

MACHINE LEARNING FOR QUANTUM AND COMPLEX SYSTEMS

Aaron Tranter

A thesis presented for the degree of
Doctor of Philosophy of The Australian National University



**Australian
National
University**

Department of Quantum Science
July 2020

© Copyright by Aaron Tranter 2020
All Rights Reserved

DECLARATION

This thesis is an account of research undertaken between February 2016 and July 2020 at The Department of Quantum Science, Research School of Physics, Faculty of Science, The Australian National University, Canberra, Australia. Except where acknowledged in the customary manner, the material presented in this thesis is, to the best of my knowledge, original and has not been submitted in whole or part for a degree in any university.

Aaron Tranter

2020

ACKNOWLEDGEMENTS

Without the support of many people over the course of my PhD, this thesis would not be of the same quality. Firstly to Ben, thank you for your time, effort and expert advice (both in physics and life). Your level of engagement across so many topics and tasks is truly impressive to witness. To Ping Koy, your tireless work in keeping our group organised and interested is key to the success enjoyed by myself and others. Thank you for the many valuable discussions and insights.

To Geoff, thank you for helping to bring my optics experience from minimal to acceptable. Your problem solving prowess continues to impress me. You have my utmost respect. I look forward to hurting myself outdoors in new and exciting ways.

To the keepers of the sacred Tim Tams, Amanda, Lynne and Siobhan, your support was crucial to my sanity, and for that I thank you. Your dedication to keeping our group functional is unparalleled.

During my PhD I was fortunate to work with many talented people in the research group. In (rough) order of appearance: Jesse, your knowledge in physics never ceases to amaze me. Anthony, thank you for being excellent lab partner, I always enjoyed our lab sessions. Oliver, you made me feel at home in the group initially, thank you. Jinyong, thanks for many insightful conversations. Giovanni (and by extension Simone), I can always count on you to take me somewhere pretty, although we may take a detour or two. Young-Wook, thanks for your help in the lab and keeping me company in Seoul. Jing-Yan, thanks for many interesting discussions and helping me get acquainted with Canberra. Mark, always polite and friendly, thanks for the interesting chats! Tobias, my motorbike bro, see you in Alps soon! Pierre, excellent wine and cheese advice and a great travel companion, what's not to love? Thibault, thanks for many glorious trips, I could always count on you. Sophie, ai-ya, need I say more? Karun, your lovable naivety always makes me laugh. Ruvi, the Baileys loving bandit, thanks for the laughs. Daniel, your opinions on specific details that I never would have thought of is great, all the best. Sarah, the single loudest 310cc QLD biker known to man, much respect. Harry, you were a pleasure to work with, Skynet will always live on. Jiri, thanks for a wide variety of conversational topics! Jiayi, I've never seen you not smile which makes you an absolute joy to be around. Assad, thanks for your expert advice (in climbing and physics), I can always count on you. Hao, you have a level of dedication I wish I could emulate. Kabilan, the most open minded of us all, you may just change the world yet. Maria, always willing to give anything a go, all the best. Biveen, my kindred-spirit in tech, stay classy my friend. Ida, quiet yet tough as nails, your determination and

tenacity continues to impress me. Lorcan, you're a top lad, together we may only fear the sun. Naomi, I couldn't ask for a better partner to muddle through code with, my best wishes for your future. Ozlem, we will get better at balancing don't worry. Ciaran, always around for a tea or a beer. Spyros, a late addition, but a welcome one to be sure. See you in Greece soon for more philosophical mind blows.

Thank you to the workshop crew: Niel, Paul and Tom. Your expert advice and fantastic skills help make the experiments what they are.

To Mum, Dad and Tara, thank you for your support throughout the years. I could not have ended up here without your tireless effort and sacrifice. Words cannot express my gratitude.

To the members of the Department of Quantum Science, past and present. Thank you for the many morning coffees, the great conversation and the advice.

めぐみ...あなたの支えと愛はいつも僕を励ましてくれているね。
あなたの助けなしではここまで辿り着くことは不可能であったし、そもそもあなたの存在がなかったらまずチャレンジしようなどとは思わなかっただろう。
またこんなに頑張ることもなかったのだろうと思う。

いつも一緒にいてくれてありがとう。

あなたはこの世で一番大切な存在です。

ありがとう。

ABSTRACT

Machine learning now plays a pivotal role in our society, providing solutions to problems that were previously thought intractable. The meteoric rise of this technology can no doubt be attributed to the information age that we now live in. As data is continually amassed, more efficient and scalable methods are required to yield functional models and accurate inferences.

Simultaneously we have also seen quantum technology come to the forefront of research and next generation systems. These technologies promise secure information transfer, efficient computation and high precision sensing, at levels unattainable by their classical counterparts. Although these technologies are powerful, they are necessarily more complicated and difficult to control.

The combination of these two advances yields an opportunity for study, namely leveraging the power of machine learning to control and optimise quantum (and more generally complex) systems. The work presented in thesis explores these avenues of investigation and demonstrates the potential success of machine learning methods in the domain of quantum and complex systems.

One of the most crucial potential quantum technologies is the quantum memory. If we are to one day harness the true power of quantum key distribution for secure transmission of information, and more general quantum computing tasks, it will almost certainly involve the use of quantum memories.

We start by presenting the operation of the cold atom workhorse: the magneto-optical trap (MOT). To use a cold atomic ensemble as a quantum memory we are required to prepare the atoms using a specialised cooling sequence. During this we observe a stable, coherent optical emission exiting each end of the elongated ensemble. We characterise this behaviour and compare it to similar observations in previous work.

Following this, we use the ensemble to implement a backward Raman memory. Using this scheme we are able to demonstrate an increased efficiency over that of previous forward recall implementations. While we are limited by the optical depth of the system, we observe an efficiency more than double that of previous implementations.

The MOT provides an easily accessible test bed for the optimisation via some machine learning techniques. As we require an efficient search method, we implement a new type of algorithm based on deep learning. We design this technique such that the artificial neural networks are placed in control of the online optimisation, rather than simply being used as surrogate models.

We experimentally optimise the optical depth of the MOT using this method, by parametrising the time varying compression sequence. We identify a new and unintuitive method for cooling the atomic ensemble which surpasses current methods.

Following this initial implementation we make substantial improvements to the deep learning approach. This extends the approach to be applicable to a far wider range of complex problems, which may contain extensive local minima and structure. We benchmark this algorithm against many of the conventional optimisation techniques and demonstrate superior capability to optimise problems with high dimensionality. Finally we apply this technique to a series of preliminary problems, namely the tuning of a single electron transistor and second-order correlations from a quantum dot source.

CONTENTS

1	INTRODUCTION	1
I FUNDAMENTALS		
2	LIGHT AND MATTER WITH QUANTUM MECHANICS	7
2.1	What is quantum mechanics anyway?	7
2.2	The basic theory	7
2.2.1	Superposition	8
2.2.2	Entanglement	8
2.2.3	Operators	9
2.2.4	Density matrix formalism	9
2.3	Classical atom light interactions	10
2.3.1	Lorentz model	10
2.3.2	Absorption cross-section	11
2.3.3	Optical depth	12
2.4	Semi-classical picture	12
2.4.1	Two level atom	13
2.4.2	Optical Bloch equations	14
2.4.3	Useful results	15
2.4.4	Three level atom	16
3	LASER COOLING	19
3.1	Doppler cooling	19
3.1.1	Doppler shift	19
3.1.2	Spontaneous emission	19
3.1.3	1D optical molasses	20
3.2	Sub-Doppler cooling	21
3.2.1	Comparison to Doppler cooling	23
3.2.2	Sisyphus cooling	24
3.3	Magneto-optical traps	24
3.3.1	Magnetic fields	25
3.3.2	Optical fields	25
3.3.3	3D potentials and polarisation gradients	26
3.3.4	Density restrictions and loss	27
3.3.5	Time dependent MOTs	27
II ATOMS AND QUANTUM MEMORIES		
4	TRAPPING RUBIDIUM	31
4.1	General setup	31
4.1.1	Vacuum system	32

4.1.2	Laser cooling systems	32
4.1.3	Saturated absorption locking	34
4.1.4	Magnetic fields	36
4.1.5	Optical pumping	37
4.2	Experimental control	38
4.2.1	Generalised experimental control	38
4.2.2	Python <i>RunBuilder</i>	38
4.2.3	Atom loading	39
4.3	Characterisation and measurement	40
4.3.1	Optical depth	40
4.3.2	Imaging	40
4.3.3	Temperature	40
5	ATOMIC EMISSION	43
5.1	Literature	43
5.2	Experimental setup	44
5.3	Characterisation	45
5.3.1	Steady state time dependence	45
5.3.2	Compression time dependence	46
5.3.3	Polarisation	48
5.3.4	Feedback	50
5.3.5	Far field profile	52
5.3.6	Physical description	52
5.4	Conclusions and outlook	53
6	QUANTUM MEMORIES	55
6.1	Quantum computing and memories	55
6.2	Quantum memory theory	57
6.2.1	Collective operators	57
6.2.2	Extension to Λ -system	57
6.2.3	Raman transitions	60
6.3	Efficacy and loss	61
6.3.1	Efficiency	62
6.3.2	Lifetime	62
6.3.3	Bandwidth	62
6.3.4	Fidelity	63
6.3.5	Multi-mode capacity	63
6.3.6	Additional considerations	64
6.4	Quantum memory literature	65
6.4.1	Delay lines and cavities	65
6.4.2	Slow light and EIT	65
6.4.3	Atomic frequency comb	66
6.4.4	Raman memory	67
6.4.5	Controlled reversible inhomogeneous broadening	68

6.4.6	Summary	69
7	BACKWARD RAMAN MEMORY	71
7.1	From scattering to memory	71
7.2	Forward Raman memory	71
7.2.1	Memory protocol	72
7.2.2	Numerical simulation	74
7.3	Backward configuration	76
7.3.1	Phase matching	78
7.4	Experimental Implementation	78
7.4.1	Memory scheme	80
7.4.2	Atom Preparation	80
7.4.3	Optical Fields	80
7.4.4	Alignment and detection	82
7.4.5	Machine learned compression and PGC sequence	83
7.5	Experimental Results	84
7.5.1	Storage efficiency	84
7.5.2	Storage decay	85
7.5.3	Losses	87
7.6	Conclusions and outlook	88
III MACHINE LEARNING		
8	INTRODUCTION TO MACHINE LEARNING	91
8.1	What is machine learning anyway?	92
8.1.1	Types of machine learning	93
8.2	Deep learning	93
8.2.1	Artificial neural networks	94
8.2.2	Function approximation	95
8.2.3	Multilayer perceptrons	95
8.2.4	Loss functions	96
8.2.5	Training neural networks	97
8.2.6	Backpropagation	98
8.3	Hyperparameters	100
8.3.1	Layers and neurons	100
8.3.2	Regularisation	100
8.3.3	Activation functions	103
8.3.4	Initialisation	104
8.3.5	Training algorithms	105
8.3.6	Choosing hyperparameters	106
8.4	Extended networks	106
8.4.1	Convolutional neural networks	106
8.4.2	Recurrent neural networks	107
8.4.3	Actor-critic and Deep Q networks	107
8.4.4	Generative adversarial networks	108

