol{y} = \operatorname{ReLU}(\mathbf{W}_2\boldsymbol{h} + \boldsymbol{b}_2), \qquad (5.6)$$

where h denotes the hidden nodes, x denotes the input nodes, y denotes the output nodes,  $\mathbf{W}_1$ ,  $\mathbf{W}_2$  are the matrices of weights, and  $b_1$ ,  $b_2$  are vectors of biases. An example topology of the Q-network is shown in Fig. 5.3.



Figure 5.3: An example topology of the network that we used in the deep Q-algorithm.

## 5.4 Component design

We now use the Q-learning framework to design the essential elements for building an interferometer, i.e. a splitter, reflectors, and a recombiner. One aspect that is worth emphasizing is that the standard Mach-Zehnder interferometer topology may not always be the optimal layout for determining an accumulated quantum phase. For example, it has been demonstrated in [57] that novel optical experimental setups that could potentially provide sensitivity surpassing that of the Mach-Zehnder configuration can be designed with reinforcement learning. However, with the aim of simplicity, and even though it constrains our solution space, we will limit the discussion here to the design of only the specified individual interferometer components. This will allow us to evaluate the efficacy of the learned protocols by direct comparison with the device characteristics of the analogous Bragg interferometer. In the future, it may prove beneficial to allow for more freedom in the design topology and to thereby ascertain whether this leads to further performance gains.

We have developed the learning terminology around words such as "state" and "environment" where they have been given a specific meaning. We are now going to apply the methodology to the design task of controlling a complex quantum system. Unfortunately we will encounter the issue that terms such as state and environment have conflicting and completely different meanings in the theory of open quantum systems. In particular, to enable an efficient machine learning algorithm with an addressable number of input nodes, the state space that we use is typically a compact representation of a quantum state or operator. The environment, which denotes the system that we wish to control and that is evolving under our imposed actions, is a completely different use of the term to that common in open quantum system literature. In the next design descriptions, we will provide the dictionary for mapping the learning terminology to the quantum variables.

#### 5.4.1 Beam splitter

The beam splitter is a target-state design problem that can be specified as follows. We initialize a quantum state in the ground state of an optical lattice potential, where the lattice depth  $V_0$  is set to 10  $E_r$  ( $E_r = \hbar k_L^2/2m$  is the recoil energy). The task is to shake the lattice in such a way as to split the atoms from the ground state into an approximate equal superposition of  $|\pm p_0\rangle$  momentum states. We will focus on  $p_0 = 4\hbar k_L$  here, which we find to balance well the tradeoff between large momentum splitting and high-frequency components necessary for the shaking function. If the target quantum state is a lattice eigenstate it is a stationary state under free evolution, avoiding the need for a free evolution protocol. As a consequence, we assign the target quantum state to be the third-excited Bloch state with zero quasimomentum. This eigenstate has the attractive property that it closely approximates the desired superposition of momenta, although not perfectly, and can be strongly coupled to the ground state by a reasonably simple shaking solution.

The state space for learning, i.e., the vector space containing s, is a compact representation of the Hilbert space. For this learning task, we define s to be a vector of populations (i.e., probability distribution) for momenta chosen from a comb of discrete values,  $\{2n\hbar k\}$  for  $n \in \mathbb{Z}$ ,  $|n| < n_{\text{max}}$ . This choice is motivated by the fact that the stimulated absorption and emission of photons only couple momenta separated by this quantized spacing. The possible actions, a, are also reduced to a discrete set of possibilities, each action corresponding to a specific constant phase to be applied during the corresponding step. This means that with these associations, the deep Q-network is specified as mapping momentum populations at its input layer to Q-values at its output layer associated with a finite set of possible lattice phases to be applied during the next update.

In order to apply the action generated by a choice of lattice phase and to thereby determine the updated momentum distribution that results, we need to specify the environmental model. In principle, this could be experimental or theoretical, and even if theoretical, could include stochastic effects that arise from dissipation and noise. However, for simplicity, we will take the most basic formulation. This means that our environmental model will update a pure quantum state through the Schrödinger equation evolution for an isolated system. The reward will be based on the fidelity between the evolved quantum state and the target quantum state, i.e., the third-excited Bloch state, as calculated from the modulus-square inner product.

Actually, this environmental model allows us to accomplish two design tasks at the same time. Once we find a protocol for splitting, the recombining task can be achieved by simply implementing the time-reversed protocol, a consequence of the time-reversal symmetry underpinning the Schrödinger evolution and the time-reversal symmetry of both the initial and the target states.

A reinforcement learning outcome for the beam splitter is shown in Fig. 5.4, which illustrates the learned shaking function,  $\phi(t)$ , and the corresponding evolution for the momentum probability distribution that results. Even though this sequence of phases does not appear to have a predictable structure, it is apparent that at the terminal time the momentum distribution has the anticipated form of two well-defined peaks at  $\pm 4\hbar k_L$ . The calculated fidelity of the state is approximately 95%. It would be possible to optimize further by expanding the set of possible phase elements, by increasing the total evolution time, or by optimizing the learning cycles. However, any real experiment may possess imperfections and other aspects that are not well described by our isolated system's Schrödinger evolution, and the level of performance of our design is sufficient for the task at hand.

## 5.4.2 Details of reinforcement learning for target state

Here we present how we formulate the control for splitting as a reinforcement learning problem. First of all, we define the input state to be a vector of the populations in the momentum eigenstates,

$$\{|2n\hbar k_L\rangle \mid n = -3, -2, -1, 0, 1, 2, 3\}.$$
(5.7)

The momentum states that are neglected here are barely populated the whole time. The actions are chosen from a set of discrete phase values,

$$\{n\pi/4 \mid n = -4, -2, -1, 0, 2\}.$$
(5.8)

The reward is a function of fidelity, defined as

$$r = \begin{cases} 0 & \text{if } d = 0 \\ \\ \frac{\mathcal{F}}{1 - \mathcal{F}} & \text{if } d = 1 \end{cases}$$
(5.9)

$$\mathcal{F} = |\langle \psi_{\text{target}} | \psi(t) \rangle|^2, \qquad (5.10)$$

where d = 1 when the maximum terminal time is reached or when the fidelity is higher than the threshold 0.95, and d = 0 otherwise. The Q-network we use for this task has an input dimension of 7 and an output dimension of 6, as a result of (5.7) and (5.8). The neural network is implemented in PyTorch. The values for the hyperparameters used for training this Q-network are listed in Table 5.1. The hyperparameters are determined from a randomized grid search carried out in the

Hyperparameters	Values
$\gamma$	0.999
au	0.999
$\alpha$	0.001
episodes	20,000
$\epsilon$ decay	0.0001
Hidden size	98
Batch size	64
Optimizer	Adam

Table 5.1: Choice of hyperparameters for the splitting task. The exploration probability  $\epsilon$  decays linearly for each episode at the rate represented by  $\epsilon$  decay. The optimizer we use, 'Adam' (which stands for adaptive momentum estimation), is a method for efficient stochastic gradient descent [42].

package Ray Tune, where a set of hyperparameters that results in the best reward at the end of the training among the 20 sets is picked out.

#### 5.4.3 Mirror

The second task is to reflect the momentum from  $|\pm p_0\rangle$  to  $|\mp p_0\rangle$ , that is, corresponding to a matter-wave mirror. We should point out at the start that this task is essentially different in character from the beam splitter, because it is not a target-state but a target-operator design problem. The mirror is defined by the desired map,

$$\alpha |p_0\rangle + \beta |-p_0\rangle \to \alpha |-p_0\rangle + \beta |p_0\rangle,$$

for any arbitrary  $\alpha$  and  $\beta$ . This implies that the target unitary is any operator  $\hat{U}_{\text{target}}$  that satisfies

$$\hat{P}\hat{U}_{\text{target}}\hat{P} = |p_0\rangle \langle -p_0| + |-p_0\rangle \langle p_0|, \qquad (5.11)$$

where  $\hat{P} = |p_0\rangle \langle p_0| + |-p_0\rangle \langle -p_0|$  is the projector onto the 2-dimensional subspace. The design goal is, therefore, to shake the lattice as a function of time, such that the corresponding unitary evolution,  $\hat{U}(t)$ , approximates any  $\hat{U}_{\text{target}}$  that satisfies Eq. (5.11) at the terminal time. The vector space containing s is defined as a compact representation, in this case, of the set of quantum unitary operators. Given the mirror behavior, a suitable choice for s for this learning task is a vector of norms of selected matrix elements of the unitary represented in momentum space, i.e.,  $|\langle \pm p_0 | \hat{U}(t) | p = 2n\hbar k \rangle|^2$  for  $n \in \mathbb{Z}$ ,  $|n| < n_{\text{max}}$ . For the set of possible actions, a, we use a different parametrization of the phase than for the beam splitter design. We found that the quality of the mirror is very sensitive to the frequency of the shaking function. The characteristic frequency,  $12\omega_r$ , where  $\omega_r = \hbar k_L^2/2m$  is known as the recoil frequency of the lattice, corresponds to the kinetic energy difference between  $|4\hbar k_L\rangle$  and  $|2\hbar k_L\rangle$ , and works especially well. Consequently, we enforce a sinusoidal  $\phi(t)$  at frequency  $12\omega_r$ , and at each half cycle, the agent chooses the action a as the amplitude of the sinusoidal shaking function selected from a small set of possible values. This means here that the deep Q-network is specified as mapping a compact representation formed from the norm of a subset of elements of the unitary matrix to Q-values associated with the amplitude of the sinusoidal shaking function for its next half cycle.

We also need to define what is meant by the environment in this case. The environmental model establishes the unitary operator update during a shaking function half cycle. If, as before, we take the environmental model to be a closed quantum system, the evolution is given by solving the Schrödinger equation of motion for the time evolution operator

$$i\hbar \frac{d\hat{U}(t)}{dt} = \hat{H}(t)\hat{U}(t).$$
(5.12)

The reward is associated with the channel fidelity in the  $d_s = 2$  dimensional subspace spanned by  $|\pm p_0\rangle$  (see [71])

$$\mathcal{F} = \frac{1}{d_s(d_s+1)} \left[ \operatorname{Tr} \left( M M^{\dagger} \right) + |\operatorname{Tr} (M)|^2 \right]$$
(5.13)

where the density operator is  $M = PU_{\text{target}}^{\dagger}U(T)P$ .

An example outcome from reinforcement learning for the mirror is shown in Fig. 5.5. This illustrates the learned shaking function and corresponding momentum probability density that this generates starting from  $-4\hbar k_L$ . From matrix elements of the resulting unitary operator, we also verify that the phase relation between the coefficients of the  $\pm 4\hbar k_L$  states is preserved. As for the beam splitter case, it would be difficult to anticipate the form of  $\phi(t)$ , and yet the resulting quantum state evolution closely approximates that expected for a mirror. The channel fidelity reaches 93% at the terminal time, which as discussed earlier could still be improved on, but nevertheless is sufficient for our purpose.

## 5.4.4 Details of reinforcement learning for target unitary

The reflection operation is defined by a target unitary operator that satisfies Eq. (5.11). Our goal is to control the phase as a function of time, such that the unitary operator

$$\hat{U}(t) = \mathcal{T}\left\{\exp\left[\int_0^t -i\hat{H}(\phi(t'))dt'/\hbar\right]\right\},$$
(5.14)

reaches the target operator at the terminal time. Here  $\mathcal{T}$  represents the time-ordering operator and  $\hat{H}(\phi)$  is defined in Eq. (5.1). To formulate the control problem as a reinforcement learning problem, we assign the state to be the matrix elements of the unitary operator. Since before the wave function enters the reflection stage, most of its population is in the  $|\pm 4\hbar k_L\rangle$  subspace, we only keep track of the operation in this subspace. Because of this, we reduce the input state to the modulus square of the matrix elements for only

$$\{|2n\hbar k_L\rangle \langle \pm 4\hbar k_L| | n = -3, -2, -1, 0, 1, 2, 3\}.$$
(5.15)

The state is chosen so that it forms an experimental observable. This state can be obtained by initializing the quantum state to one of the  $|\pm 4\hbar k_L\rangle$  states, applying the control and measuring the populations for all the momentum states through time-of-flight imaging.

As discussed in the main text, we found that the frequency of the shaking function is important for the performance of the mirror. Simply by modulating at the characteristic frequency,  $12\omega_r$ , which corresponds to the energy difference between  $|4\hbar k_L\rangle$  and  $|2\hbar k_L\rangle$ , with a fixed amplitude for a duration of  $\sim 3\omega_r^{-1}$ , the channel fidelity can reach as high as 0.8. We take advantage of this knowledge, and try to improve the fidelity on the fixed-amplitude modulation. We choose our actions to be the amplitude of the sinusoidal modulation to the phase, so the resulting phase is  $\phi(t) = \text{Amp}(t) \times \sin(12\omega_r t)$ . The amplitudes are chosen from a discrete set of values,

$$Amp = \{0.4, 0.6, 0.8, 1.0, 1.2\}.$$
(5.16)

The time interval between each decision point is  $\pi/12 \approx 0.26 \,\omega_r^{-1}$ , and in each interval the amplitude is held constant.

The reward function is the same as Eq. (5.9), except that the fidelity  $\mathcal{F}$  is now defined with the channel fidelity in the relevant subspace [see Eq. (5.13)].

With the states, actions, rewards specifically designed for reflection, we use the double deep Q-learning algorithm to learn the strategy for controlling the amplitude of the sinusoidal modulation to the phase. The deep Q-network used for learning the mirror has an input size of 14, and an output size of 5, as a result of (5.15) and (5.16). The hyperparameters are listed in Table 5.2.

Hyperparameters	Values
$\gamma$	0.99
au	0.99
$\alpha$	0.001
$\epsilon$ decay	0.0005
episodes	8,000
Hidden size	128
Batch size	32
Optimizer	Adam

Table 5.2: Choice of hyperparameters for the reflection task.

## 5.4.5 Matter-wave interferometer

We can now cascade together the learned components to make a shaken lattice interferometer. In between the components, we allow the system to propagate freely with the lattice present but with  $\phi(t) = 0$ . This gives a time evolution of the wave function in real space as shown in Fig. 5.6, where the state has been initialized to be in the ground state of the lattice multiplied in real space by a Gaussian envelope with a width of a few lattice sites. The free-propagation time is set to  $10 \ \omega_r^{-1}$ . During free propagation, we can see the wave packets propagating at the anticipated  $\pm 4\hbar k_L/m$  group velocity. The anticipated features are evident in the observed splitting, reflection and recombination of the matter wave.

## 5.5 Statistical analysis

We compute the final momentum distribution on a grid of acceleration values in order to see how useful our device is for inertial sensing.

To interpret the outcome of the interferometer, we use Bayes theorem to derive information about the acceleration from the momentum distribution. This approach is always more effective than curve-fitting when there is knowledge of the probability distribution generator, and the difference is particularly notable in situations like this when the distributions are multimodal. The probability distribution of the acceleration after measurement of a particle in a particular momentum state, P(a|p), is given by,

$$P(a|p) \propto P(p|a)P(a), \qquad (5.17)$$

where P(a) is the prior distribution of the acceleration, and the proportionality is resolved by normalizing the total probability to unity. The probability for measuring a particular momentum p conditioned on the acceleration a, P(p|a) is directly calculated by propagating the wave function in time. We combine multiple measurements by iterating Eq. (5.17) to formulate the conditional probability distribution based on the entire measurement record [35]. For a measurement record  $\{p_1, p_2, \ldots, p_N\}$ , where N is the number of measurements, the probability distribution becomes

$$P(a|p_1,\ldots,p_N) \propto P(p_N|a)\ldots P(p_1|a)P(a).$$
(5.18)

normalized to unity. Since we assume that each atom is independent of each other, N is the total number of atoms we observe. We combine the distributions for each atom using Eq. (5.18) to accumulate the distribution for a conditioned on the measurement record, and estimate the acceleration by taking the expectation value.

In Fig. 5.7(a), we show an example of the parameter estimation process, where the actual acceleration is  $-3 \times 10^{-4} \omega_r v_r$  ( $v_r = \hbar k_L/m$  is the recoil velocity). The corresponding momentum probability distribution for each atom is shown in the inset. What is important about this distribution is that it is almost unique for each acceleration value and therefore acts as a fingerprint. In Figure 5.7(b), we show the standard deviation of the acceleration as a function of the number of measurements. A typical number of atoms for an ultracold gas experiment may be of order  $\sim 10^4$ , and with  $10^4$  measurements, the estimated acceleration agrees extremely well with the actual value. We observe that the standard deviation approaches the standard quantum limit, which scales as  $\sim 1/\sqrt{N}$ . This result is consistent with the fact that we are not considering atom interactions in the system, so each atom may be thought of as making an independent measurement.

We demonstrate that reinforcement learning provides additional capability over traditional experimental techniques by comparing the sensitivity of the shaken lattice interferometer with the conventional Bragg interferometer. The Bragg interferometry sequence is shown in Fig. 5.1(b). We use the Bayesian parameter estimation result with the Bragg interferometer as a baseline for benchmarking the shaken lattice interferometer. To make a fair comparison, we consider the case where the Bragg interferometry sequence takes the same time as the shaken lattice interferometer The protocol found by our reinforcement learning algorithm generates a momentum splitting of  $8\hbar k_L$  (between  $-4\hbar k_L$  and  $+4\hbar k_L$ ) at the beam splitter, four times larger than the  $2\hbar k_L$  (between  $-1\hbar k_L$  and  $+1\hbar k_L$ ) splitting from Bragg diffraction, and therefore this results in higher sensitivity. One may argue that it is possible to generate larger momentum splitting with Bragg diffraction, and while this is true, higher-order transitions typically require much more time than than the short pulses applied here. For these reasons, we have observed from our simulations that the standard deviation of the shaken lattice interferometer can be approximately four times lower than that of the corresponding Bragg interferometer given the time constraint, implying that we have realized an approximate factor of four in sensitivity gain. The results are shown in Fig. 5.7(b).

We verify that the standard deviations  $\sigma_a$  for both the shaken lattice interferometer and the Bragg interferometer are close to the limits set by the Cramér-Rao bound, which can be calculated from the classical Fisher information as,

$$\sigma_a \geq 1/\sqrt{NI_1(a)},$$

$$I_1(a) = \sum_p \frac{1}{P(p|a)} \left[\frac{\partial P(p|a)}{\partial a}\right]^2,$$
(5.19)

where  $I_1(a)$  is the Fisher information for one independent measurement. The Cramér-Rao bound for the shaken lattice interferometer is determined by numerically calculating the classical Fisher information using the pre-calculated probability P(p|a), and the Cramér-Rao bound for the Bragg interferometer is determined by the analytical form  $\sigma_a \geq (2k_LT^2)^{-1}N^{-1/2}$ , where T is the free propagation time.

Note that the sensitivity presented here is not the ultimate achievable sensitivity since we are constraining system parameters. It would be possible to improve the sensitivity of our interferometer, for example, by increasing the free-propagation time, or by further accelerating the matter-wave components in opposite directions after applying our splitting protocol, as discussed

in Refs. [26, 69].

## 5.6 Conclusion

In summary, we have demonstrated a machine learning methodology to control a complex quantum system for the purpose of performing a quantum metrology task. Specifically, we demonstrated how to utilize reinforcement learning for the design of a lattice-based matter-wave interferometer. We showed that by shaking the lattice with protocols derived from deep Q-learning, matter-wave analogs of optical components are realized. We showed that these can be concatenated together to build a high-precision interferometric device. The multipath interferometer that was constructed in this way is capable of the measurement of acceleration with higher sensitivity than that achieved by a conventional matter-wave interferometer.

While we have assumed that atom interactions are negligible, they could potentially bring more quantum advantage if harnessed appropriately. For example, if we could discover a protocol for shaking the lattice to generate momentum-squeezed states, the sensitivity could potentially be higher than the standard quantum limit that is generated by the statistical averaging of independent atoms. In fact, the quantum entanglement in such squeezed states could allow the resulting interferometer to approach the ultimate limit to sensitivity set by the Heisenberg uncertainty relation between number and phase. We have focused on demonstrating a specific solution to the problem of designing a Mach-Zehnder interferometer, but our main outcome is actually to illustrate an effective design methodology. In the future, this approach could potentially lead to the design of alternate protocols that construct completely new types of metrological devices with a performance level surpassing conventional experimental methods.



Figure 5.4: Shaking function for splitting. The learned lattice phase,  $\phi(t)$ , (top) is allowed to take on one of a discrete set of five possible values that span the range shown. The momentum probability distribution (bottom) is initialized to the ground state of the lattice and at the terminal time well approximates the desired superposition for a beam splitter.


Figure 5.5: Shaking function for reflection. The learned lattice phase,  $\phi(t)$ , (top) is sinusoidal and the amplitude is allowed to take on any one of a discrete set of five values at each half-cycle. The momentum distribution (bottom) is prepared in  $-4\hbar k_L$  and well approximates  $4\hbar k_L$  at the terminal time.



Figure 5.6: Time evolution of the matter-wave density throughout the entire interferometry sequence. The white dotted lines separate the plot into five regions. In region I, we apply the splitting protocol. In region II, we allow the matter-wave to propagate freely in the lattice. The appearance of two wave packets traveling in opposite directions shows that the beam splitter operates as expected. In region III, we apply the mirror shaking function. The matter-wave undergoes free propagation in region IV again, and the two wave packets switch directions, demonstrating the functionality of the mirror. Last, we apply the recombining protocol in region V. Apart from the main closed diamond-shaped paths, we observe that there are auxiliary paths that are fainter but still clearly evident. They arise due to the imperfection of the components, and also due to the side peaks arising from the third-excited Bloch state.



Figure 5.7: (a) Posterior probability density of the acceleration for the first 100 atoms. The true acceleration that we aim to reveal by the measurements is  $-3 \times 10^{-3} \omega_r v_r$  (red line), and the measurements are sampled from the momentum distribution at the end of the interferometry sequence, as shown in the inset. (b) Standard deviation of the acceleration estimated using Bayes theorem for up to  $10^4$  atoms. We show the results from both the shaken lattice interferometer (SLI) and the Bragg interferometer and conclude that the shaken lattice interferometer has a higher sensitivity. The standard deviations  $\sigma_a$  are roughly inversely proportional to the square-root of the number of measurements N, for  $N > 10^2$ . The the black dashed lines are the Cramér-Rao (CR) lower bounds and scale exactly as  $1/\sqrt{N}$ .