9	DEEP LEARNING OPTIMISATION	109
9.1	Convex vs non-convex	109
9.2	Solving non-convex problems: a review	110
9.2.1	Global optimisation	110
9.2.2	Evolutionary algorithms	111
9.2.3	Stochastic methods	111
9.2.4	Metaheuristics	111
9.2.5	Neural net surrogates	112
9.2.6	Summary and issues	112
9.3	Deep learning approach	113
9.3.1	Multiple networks	114
9.3.2	Generating predictions	114
9.3.3	Pseudo-random learner	115
9.3.4	Initial sampling	115
9.3.5	In-loop optimisation	115
9.3.6	Function evaluations	116
9.3.7	Testing efficacy	116
10	OPTIMISING A MAGNETO-OPTICAL TRAP	119
10.1	The problem	119
10.1.1	Parametrisation	119
10.1.2	Online approach	120
10.2	Experimental implementation	121
10.2.1	(a) Cold atoms and measurement	122
10.2.2	(b) SANN training and minimisation	123
10.2.3	(c-d) Predictions and implementation	124
10.3	Experimental results	124
10.3.1	Human and machine ramps	125
10.3.2	Quantifying success	125
10.3.3	Time dependence	127
10.3.4	Temperature	127
10.3.5	Monitoring convergence	127
10.3.6	Solution efficacy and local minima	130
10.3.7	Physical mechanisms	132
10.4	Summary and future directions	133
11	OPTIMISATION UPGRADES	135
11.1	Previous issues	135
11.2	Method improvements	136
11.2.1	General cost functions	136
11.2.2	The sampling problem	137
11.2.3	Local minima methods	138
11.2.4	Exploration by overfitting	139
11.2.5	Local exploitation	139

11.2.6	New data generation	140
11.2.7	Improved performance	140
11.3	Optimisation benchmarks	141
11.4	Real world applications	144
11.4.1	Single electron transistor	144
11.4.2	Quantum dot frequency shaping	147
11.5	Further machine learning applications	148
11.5.1	Actor-critic networks	148
11.5.2	Learning transfer functions	149
11.6	Summary and outlook	150
12	CONCLUSION	151
IV APPENDIX		
A	ACTOR-CRITIC MODEL FOR AUTO ALIGNMENT	155
A.1	Reinforcement learning platform	155
A.2	General actor-critic	156
A.2.1	Temporal difference learning	156
A.2.2	Network losses	157
A.2.3	Network design	157
A.3	Physical problem	160
A.4	Algorithmic approaches	160
A.4.1	Data acquisition and reward	161
A.5	Preliminary results	162
A.5.1	Acuminata	162
A.5.2	Actor-Critic	162
A.6	Continuing work	164
B	FRUIT FLY OPTIMISATION	167
B.1	Raman memory simulation	168
C	LEARNING TRANSFER FUNCTIONS	169
C.1	Machine learning approach	170
C.1.1	Function parametrisation	170
C.2	Method testing	171
C.2.1	Cost function	172
C.2.2	Avoiding overfitting	174
C.3	Continuing work	174

LIST OF FIGURES

Figure 2.1	Three level Λ -scheme with two applied fields, a strong coupling field Ω and weak probe field \mathcal{E} , with frequencies ω_{es} and ω_{eg} respectively.	17
Figure 3.1	Level structure for a fictitious atom with two levels with $J=1$ and 2 respectively. The Clebsch-Gordan coefficients for each transition are indicated.	21
Figure 3.2	The resulting spatially varying polarisation as a result of counter propagating fields in (a) linear- \perp -linear configuration and (b) σ_+ - σ_- configuration. As can be seen the polarisation modulates between linear and circular in the linear- \perp -linear case whereas for circularly polarised fields the resulting polarisation is a linear polarisation, the axis of which rotates about the propagation axis.	22
Figure 3.3	Coil geometries for creating a linear magnetic field across the central region. (a) corresponds to Helmholtz configuration, whereas when one of the currents are reversed as in (b) anti-Helmholtz is achieved with 0 field in the central region.	25
Figure 3.4	Schematic showing the orientation of the cooling beams relative to the magneto-optical trap (MOT) coils. Trapping and repump co-propagate after being combined on a polarising beam splitter. Each pair of counter propagating beams are orthogonally polarised to establish the needed standing wave. The direction of gravity is also noted in the second orientation.	26
Figure 3.5	Timing and coordination of the various control channels during static loading followed by a temporal dark SPOT transient compression sequence.	28
Figure 4.1	Current experimental setup depicting a ^{87}Rb MOT illuminated by the cooling beams used for trapping. Also depicted is the glass cell attached to the vacuum chamber and the magnetic field coils.	32
Figure 4.2	Schematic of the ^{87}Rb MOT vacuum chamber and associated glass cell. The cell is coated on the external faces with anti-reflection coatings for 780 – 795nm. Continuous operation of the ion pump ensures the vacuum pressure remains on the order of $\sim 10^{-9}$ Torr.	33

Figure 4.3	D ₂ transition of ⁸⁷ Rb showing the detuned trapping and repump transitions used when loading the MOT. The hyperfine splitting arises due to coupling between the nuclear spin ($I=3/2$ for Rb) and the total angular momentum of the electron.	33
Figure 4.4	Cooling beam geometry relative to the glass cell. The trapping and repump beams co-propagate and are retro-reflected to create the standing waves in all 3-dimensions. The trapping region is defined by the intersection of the three beams.	34
Figure 4.5	Saturated absorption measurement of the D ₁ line of Rb. Both ⁸⁷ Rb and ⁸⁵ Rb are present. Each set of hyperfine transitions are highlighted with the transition frequencies relative to the leftmost transition for ⁸⁷ Rb and ⁸⁵ Rb respectively. Note that the ground state splittings are different for ⁸⁷ Rb and ⁸⁵ Rb.	35
Figure 4.6	Optical setup used to lock the diode laser onto the cooling transition. The saturated absorption measurement provides a locking signal which can be locked to a particular hyperfine transition. This setup is cloned with the exception of the lock offset block for locking a second diode laser onto the repump transition.	36
Figure 4.7	Geometry of the race track coils used to generate the elongated MOT. The dashed box in the centre represents the location of the vacuum cell with the storage axis for memory operations labelled as the z-axis. The field profiles are simulated using a finite difference time domain (FDTD) simulation in COMSOL. The field magnitude is given at a slice through the centre axis which possesses radial symmetry about the z-axis. The vector field arrows are logarithmically scaled to the magnitude of the field. For clarity the field from the capping coils is not shown which would add additional confinement along the z-axis.	37
Figure 4.8	Schematics of the control scheme used for experimental control. The modular setup allows for automated experimental control which is useful for optimisation purposes. Extensions can be included as add-ons to each system. Field programmable gate array (FPGA) code is contributed by Geoff Campbell.	39
Figure 4.9	Experimental schematic for transverse imaging of the MOT. An example of both a fluorescence emission image and absorption image are shown. Note that the intensity of the absorption image has been inverted.	41

Figure 4.10	Ballistic drop measurement of the ensemble temperature with acceleration due to gravity. The images are captured under absorption measurement on the $F = 1 \rightarrow F = 2$ transition of the D_2 line. The measured temperature of the ensemble above is on the order of $\sim 240 \mu\text{K}$	42
Figure 5.1	Experimental setup used when observing the atomic emission. Unlike the previous examples in the literature we do not require any additional beams other than those needed for laser cooling. The emission exits the atomic ensemble parallel to the axis with the highest optical depth. The inset shows a typical spatial pattern associated with the emitted mode.	45
Figure 5.2	Observed time dependence of the atomic emission for two different detunings. The dominant oscillation frequency is given in the inset of each plot. A larger detuning also corresponds to a lower steady state atomic density.	47
Figure 5.3	Measurement of the emission power showing the dependence on atom number during a transient compression sequence. The top row shows the evolution of the emission optical power over time. The bottom row shows the transient compression sequence parameters where the magnetic field is ramped from 6 Gs/cm to 28 Gs/cm and the repump frequency is detuned to induce a temporal dark SPOT. The highlighted area is expanded in the second column where the characteristic oscillation is once again observed. The compression sequence lasts for 20 ms however only the final 10ms is shown as the emission does not start until a critical density is reached for this lower trapping detuning of -32 MHz. The time dependence of the density for this sequence is depicted in Fig. 5.4.	48
Figure 5.4	MOT density and absorption as a function of compression time. Here the total absorption is not a strict measure of atom number but is $\propto N$. The right plot shows the evolution of the spatial distribution of the MOT as the transient compression sequence continues. Here, absorption images have been integrated to show the distribution of the ensemble along the x-axis (see Fig. 5.1). The shift towards the right of the distribution is a result of imbalanced trapping beam powers due to the retro-reflecting geometry. As expected the density increases leading to a narrow spatial distribution, which is radially symmetric around the z-axis, similar to a cigar.	49

Figure 5.5	Polarimeter setup for determining the Stokes parameters. Either α or β may be swept and the intensity measured. This does necessitate the ability to measure faster than physical changes in the system that may manifest separate polarisations.	49
Figure 5.6	Change in measured intensity as a function of polariser angle. From these measurements an approximate polarisation can be obtained by fitting the Stokes parameters. However the dynamic nature of the emission precludes an accurate measurement as evidenced by the large errors associated with each measurement.	50
Figure 5.7	Measured oscillation patterns observed with a photodetector upon retro reflection of the emission from one side of the MOT. It can be seen that the retro reflected mode helps to reinforce the oscillations seen previously. The trapping detuning used is -35MHz which can be directly compared to Fig. 5.2 without retro reflection. The respective periodogram is given for each trace in the second column, often multiple spectral components are present which are distinct from harmonics.	51
Figure 5.8	(a)-(d) Far field profiles of the emitted modes from the atomic ensemble. The intensity is measured relative to the highest emitted intensity. The mode structure shown in (a) corresponds to the high intensity pulsing from Fig. 5.7. The trap detuning for (a)-(d) is constant at -35MHz , with each image representing a different oscillatory behaviour. We have observed similar patterns are present for the full range of detunings at which this phenomena occurs.	52
Figure 6.1	Depictions of various three level schemes. For the present work we are most concerned with (a) the Λ -scheme, however there also exists the (b) Ξ and (c) V configurations.	58
Figure 6.2	Preparation of an inhomogeneously broadened ensemble into an atomic frequency comb with comb spacing Δ . Once the atomic ensemble has been prepared a signal field may be stored and additionally shelved on a metastable state to increase the storage lifetime.	67
Figure 6.3	Storage of an input optical pulse in an ensemble with a spatially dependent inhomogeneous broadening. After some storage time the sign of the broadening is reversed and the stored pulse is re-emitted. The relative frequency of the transition for a given spatial location is depicted by the colouring.	69

Figure 6.4	Summary of the maximum results for four parameters of interest for each of the main quantum memory schemes currently being pursued. Note that this is independent of platform and as such concurrent access to the maximum of each metric may not necessarily be possible. However this does serve as a useful guide to illustrate the successes of each scheme.	70
Figure 7.1	(a) Stimulated Raman scattering process with 2 optical fields applied to a 3-level system. The anti-Stokes photon has gained energy from the difference between the two ground state energies. The transition of an atom between the ground states can be used as a quantum memory protocol by considering the timings shown in (b). The S operator is the spin-wave coherence that is generated as a result of this storage operation.	72
Figure 7.2	$\chi(\omega)$ for different probe detunings and control configurations. For each subplot the blue curve represents the imaginary (absorptive) part of χ and the green dashed curve gives the real (dispersive) part. (a) shows the case where $\Omega_c = 0$ which returns the Lorentzian profile for a 2-level atom. (b) shows the two-photon resonance case where $\delta = 0$. Here the electromagnetically induced transparency (EIT) window is observed. (c) shows the emergence of the narrow Raman absorption line where the two-photon resonance $\delta = 10\Gamma$. The inset shows the narrow absorption line starting at 10.0γ . For all plots the parameters used are $\Gamma = \Omega_c = 5$ and $\gamma_{gs} = 0$	73
Figure 7.3	Numerical simulation of a Raman memory. Plots (a) and (b) show the amplitude of the electric field (probe) and spin-wave operators respectively. The atoms are assumed to extend uniformly over the normalised distance $-1 \rightarrow 1$. Intuitively it can be seen that the incident probe pulse is mapped onto the atomic spin-wave (write + storage) and converted back to a propagating field some time later (read). (c) shows the time trace that could be measured experimentally with a photo-detector. The recall efficiency is $\approx 26\%$	75
Figure 7.4	Numerical simulation of a Raman memory operating on two-photon resonance. The control field has been optimally placed and the Rabi frequency optimally tuned to produce a maximum efficiency of $\approx 84\%$	75

Figure 7.5	Numerical simulation of a Raman memory operating on two-photon resonance showing (a) the electric field operator amplitude, (b) the spin-wave operator magnitude and (c) the time trace of the input, recall and control fields. The control field has been optimally shaped as a series of two Gaussian pulses to give an efficiency of $\approx 91\%$. Note, the control fields have been normalised to the read pulse Rabi frequency.	77
Figure 7.6	Numerical simulation of a Raman memory operating on two-photon resonance using the simplified two level model. The forward retrieval time trace is shown in (a) with a efficiency of 75%. (c) - (e) shows the magnitude of the operators. The backward retrieval time trace in (b) yields an efficiency of 94% with (f) - (h) showing the magnitude of the operators. The simulation parameters used are optical depth (OD) = 1000, $\Delta_{\pm} = \pm 33\Gamma$, $\delta = 0$ and $[\Omega_{write}, \Omega_{read}] = [3.31, 11.54]$ and $[4.0, 10.7]$ for the forward and backward retrieval respectively. In both cases the Gaussian read/write pulses were optimised for Rabi frequency, write input time and pulse bandwidth, with the bandwidth being kept symmetric for simplicity. The respective bandwidths were Γ for the input pulse, 1.8Γ and 2.2Γ for the forward and backward retrieval respectively.	79
Figure 7.7	Phase matching conditions required to satisfy the co-linear storage and retrieval of the probe fields. The plotted vectors correspond to the wave vectors of each respective field. The difference in their wave vectors determine the wave vector of the atomic spin-wave.	80
Figure 7.8	Optical setup for generating the correct frequencies used in the experimental implementation of the backward Raman memory. Beams are fibre coupled to the MOT table where the main vacuum chamber assembly is located.	81
Figure 7.9	Simplified schematics of the optical path length and detection setup used for detecting the backward recalled probe. (a) shows the folded beam path used to elevate the optical fields to the MOT height. Phase matching and filtering requires that the total optical path length is 6m from when the probe and control are combined. (b) shows the detection scheme used during the quantum memory experiment. Note that for simplicity the angle is exaggerated and that normally the entire probe interaction region is illuminated by the control fields.	83

Figure 7.10	Relative timings of the experimental run for loading the atomic ensemble, compression, PGC, optical pumping and the experimental memory run. The dead time is included to allow the dissipation of eddy currents induced from fast switching of the magnetic field coils. Figure is adapted from [136].	84
Figure 7.11	Experimentally observed storage and backward recall of a 5 μ s width pulse using the backward Raman memory protocol. The efficiency of the storage operation shown here is $63 \pm 6\%$ for a storage time of 5 μ s. The time axis is zeroed to the centre of the forward control pulse as this is the mechanism that controls the timing of the memory operation. The backward control field is centred on 15 μ s.	85
Figure 7.12	High bandwidth storage operation yielding a time-bandwidth product of 160. The pulse stored has a $1/e^2$ -width of 360ns and is stored for one pulse width for a retrieval efficiency of 40%. The characteristic decay time is found to be $\tau = 60 \mu$ s. Figure is adapted from [136].	86
Figure 7.13	Experimentally observed efficiency decay as a function of storage time. Atomic motion and extraneous magnetic fields are responsible for the decay in efficiency as described by Eq. 7.4 which determines the fit shown in green. Errors in the efficiencies are determined from the standard deviation of the experimental acquisitions.	87
Figure 8.1	Comparison between the different model types during the training stage. Once each model is considered trained, the prediction stage would simply consist of passing the input for each model to receive an output.	94
Figure 8.2	(a) Approximation of a Gaussian function using two network structures with 1 layer (depth 1). The first structure consists of a single neuron and poorly approximates the function. With the addition of a second neuron we can already begin to approximate a Gaussian with high efficiency. The labelled parts of the neuron show the construction of the operations from input to output. The weights and biases for optimal overlap are given for both networks. Weights are depicted as connecting lines, whereas biases are shown in or above the circles. Each circle represents either neurons or input/outputs. (b) depicts a network with two layers that are fully connected. As the input propagates through the graph, from left to right, the input is constantly rescaled and offset with the weights and biases. Additionally at each neuron the non-linear activation function A is applied.	96

Figure 8.3	Structure and connected weights and biases to form a 3×3 MLP. The indicated neurons have been chosen to display how information flows through the network as a function of the weights and biases. The central neurons ($n=2$) have been highlighted to demonstrate which weights feed into the neuron from each respective layer. It is assumed that all neurons have the same activation function $a(x)$. Two neuron indices are required to distinguish the weights, however we have dropped the second index for clarity since we will only deal with the central neurons. Note that the superscript is not an exponent, but instead are neuron indices.	98
Figure 8.4	Learned fits for $f(t)$ given noisy sampled data, for different network topologies. Depths of 1 to 5 are shown with 8 and 32 neurons per layer. The dashed line indicates the true relationship while the purple points indicate the noisy samples on which the models are trained. The solid line is the predicted fit from the model. The training loss for each network as a function of epoch is also shown.	101
Figure 8.5	Different regularisation strategies applied to a multilayer perceptron (MLP) of depth 4 with 32 neurons per layer. The overfitting tendency is mitigated via the use of dropout and L2 regularisation. Note that the loss cannot be compared directly as the L2 regularisation includes a term into the loss function. In general the choice of activation function is decided empirically.	103
Figure 8.6	Different activation functions applied to the neurons of a MLP of depth 4, with 32 neurons per layer. The activation function is plotted below the fit of each model to the random walk data.	104
Figure 8.7	Demonstration of different convolutional filters that can be applied to images. (a) shows the original image of a house, (b) a horizontal edge detect kernel and (c) a vertical edge detect kernel. (d) Shows a standard convolutional neural network (CNN) architecture that may be used for image classification. The pooling layers reduce the dimensionality of each layer, while the 'conv' layers perform a convolution. Original image in (a) by Luke Stackpole.	107
Figure 9.1	Two different objective function topologies demonstrating the difference between convex and non-convex functions. The red section of the non-convex landscape represents a violation of the convex constraint. The presence of local minima in the landscape breaks prevents the use of efficient convex optimisers to consistently find the global minimum.	110

- Figure 9.2 Different predicted landscapes from the artificial neural networks (ANNs) trained on the sampled (grey) data. With a sparse sampling of the landscape the neural networks will disagree as to the location of the best minima, facilitating exploration. . . . 114
- Figure 9.3 Benchmark of different optimisations on the random polynomial problem. The best performing algorithm in terms of best minima found is the neural net algorithm (NN), followed by Gaussian process (GP), fruit fly (FF) and differential evolution (DE). Algorithms are restricted to 200 function evaluations (including training for the neural nets). 117
- Figure 10.1 Experimental implementation of the online deep learning algorithm for optimisation of the OD. The feedback like loop starts by (a) sampling the response of the MOT via a probe beam for different parameter sets. These parameter, cost pairs are then used by (b) the stochastic artificial neural network (SANN) to train a representative model of the cost landscape. (c) shows the mapping between parameters and cost indicated by the line colour, observed over the experimental acquisition (lower cost is an improvement). The SANN predicts the next best set of parameters to examine. Parameters that are found by the SANN during a minima search are passed to the experimental controls (d) to be implemented experimentally. 121
- Figure 10.2 Results of the 63 dimensional optimisation using the in-loop artificial intelligence (AI) algorithm. (a) Shows the original human ramp which was optimised using a Gaussian process learner. (b) the progression of the measure cost as a function of run number. The neural nets significantly outperform the differential evolution (DE) learner with the optimal solution identified at run 583. The bottom plot of (b) shows the historical minimum. The resulting optimal ramp is shown in (c) displaying features which are neither continuous nor monotonic. The corresponding OD measurements for each run are shown in (d) with an OD of 580 ± 8 and 960 ± 20 for the human and machine runs respectively. The orange line represents the probe detuning during the optimisation. (e) shows a quantitative comparison between the human and machine learned ramps effect on the atomic ensemble. Side profile absorption imaging shows a distinct difference in density in the centre region with a halo of atoms surrounding the machine learned ramp. The integration of these absorption plots is shown below the side images on a logarithmic scale, with the dashed boxes indicating the integration region that is chosen to be free of image distortions. 126

Figure 10.3	Side absorption images of the MOT at chosen times, to illustrate the difference in compression. While the human optimised sequence compresses the ensemble monotonically, the machine learned sequence opts for a multimodal technique, releasing the compression partially during the sequence. Note these images are not calibrated for OD.	128
Figure 10.4	Evolution of the atom spatial distribution, as a function of compression time, for both the human and machine learned sequences. The human sequence is monotonic whereas the machine learned sequence performs multiple compressions. The absorption traces are normalised to the imaging beam without atoms present, and the maximum observed absorption of each sequence.	129
Figure 10.5	Evolution of ensemble temperature as a function of compression time. Each point is determined by performing a ballistic drop measurement after applying the sequence up until that time.	129
Figure 10.6	Outputs from the SANN network models. (a) shows the optimal run with the highlighted points being used to generate the similarly coloured slices in (b). The blue curves represent every other parameter slice not highlighted. (c) shows the convergence between the measured and predicted cost according to the SANN with the red region corresponding to experimental noise. The difference between these two values is shown in the middle plot of (c) with a moving average. The lower plot of (c) shows the distance of a given parameter set from the best observed parameter set which can be used as a measure of exploration within the space.	131
Figure 10.7	Other identified local minima with their corresponding OD value. Multiple local minima can be identified via repeated application of the optimisation algorithm. Due to the large nature of the space there are classes of local minima with similar values.	132
Figure 10.8	Absorption images of the MOT captured during the compression sequence without pumping into the imaging state. The population will sometimes occupy the dark state for a period before finally ending up in that state towards the end of the run.	133

- Figure 11.1 Optimisation over an eggshell benchmark function. (a) shows the general landscape shape for a 1 dimensional parameter slice, the total function for a multidimensional problem is simply the addition of each of these profiles. (b) shows the results of a 50 dimensional optimisation run using the previous deep learning algorithm, fruit fly, differential evolution and Gaussian process. As can be seen the fruit fly algorithm is now outperforming the neural nets. 136
- Figure 11.2 Two separate objective function landscapes with different characteristics that pose a difficult optimisation problem. (a) presents a sampling problem, as knowledge of different parts of the landscape does not facilitate inference as to new local minima. (b) is a complex landscape that presents a more friendly approach, while there are local minima, an approximate solution to the problem can help to infer the location of better parameters (shown in green). 137
- Figure 11.3 Sampling method, where samples are determined by taking a uniform sampling of each parameter, to determine the coordinates in the objective function space. The scatter plot shows the distribution of 1000 samples. The histogram shows the frequency of the distance from the centre of the space $(0, 0, 0)$, which is not a uniform distribution. 138
- Figure 11.4 Different regularisations applied to the neural networks of the [SANN](#) to produce different exploration results. For the higher regularisation a lack of overfitting reduces exploration and thus the effectiveness of the algorithm. Conversely, too much overfitting also hampers to the effectiveness of the algorithm which avoids exploitation. To perform effectively, the algorithm is required to balance exploration and exploitation, as is the case for the middle plot which has identified a better minimum. The L2 regularisation coefficient for each optimisation run is noted above each plot. 140
- Figure 11.5 Ackley test problem landscape shown as both a 3-dimensional and contour plot, for the 2 parameter problem. The landscape exhibits a number of local minima with a global minimum at the centre of the landscape. The contour plot is shown for a smaller region of $[-5, 5]$ around the global minimum, illustrating the different length scales of the eggshell structure. 141