# Chapter 6

### Practical considerations for the shaken lattice accelerometer

In the previous chapter, we demonstrated the construction of a matterwave accelerometer with individual components designed through reinforcement learning. The sensitivity of the accelerometer shown was obtained from a relatively small enclosed area, with the matterwaves propagating over only about 25 lattice sites. While this was sufficient as a proof of principle demonstration to show the effectiveness of the matterwave-optics components, there is still a lot of potential to gain in the performance of the shaken lattice accelerometer by increasing the dimensions. In fact, the experiment built by the Anderson group and ColdQuanta Inc. can support an enclosed area that is of order 100 times larger. On the other hand, we have not taken into account various noise sources that are unavoidable in a real experiment, which could potentially have detrimental effects on the performance of the interferometer. In this chapter, we will explore the ultimate limits of the shaken lattice accelerometer in terms of sensitivity and dynamic range. Furthermore, we will take into consideration limitations on the engineering design parameters, including laser intensity noise, laser phase noise, dark counts on the CCD camera, and the finite quantum efficiency of the CCD camera. These simulations provide a guide as to the requirements for equipment specifications to achieve a desired sensitivity.

# 6.1 From natural units to experimental units

To connect our theory to real-world scenarios, the first step is to convert the natural units used in the theory to experimental units. The system we consider is made up of an optical lattice formed by counter-propagating lasers with wavelength  $\lambda_L = 852$  nm, and the atoms trapped in the lattice are rubidium atoms <sup>87</sup>Rb, whose individual masses are m = 87 amu. With  $\hbar = 1.05 \times 10^{-34} \text{ m}^2 \cdot \text{kg/s}$ and 1 amu =  $1.66 \times 10^{-27}$  kg, we can calculate the recoil frequency  $\omega_r$  and the recoil velocity  $v_r$ ,

$$\omega_r = \frac{\hbar k_L^2}{2m} = 1.98 \times 10^4 \,\mathrm{s}^{-1} \,, \tag{6.1}$$

$$v_r = \frac{\hbar k_L}{m} = 0.0054 \text{ m/s},$$
 (6.2)

where we have used  $k_L = 2\pi/\lambda_L$ . Note that the recoil frequency is expressed in radians per second above and converts to 3.15 kHz in hertz. The natural unit we have used for acceleration is  $\omega_r v_r$ , and from Eqs. (6.1-6.2), we get

$$1 \omega_r v_r = 106 \text{ m/s}^2 \approx 10 g,$$
 (6.3)

where g is the gravitational acceleration on the surface of Earth. Time has been presented in  $\omega_r^{-1}$ , where 1  $\omega_r^{-1} = 50 \ \mu$ s, and the unit for spatial dimension,  $k_L^{-1}$ , converts to 1  $k_L^{-1} = 136 \ \text{nm}$ .

In Chapter. 5, the sensitivity for the acceleration with  $10^4$  measurements in the shaken lattice interferometer reached  $10^{-5}\omega_r v_r = 1.1 \times 10^{-3} \text{ m/s}^2 \approx 10^{-4} \text{ g}$ . The total time for one run of the interferometry sequence was  $T_{\text{tot}} \approx 27 \ \omega_r^{-1} = 1.36 \text{ ms}$ , during which the wave function traveled across a distance of about  $L = 160 \ k_L^{-1} = 0.0217 \text{ mm}$ .

## 6.2 Noise sources in the experiment

#### 6.2.1 Laser linewidth

A laser with non-zero linewidth has phase diffusion that can be modeled as white noise. This means the phase undergoes a random walk in time, and the phase noise can be described by a Wiener process,

$$\phi_{t+dt} - \phi_t \sim \mathcal{N}(0, \Delta \, dt) \,, \tag{6.4}$$

where  $\mathcal{N}(\mu, \sigma^2)$  describes a Gaussian distribution with mean  $\mu$  and standard deviation  $\sigma$ ,  $\Delta$  is the laser linewidth, and dt is the time step used in the simulation. If the two counter-propagating lasers have common phase noises, which is the case when the counter-propagating light is generated from

a common source, e.g., by a single retro-reflected laser, then the relative phase between them, i.e.,  $\phi$  in Eq. (5.1), will not suffer from the noise from the laser linewidth. However, if the laser noises are independent, which is the case when the lattice is formed by two separate light sources, then we need to take the phase noise into account. The standard deviation of the noise in the relative phase will be  $\sqrt{2}$  times the standard deviation in a single laser.

To incorporate the laser noise in our simulations, we assume that

$$\phi(t) = \phi_C(t) + \phi_r(t), \qquad (6.5)$$

where  $\phi_C$  is the controlled phase, and  $\phi_r$  is the phase noise that fluctuates as described in Eq. (6.4). The controlled phase functions are the ones the reinforcement learning agent learns in a noiseless environment, as shown in Chapter 5. We simulate the output momentum distributions from the interferometer with  $\Delta = 2 \times 10^{-4} \omega_r$ , which corresponds to a linewidth of ~ 16Hz, and use a Bayesian approach to reconstruct the acceleration. The value of  $\Delta$  is chosen to give the same accumulated phase during the interferometer sequence that would be generated by the smallest detectable acceleration in the simulation with monochromatic lasers. Fig. (6.1)(a) shows the reconstruction process for the acceleration in this case. It may be observed that even though the distribution becomes narrower as the number of measurements increases, the center of the peak is shifted from the true value, meaning that even though the precision converges, the estimation is biased, compromising accuracy. This can also be concluded from Fig. (6.1)(b) and (c).

Fig. (6.1)(b) shows the standard deviation of the acceleration, which goes down as the standard quantum limit. Fig. (6.1)(c) shows the mean squared error of the estimated acceleration. The value where  $\sigma_a$  saturates corresponds to the bias of the estimation.

The non-zero laser linewidth also results in a spread in the wave vector  $k_L$ . However, since  $\Delta \nu / \nu \sim \Delta k_L / k_L \sim 3 \text{ MHz}/300 \text{ THz} \sim 10^{-8}$ , it is expected that the fluctuation in the wave vector will not be a main source of error.



Figure 6.1: (a) Bayesian reconstruction of the true acceleration in the presence of a phase noise that corresponds to a linewidth  $\Delta = 2 \times 10^{-4} \omega_r$ . The red line indicates the true acceleration. (b) Standard deviation of the reconstructed acceleration. (c) Mean squared error of the acceleration.

### 6.2.2 Laser intensity fluctuation

To model the laser intensity noise, we assume that the lattice depth fluctuates as

$$V_0(t) = \overline{V_0} \, \left[ 1 + W(t) \right] \,, \tag{6.6}$$

where  $\overline{V_0}$  represents the average lattice depth, and W(t) is a white noise that fluctuates as

$$W_{t+dt} - W_t \sim \mathcal{N}(0, \xi \, dt) \,. \tag{6.7}$$

We simulate the shaken lattice interferometer with laser intensity fluctuation rate  $\xi = 2 \times 10^{-4} \omega_r$ , and reconstruct the acceleration signal using a Bayesian approach. This rate can be interpreted as an empirical upper bound, and can be expected to be much larger than what would be typical in experiment. Fig. (6.2)(a) shows the reconstruction process for the first 100 measurements (atoms). Fig. (6.2)(b) shows the standard deviation, which still scales as the standard quantum limit at large N despite the large level of noise at small number of measurements. However, in Fig. (6.2)(c), we can see that the level of the mean squared error saturates at a relatively large value. The explanation is that during the first 100 measurements, there the distribution of the acceleration is multi-peaked due to shot noise. After we make more measurements, the distribution becomes single-peaked, which is a result of the maximum likelihood principle, and therefore, the standard deviation starts to decrease. However, the peak may not coincide with the true value due to the presence of the intensity noise, and therefore, we observe the high level of mean squared error.

In addition to the white noise that is uniform for all frequencies, we also simulate the intensity noise with a certain pattern in the power spectral density, which can be captured on a photodetector connected to a spectrum analyzer. The resulting time series is shown in Fig. 6.3(a). We only use a segment of the time series for each acceleration value, and for the next value we use the next segment, and so on. When we reach the end of the series, we go back to the beginning. This ensures a certain amount of randomness. The momentum distribution of the interferometer is shown in Fig. 6.3(a). The maximum difference between the noisy distribution and the noiseless distribution is only about 1%.



Figure 6.2: (a) Reconstruction of the true acceleration in the presence of intensity noise  $\xi = 2 \times 10^{-4} \omega_r$ . (b) Standard deviation of the acceleration. (c) Mean squared error of the acceleration.



Figure 6.3: (a) Reconstruction of the acceleration from a system with laser intensity noise as shown in (c). (b) The standard deviation of the acceleration as calculated from the reconstruction. (c) Time series of the noise. The noise is defined as the ratio between the intensity fluctuation and the mean intensity.

### 6.2.3 Quantum efficiency

An important property of the photodetector is that it may not detect all the atoms. This effect is known as the quantum efficiency and is considered to be an imaging noise. Suppose the detector has a quantum efficiency of  $\mathcal{E}$ , then each atom has a probability of  $1 - \mathcal{E}$  to be not counted. We simulate this effect by ignoring each sample drawn from the momentum distribution with probability  $1 - \mathcal{E}$ . The result of this is summarized in Fig. (6.4).

### 6.2.4 Dark counts

Dark counts come from the background noise on the CCD image when there is actually no atom. We simulate this effect by adding uniform random noise on top of the original samples. The noise level is defined by the ratio between the number of random samples and the number of atoms there actually is. The result of this is summarized in Fig. (6.5).

# 6.3 Dynamic range

To study the limit for the dynamic range, we simulate the interferometry outcome for a wide range of acceleration. We observe that within a range of acceleration close to 1 g, the momentum distribution for each acceleration value is still unique, and therefore acts as a fingerprint that allows us to make unbiased estimations. The results are shown in Fig. 6.6.

We then estimate the acceleration from the interferometry outcome simulated using the true acceleration. The estimation is calculated by taking the mean from the reconstructed probability distribution of the acceleration. We do this for all the acceleration values in the range presented in Fig. 6.6, and look at the relation between the true acceleration and the estimated acceleration. The fact that the curve lies on x = y means that the estimation is unbiased in this range.



Figure 6.4: (a) Reconstruction of acceleration with  $\mathcal{E} = 0.9$ . (b) Standard deviation. We assume that there are originally 10000 atoms, but since  $\mathcal{E} \neq 1$ , we only get 9035 counts on the detector. The inset shows a zoom-in of N = 8000 to N = 9035.



Figure 6.5: Error and standard deviation in the presence of dark counts. The absolute value of the error grows larger as the noise level increases, whereas the standard deviation remains almost the same.



Figure 6.6: Momentum distribution for a wide range of acceleration. The free-propagation time is set to 10  $\omega_r^{-1}.$ 



Figure 6.7: (a) A plot of the estimated acceleration versus the true acceleration. Each acceleration is reconstructed from 10000 samples. The linear relation verifies that the estimations are unbiased. (b) A scatter plot of the error. It can be seen that the error is uniformly distributed across the whole range of true acceleration. (c) A histogram of the error, where the width of each bin is  $4 \times 10^{-6} \ (\omega_r v_r)$ .

## 6.4 The ultimate limit for sensitivity

Now let us extend the free propagation time to  $T = 1000 \ \omega_r^{-1}$ . This corresponds to a free propagation distance of ~ 1 mm, which fits into a typical optical lattice experiment in a glass cell. We found that at this limit, the sensitivity reaches  $10^{-9} \ \omega_r v_r = 10^{-8} \ g$ , and if we repeat the interferometry 100 times, which takes about 20 min due to the finite loading time for the BEC, we could push the sensitivity to  $10^{-9} \ g$ . At this level, there are already some interesting applications for space navigation, such as detecting the sudden change in acceleration caused by the solar radiation pressure on a spacecraft when it emerges from a region of shadow.

### 6.5 Outlook

Since the noises could impose a bias to the parameter estimation as shown in our analysis, it is crucial to resolve this problem, and machine learning could potentially mitigate this adverse effect. One direction is to employ a supervised learning algorithm that learns the relation between the momentum distribution and the acceleration from the noisy data. The data are labelled by the true acceleration, which is considered a known parameter under a well-controlled experimental setup, and after the supervised-learning model is trained, we could test its efficacy with distributions that come from an unknown acceleration. The other direction would be to develop a reinforcement learning agent that is capable of learning in a noisy environment and finding the protocols that are the most robust against the noises. In this case, we could add noise robustness to the reward function, so that the reward is large for small variations of the distribution with respect to laser intensity or frequency. The learning cycle would then result in the agent optimizing with respect to this quantity.



(a)









Figure 6.8: (a) Final momentum distribution for an interferometer where the free propagation time is 1000  $\omega_r^{-1}$ . (b) A zoom-in of (a) near a = 0. (c) The error and the standard deviation of the estimation of the acceleration from plot (b) with 10000 samples. (d) Histogram of the estimation error.

# Chapter 7

# How to train your gyro: Reinforcement learning for rotation sensing with a shaken optical Lattice

# 7.1 Introduction

If one asks an expert for advice about how to measure tiny displacements or forces over large distances, they will likely point to the very large body of literature on interferometry, perhaps highlighting LIGO [1] and other such remarkable interferometry systems. Optical interferometers are not the only ones useful in high precision measurements, and indeed matter-wave interferometers have also offered the potential for exquisite sensitivity to inertial forces. In the work by Weidner and Anderson, a new kind of matter-wave interferometric accelerometer that utilizes an optical lattice was proposed and demonstrated experimentally [102, 100, 101]. The purpose of the optical lattice is to provide robustness of the system in the face of a harsh dynamical environment typical of real-world applications. It has been demonstrated that quantum design of optical-lattice-based interferometers can benefit greatly from machine learning approaches, since the machine can explore an extensive number of possible control protocols and come up with one that is not intuitive but very effective. The previous chapter established that one can 'teach' the system to become sensitive to acceleration by learning how to modulate the phase of the lattice with unconventional patterns. that is, how to shake the lattice such that the system carries out an interferometric protocol. While this work demonstrated the power of machine learning to construct an interferometer to surpass conventional Bragg interferometry, the machine learning was constrained, let us say, by an expert's understanding of how interferometry is to take place. Indeed, nearly every kind of interferometry, optical or matter wave, takes place as a sequence of wavefront splitting, reflection, and recombination, with wave propagation in-between, and this is known as the Mach-Zehnder configuration [54, 107].

If we free the learning agent from the constraints set by conventional configurations, it could potentially learn beyond what human has explored. For example, Ref. [57] has shown that an active-learning agent can predict novel configurations for quantum optics experiments that generate highly-entangled states. This chapter presents results from the application of machine learning to train a two-dimensional optical lattice to sense rotation, i.e., to form a gyroscope, as depicted in Figure 7.1. In contrast to the accelerometer, in this case we utilize reinforcement learning in a manner that removes the prejudice of an expert about how interferometry is supposed to be done. Instead, learning is driven by a focus on performance metrics — that is, what it means to be a good gyroscope rather than what it means to be a proper interferometer. We show that the solution looks nothing like standard interferometry and that the system, in fact, exhibits a sensitivity to rotation that is multiple factors better than a standard interferometer. The fact that there is a numerical improvement *per se* is perhaps less important than the fact that the lattice interferometer exhibits the complexity enabled by quantum mechanics; one associated with a large Hilbert space dimensionality. The learned solution outperforms the expert's preconceived notions about how one might use that complexity to advantage. Moreover, the machine learning can cater to a set of performance metrics, such as a combination of sensitivity, dynamic range, and tolerance to noise and experimental drifts. The learning algorithm navigates a trade-space that may well be opaque to the expert designer.

# 7.2 Physical model

The system we consider is one where the atoms are loaded into a two-dimensional optical lattice. The lattices in both directions can be 'shaken' by varying the phases of the corresponding interfering laser beams. The entire system, shown in Figure 7.1, is considered in a rotating non-



Figure 7.1: A two-dimensional optical lattice potential in a rotating frame, where atoms are trapped and their dynamics can be controlled by the shaking of the lattice. The blue arrows indicate possible trajectories associated with the splitting of the matterwaves.

inertial frame, and therefore described by the Hamiltonian

$$\hat{H} = \frac{\hat{p}^2}{2m} - \boldsymbol{\Omega} \cdot \hat{\boldsymbol{L}} - \frac{V_x}{2} \cos(2k_L \hat{x} + \phi_x) - \frac{V_y}{2} \cos(2k_L \hat{y} + \phi_y), \qquad (7.1)$$

where m is the mass of a single atom,  $k_L$  is the laser wavenumber,  $\phi_x$  and  $\phi_y$  are the phase differences between the two counter-propagating lasers in the x and y directions respectively, and  $V_x$  and  $V_y$ are the corresponding lattice potential strengths. The term  $\boldsymbol{\Omega} \cdot \hat{\boldsymbol{L}}$  describes the rotational kinetic energy of the system, with  $\hat{\boldsymbol{L}} = \hat{\boldsymbol{x}} \times \hat{\boldsymbol{p}}$  the angular momentum, and  $\boldsymbol{\Omega}$  being the angular velocity. It is the magnitude,  $\Omega$ , that we aim to measure with high precision.

# 7.3 Two-path Sagnac matter-wave interferometer

To measure the angular velocity, we first consider a two-path loop configuration [32, 50] that resembles the construction of the fiber-optic gyroscope [93], which directly utilizes the Sagnac effect [77, 6]. The Sagnac effect describes the physical process that generates a phase difference between the two waves that propagate in the opposite directions in a loop when in a non-inertial rotating frame of reference. In general, this predicts the property that the phase difference,  $\Delta\phi$ , is proportional to both the angular velocity and to the area that the two paths enclose,  $\boldsymbol{A}$ . In the case of matter-waves, this relationship can be derived to be  $\Delta\phi = 4m\boldsymbol{\Omega}\cdot\boldsymbol{A}/\hbar$ .

A two-path Sagnac matter-wave interferometer can be constructed in the lattice using the beam-splitting and reflecting protocols presented in [17], based on components optimized by reinforcement learning. In that case, the splitting protocol was demonstrated to transfer the ground state to a superposition of the  $|\pm 4\hbar k_L\rangle$  states, and the reflecting protocol to map any linear combination of the  $|\pm 4\hbar k_L\rangle$  states to the corresponding combination of the  $|\mp 4\hbar k_L\rangle$  states.

In order to build a gyroscope out of these components, the sequence of operations is the following. Initially, the wavefunction is prepared in the ground state of the 2D lattice. In the x-direction, the 1D beam-splitting protocol is applied, allowing free propagation for a duration, denoted as T, and then the reflecting protocol is applied. Following this, free propagation occurs for a duration 2T, and then the reflecting protocol is applied, followed by free propagation for

a duration T, and finally the recombining protocol. In the y-direction, the lattice operates as a conveyor belt for the trapped atoms. The lattice is smoothly accelerated to the velocity  $4\hbar k_L/m$ according to the adiabaticity criterion, and then translated at constant velocity. During the halfway point of the sequence, when the two-paths in the x-directions cross, the velocity of the y-lattice is symmetrically decelerated through zero to  $-4\hbar k_L/m$ , so that the lattice can be translated backward at the constant velocity from that point on. The final step is to adiabatically accelerate the y-lattice back to zero velocity for the final recombination protocol.

To ensure that the acceleration of the y-lattice is indeed adiabatic, let us investigate the adiabaticity condition for the wavefunction to stay in the ground state of an accelerating lattice. The Hamiltonian in the y direction in the lab frame is

$$\hat{H}_y = \frac{\hat{p}_y}{2m} - \frac{V_y}{2} \cos\left(2k_L \hat{y} + \phi(t)\right).$$
(7.2)

With the unitary transform by the operator  $\hat{U} = \exp(-i\hat{p}_y\phi/2\hbar k_L)$ , we move into the frame co-moving with the lattice. The Hamiltonian in the co-moving frame is

$$\hat{\tilde{H}}_{y} = \frac{\hat{p}_{y}}{2m} - \frac{V_{y}}{2} \cos\left(2k_{L}\hat{y} + \phi(t)\right) + \frac{\dot{\phi}}{2k_{L}}\hat{p}_{y}.$$
(7.3)

The instantaneous eigenstates are  $\left|\tilde{\psi}_{N,q}\right\rangle = \exp\left(im\hat{y}\dot{\phi}/2\hbar k_L\right)\exp\left(-i\hat{p}\phi/2\hbar k_L\right)\left|\psi_{N,q}\right\rangle$ , where  $\left|\psi_{N,q}\right\rangle$  are the Bloch eigenstates with band index N and quasimomentum q. If we wish the wave-function stays in the ground state (N=0), then the quantum adiabatic theorem requires that [19]

$$\sum_{m \neq 0} \left| \frac{\hbar \langle m | \tilde{H} | 0 \rangle}{(E_m - E_n)^2} \right| = \sum_{m \neq 0} \left| \frac{\hbar \langle m | \hat{p} \, \ddot{\phi} / 2k_L | 0 \rangle}{(E_m - E_n)^2} \right| \ll 1.$$
(7.4)

For  $V_y = 10 E_r$ , numerical calculations show that the adiabaticity condition is  $|\ddot{\phi}| \ll 60 \omega_r^2$ . In our simulations, we use  $|\ddot{\phi}| = 6.4\omega_r^2$ , which satisfies the adiabaticity condition.

If higher sensitivity is desired, one could trivially extend this sequence to allow the matter waves to travel around in multiple loops before applying the last recombination step, although in practice, one needs to balance this against the resulting reduced fringe visibility. A schematic to illustrate the paths that the wave functions traverse is shown in Fig. 7.2. The point here is to



Figure 7.2: The clockwise path (blue) and the counter-clockwise path (red) taken by the wave functions in the two-path shaken lattice gyroscope. The vertical back line is where the wave functions were initially split and are eventually recombined.

demonstrate that a gyroscope can indeed be constructed in a two-dimensional lattice, and to set a baseline for sensitivity for the purpose of making comparison with the gyroscope designed by end-to-end reinforcement learning that will follow later.