Figure 11.6	Benchmark results of multiple optimisation algorithms on a 100 dimensional Ackley test problem. The costs observed by other optimisation algorithms indicate that they are not yet within the main central region of the global minimum and are still stuck on the local minima located on the initial edge. The benchmark algorithms are the old neural net algorithm (NN), fruit fly (FF), differential evolution (DE), Gaussian process (GP) and Acuminata (AC).	142
Figure 11.7	Results of applying the fruit fly (blue) and Acuminata (green) to different benchmark problems. The first plot shows the representative landscapes for a 2D version of each problem. The scatter plots represent the best cost found for each problem with increasing dimensions.	143
Figure 11.8	Results of the optimisation of the single electron transistor operating point using Acuminata. While only 4 dimensional, the problem has a complex structure with experimental drift. The blue line is the observed costs and the green line is the running best. The 4 parameters correspond to sensor top gate, sensor left barrier gate, sensor right barrier gate and source drain bias of the single electron transistor.	145
Figure 11.9	10 dimensional Ackley problem with a drift that increases over time. Without the compensation the algorithm struggles to follow the minimum. With compensation Acuminata follows the drift in the landscape over time. The first 200 runs of each optimisation constitute training.	146
Figure 11.10	General experimental setup for performing spectral shaping using a spatial light modulator (SLM). A diffraction grating is used to spatially separate the spectral components of an incident optical pulse. The spatial light modulator (SLM) is positioned at the Fourier plane and can be used to modify the phase and amplitude of the different spectral components, by applying masks to the SLM pixels. The filtered beam is then recombined on a second diffraction grating.	148
Figure 11.11	Actor critic schematic for reinforcement learning. The actor enacts an action on the environment and is provided feedback by the critic in terms of the temporal difference (TD) error. The critic learns the value of each state which provides the basis for advice provided to the actor network.	149
Figure A.1	actor-critic (AC) agent episode history when attempting to learn the problem. The only difference between the two agents is the activation function used in the value network, RELU in (a) and GELU in (b).	159

Figure A.2	AC agent applied to the 1D Ackley function. (a) shows the episodic performance of the agent applied to this problem. (b) shows the structure inherent in the Ackley function that determines the reward structure. Clearly it is no longer monotonic.	160
Figure A.3	Optical setup for the auto alignment experiment into a cavity. 6 actuators are used to walk the beam, to maximise a particular mode. The cavity is scanned using a piezo to determine the present modes.	161
Figure A.4	Preliminary alignment using Acuminata. (a) shows the observed costs during the optimisation which lasted for 751 runs. (b) shows the position of the actuators as determined by Acuminata. The four actuators being controlled correspond to two separate steering mirrors with x and y control.	163
Figure A.5	Preliminary alignment using the Actor-Critic model. (a) shows the cumulative reward and final state for each episode. (b) shows the cumulative reward and final state for episodes which achieved a cumulative reward greater than 400.	164
Figure C.1	First five Legendre polynomials. The polynomials form an orthogonal basis from which arbitrary functions can be constructed.	171
Figure C.2	Original generated spectrum centered around 50Hz, filtered with a narrow band filter. This test set is used to test recovery of the original spectrum via the proposed method.	172
Figure C.3	Results of a test optimisation aimed at recovering the original signal. (a) shows the overlap between the original and recovered spectrum. (b) shows the same overlap but on a logarithmic scale. The edges of the spectrum show most relative divergence. (c) shows a comparison between the time series for the original, recovered and filtered signals. (d) shows the progression of the optimisation loss as a function of function evaluations.	173

LIST OF TABLES

Table 4.1	Convenient optical and atomic properties for the D_1 and D_2 transitions of ^{87}Rb . Trapping is generally performed on the D_2 line leaving the D_1 for memory operations.	31
Table 11.1	Definitions, domains and minima of the various benchmark functions used. All benchmark functions are found within optimisation literature as standard test problems.	144

ACRONYMS

AC	actor-critic	GP	Gaussian process
AFC	atomic frequency comb	MLP	multilayer perceptron
AI	artificial intelligence	ML	machine learning
ANN	artificial neural network	MOT	magneto-optical trap
AOM	acousto-optic modulator	MSE	mean squared error
APD	avalanche photo-detector	OD	optical depth
BEC	Bose-Einstein condensate	PDH	Pound-Drever-Hall
CNN	convolutional neural network	PD	photo detector
CRIB	controlled reversible inhomogeneous broadening	PGC	polarisation gradient cooling
DE	differential evolution	QD	quantum dot
EA	evolutionary algorithm	RL	reinforcement learning
EIT	electromagnetically induced transparency	RNN	recurrent neural network
EOM	electro-optic modulator	ROSE	revival of silenced echo
FDTD	finite difference time domain	SANN	stochastic artificial neural network
FF	fruit fly	SGD	stochastic gradient descent
FPGA	field programmable gate array	SLM	spatial light modulator
GAN	generative adversarial network	SO	stochastic optimisation
GD	gradient descent	SPDC	spontaneous parametric down conversion
GEM	gradient echo memory	TD	temporal difference
		UHV	ultra-high vacuum

1

INTRODUCTION

Almost everything in the natural world can be thought of as some sort of optimisation problem. Nature inherently optimises, with physical systems tending towards the lowest energetic state. Evolutionary biology demonstrates to us that organisms which show a natural tendency to adapt, will outperform and supersede their peers. Certainly the need for optimisation comes from the finite reality of the natural world. That is to say without restrictions such as resource scarcity and finite lifetimes, there would be no need to bother with optimisation at all, outside of curiosity. We do, however, live in a world which is confronted with scarcity and flux, requiring us to seek ever more efficient solutions to problems. This is most evident with the meteoric rise of communications and connected societies, presenting us with complex networking and distribution problems requiring new, efficient methods to solve, even with ever increasing computational power.

Complex quantum systems are now beginning to appear at the forefront of mainstream technology, with promises of superior computational power, secure communications and solutions to so far intractable problems. It is still unknown as to whether quantum computing will provide answers to intractable optimisation problems in the general sense. In either case quantum computing and communication is potentially interesting and useful enough on its own, to warrant further study. The design and operation of such devices is however no trivial task, with sources of decoherence, noise and extreme sensitivity to environmental conditions common to many quantum systems. While quantum systems promise powerful capabilities, they are also fragile. In the wake of this realisation, more attention has turned to issues regarding error correction, control, and optimisation. The challenges associated with these causes are the high dimensionality and inaccessibility of many quantum properties. Clearly in the absence of greater computational power, one will have to conceive smarter ways of dealing with such problems.

This thesis is essentially split into three parts, with the broader aim of addressing optimisation in the context of quantum and high dimensional systems. The first part introduces the concepts of quantum mechanics and the relevant cold atom theory for understanding the second part, which details the behaviour of cold atoms in the context of quantum information processing.

In the second part I detail some of the odd behaviours that can be observed in cold atoms as well as the operation of a quantum memory protocol which aims to improve on existing protocols. The work with cold atoms details the first attempts to improve this system using some basic machine learning approaches. It also serves to highlight the complexities of dealing with quantum information processes in the real world.

The third part presents a new method for high dimensional optimisation, leveraging the computational power of deep learning approaches. This algorithm is applied to the optimisation of the cold atom ensemble, with the aim of improving future experiments, the performance of which relies on the characteristics of the atom trap. This marks the first instance of an online deep learning approach applied in an experimental setting. I then go on to expand and improve the algorithmic approach, substantially improving the performance of the algorithm on different benchmark problems. The improvement of this algorithm leads to the application on additional quantum systems, highlighting the powerful ability of this algorithm to operate efficiently in high dimensional spaces.

PUBLICATIONS AND CONTRIBUTIONS During the course of my PhD, I was involved in a number of publications which are covered in this thesis. The first experimental component of this thesis is covered in Ch. 5. I was responsible for the experimental characterisation and was assisted by Anthony Leung. This work will be prepared for publication following further characterisation of this effect.

Ch. 7 covers the first published work on a backwards Raman memory scheme:

- Vernaz-Gris, P., Tranter, A.D., Everett, J.L., Leung, A.C., Paul, K.V., Campbell, G.T., Lam, P.K. and Buchler, B.C., 2018. High-performance Raman memory with spatio-temporal reversal. *Optics express*, 26(10), pp.12424-12431.

This chapter covers a number of theoretical simulations which I developed with contributions from Jesse Everett to the base code. I was also responsible for the development and integration of the optimisation code with the simulations. The experimental component of this chapter and the associated main result was performed in conjunction with Pierre Vernaz-Gris. I was directly involved in the experimental design, construction and measurement. Additionally I was responsible for the development and integration of the machine learning component of the experiment.

The second published work is covered in Ch. 10, detailing the experimental optimisation of the magneto-optical trap:

- Tranter, A.D., Slatyer, H.J., Hush, M.R., Leung, A.C., Everett, J.L., Paul, K.V., Vernaz-Gris, P., Lam, P.K., Buchler, B.C. and Campbell, G.T., 2018. Multiparameter optimisation of a magneto-optical trap using deep learning. *Nature communications*, 9(1), pp.1-8.

In this work I was directly involved in the experimental design, construction and measurement. Co-authors Harry Slatyer and Michael Hush were directly involved in the development of the deep learning algorithm. I was also responsible for the integration of the algorithm into the experimental setup and modifications to the algorithm.

The work covered in Ch. 11 has been withheld from publication due to the opportunity for commercialisation. I was independently responsible for the complete re-design and subsequent improvements to the deep learning approach. The following works are currently in preparation relating to the machine learning extensions:

- Auto-Alignment of an experimental quantum optics experiment using actor-critic algorithms. Other authors: Naomi Chaix-Eichel, Thibault Michel.
- Improving the sensitivity of gravimeter measurements with machine learning. Other authors: Christian Freier, Kyle Hardman, Paul Wigley.

Information on these works can be found in the appendices. Additional works outlined in Ch. 11 are also expected to yield publications, however the projects are still in the preliminary stages and have not been mentioned above.

The final work I was involved in is not covered in this thesis, however the interested reader may consult:

- Everett, J.L., Vernaz-Gris, P., Campbell, G.T., Tranter, A.D., Paul, K.V., Leung, A.C., Lam, P.K. and Buchler, B.C., 2018. Time-reversed and coherently enhanced memory: A single-mode quantum atom-optic memory without a cavity. *Physical Review A*, 98(6), p.063846.

Part I

FUNDAMENTALS

"Exactly!" said Deep Thought. "So once you do know what the question actually is, you'll know what the answer means."

- *Douglas Adams, The Hitchhiker's Guide to the Galaxy*

2

LIGHT AND MATTER WITH QUANTUM MECHANICS

In order to reach a broader audience, initially we will cover the basics of quantum mechanics to facilitate the reader in understanding the broader field and applications of the work at hand. The interested reader may consider following [1] and [2] for a more in depth discussion of the concepts introduced here.