The sensitivity of a gyroscope constructed this way is presented in terms of the classical Fisher information in Fig. 7.3. The classical Fisher information is a quantitative measure of the sensitivity of the output distribution on the metrological parameter. Note that the value of the Fisher information presented here is normalized to unity for an ideal two-path gyroscope with the same momentum splitting and free-propagation duration as the one we construct, i.e.,  $(4p_0T^2)^2$ , where  $p_0 = 4\hbar k_L$ . The shaken lattice gyroscope gives a Fisher information of order unity as can be anticipated, and the slight deviation is due to imperfections in the components and the fact that the shaken lattice is does not fit perfectly within the two-path framework.

## 7.4 Effective 1D model

One challenge of simulating the two-dimensional rotating lattice is that, in general, the computational complexity scales quadratically with the number of momentum states. To enable efficient simulations, which is especially important when running many machine-learning epochs, we reduce the system to a linear complexity by introducing an approximation of a deep potential in the y-direction that reduces the system to an effective one-dimensional model. Explicitly, we take advantage of the fact that when the deep lattice in the y-direction is adiabatically translated, the shape of the wavepacket in the y-direction transported along with the lattice remains approximately invariant. This approximation is encapsulated by a separable ansatz, where we assume the two-dimensional wavefunction is a simple product state as represented by the tensor product of the wavefunctions in the x and y directions, i.e.,



Figure 7.3: The Fisher information for the machine-learned shaken-lattice gyroscope normalized with respect to the Fisher information that results from the gyroscope in which the matterwaves follow the ideal paths of the conventional two-path inteferometer.

$$\begin{aligned} |\psi\rangle &\approx |\psi_x\rangle \otimes |\psi_y\rangle ,\\ \hat{p}_y &\approx \langle p_y\rangle ,\\ \hat{y} &\approx \langle y\rangle . \end{aligned}$$
(7.5)

With this ansatz, the time-dependent solution and the interference that is formed is determined by the x-direction evolution, modeled from,

$$\widetilde{H}_x = \frac{\widehat{p}_x}{2m} - \frac{V_0}{2}\cos(2k_L\hat{x} + \phi_x) - \Omega\widetilde{L}_z, \qquad (7.6)$$

where  $\tilde{L}_z = (\langle p_y \rangle \hat{x} - \langle y \rangle \hat{p}_x)$ . The mean position and momentum in the y-direction for the sequence described in Sec. 7.3 is shown in Fig. (7.4). The oscillations in the momentum plot are due to the atoms being trapped in the bottom of the lattice sites, which are close to harmonic traps, and the accleration phase being not perfectly adiabatic.

Let us further discuss the validity of the separation ansatz. We can write the 2D wavefunction in the form of its Schmidt decomposition,

$$\psi(x,y) = \sum_{i} \sqrt{p^{(i)}} \ \psi_x^{(i)}(x) \ \psi_y^{(i)}(y) \tag{7.7}$$

This decomposition can be obtained numerically from singular value decomposition. If the highest singular value  $p^{(0)}$  is close to identity, we can drop the other terms in the summation, and approximate the 2D wavefunction as the separable form, i.e.,

$$\psi(x,y) \approx \psi_x^{(0)}(x) \,\psi_y^{(0)}(y)$$
(7.8)

To verify that in the case where the y-wavepacket is adiabatically transported along the lattice, the approximation shown above is valid, we compute the Schmidt decomposition of the 2D wavefunction evolved from the full 2D Hamiltonian. The highest singular value is 0.95. We compare the population distribution in x from the full 2D wavefunction,  $|\psi_{2D}(x)|^2 = \int |\psi(x,y)|^2 dy$ , and that from the Schmidt decomposition,  $|\psi_x^{(0)}(x)|^2$ , and as shown in Fig. 7.5, the two populations agree well.



Figure 7.4: (a) Average position of the wavepacket in the y-direction. (b) Average momentum of the wavepacket in the y-direction. The lattice depth  $V_y$  is set to 10  $\omega_r$ .



Figure 7.5: Comparison between the simulations results of the probability distributions in the x direction from the effective 1D model (blue), the full 2D model (orange), and the highest singular value component from the 2D model (green).

The validity of the approximation allows us to use the separable ansatz, i.e.,

$$\psi_{\rm sep}(x,y) = \psi_x(x)\psi_y(y) \tag{7.9}$$

and evolve the wavefunctions in the x and y directions separately using the effective 1D Hamiltonians. From now on, we will work under the effective 1D assumption, where only the dynamics in x-direction is investigated.

# 7.5 End-to-end design with reinforcement learning

We have shown how to build a two-path Sagnac gyroscope from individual interferometer components. In this section, we consider applying reinforcement learning to the design of the gyroscope to go beyond the traditional Mach-Zehnder configuration. The details of the principles of reinforcement learning and deep-Q learning can be found in [17] and will not be repeated here. Our purpose is to demonstrate how to formulate an end-to-end design problem for a matter-wave gyroscope within the scope of the reinforcement learning framework.

A basic reinforcement learning framework consists of a agent that select actions based on the observed state of the environment, and an environment that interacts with the action and generates a reward based on the observed quality of the current state. In our design framework, the state (not the quantum state, but the state of the environment within the context of machine learning) contains information about the population in the discrete momentum basis in x and the average position and momentum in y. The actions are defined as a finite set of discrete choices for the time derivative of the phase, i.e.,  $\dot{\phi}$ . This choice is motivated by the experimental implementation for the control of the lattice phase, where an acoustic-optical modulator is employed and the frequency of the modulation is proportional to  $\dot{\phi}$ . Let us now elaborate more on the states and actions we have chosen. The state given to the reinforcement learning agent consists of information about the populations in  $\{|2n\hbar k_L\rangle | \pm 3, \pm 2, \pm 1, 0\}$  in the x-direction and the position and momentum in the y-direction. It is represented by the vector,

$$s = \left\{ P_{-3}, P_{-2}, P_{-1}, P_{-0}, P_{-1}, P_{-2}, P_{-3}, \frac{\langle \hat{y} \rangle}{40}, \frac{\langle \hat{p}_y \rangle}{4} \right\},$$
(7.10)

where  $P_n$  represent the population in the  $|2n\hbar k_L\rangle$  state. The action is the time derivative of the laser phase in the x-direction selected from a discrete set of values,

$$\mathcal{A} \in \{-\dot{\phi} = -8, \ \dot{\phi} = 0, \ \dot{\phi} = 8\}.$$
(7.11)

In order to compute the reward, we construct the environment as three copies of the quantum system with different rotation rates,  $\Omega = [\Omega_0 - d\Omega, \Omega_0, \Omega_0 + d\Omega]$ , for small  $d\Omega$  and with each evolving according to Eq. (7.6). This choice is motivated by the requirement to compute the derivative with respect to  $\Omega$  by numerical symmetric finite differencing. Explicitly, the reward is defined as the classical Fisher information computed with respect to  $\Omega$  using the following formula,

$$I_1(\Omega) = \sum_p \frac{1}{\Pr(p|\Omega)} \left[ \frac{\partial \Pr(p|\Omega)}{\partial \Omega} \right]^2.$$
(7.12)

The goal of learning is to find a sequence of  $\dot{\phi}$  that optimizes the Fisher information at the terminal time.

We point out that the Fisher information is related to the theoretical bound (Cramér-Rao bound) for the standard deviation of the measurement for  $\Omega$  as follows,

$$\sigma_{\Omega} \ge 1/\sqrt{NI_1(\Omega)},\tag{7.13}$$

so by optimizing the Fisher information with reinforcement learning, we are effectively optimizing the measurement sensitivity with respect to the generated quantum state at the unitary output, but based only on a specific choice of the angular frequency,  $\Omega_0$ . A variety of extensions to what we present here are clearly possible, based on alternative design conditions.

## 7.6 Results for end-to-end design

Within the context of the design framework described in Sec. 7.5, we are able to obtain lattice shaking protocols that outperform the conventional two-path gyroscope demonstrated in Sec. 7.3 in terms of sensitivity. The protocol varies for every instantiation due to the randomized initialization of the agent, and Fig. 7.6 shows one of the machine-learned protocols. The time-evolution of the resulting machine-learned multi-path interferometer is shown in Fig. 7.6. The pattern is reminiscent of the speckle patterns that emerge from a multi-mode fiber (fiber specklegram), which are known to be sensitive interference detectors of inertial phase. On the other hand, in the case of a multi-mode fiber, that patterns are scrambled by temperature changes or strain on the fiber. For the multimode matter-wave interferometer we consider here, the situation is quite different, and in spite of the irregular pattern it generates, we show that this device can, in principle, measure rotational signals in an unbiased fashion and reach the sensitivity predicted by the theoretical bound.

The momentum distributions that are produced by the multi-path interferometer under conditions of different values of the rotation rate are shown in Fig. 7.7. The larger the variation in the momentum distribution is as the rotation rate is varied, the more sensitive is the interferometer. The Fisher information calculated from the distribution is shown in Fig. 7.7. The value on the y-axis is the ratio between the Fisher information from the multi-mode interferometer constructed by reinforcement learning and the one shown in in Sec. 7.3. We optimize the Fisher information around  $\Omega_0 = 0$ , and one can see that the Fisher information is maximum at  $\Omega = 0$ , with the dynamic range being sacrificed for the high sensitivity. The machine-learned gyroscope achieved a Fisher information of around 25 times higher than the two-path gyroscope, meaning that the sensitivity is improved by an extremely impressive factor of 5. If we make a comparison to a gyroscope constructed from Bragg diffraction, the improvement is  $4 \times 5 = 20 - \text{fold.}$ 

To extract the rotational signal from the distribution shown in Fig. 7.7, we apply Bayesian reconstruction. This means that we iteratively update the prior distribution for  $\Omega$  from each sample. The reconstruction of the rotational signal is shown in Fig. 7.8(a), and the estimation for  $\Omega$  is unbiased. In Fig. 7.8(b), we plot the standard deviation of  $\Omega$  with the number of samples taken, and see that the sensitivity scales inversely with the square root of the number of atoms, as expected from independent measurements.



Figure 7.6: Phase and the spatial wavefunction as a function of time. The shaking protocol is generated from the reinforcement learning agent with  $\dot{\phi}$  as the actions. There are 125 steps in this protocol. The wavefunction is initialized in the ground state with a momentum width  $\sigma_p = 0.1$ , and evolve with the shaking protocol shown above.



Figure 7.7: Population distribution in the momentum basis and the resulting classical Fisher information for the machine-learned shaken-lattice gyroscope as a function of the rotation rate  $\Omega$ . The value of the Fisher information is normalized with the Fisher information of a conventional two-path interferometer with splitting  $\Delta p = 8\hbar k_L$ .



Figure 7.8: (a) A reconstruction of the rotation rate from 1 sample to 100 samples. (b) The standard deviation of the rotation rate as a function of the atom number. The result at large atom number is consistent with the Cramér-Rao bound, which scales as  $1/\sqrt{N}$ .

# 7.7 Conclusion and outlook

In this chapter, we have proposed and analyzed the efficacy of the use of deep reinforcement learning for the optimization of a quantum sensor. We have demonstrated the explicit application of this concept through an example of the design of a shaken-lattice matter-wave gyroscope. We first showed that a Mach-Zehnder-type matter-wave gyroscope can be constructed in a 2D lattice from a conveyor belt in the *y*-direction and the previously developed beam-splitting and the reflecting shaking functions in the *x*-direction. We then applied reinforcement learning in an end-to-end fashion to find a shaking protocol that improves the sensitivity of the gyroscope by a factor of 5 compared to the conventional Mach-Zehnder interferometry. These impressive results open up the possibility for applying the same principal machinery to a variety of quantum sensing problems with more complex landscapes, such as multi-parameter estimation and entanglement-enhanced metrology. However, the implications of our results go beyond this framework, and the same ideas of reinforcement learning may be applied to a variety of quantum and classical systems where design of complex circuits is needed.

# Chapter 8

### Summary and outlook

In this thesis, we applied the basic concepts of quantum metrology and machine learning, presented in Chapter 2 and 3 respectively, to the design of quantum systems for extreme sensing. We gave three examples of quantum design in this thesis corresponding to three levels of design philosophy. First, we showed that the periodic modulation of the scattering length in a Bose-Einstein condensate at the resonant frequency can generate squeezed states in momentum, and proposed a protocol to generate Ramsey fringes to probe the coherence in this system. Next, we demonstrated the application of reinforcement learning to the design shaking functions that generate the essential components of Mach-Zehnder atom interferometry in a one-dimensional lattice. We were able to construct an interferometer with a larger enclosed space-time area than the conventional Bragg interferometer, and therefore achieved a higher sensitivity to acceleration. We further investigated how well this interferometer would work in a hypothetical experimental setting. Finally, we moved forward and used reinforcement learning in an end-to-end manner, with the aim to design a gyroscope in a two-dimensional shaken lattice that is not constrained by our prior knowledge about how interferometers should work. We discovered the shaking function that led to a device that performs better than the conventional Sagnac gyroscope in terms of sensitivity to rotation.

The reinforcement learning framework for quantum design is general, and has inspired proposals for many potential future research works, including:

• Using the formalism established in Chapter 4 to simulate an environment that can be controlled by a reinforcement learning agent, with the goal of exploring methods to generate squeezed states in a Bose-Einstein condensate that can potentially be used for entanglementenhanced atom interferometry.

- Extending the design of a unitary operator demonstrated in Chapter 5 to the design of a general trace-preserving quantum channel in the context of an open quantum system, which could be useful for tackling noise in a quantum processor.
- Designing the optimal encoding and decoding protocols for a multi-qubit programmable quantum sensor using natively available gates on platforms such as ion or neutral atom arrays, for the purpose of achieving entanglement-enhanced metrology.
- Finding the protocol to prepare the optimal quantum state for a multi-parameter estimation scheme, in particular for a set of non-commuting observables.
- Improving the momentum capture range in laser cooling and particle slowing, which leads to a larger number of particles that can be used for performing a metrology task, and therefore increasing the sensitivity.

Even though the central focus of this thesis has been on metrology, the design philosophy is general, and can now be applied to a wide variety of problems relevant to engineering quantum systems to perform desired tasks. These include algorithmic operations that are relevant for quantum computation, quantum communication applications, quantum simulation ideas and concepts, and quantum control, and many of these contain quantum entanglement and many-body correlations that will go beyond what we have presented here.

# Bibliography

- B. P. Abbott et al. Observation of gravitational waves from a binary black hole merger. <u>Phys.</u> Rev. Lett., 116:061102, Feb 2016.
- [2] Z. An and D. L. Zhou. Deep reinforcement learning for quantum gate control. <u>EPL</u> (Europhysics Letters), 126(6):60002, jul 2019.
- [3] G. B. Arfken, H. J. Weber, and F. E. Harris. Chapter 14 bessel functions. In G. B. Arfken, H. J. Weber, and F. E. Harris, editors, <u>Mathematical Methods for Physicists (Seventh Edition</u>), pages 643 713. Academic Press, Boston, seventh edition edition, 2013.
- [4] M. Arratia. On the jets emitted by driven bose–einstein condensates. Journal of Physics B: Atomic, Molecular and Optical Physics, 52(5):055301, feb 2019.
- [5] B. Barrett, A. Bertoldi, and P. Bouyer. Inertial quantum sensors using light and matter. Physica Scripta, 91(5):053006, apr 2016.
- [6] B. Barrett, R. Geiger, I. Dutta, M. Meunier, B. Canuel, A. Gauguet, P. Bouyer, and A. Landragin. The sagnac effect: 20 years of development in matter-wave interferometry. <u>Comptes Rendus Physique</u>, 15(10):875–883, 2014. The Sagnac effect: 100 years later / L'effet Sagnac : 100 ans après.
- [7] R. Bellman. On the theory of dynamic programming. <u>Proceedings of the National Academy</u> of Sciences, 38(8):716–719, 1952.
- [8] T. Bergeman, M. G. Moore, and M. Olshanii. Atom-atom scattering under cylindrical harmonic confinement: Numerical and analytic studies of the confinement induced resonance. Phys. Rev. Lett., 91:163201, Oct 2003.
- [9] J.-P. Blaizot and G. Pipka. Quantum theory of finite systems. MIT Press, 1986.
- [10] N. N. Bogolyubov. On the theory of superfluidity. J. Phys. (USSR), 11:23–32, 1947.
- [11] C. Bordé. Atomic interferometry with internal state labelling. <u>Physics Letters A</u>, 140(1):10 - 12, 1989.
- [12] T. Bothwell, C. J. Kennedy, A. Aeppli, D. Kedar, J. M. Robinson, E. Oelker, A. Staron, and J. Ye. Resolving the gravitational redshift across a millimetre-scale atomic sample. <u>Nature</u>, 602(7897):420–424, Feb 2022.
- [13] M. Bukov. Reinforcement learning for autonomous preparation of floquet-engineered states: Inverting the quantum kapitza oscillator. Phys. Rev. B, 98:224305, Dec 2018.
- [14] M. Bukov, A. G. R. Day, D. Sels, P. Weinberg, A. Polkovnikov, and P. Mehta. Reinforcement learning in different phases of quantum control. Phys. Rev. X, 8:031086, Sep 2018.
- [15] J. Carrasquilla. Machine learning for quantum matter. <u>Advances in Physics: X</u>, 5(1):1797528, 2020.
- [16] T. Chen and B. Yan. Characterization of stimulated excitations in a driven bose-einstein condensate. Phys. Rev. A, 98:063615, Dec 2018.
- [17] L.-Y. Chih and M. Holland. Reinforcement-learning-based matter-wave interferometer in a shaken optical lattice. Phys. Rev. Research, 3:033279, Sep 2021.
- [18] L. W. Clark, A. Gaj, L. Feng, and C. Chin. Collective emission of matter-wave jets from driven bose-einstein condensates. Nature, 551:356 EP –, Nov 2017.
- [19] D. Comparat. General conditions for quantum adiabatic evolution. <u>Phys. Rev. A</u>, 80:012106, Jul 2009.
- [20] G. M. D'Ariano and M. G. A. Paris. Arbitrary precision in multipath interferometry. <u>Phys.</u> Rev. A, 55:2267–2271, Mar 1997.
- [21] E. A. Donley, N. R. Claussen, S. L. Cornish, J. L. Roberts, E. A. Cornell, and C. E. Wieman. Dynamics of collapsing and exploding bose-einstein condensates. <u>Nature</u>, 412(6844):295–299, 2001.
- [22] L. Feng, J. Hu, L. W. Clark, and C. Chin. Correlations in high-harmonic generation of matter-wave jets revealed by pattern recognition. Science, 363(6426):521–524, 2019.
- [23] T. Fösel, P. Tighineanu, T. Weiss, and F. Marquardt. Reinforcement learning with neural networks for quantum feedback. Phys. Rev. X, 8:031084, Sep 2018.
- [24] S. Fray, C. A. Diez, T. W. Hänsch, and M. Weitz. Atomic interferometer with amplitude gratings of light and its applications to atom based tests of the equivalence principle. <u>Phys.</u> Rev. Lett., 93:240404, Dec 2004.
- [25] H. Fu, L. Feng, B. M. Anderson, L. W. Clark, J. Hu, J. W. Andrade, C. Chin, and K. Levin. Density waves and jet emission asymmetry in bose fireworks. <u>Phys. Rev. Lett.</u>, 121:243001, Dec 2018.
- [26] M. Gebbe, J.-N. Siemß, M. Gersemann, H. Müntinga, S. Herrmann, C. Lämmerzahl, H. Ahlers, N. Gaaloul, C. Schubert, K. Hammerer, S. Abend, and E. M. Rasel. Twin-lattice atom interferometry. Nature Communications, 12(1):2544, 2021.
- [27] D. M. Giltner, R. W. McGowan, and S. A. Lee. Theoretical and experimental study of the bragg scattering of atoms from a standing light wave. Phys. Rev. A, 52:3966–3972, Nov 1995.
- [28] S. Giorgini. Collisionless dynamics of dilute bose gases: Role of quantum and thermal fluctuations. Phys. Rev. A, 61:063615, May 2000.

- [29] V. Giovannetti, S. Lloyd, and L. Maccone. Advances in quantum metrology. <u>Nature</u> Photonics, 5(4):222–229, Apr 2011.
- [30] E. Greplova, C. K. Andersen, and K. Mølmer. Quantum parameter estimation with a neural network, 2017.
- [31] A. Griffin. Conserving and gapless approximations for an inhomogeneous bose gas at finite temperatures. Phys. Rev. B, 53:9341–9347, Apr 1996.
- [32] T. L. Gustavson, P. Bouyer, and M. A. Kasevich. Precision rotation measurements with an atom interferometer gyroscope. Phys. Rev. Lett., 78:2046–2049, Mar 1997.
- [33] M. N. Holger B. Nielsen. Dirac sea for bosons. Aug 1998.
- [34] M. Holland, J. Williams, and J. Cooper. Bose-einstein condensation: Kinetic evolution obtained from simulated trajectories. Phys. Rev. A, 55:3670–3677, May 1997.
- [35] M. J. Holland and K. Burnett. Interferometric detection of optical phase shifts at the heisenberg limit. Phys. Rev. Lett., 71:1355–1358, Aug 1993.
- [36] J. Hu, L. Feng, Z. Zhang, and C. Chin. Quantum simulation of unruh radiation. <u>Nature</u> Physics, 2019.
- [37] N. M. Hugenholtz and D. Pines. Ground-state energy and excitation spectrum of a system of interacting bosons. Phys. Rev., 116:489–506, Nov 1959.
- [38] C. Jekeli. Navigation error analysis of atom interferometer inertial sensor. <u>NAVIGATION</u>, 52(1):1–14, 2005.
- [39] C. Jekeli. Precision free-inertial navigation with gravity compensation by an onboard gradiometer. Journal of Guidance, Control, and Dynamics, 29(3):704–713, 2006.
- [40] M. Kasevich and S. Chu. Atomic interferometry using stimulated raman transitions. <u>Phys.</u> Rev. Lett., 67:181–184, Jul 1991.
- [41] M. Khanahmadi and K. Mølmer. Time-dependent atomic magnetometry with a recurrent neural network, 2020.
- [42] D. Kingma and J. Ba. Adam: A method for stochastic optimization. <u>International Conference</u> on Learning Representations, 12 2015.
- [43] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. Optimization by simulated annealing. <u>Science</u>, 220(4598):671–680, 1983.
- [44] M. Kitagawa and M. Ueda. Squeezed spin states. Phys. Rev. A, 47:5138–5143, Jun 1993.
- [45] S. J. J. M. F. Kokkelmans and M. J. Holland. Ramsey fringes in a bose-einstein condensate between atoms and molecules. Phys. Rev. Lett., 89:180401, Oct 2002.
- [46] S. J. J. M. F. Kokkelmans, J. N. Milstein, M. L. Chiofalo, R. Walser, and M. J. Holland. Resonance superfluidity: Renormalization of resonance scattering theory. <u>Phys. Rev. A</u>, 65:053617, May 2002.