2.1 WHAT IS QUANTUM MECHANICS ANYWAY?

While many physicists will often remark that no-one really "understands" quantum mechanics, it is fortunate for us that the ability to perform quantum mechanical calculations lies within our reach. In fact it is the language of quantum mechanics that allows us to describe some of the most fundamental interactions in the natural world. Until 1900, the physics of so called quantum mechanical systems (a definition which we will limit to the interaction between a quantised object and some other system, be it semi-classical or not) had been limited in their description by classical physics. The origin of quantum mechanics is fuzzy in its conception, however, following the first resolution to the "ultra-violet catastrophe" by Max Planck, the field of quantum mechanics exploded into a frenzy of study and controversy with many out-right rejecting the notion of lights quantised nature following Albert Einstein's quantum explanation of the photoelectric effect [3].

Fast forward to the present and still we are continuing to discover the rich and often strange physics associated with quantum systems. This can be partly attributed to the common language that quantum mechanics has provided, allowing the description of the interaction between what would otherwise be distinct physical systems. Nowhere has this become more apparent than the quest for quantum computing and communication infrastructures, where this language provides a common overlay covering single donor semi-conductor and solid state physics all the way through to nuclear magnetic resonance and quantum optics.

2.2 THE BASIC THEORY

There are a number of equivalent descriptions of quantum mechanics mainly differing in their mathematical representation for the sake of convenience or conceptual understanding. While one of the earliest conceptual methods was to construct a wave function Ψ we will restrict our discussion to the more computationally convenient

Dirac notation. We wish to consider a state vector $|\psi\rangle$ which will represent our system's state. If we consider our state vector to reside in a complex vector space, then it must be true that any given pure state can be decomposed into a linear combination of the basis vectors that span our complex vector space. That is to say

$$|\psi\rangle = \sum_i c_i |e_i\rangle, \quad (2.1)$$

where c_i are the complex coefficients and $|e_i\rangle$ are orthonormal basis vectors of this space. The complex vector space generally used in quantum mechanics is referred to as a Hilbert space when possessing certain properties (see [1] for a more complete description). In this formalism the probability of finding the system in the state $|e_i\rangle$ is simply given by $|c_i|^2$, where $\sum_i |c_i|^2 = 1$.

2.2.1 Superposition

From this construction it follows that we can then have a state represented as a linear combination of two basis states. In the example of a spin- $\frac{1}{2}$ particle we can prepare the state

$$|\psi\rangle = \frac{1}{\sqrt{2}}(|\uparrow\rangle + |\downarrow\rangle), \quad (2.2)$$

where $|\uparrow\rangle$ and $|\downarrow\rangle$ represent the spin-up and spin down-states respectively. This linear combination represents a coherent superposition where the system occupies both states simultaneously. In this case upon measurement we find spin-up $|\langle\uparrow|\psi\rangle|^2 = \frac{1}{2}$ of the time and similarly for spin-down. Superposition is a fundamental property of quantum mechanics that forms the basis for many quantum computing algorithms as it may be leveraged as a form of parallel computation in certain contexts.

2.2.2 Entanglement

In general entanglement in quantum mechanics refers to the phenomenon in which the correlation between two or more systems cannot be explained by classical physics [4]. Consider the Bell state defined between two spin- $\frac{1}{2}$ systems as

$$|\psi_{Bell}\rangle = \frac{1}{\sqrt{2}}(|\uparrow\rangle_A |\downarrow\rangle_B - |\downarrow\rangle_A |\uparrow\rangle_B), \quad (2.3)$$

where $|i\rangle_k$ describes the state of system k . As this state is superposed we expect that upon measurement we should find the state in one of two states with equal probability. However this state has the additional property that a measurement performed on system A will determine the state of system B and vice versa. This property holds true irrespective of the distance between A and B and led to the now famous Bell's theorem

[5] and associated experimental confirmations [6–8] that ruled out local realism theory as an explanation for such phenomena.

2.2.3 Operators

An operator can be thought of as a transformation from one state to another. For example, if we have a state $|\psi\rangle$ then we can define an operator \hat{O}

$$\hat{O}|\psi\rangle = |\psi'\rangle, \quad (2.4)$$

where $|\psi'\rangle$ is the target state of the operation. A given state vector is said to be an eigenstate of an operator if it satisfies the relation

$$\hat{O}|\psi\rangle = c|\psi\rangle, \quad (2.5)$$

where the eigenvalue c is in general a complex value, however if \hat{O} is Hermitian then c is real. Hermitian operators are also called observables, and a typical example of those is

$$\mathcal{H}|a\rangle = E_a|a\rangle, \quad (2.6)$$

which will return the total energy of the system E_a .

2.2.4 Density matrix formalism

A state $|\psi\rangle$ is assumed to contain all the information regarding a system, and thus it is called a pure state. The mean value of a given operator \hat{O} is given by

$$\langle\hat{O}\rangle = \langle\psi|\hat{O}|\psi\rangle, \quad (2.7)$$

where $\langle\psi|\psi\rangle = 1$. Alternatively, we can describe a pure system in what is called density formalism which is given by

$$\rho = |\psi\rangle\langle\psi| = \sum_{i,j} c_{ij} |\psi_i\rangle\langle\psi_j|, \quad (2.8)$$

where $c_{ij} = c_i \cdot c_j^*$. The expectation value of the operator now is

$$\langle\hat{O}\rangle = \text{Tr}(\rho\hat{O}), \quad (2.9)$$

where Tr denotes the trace.

A statistical mixture of multiple pure states is also a valid state, and it is called a mixed state. Utilizing the density matrix formalism, mixed states are represented as

$$\rho = \sum_i p_i |\psi_i\rangle \langle \psi_i|, \quad (2.10)$$

with $\text{Tr}(\rho) = 1$. Here, p_i , represent probabilities and therefore $\sum_i p_i = 1$. Thus, pure states are only a subset of all possible quantum states. A straightforward way to differentiate between pure and mixed states is the following: if $\rho = \rho^2$ the state is pure, otherwise it's mixed.

2.3 CLASSICAL ATOM LIGHT INTERACTIONS

It is useful to consider the most basic picture of an atom with a single electron interacting with a classical electric field. We can initially consider this in the semi-classical picture and arrive at a number of useful quantities that are valid when the electromagnetic field contains many photons with a weak atom-field coupling. These assumptions are justified for many of the cases we will deal with.

2.3.1 Lorentz model

Consider an electron bound to an atom modelled as a damped harmonic oscillator. We can consider a monochromatic electric field interacting with our atom given by

$$\mathbf{E}(t) = \hat{\mathbf{e}} E_0 e^{-i\omega t}, \quad (2.11)$$

where $\hat{\mathbf{e}}$ is the polarisation unit vector, E_0 is the amplitude and ω is the driving frequency. We have used the dipole approximation (the atom is much smaller than the wavelength) to remove any spatial dependence. From this, it follows that the force on the electron due to the oscillating field will be

$$\mathbf{F} = -e\mathbf{E}, \quad (2.12)$$

where e is the fundamental charge. We represent the average position of the electron x

$$m\ddot{x} + m\gamma\dot{x} + m\omega_0^2 x = -e\hat{\mathbf{e}} E_0 e^{-i\omega t}, \quad (2.13)$$

where m is the reduced mass of the electron, ω_0 is the resonant frequency of the harmonic potential, and Γ the damping rate is equivalent to the excited state decay rate. If we assume a solution that has the same form as the driving field

$$\mathbf{x}(t) = \hat{\mathbf{e}} x_0 e^{-i\omega t}, \quad (2.14)$$

then substituting into Eq. 2.13 we can obtain

$$x_0 = \frac{eE_0/m}{\omega^2 - \omega_0^2 + i\Gamma\omega}. \quad (2.15)$$

Now we have a complete expression for the electron displacement we can think about the induced dipole moment

$$\mathbf{d} = -ex = \alpha(\omega)\mathbf{E}, \quad (2.16)$$

where α is the polarizability describing the response of the atom to the applied electric field. From this we find the polarizability to be,

$$\alpha(\omega) = \frac{e^2/m}{\omega_0^2 - \omega^2 - i\Gamma\omega} \quad (2.17)$$

and through the relationship $\chi(\omega) = N\alpha(\omega)/\epsilon_0$ we also find the susceptibility

$$\chi(\omega) = \frac{e^2N}{m\epsilon_0(\omega_0^2 - \omega^2 - i\Gamma\omega)}, \quad (2.18)$$

where N is the number density and ϵ_0 is the vacuum permittivity. The complex susceptibility describes the response of the atom to the applied field, with the imaginary part describing the absorptive properties and the real part describing the dispersive properties. We can define the complex refractive index in terms of this susceptibility as

$$\tilde{n}^2 = 1 + \chi(\omega). \quad (2.19)$$

2.3.2 Absorption cross-section

As the atomic vapours that we are dealing with have a low vapour pressure we can take the complex refractive index to be

$$\tilde{n}(\omega) \approx 1 + \frac{\chi(\omega)}{2}, \quad (2.20)$$

where the imaginary part of $\tilde{n}(\omega)$ gives us the absorption coefficient

$$a(\omega) = \frac{2\omega}{c} \mathbf{Im}[\tilde{n}(\omega)], \quad (2.21)$$

where c is the speed of light. If we take our detuning from resonance ω_0 to be small then we recover the familiar Lorentzian absorption profile

$$a(\omega) \approx \frac{Ne^2}{m\epsilon_0c\Gamma} \frac{(\Gamma/2)^2}{(\omega_0 - \omega)^2 + (\Gamma/2)^2}. \quad (2.22)$$

We can use this to find the absorption cross section, $\sigma(\omega)$, via the relationship $a(\omega) = \sigma(\omega)N$ giving

$$\sigma(\omega) = \frac{\sigma_0}{1 + 4(\Delta/\Gamma)^2}, \quad (2.23)$$

where $\sigma_0 = \sigma(\omega_0)$ and $\Delta = \omega_0 - \omega$. However this expression will only apply in the case where the intensity of our driving field is much less than that of the saturation intensity of our given transition.

2.3.3 Optical depth

A useful quantity when working with ensembles of atoms is optical depth. Here we define optical depth as the quantity, d , for which an incident resonant field is attenuated by a factor e^{-d} . Such a measurement is challenging as ensembles with high optical depth, such as those found in the present work, will attenuate resonant beams beyond measurement capacity. Thus, it is useful to define resonant optical depth in terms of a transmitted off-resonant beam. It follows from Eq. 2.23 that an off resonant beam will instead experience an attenuation of

$$\exp\left(-\frac{d}{1 + 4(\Delta/\Gamma)^2}\right). \quad (2.24)$$

Thus we can define optical depth through a measurement of off resonant transmittance as

$$d = (1 + 4(\Delta/\Gamma)^2) \cdot \ln\left(\frac{I_0}{I_T}\right), \quad (2.25)$$

where I_T is the intensity of the transmitted probe attenuated from a reference intensity I_0 .

2.4 SEMI-CLASSICAL PICTURE

Unfortunately the classical picture of atom-light interactions will only take us so far. For example, our previous discussion is limited to the case where we have a bright classical field, but not so bright that we begin to saturate the populations of particular transitions. Things start to become more complicated when dealing with quantised states corresponding to orbital angular momentum and spin angular momentum. Restricting our discussion to Hydrogen-like atoms¹ we can begin to think about a more quantum model for our atom.

¹ The present work deals with alkali atoms such as ⁸⁷Rb which can be thought of as Hydrogen-like as they have one valence electron.