- [47] T. Kovachy, P. Asenbaum, C. Overstreet, C. A. Donnelly, S. M. Dickerson, A. Sugarbaker, J. M. Hogan, and M. A. Kasevich. Quantum superposition at the half-metre scale. <u>Nature</u>, 528(7583):530–533, Dec 2015.
- [48] S. Kuang and S. Cong. Lyapunov control methods of closed quantum systems. <u>Automatica</u>, 44(1):98 – 108, 2008.
- [49] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. nature, 521(7553):436, 2015.
- [50] A. Lenef, T. D. Hammond, E. T. Smith, M. S. Chapman, R. A. Rubenstein, and D. E. Pritchard. Rotation sensing with an atom interferometer. <u>Phys. Rev. Lett.</u>, 78:760–763, Feb 1997.
- [51] E. H. Lieb and W. Liniger. Exact analysis of an interacting bose gas. i. the general solution and the ground state. Phys. Rev., 130:1605–1616, May 1963.
- [52] L.-J. Lin. Self-improving reactive agents based on reinforcement learning, planning and teaching. Machine Learning, 8(3):293–321, May 1992.
- [53] O. J. Luiten, M. W. Reynolds, and J. T. M. Walraven. Kinetic theory of the evaporative cooling of a trapped gas. Phys. Rev. A, 53:381–389, Jan 1996.
- [54] L. Mach. Ueber einen interferenzrefraktor, 1892.
- [55] J. Mackeprang, D. B. R. Dasari, and J. Wrachtrup. A reinforcement learning approach for quantum state engineering. Quantum Machine Intelligence, 2(1):5, May 2020.
- [56] P. J. Martin, B. G. Oldaker, A. H. Miklich, and D. E. Pritchard. Bragg scattering of atoms from a standing light wave. Phys. Rev. Lett., 60:515–518, Feb 1988.
- [57] A. A. Melnikov, H. Poulsen Nautrup, M. Krenn, V. Dunjko, M. Tiersch, A. Zeilinger, and H. J. Briegel. Active learning machine learns to create new quantum experiments. <u>Proceedings</u> of the National Academy of Sciences, 115(6):1221–1226, 2018.
- [58] A. A. Melnikov, P. Sekatski, and N. Sangouard. Setting up experimental bell tests with reinforcement learning. Phys. Rev. Lett., 125:160401, Oct 2020.
- [59] J. N. Milstein, C. Menotti, and M. J. Holland. Feshbach resonances and collapsing boseeinstein condensates. New Journal of Physics, 5:52–52, may 2003.
- [60] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller. Playing atari with deep reinforcement learning, 2013.
- [61] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, et al. Human-level control through deep reinforcement learning. Nature, 518(7540):529, 2015.
- [62] H. P. Nautrup, N. Delfosse, V. Dunjko, H. J. Briegel, and N. Friis. Optimizing Quantum Error Correction Codes with Reinforcement Learning. Quantum, 3:215, Dec. 2019.
- [63] J. A. Nelder and R. Mead. A Simplex Method for Function Minimization. <u>The Computer</u> Journal, 7(4):308–313, 01 1965.

- [64] J. Niederle, V. Popov, and L. Hlavatỳ. <u>Functional Integrals in Quantum Field Theory and Statistical Physics</u>. Mathematical Physics and Applied Mathematics. Springer Netherlands, 1983.
- [65] M. Y. Niu, S. Boixo, V. N. Smelyanskiy, and H. Neven. Universal quantum control through deep reinforcement learning. npj Quantum Information, 5(1):33, Apr 2019.
- [66] S. P. Nolan, A. Smerzi, and L. Pezzè. A machine learning approach to bayesian parameter estimation, 2020.
- [67] E. Oelker, R. B. Hutson, C. J. Kennedy, L. Sonderhouse, T. Bothwell, A. Goban, D. Kedar, C. Sanner, J. M. Robinson, G. E. Marti, D. G. Matei, T. Legero, M. Giunta, R. Holzwarth, F. Riehle, U. Sterr, and J. Ye. Demonstration of 4.8×10<sup>(-17)</sup> stability at 1 s for two independent optical clocks. Nature Photonics, 13(10):714–719, Oct 2019.
- [68] M. Olshanii. Atomic scattering in the presence of an external confinement and a gas of impenetrable bosons. Phys. Rev. Lett., 81:938–941, Aug 1998.
- [69] Z. Pagel, W. Zhong, R. H. Parker, C. T. Olund, N. Y. Yao, and H. Müller. Symmetric bloch oscillations of matter waves. Phys. Rev. A, 102:053312, Nov 2020.
- [70] R. H. Parker, C. Yu, W. Zhong, B. Estey, and H. Müller. Measurement of the fine-structure constant as a test of the standard model. Science, 360(6385):191–195, 2018.
- [71] L. H. Pedersen, N. M. Møller, and K. Mølmer. Fidelity of quantum operations. <u>Physics</u> Letters A, 367(1):47 – 51, 2007.
- [72] A. P. Peirce, M. A. Dahleh, and H. Rabitz. Optimal control of quantum-mechanical systems: Existence, numerical approximation, and applications. Phys. Rev. A, 37:4950–4964, Jun 1988.
- [73] M. E. Peskin and D. V. Schroeder. <u>An Introduction to quantum field theory</u>. Addison-Wesley, Reading, USA, 1995.
- [74] A. Peters, K. Y. Chung, and S. Chu. Measurement of gravitational acceleration by dropping atoms. Nature, 400(6747):849–852, Aug 1999.
- [75] D. S. Petrov, M. Holzmann, and G. V. Shlyapnikov. Bose-einstein condensation in quasi-2d trapped gases. Phys. Rev. Lett., 84:2551–2555, Mar 2000.
- [76] H. F. Rice, V. Benischek, and L. Sczaniecki. Application of atom interferometric technology for gps independent navigation and time solutions. In <u>2018 IEEE/ION Position, Location</u> and Navigation Symposium (PLANS), pages 1097–1106, 2018.
- [77] F. Riehle, T. Kisters, A. Witte, J. Helmcke, and C. J. Bordé. Optical ramsey spectroscopy in a rotating frame: Sagnac effect in a matter-wave interferometer. <u>Phys. Rev. Lett.</u>, 67:177–180, Jul 1991.
- [78] G. Rummery and M. Niranjan. On-line q-learning using connectionist systems. <u>Technical</u> Report CUED/F-INFENG/TR 166, 11 1994.
- [79] P. A. Ruprecht, M. J. Holland, K. Burnett, and M. Edwards. Time-dependent solution of the nonlinear schrödinger equation for bose-condensed trapped neutral atoms. <u>Phys. Rev. A</u>, 51:4704–4711, Jun 1995.

- [80] L. Salvi, N. Poli, V. Vuletić, and G. M. Tino. Squeezing on momentum states for atom interferometry. Phys. Rev. Lett., 120:033601, Jan 2018.
- [81] J. F. Schaff, T. Langen, and J. Schmiedmayer. Interferometry with atoms. 2015.
- [82] C. Schubert, S. Abend, M. Gersemann, M. Gebbe, D. Schlippert, P. Berg, and E. M. Rasel. Multi-loop atomic sagnac interferometry. Scientific Reports, 11(1):16121, Aug 2021.
- [83] J. Schuff, L. J. Fiderer, and D. Braun. Improving the dynamics of quantum sensors with reinforcement learning. New Journal of Physics, 22(3):035001, mar 2020.
- [84] A. Shankar, L. Salvi, M. L. Chiofalo, N. Poli, and M. J. Holland. Squeezed state metrology with bragg interferometers operating in a cavity. <u>Quantum Science and Technology</u>, 4(4):045010, oct 2019.
- [85] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. van den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner, I. Sutskever, T. Lillicrap, M. Leach, K. Kavukcuoglu, T. Graepel, and D. Hassabis. Mastering the game of go with deep neural networks and tree search. <u>Nature</u>, 529(7587):484–489, Jan 2016.
- [86] D. Silver, T. Hubert, J. Schrittwieser, I. Antonoglou, M. Lai, A. Guez, M. Lanctot, L. Sifre, D. Kumaran, T. Graepel, T. Lillicrap, K. Simonyan, and D. Hassabis. A general reinforcement learning algorithm that masters chess, shogi, and go through self-play. <u>Science</u>, 362(6419):1140–1144, 2018.
- [87] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton, et al. Mastering the game of go without human knowledge. Nature, 550(7676):354, 2017.
- [88] K. Staliunas, S. Longhi, and G. J. de Valcárcel. Faraday patterns in bose-einstein condensates. Phys. Rev. Lett., 89:210406, Nov 2002.
- [89] R. Sutton, D. A. McAllester, S. Singh, and Y. Mansour. Policy gradient methods for reinforcement learning with function approximation. In NIPS, 1999.
- [90] R. S. Sutton and A. G. Barto. Reinforcement learning: An introduction. MIT press, 2018.
- [91] J. Taylor. Scattering theory: the quantum theory on nonrelativistic collisions. Wiley, 1972.
- [92] G. Tesauro. Temporal difference learning and td-gammon. <u>Commun. ACM</u>, 38(3):58–68, Mar. 1995.
- [93] V. Vali and R. W. Shorthill. Fiber ring interferometer. <u>Appl. Opt.</u>, 15(5):1099–1100, May 1976.
- [94] H. Van Hasselt, A. Guez, and D. Silver. Deep reinforcement learning with double q-learning. In Thirtieth AAAI conference on artificial intelligence, 2016.
- [95] J. Wallnöfer, A. A. Melnikov, W. Dür, and H. J. Briegel. Machine learning for long-distance quantum communication. PRX Quantum, 1:010301, Sep 2020.
- [96] D. F. Walls. Squeezed states of light. Nature, 306(5939):141–146, 1983.