2.4.1 Two level atom

Let us start by assuming we have a single atom with an excited and ground state, labelled $|e\rangle$ and $|g\rangle$ respectively, with resonant frequency ω_{eg} . We also wish to define some atomic operators in the $\{|e\rangle, |g\rangle\}$ basis. We will define three atomic operators, $\hat{\sigma}_z$, $\hat{\sigma}_+$ and $\hat{\sigma}_-$ which can be constructed as such

$$\hat{\sigma}_z = |e\rangle\langle e| - |g\rangle\langle g| \quad \hat{\sigma}_+ = |e\rangle\langle g| \quad \hat{\sigma}_- = |g\rangle\langle e|, \quad (2.26)$$

where $\hat{\sigma}_z$ is the Pauli spin matrix and $\hat{\sigma}_\pm$ are effectively the atomic raising and lower operators, respectively. The matrix representation of these operators corresponds to

$$\hat{\sigma}_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \quad \hat{\sigma}_+ = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix} \quad \hat{\sigma}_- = \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}. \quad (2.27)$$

We wish to find the Hamiltonian of this system which should take the form $\mathcal{H} = \mathcal{H}^A + \mathcal{H}^I$, where \mathcal{H}^A is the free atom Hamiltonian and \mathcal{H}^I describes the atom-light interaction.

To construct \mathcal{H}^A we consider the energy of the two states $|e\rangle$ and $|g\rangle$. Given the resonant frequency ω_{eg} , the energy difference between the two states is $\hbar\omega_{eg}$. If we define a zero point energy as being half way between these two states we then have

$$E_e = \frac{1}{2}\hbar\omega_{eg} \quad E_g = -\frac{1}{2}\hbar\omega_{eg}, \quad (2.28)$$

where E_i is the energy of state i . From this construction it follows that we expect the free atom Hamiltonian to take the form

$$\mathcal{H}^A = \frac{1}{2}\hbar\omega_{eg} |e\rangle\langle e| - \frac{1}{2}\hbar\omega_{eg} |g\rangle\langle g| = \frac{1}{2}\hbar\omega_{eg}\hat{\sigma}_z. \quad (2.29)$$

We will consider a classical field $\mathbf{E}(t)$ decomposed into its counter rotating components of frequency ω which is detuned Δ from ω_{eg}

$$\mathbf{E}(t) = \frac{\hat{\epsilon}E_0}{2} \left(e^{-i\omega t} + e^{i\omega t} \right) \quad (2.30)$$

$$= \mathbf{E}^+ + \mathbf{E}^-. \quad (2.31)$$

To determine \mathcal{H}^I we consider the interaction between the atom and field to be the dipole interaction similar to that of the classical case presented previously

$$\mathcal{H}^I = -\mathbf{d} \cdot \mathbf{E}. \quad (2.32)$$

Due to parity ($\langle g|\mathbf{d}|g\rangle = \langle e|\mathbf{d}|e\rangle = 0$) we can consider only the cross terms in the atomic dipole operator,

$$\mathbf{d} = \langle g|\mathbf{d}|e\rangle |g\rangle\langle e| + \langle e|\mathbf{d}|g\rangle |e\rangle\langle g|. \quad (2.33)$$

By choosing the phase appropriately we can simplify this expression further in terms of the operators in Eq. 2.26. Similar to the electric field we may decompose the dipole operator into its counter rotating components

$$\mathbf{d} = \langle g | \mathbf{d} | e \rangle (\hat{\sigma}_+ + \hat{\sigma}_-) \quad (2.34)$$

$$= \mathbf{d}^+ + \mathbf{d}^-, \quad (2.35)$$

where $\mathbf{d}^\pm \sim e^{\mp i\omega_{eg}t}$. If we expand our expression from Eq. 2.32 we get the expansion

$$\mathcal{H}^I = -\mathbf{d}^+ \cdot \mathbf{E}^+ - \mathbf{d}^+ \cdot \mathbf{E}^- - \mathbf{d}^- \cdot \mathbf{E}^+ - \mathbf{d}^- \cdot \mathbf{E}^-. \quad (2.36)$$

The first and last term in this expression will be oscillating much faster ($\omega + \omega_{eg}$) whereas the central terms will be oscillating much more slowly (Δ). We can neglect the fast oscillating terms under the rotating wave approximation which reduces the above to

$$\mathcal{H}^I = -\mathbf{d}^+ \cdot \mathbf{E}^- - \mathbf{d}^- \cdot \mathbf{E}^+. \quad (2.37)$$

We can form our Hamiltonian in terms of the Rabi frequency, Ω

$$\mathcal{H}^I = \frac{\hbar\Omega}{2} (\hat{\sigma}_- e^{i\omega t} + \hat{\sigma}_+ e^{-i\omega t}), \quad (2.38)$$

where $\Omega = -\langle g | \mathbf{d} | e \rangle \cdot \hat{\mathbf{e}} E_0 / \hbar$. The Rabi frequency describes the coupling strength between the two level atom and the field, as the inner product gives the overlap between the atomic dipole and the electric field.

2.4.2 Optical Bloch equations

Until this point we have assumed a simple semi-classical model which has not characterised certain processes we are familiar with such as spontaneous emission, in which an excited atom will tend to decay to the ground state. Considering the Hamiltonian that we have just constructed, we can think about the evolution of the density matrix elements and how our system will evolve. A full derivation of these elements is presented in [2, 9], however here for brevity we will simply present them:

$$\begin{aligned} \partial_t \rho_{ee} &= i\frac{\Omega}{2} (\tilde{\rho}_{eg} - \tilde{\rho}_{ge}) \\ \partial_t \rho_{gg} &= -i\frac{\Omega}{2} (\tilde{\rho}_{eg} - \tilde{\rho}_{ge}) \\ \partial_t \tilde{\rho}_{ge} &= -i\Delta \tilde{\rho}_{ge} - i\frac{\Omega}{2} (\rho_{ee} - \rho_{gg}) \\ \partial_t \tilde{\rho}_{eg} &= -i\Delta \tilde{\rho}_{eg} + i\frac{\Omega}{2} (\rho_{ee} - \rho_{gg}). \end{aligned} \quad (2.39)$$

Eq. 2.39 describes the evolution of the populations and coherences of our two level system. Here the tilde is used to denote the elements of the density matrix in the rotating

frame (assumed to be rotating at the frequency ω) for mathematical convenience. If we consider the case of being on resonance with no external field (this is equivalent to setting Δ and Ω to 0) then we can consider adding some terms to account for a decay rates with the following forms

$$\begin{aligned}
\partial_t \rho_{ee} &= -\Gamma \rho_{ee} \\
\partial_t \rho_{gg} &= \Gamma \rho_{ee} \\
\partial_t \tilde{\rho}_{ge} &= -\gamma \tilde{\rho}_{ge} \\
\partial_t \tilde{\rho}_{eg} &= -\gamma \tilde{\rho}_{eg},
\end{aligned} \tag{2.40}$$

where Γ is simply our spontaneous decay rate from earlier and γ accounts for additional coherence decay. Physically it is clear that the Γ terms are facilitating population shuffling from the excited to ground state thanks to spontaneous emission. The additional decay rate γ here corresponds to dephasing that doesn't affect the populations, for example atom-atom collisions. In the limit of homogeneous broadening where there are no additional dephasing effects $\gamma = \Gamma/2$. Combining 2.39 and 2.40 we arrive at the optical Bloch equations

$$\begin{aligned}
\partial_t \rho_{ee} &= i \frac{\Omega}{2} (\tilde{\rho}_{eg} - \tilde{\rho}_{ge}) - \Gamma \rho_{ee} \\
\partial_t \rho_{gg} &= -i \frac{\Omega}{2} (\tilde{\rho}_{eg} - \tilde{\rho}_{ge}) + \Gamma \rho_{ee} \\
\partial_t \tilde{\rho}_{ge} &= -(\gamma + i\Delta) \tilde{\rho}_{ge} - i \frac{\Omega}{2} (\rho_{ee} - \rho_{gg}) \\
\partial_t \tilde{\rho}_{eg} &= -(\gamma - i\Delta) \tilde{\rho}_{eg} + i \frac{\Omega}{2} (\rho_{ee} - \rho_{gg}).
\end{aligned} \tag{2.41}$$

2.4.3 Useful results

Now that we have determined the evolution of our system via the optical Bloch equations there are a number of useful properties we can define.

SATURATION A fundamental difference between the classical and quantum treatment of the atom is the excitation. While our harmonic oscillator may be excited to arbitrarily high amplitudes, this is not true for our quantum mechanical system which will undergo saturation. We can define a saturation parameter s given by

$$s = \frac{\Omega^2 / \gamma \Gamma}{1 + \Delta^2 / \gamma^2}. \tag{2.42}$$

Framing the steady state solution in terms of this parameter we can show that the steady state population of the excited state tends to

$$\rho_{ee} = \frac{s/2}{1+s}. \quad (2.43)$$

Thus a small intensity field will increase the excitation linearly with s , tending to $\frac{1}{2}$ as $s \rightarrow \infty$. This effect is a major departure from our classical modelling and becomes important when considering the properties of laser cooling (see Ch. 3).

SCATTERING CROSS SECTION As we have altered the mechanism for illuminating a quantum mechanical atom, so must we modify our expression for the scattering cross section. We now have the updated expression

$$\sigma(\omega) = \frac{\sigma_0}{1 + 4\Delta^2\Gamma^2 + I/I_{sat}}, \quad (2.44)$$

where I/I_{sat} is a rescaling of the incident intensity to the saturation intensity and $\sigma_0 = \hbar\omega_{eg}\Gamma/2I_{sat}$.

POWER BROADENING If we consider the steady state population of the excited state in the limit of a strong field we find the Lorentzian

$$\rho_{ee} = \frac{\Omega^2/4}{\Delta^2 + \Omega^2/2}. \quad (2.45)$$

As stated previously this value should tend to a maximum of $\frac{1}{2}$ in this limit, however the full-width half-maximum of this Lorentzian is also $\sqrt{2}\Omega$. This dependence on the power of the driving field is known as power broadening, where the effective linewidth of the transition is increased due to strong coupling with the field.

INCOHERENT SCATTERING An issue when considering atom-light interactions in the context of coherent processes, is losses due to incoherent scattering. In the homogeneous broadening case ($\gamma = \Gamma/2$), the incoherent scattering rate can be calculated in terms of the saturation parameter as

$$R_{ISc} = \frac{\Gamma s^2}{2(1+s)^2}. \quad (2.46)$$

It is important to note that while the incoherent scattering rate increases quadratically with s , the coherent scattering rate only increases linearly.

2.4.4 Three level atom

A simple extension to the two level atom case is to consider a three level atom. The structure of these transitions can take many forms however here we will consider a




Figure 2.1: Three level Λ -scheme with two applied fields, a strong coupling field Ω and weak probe field \mathcal{E} , with frequencies ω_{es} and ω_{eg} respectively.

Λ type scheme (see Fig. 2.1). Here we will assume there exists three states, a ground state $|g\rangle$, an excited state $|e\rangle$ and a meta-stable state (or second ground state) $|s\rangle$. Additionally we will assume that there are optical transitions $|g\rangle \rightarrow |e\rangle$ and $|s\rangle \rightarrow |e\rangle$ with detuned frequencies ω_{eg} and ω_{es} , respectively. A similar treatment to the two level system above can be taken. If we assign the excited state to have zero energy we can determine the interaction Hamiltonian in the rotating frame as

$$\mathcal{H}_I = \frac{\hbar}{2} \left(\Omega_{eg} |e\rangle \langle g| e^{-i\omega_{eg}t} + \Omega_{es} |e\rangle \langle s| e^{-i\omega_{es}t} + c.c. \right), \quad (2.47)$$

where Ω_{ji} represents the Rabi frequency for transition $|i\rangle \rightarrow |j\rangle$. This system will behave like an effective two level system in the regime of large detunings, driving populations between the two ground states without any significant population occupying the excited state. This effect is a result of the coherence between the dipole moments of each transition. This convenient approximation will be used in Ch. 7.

3 | LASER COOLING

A particularly useful tool for the study of atom-light interactions is some form of atomic trap. While there are many types of traps, for the present work we are mainly concerned with a magneto-optical trap. For this we will need to introduce concepts related to laser cooling and trapping.

3.1 DOPPLER COOLING

We require the ability to cool atomic species down to temperatures at which point the atomic motion is either negligible or at the very least, conveniently small. A method which was proposed simultaneously by Hänsch and Schawlow, and, Wineland and Dehmelt [10, 11] in 1975, is to use laser light as a method for cooling gaseous media. This method is known as Doppler cooling.

3.1.1 Doppler shift

Let us consider an atom illuminated by a single propagating field at some frequency ω . In the case where the atom has non-zero or at least non-negligible velocity, we need to take into account the Doppler shift associated with this velocity. This can be expressed via the relationship

$$\omega_d = -\mathbf{k} \cdot \mathbf{v} = -kv \cos(\theta), \quad (3.1)$$

where ω_d is the Doppler shift, \mathbf{v} is the atomic velocity and \mathbf{k} is the wavevector with magnitude $2\pi/\lambda$. This can have a significant effect when addressing optical transitions, for example, for Rubidium at room temperature this can be on the order of 500 MHz. This effect is responsible for the Doppler broadened transitions observed in atomic vapours during an absorption measurement (see Fig. 4.5).

3.1.2 Spontaneous emission

An atom absorbing a photon results in a transfer of momentum from the optical field to the atom. Similarly the process of spontaneous emission also results in a momentum kick. This reshuffling of momenta is what allows the process of laser cooling. The optical field will always have a momentum vector in the direction of propagation, however spontaneous emission will occur in all directions, effectively averaging to zero

net effect over many events. In this way we can define the force associated with such a process as

$$F = \hbar k R_{Sc} = \hbar k \Gamma \rho_{ee}, \quad (3.2)$$

where R_{Sc} is the scattering rate derived in the steady state. We can thus frame this in terms of $s_0 = s(\Delta = 0)$, the saturation parameter at zero detuning as

$$F = \frac{\hbar k \Gamma}{2} \frac{s_0}{1 + s_0 + (2(\Delta - \omega_d)/\Gamma)^2}, \quad (3.3)$$

where we have also included the Doppler shift from Eq. 3.1. For higher intensities this force will saturate at a maximum value $\sim \hbar k \Gamma / 2$ as $\rho_{ee} \rightarrow 1/2$. However a discussion about high intensities will become more complicated as stimulated emission begins to play a role in momentum transfer.

3.1.3 1D optical molasses

Considering the 1-dimensional case of two counter-propagating optical fields we can determine the force experienced by the atom in the low intensity limit. In this case we consider a standing wave constructed from the two counter-propagating fields. We can thus derive an average force $F_{av} = -\beta_D \cdot v$ acting against the atomic velocity, where β_D is the damping coefficient and is given by

$$\beta_D = -\frac{8\hbar k^2 \Delta s_0}{\Gamma(1 + s_0 + (2\Delta/\Gamma)^2)}. \quad (3.4)$$

From Eq. 3.4 we can see that if $\Delta < 0$ we end up with a situation where the atomic motion is viscously damped. This resulting viscous force is referred to as optical molasses.

DOPPLER LIMIT Considering the viscous damping force, in the absence of other forces we would expect the temperature of our system to arrive at 0K. Clearly this result is unphysical as there will be a number of heating effects that contribute to a non-zero temperature. When undergoing spontaneous emission our sample will receive momentum kicks of $\hbar k$ per event in random directions. This therefore must contribute to heating of the system. The minimum energy we can expect to achieve is then simply when the cooling rate and heating rate are in equilibrium. This minimum energy is characterised as E_D , the Doppler energy, with an associated Doppler temperature T_D

$$E_D = \frac{\hbar}{8} \left(\frac{\Gamma + 4|\Delta|^2}{2|\Delta|} \right) \quad (3.5)$$

$$T_D = \frac{\hbar \Gamma}{2k_B},$$




Figure 3.1: Level structure for a fictitious atom with two levels with $J=1$ and 2 respectively. The Clebsch-Gordan coefficients for each transition are indicated.

where k_B is Boltzmann's constant. This turns out to be a curious result as it is independent of the atomic mass, optical wavelength or optical intensity (providing we stay below the threshold for the low intensity regime). Experimental demonstrations initially agreed with this limit, however, there quickly became mounting evidence that this limit could be exceeded.

3.2 SUB-DOPPLER COOLING

Experimental demonstrations began to show that sub-Doppler cooling was possible [12] and it was theorised this arose from evolution of the internal atomic state [13]. Here we will present the approach for determining the origin of this increased cooling capacity.

For simplicity it is instructive to consider a fictitious atom with ground states $|g\rangle$ and $|e\rangle$ with total angular momentum $J_g = 1$ and $J_e = 2$ respectively¹. We will also assume the Clebsch-Gordan coefficients² shown in Fig. 3.1. A key component of sub-Doppler cooling mechanisms is the light shift experienced by ground magnetic substates. In the low intensity limit with two laser beams of intensity $s_0 I_s$ the energy shift δ_g is given by

$$\delta_g = \frac{\hbar \Delta s_0 C_{eg}^2}{1 + (2\Delta/\Gamma)^2}, \quad (3.6)$$

where C_{eg} is the relevant Clebsch-Gordan coefficient for the transition with a given polarisation. I will only consider the case of two counter-propagating circularly polarised beams as this is what is used in the present work. If we consider the standing wave created by these beams propagating in the z -direction we find that the electric field amplitude is constant along z . The polarisation on the other hand is linear, with an axis that varies helically along the z -direction. This is an important distinction from the linear- \perp -linear case which has a spatially varying polarisation that modulates between linear and circular as shown in Fig. 3.2. When this is the case the origin of the damping

¹ While the atom is fictitious, the level scheme is still physical.

² Clebsch-Gordan coefficients arise from angular momentum coupling. More information can be found in [9]




Figure 3.2: The resulting spatially varying polarisation as a result of counter propagating fields in (a) linear- \perp -linear configuration and (b) σ_+ - σ_- configuration. As can be seen the polarisation modulates between linear and circular in the linear- \perp -linear case whereas for circularly polarised fields the resulting polarisation is a linear polarisation, the axis of which rotates about the propagation axis.

force can be explained in terms of Sisyphus cooling mechanisms as there is a spatially varying light shift on the magnetic ground states [14]. However for the σ_+ - σ_- case there is no such spatial dependence of the light shift.

In the absence of such a light shift we require a new mechanism for the origin of this damping force. Similar to Sec. 2.4.1 we can construct the coupling between the atomic states and optical fields in terms of the Clebsch-Gordan coefficients and find the average force acting on the atom in terms of the spatial dependence of this interaction. This is to say that

$$F = \left\langle \frac{dV}{dz} \right\rangle, \quad (3.7)$$

where V is the coupling between the atom and optical fields. It can be shown that the average force is given by [13]

$$F = \hbar k \frac{\Gamma}{2} \left[\rho_{gg_1} \left(s_+ - \frac{s_-}{6} \right) + \rho_{gg_0} \left(\frac{s_+ - s_-}{2} \right) + \rho_{gg_{-1}} \left(\frac{s_+}{6} - s_- \right) + C_r \left(\frac{s_+ - s_-}{6} \right) - \frac{1}{3} C_i \left(s_+ \frac{\Delta - kv}{\Gamma} + s_- \frac{\Delta + kv}{\Gamma} \right) \right], \quad (3.8)$$

where ρ_{ggi} is the population of magnetic substate i and s_{\pm} are the saturation parameters for the σ_{\pm} fields given by

$$s_{\pm} = \frac{\Omega^2/2}{(\Delta \mp kv)^2 + (\Gamma/2)^2}. \quad (3.9)$$

The first 3 terms in Eq. 3.8 are effectively the radiation pressure forces from each beam scaled by the populations of the magnetic substates. The final two terms correspond to a force from an induced coherence between the $|g_1\rangle$, and $|g_{-1}\rangle$ magnetic substates which undergo population redistribution through absorption-stimulated emission cycles. The coefficients at the front of these terms are given by

$$C_r = \text{Re} \left[\langle g_1 | \rho | g_{-1} \rangle e^{-2ikvt} \right] \quad C_i = \text{Im} \left[\langle g_1 | \rho | g_{-1} \rangle e^{-2ikvt} \right]. \quad (3.10)$$

From Eq. 3.8 it follows that an atom in the $|g_1\rangle$ will preferentially scatter from the σ_+ field by a factor of 6, and conversely so for the $|g_{-1}\rangle$ state.

The final piece to this mechanism is the population imbalance induced by atomic motion. If we consider an atom to be moving along the z-direction with a velocity v , then as mentioned previously the atom will see a linear polarisation rotating about the axis of propagation. As the atom moves through this polarisation gradient the quantisation axis, which is defined by the polarisation, also rotates. This rotation forces optical pumping to occur as the population of the ground states lags behind the steady state population defined by the rotating quantisation axis. The difference between the populations can be quantified by

$$\rho_{gg_1} - \rho_{gg_{-1}} = \frac{40kv}{17\delta'_{g_0}}, \quad (3.11)$$

where δ'_{g_0} is the mean light shift of the $|g_0\rangle$ state. Here we can see the explicit dependence of the population imbalance on the velocity of our atom. As the atom moves through the polarisation gradient, the population is optically pumped into the state which is more likely to scatter counter-propagating photons. In this manner an atom not at rest will experience an unbalanced radiation pressure force acting against its motion. We expect that for a stationary atom the population will arrive at the steady state with $|g_0\rangle$ containing 9/17 of the population and the remaining population spread evenly amongst the $|g_{\pm 1}\rangle$ states. We can calculate the damping coefficient for this polarisation gradient cooling β_P in the low-velocity regime as

$$\beta_P = -\frac{120}{17} \frac{\Delta\Gamma\hbar k^2}{5\Gamma^2 + 4\Delta^2}. \quad (3.12)$$

3.2.1 Comparison to Doppler cooling

If we compare our new damping coefficient to that of Eq. 3.4 for Doppler cooling we find some stark differences. For one this new factor is independent of laser power. Secondly, for low velocity our new force is stronger than Doppler cooling by a factor $2|\Delta|/\Gamma$. As we tend to higher velocities however this force will become approximately equal to that of Doppler cooling. Given this new mechanism we expect that the final

temperature should be lower than that of Doppler cooling. Indeed it can be shown that we should expect the minimum energy to approach the recoil energy E_r

$$E_r = k_B T_r = \frac{\hbar k^2}{2M}. \quad (3.13)$$

However in reality the lowest achievable energies will be a few times larger than the recoil energy due to assumptions we have initially made about the atoms being localisable in the polarisation gradient such that they only see one linear polarisation at any given time. To achieve temperatures lower than the recoil limit one must resort to techniques such as evaporative cooling, commonly used in the creation of Bose-Einstein condensates (BECs).

3.2.2 Sisyphus cooling

An additional mechanism that is also responsible for cooling beyond the Doppler limit is known as Sisyphus cooling and as mentioned previously, is a result of spatially varying light shifts such as that found in a linear- \perp -linear configuration. Conceptually if one considers an atom moving in the z -direction, then where the potential is spatially dependent there is a net transfer between kinetic and potential energy as the atom essentially "climbs" the potential. In the case of polarisation gradients when atoms are at the top of the hill they are most likely to be optically pumped into the ground state with a lower energy thanks to the spatially dependent light shift. This in turn means that over many cycles the atom will experience a net loss of kinetic energy, thus cooling the atom. In general Sisyphus cooling can be performed on an atomic system with multiple levels, where kinetic energy can be exchanged as the atom traverses the increasing potential. A choice of pumping is also required to switch the atomic system back to the lower state with a net loss of kinetic energy.