- [97] R. Walser. Ground state correlations in a trapped quasi one-dimensional bose gas. <u>Optics</u> Communications, 243(1):107 – 129, 2004. Ultra Cold Atoms and Degenerate Quantum Gases.
- [98] Z. T. Wang, Y. Ashida, and M. Ueda. Deep reinforcement learning control of quantum cartpoles. Phys. Rev. Lett., 125:100401, Sep 2020.
- [99] C. J. C. H. Watkins and P. Dayan. Q-learning. In Machine Learning, pages 279–292, 1992.
- [100] C. A. Weidner and D. Z. Anderson. Experimental demonstration of shaken-lattice interferometry. Phys. Rev. Lett., 120:263201, Jun 2018.
- [101] C. A. Weidner and D. Z. Anderson. Simplified landscapes for optimization of shaken lattice interferometry. New Journal of Physics, 20(7):075007, jul 2018.
- [102] C. A. Weidner, H. Yu, R. Kosloff, and D. Z. Anderson. Atom interferometry using a shaken optical lattice. Phys. Rev. A, 95:043624, Apr 2017.
- [103] R. J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. Mach. Learn., 8(3–4):229–256, May 1992.
- [104] D. J. Wineland, J. J. Bollinger, W. M. Itano, F. L. Moore, and D. J. Heinzen. Spin squeezing and reduced quantum noise in spectroscopy. Phys. Rev. A, 46:R6797–R6800, Dec 1992.
- [105] L.-A. Wu, M. Xiao, and H. J. Kimble. Squeezed states of light from an optical parametric oscillator. J. Opt. Soc. Am. B, 4(10):1465–1475, Oct 1987.
- [106] H. Xu, J. Li, L. Liu, Y. Wang, H. Yuan, and X. Wang. Generalizable control for quantum parameter estimation through reinforcement learning. <u>npj Quantum Information</u>, 5(1):82, Oct 2019.
- [107] L. Zehnder. Ein neuer interferenzrefraktor, 1891.
- [108] X.-M. Zhang, Z. Wei, R. Asad, X.-C. Yang, and X. Wang. When does reinforcement learning stand out in quantum control? a comparative study on state preparation. <u>npj Quantum</u> Information, 5(1):85, Oct 2019.
- [109] H. Z. Zhigang Wu. Dynamics and density correlations in matter wave jet emission of a driven condensate. Jun 2019.

# Appendix A

#### Numerical methods for the Bose-Einstein condensate

#### A.1 Imaginary-time propagation

The imaginary-time propagation method is a numerical algorithm for finding the ground state of the Schrödinger equation. It is particularly useful when the Schrödinger equation is nonlinear and exact diagonization of the Hamiltonian is difficult or not possible. This method involves replacing the time t with the imaginary time  $-i\tau$  in the time-dependent Schrödinger equation, i.e.,

$$-\frac{\partial\Psi}{\partial\tau} = H\Psi.$$
(A.1)

The solution to the equation above (we have set  $\hbar$  to 1 here) is  $\Psi(\tau) = \sum_i c_i e^{-E_i \tau} \psi_i$ , where  $E_i$ and  $\psi_i$  are the eigenvalues and the eigenvectors respectively. At large  $\tau$ , the state with the lowest energy, i.e., the ground state, dominates the solution.

In practice, we start with an initial guess of the wavefunction,  $\Psi(0) = \sum_i c_i \psi_i$ , and we need to make sure that the ground state amplitude is non-zero, i.e., the guessed wavefunction is not orthogonal to the ground state. We use numerical integration methods such as second order Runge-Kutta (RK2) to advance in time, and the wavefunction needs to be renormalized every step to avoid the loss of numerical precision. The total time it takes to reach the ground state is related to the inverse of the energy gap between the ground state and the first excited state.

This method is analogous to the gradient descent method introduced in Chapter 3, with the cost function being the energy.

# A.2 Spatial derivative in 1D

To calculate the spatial derivative in the kinetic term in the one dimensional Hamiltonian,  $-\frac{\hbar^2}{2m}\frac{\partial^2}{\partial x^2}$ , we use the finite differencing method, i.e.,

$$\frac{\partial^2 \psi}{\partial x^2} \approx \frac{\psi(x - \Delta x) + \psi(x + \Delta x) - 2\psi(x)}{(\Delta x)^2} \,. \tag{A.2}$$

The numerical error of this approach (central difference approximation) is  $O(\Delta x^2)$ . With the infinite square well potential,

$$V(x) = \begin{cases} \infty & \text{if } x \le 0 \text{ or } x \ge 1 \\ 0 & \text{if } 0 < x < 1 \end{cases}$$
, (A.3)

we can impose the boundary condition,  $\psi(0) = 0$  and  $\psi(1) = 0$ , and the wavefunction amplitudes at these points stay zero throughout the simulation.

#### A.3 Spatial derivative in 2D

In Chapter 4, we considered a system in a disk geometry, and assumed that this system possess cylindrical symmetry in two dimension, i.e.,  $\Psi(\mathbf{r}) = \psi(r)$ . In this case, the spatial derivative in the kinetic term in the Hamiltonian can be derived to be

$$\nabla^2 \Psi = \frac{\partial^2 \psi}{\partial r^2} + \frac{1}{r} \frac{\partial \psi}{\partial r} \,. \tag{A.4}$$

We may use the finite differencing method outlined in the previous section to calculate the second derivative. The central difference approximation for the first derivative is

$$\frac{\partial \psi}{\partial r} \approx \frac{\psi(r + \Delta r) - \psi(r - \Delta r)}{2\Delta r} \,. \tag{A.5}$$

We may also impose the boundary condition  $\psi(r_{\text{max}}) = 0$  at some cutoff radius. It is obvious that at r = 0, Eq. (A.4) is ill-behaved with the factor 1/r, but physically one can expect that the kinetic energy is still finite and remains a continuous function in space. Therefore, we need to find the appropriate boundary condition to mitigate the issue. The boundary condition at the center arises from the symmetric assumption,  $\psi(r) = \psi(-r)$ . This means that the first derivative is zero, and that we can use the L'Hôpital's rule for the second term in Eq. (A.4), which gives

$$\lim_{r \to 0} \frac{1}{r} \frac{\partial \psi}{\partial r} = \frac{\partial^2 \psi}{\partial r^2}.$$
 (A.6)

With this boundary condition, the derivative in the kinetic term at the grid point r = 0 can be calculated as

$$\nabla^2 \Psi = 2 \frac{\partial^2 \psi}{\partial r^2} \approx 2 \frac{2 \psi(\Delta r) - 2 \psi(0)}{(\Delta r)^2} = 4 \frac{\psi(\Delta r) - \psi(0)}{(\Delta r)^2} \,. \tag{A.7}$$

# Appendix B

#### Simulation techniques for the shaken lattice

### B.1 Bloch's theorem

Bloch's theorem states that the wavefunction for each eigenstate in a spatially periodic system with periodicity a can be written as

$$\psi(x) = e^{iq \cdot x} u_q(x) \,, \tag{B.1}$$

where q is the so-called quasi-momentum, and u(x) is a periodic function, i.e., u(x) = u(x+a).

### B.2 Momentum basis

In the momentum basis, the Hamiltonian for the atoms in the optical lattice (defined in Chapter 5 can be written as

$$\hat{H} = \sum_{p} \frac{p^2}{2m} \left| p \right\rangle \left\langle p \right| - \frac{V_0}{4} \left( e^{i\phi} \left| p + 2\hbar k_L \right\rangle \left\langle p \right| + e^{-i\phi} \left| p \right\rangle \left\langle p + 2\hbar k_L \right| \right).$$
(B.2)

Note that the Hamiltonian only couples the momentum states with  $2\hbar k_L$  spacing. As a result, the manifolds with different quasi-momentum q can be simulated independently. This is reminiscent of the Bloch's theorem B.1 in the momentum representation. Assume that the initial population distribution is peaked around the momenta  $p \in \{2n\hbar k_L | n \in \mathbb{Z}\}$ , and that the width of the peak is smaller than  $2\hbar k_L$ , as can be expected from a wavefunction that spans across many lattice sites, then a convenient representation to use is  $|n, q\rangle$ , which stands for the momentum state  $|p = 2n\hbar k_L + \hbar q\rangle$ . The narrow peak assumption implies that  $\langle n', q' | n, q \rangle = \delta_{nn'} \delta(q - q')$ .

The trick mentioned above reduces the computation complexity from  $O(P^2)$  to  $O(N^2Q)$ , where P represents the length of the momentum grid across the entire range that we are interested in simulating,  $[2n_{\min} + q_{\min}, 2n_{\max} + q_{\max}]$  (factors of  $\hbar$  and  $k_L$  are dropped), N is the length of the grid for the momentum comb,  $N = n_{\max} - n_{\min} + 1$ , and Q is the length of the quasi-momentum grid.

#### B.3 Accelerating reference frame

The Hamiltonian for a one-dimensional system in a reference frame accelerating at rate a in the positive direction can be written as

$$\hat{H} = \hat{H}_0 + ma\hat{x} \,, \tag{B.3}$$

where  $\hat{H}_0$  is the original Hamiltonian in the rest frame. The second term acts as a fictitious potential on an atom with mass m.

If the quantum state and  $\hat{H}_0$  is represented in the momentum basis, as is the case in the previous section, then we need to deal with the additional term with the position operator using the split-step method. The approach we take is the symmetric Trotter decomposition

$$e^{(A+B)\Delta t} \approx e^{A\Delta t/2} e^{B\Delta t} e^{A\Delta t/2} \,. \tag{B.4}$$

where the error is in fractions of  $([A, [A, B] + 2[B, [A, B])\Delta t^3/24$ . We associate A and B with the operators  $-i\hat{H}_0/\hbar$  and  $-ima\hat{x}/\hbar$  respectively. The evolution  $e^{-i\hat{H}_0\Delta t/\hbar}$  can be simulated using the RK2 method with the formulation we set up in the previous section. The evolution that involves the position operator can be simulated as a shift operator, i.e.,

$$e^{-ima\hat{x}\Delta t/\hbar} \left| p \right\rangle = \left| p - ma\Delta t \right\rangle \,. \tag{B.5}$$

Since this shift operator acts the same on all  $|n,q\rangle$  states, we can simply shift the whole quasimomentum grid,  $q \rightarrow q - ma\Delta t/\hbar$ . This trick avoids the need to perform fast Fourier transforms when switching between the momentum and position evolutions, which are computationally expensive. ProQuest Number: 29258937

INFORMATION TO ALL USERS The quality and completeness of this reproduction is dependent on the quality and completeness of the copy made available to ProQuest.



Distributed by ProQuest LLC (2022). Copyright of the Dissertation is held by the Author unless otherwise noted.

This work may be used in accordance with the terms of the Creative Commons license or other rights statement, as indicated in the copyright statement or in the metadata associated with this work. Unless otherwise specified in the copyright statement or the metadata, all rights are reserved by the copyright holder.

> This work is protected against unauthorized copying under Title 17, United States Code and other applicable copyright laws.

Microform Edition where available © ProQuest LLC. No reproduction or digitization of the Microform Edition is authorized without permission of ProQuest LLC.

ProQuest LLC 789 East Eisenhower Parkway P.O. Box 1346 Ann Arbor, MI 48106 - 1346 USA