3.3 MAGNETO-OPTICAL TRAPS

Atom-light interactions have long been a field of fundamental study, however as atom trapping became common place the technology developed in the process of these experiments has found its way into applications such as sensing [15] and quantum memories [16]. While there are many types of atom traps we will restrict our discussion to the MOT and instead direct the reader to [14] for a more detailed discussion regarding other types of traps. In the case of the MOT we are trapping a large ensemble of atoms with the aim of using the entire optical depth of our atomic cloud.

The first MOT was demonstrated in 1987 by Raab et al. [17] with Sodium atoms. MOTs use a combination of the methods described in Secs. 3.1 and 3.2 as well as magnetic fields to localise the atoms to a particular interaction region. This idea will become important for our quantum memory applications (see Ch. 7).




Figure 3.3: Coil geometries for creating a linear magnetic field across the central region. (a) corresponds to Helmholtz configuration, whereas when one of the currents are reversed as in (b) anti-Helmholtz is achieved with 0 field in the central region.

3.3.1 Magnetic fields

Let us consider two circular loops of wire with radius r aligned parallel to each other and separated by the same distance s as shown in Fig. 3.3a. For the case when $r = s$ then we define this particular geometry as Helmholtz coils. The point of this particular geometry is to create a uniform magnetic field when current is passed through the loops. An alternate configuration is to run the current through the coils in opposite directions as shown in Fig. 3.3b. This produces a linearly increasing magnetic field from the centre of the trap which has $B = 0$. This is particularly useful as it provides an easy way to define a central region to push our trapped atoms towards. In this configuration the magnetic states of the atoms in the trap will undergo a position dependent Zeeman shift. As the atoms move further away from the trap the detuning of the magnetic states will shift relative to the detuning of the magnetically insensitive state $M_f = 0$. This asymmetry in the detuning will cause the atoms to scatter light preferentially from the trapping beams that are opposed to the direction of motion away from the trap, hence pushing the atoms towards the centre of the trap where there is no preferential scattering.

3.3.2 Optical fields

In conjunction with the magnetic fields, a MOT requires optical cooling beams to trap atoms. While 2D geometries exist we are mainly concerned with a 3D geometry such as that shown in Fig. 3.4. Here there are 3 sets of $\sigma_+ - \sigma_-$ beam pairs which provide the




Figure 3.4: Schematic showing the orientation of the cooling beams relative to the MOT coils. Trapping and repump co-propagate after being combined on a polarising beam splitter. Each pair of counter propagating beams are orthogonally polarised to establish the needed standing wave. The direction of gravity is also noted in the second orientation.

3D trapping potential. Given these 6 beams we have an intensity profile of the form [18]

$$I(\mathbf{r}) = \frac{c\epsilon_0 E_0^2}{2} [6 - 4(\cos(kz) \sin(ky) + \cos(ky) \sin(kx) - \sin(kz) \cos(kx))]. \quad (3.14)$$

In addition to the cooling beams for optical molasses an additional beam is employed known as the repump beam. So far we have considered atoms with a single transition, where in reality atoms will have many addressable optical transitions. For example the D2 line of Rubidium 87 has an excited state with 4 hyperfine states that are split by at most 266 MHz. In this case if we were to apply only the cooling beam we would simply optically pump our atoms into another state via off-resonant transitions. After this had occurred our atoms would no longer be resonant to our cooling light and we would lose the ability to trap. To mitigate this effect the repump beam is used to pump the atoms out of this dark state and back into the state resonant with the cooling beams.

3.3.3 3D potentials and polarisation gradients

In Sec. 3.2 we described a mechanism for sub-Doppler cooling of atoms. Unfortunately the extension from 1D to 3D is non-trivial and fails an intuitive explanation. For a 3D potential there are polarisation gradients in all directions due to the additional beams, leading to spatially dependent light shifts and thus, Sisyphus cooling (see 3.2.2). Experiments conducted have however observed temperatures below that of the Doppler limit and thus are undergoing some form of polarisation gradient cooling [19, 20].

3.3.4 Density restrictions and loss

For the applications at hand we will often be concerned with increasing the optical depth of the atomic ensemble, thus increasing the atom-light coupling. For a small number of atoms the density of the atomic cloud effectively scales linearly with the number of trapped atoms. However as the density increases the trap will enter what is called the multiple scattering regime. In this regime we find that the density will be constrained to a maximum value independent of the number of atoms. This effect originates from the fact that scattered light from the cooling process will be reabsorbed by nearby atoms which introduces an interatomic repulsive force. Effectively one can derive the steady state number of atoms by reconciling the capture rate and loss rate [21, 22] given by

$$N = 0.1 \frac{A}{\sigma_{col}} \left(\frac{v_c}{v_T} \right)^4, \quad (3.15)$$

where σ_{col} is the collisional cross section, A is the surface area of the trap volume, v_c is the maximum capture velocity and v_T is the thermal velocity of background atoms. Collisions here are assumed to be between cold atoms and thermal background atoms which will impart enough kinetic energy to kick atoms from the trap.

3.3.5 Time dependent MOTs

The steady state density of a MOT is fixed by the relationship between the loading (addition of atoms to the trap) and loss rates, however it is advantageous to increase the number of trapped atoms by varying some of these quantities in time. One such method is known as a transient compression sequence. During a short period, after some loading steady state has been reached, it is possible to increase the atom density transiently by increasing the magnetic field strength and detuning the repump frequencies. This technique is often referred to as a temporal dark-SPOT. The purpose of ramping the magnetic field strength is conceptually straight forward, reducing the volume occupied by the cold atomic cloud, thus increasing the density in the interaction region of interest.

The detuning of the repump frequency is a little more subtle [23]. The main idea is that the limiting factor in density in the multiple scattering regime is the repulsive force between atoms from the scattering of trapping light. Thus by increasing either the detuning or intensity of the repump light, one can shelve the atoms in the centre of the atomic cloud in a dark state where they will no longer participate in cooling, however they will not contribute to any repulsive force either. Obviously this ramp must be done over a timescale which does not facilitate a major loss of atoms from the trap. In the current work we find that to be on the order of 10-20ms although this will depend on the particular characteristics of the trap.

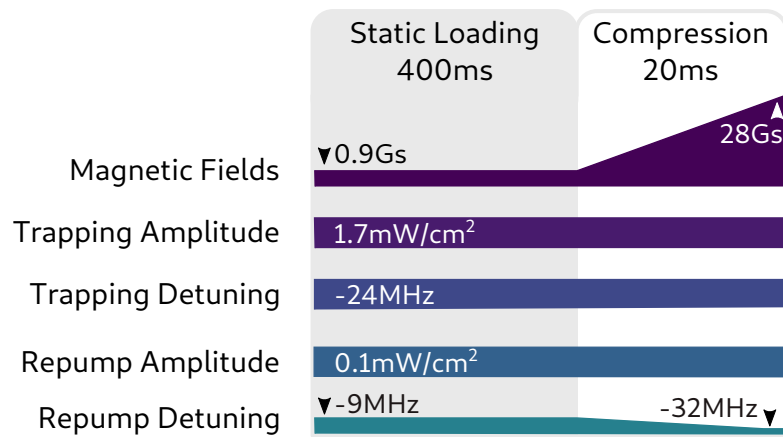


Figure 3.5: Timing and coordination of the various control channels during static loading followed by a temporal dark SPOT transient compression sequence.

Using these techniques experiments have been performed which have successfully trapped $2.6 \cdot 10^{10}$ Rubidium 87 atoms [24] and $\sim 10^{10}$ atoms in Sodium [23]. An example of such a time dependent ramp is shown in Fig. 3.5.

Part II

ATOMS AND QUANTUM MEMORIES

“If people do not believe that mathematics is simple, it is only because they do not realize how complicated life is.”

- *John von Neumann*

4

TRAPPING RUBIDIUM

As mentioned in Ch. 3, the study of light and matter generally requires confining the atomic systems in some way, either singularly or as an ensemble. For the study of quantum memories, we are more interested in trapping large numbers of atoms to collectively enhance the interaction, and hence the memory efficiency. Building on the techniques described in Ch. 3, we can now describe the operation of the MOT that enables these experiments. I will describe the experimental setup and general operating principles required to trap and form a cold atomic ensemble, however a more efficient approach will be discussed in Pt. iii.

4.1 GENERAL SETUP

Experiments carried out in the present work involve ^{87}Rb , a non-stable isotope of Rb, however its decay rate is extremely slow [18]. Fig. 4.1 depicts the experimental setup and the formation of a ^{87}Rb cold atomic ensemble. Alkali metals are often chosen as suitable candidates for atom light experiments mainly owing to their single outer electron which can be thought to behave as a hydrogen-like atom. ^{87}Rb has two D-lines: the D_1 transition $5^2S_{1/2} \rightarrow 5^2P_{1/2}$ and the D_2 transition $5^2S_{1/2} \rightarrow 5^2P_{3/2}$. Some of the more useful properties of these lines can be found in Tab. 4.1 with a comprehensive list found in [18]. ^{133}Cs and ^{23}Na are also commonly used in MOTs, however ^{87}Rb has a few advantages for our application. Firstly, it has a smaller number of hyperfine states which precludes the use of difficult pumping schemes as compared to Cs. Additionally, the availability of low noise Ti:Sapph lasers and diode lasers at the required wavelengths facilitate experimentation.

Property	D_1	D_2
Wavelength (λ)	≈ 794.979 nm	≈ 780.241 nm
Linewidth (Γ)	$2\pi \cdot 5.746$ MHz	$2\pi \cdot 6.065$ MHz
Recoil Temperature (T_r)	348.66 nK	361.96 nK
Doppler Temperature (T_D)	–	146 μK

Table 4.1: Convenient optical and atomic properties for the D_1 and D_2 transitions of ^{87}Rb . Trapping is generally performed on the D_2 line leaving the D_1 for memory operations.




Figure 4.1: Current experimental setup depicting a ^{87}Rb MOT illuminated by the cooling beams used for trapping. Also depicted is the glass cell attached to the vacuum chamber and the magnetic field coils.

4.1.1 Vacuum system

To achieve a suitable trap lifetime and atomic density for the atom light experiments we have designed, a MOT must be formed under ultra-high vacuum (UHV). As noted in Eq. 3.15, the final steady state atom number will depend on the number of collisions occurring with background atoms, thus a higher vacuum will allow us to form high density ensembles. In general, the vacuum is deemed to be sufficient provided that the background pressure remains $\leq 1 \cdot 10^{-8}$ torr. For the present work, a background pressure of $1 - 5 \cdot 10^{-9}$ torr is generally observed. To achieve such a vacuum, the entire system must undergo a high temperature bakeout on the order of 120°C for several hours. A schematic of the vacuum chamber is presented in Fig. 4.2. Here the glass cell is connected to the vacuum chamber which is pumped down to UHV. The ion-pump continues to remove background ions liberated from the walls of the cell and other surfaces. The dimensions of the cell allow us to form a MOT which is approximately 5 cm in length. Rb atoms are loaded into the vacuum chamber via a dispenser located on the top side of the cell. The dispenser is a Rb salt which heats up under the application of electrical current, this in turn liberates Rb atoms from the salt with isotopes at natural abundance. Typically under normal operations, a dispenser will last upwards of 5 years.

4.1.2 Laser cooling systems

Once the chamber is filled with ^{87}Rb , we can now go about coalescing these atoms into a cold atomic ensemble. This is done via the proper application of laser cooling beams and magnetic fields. Generally the D_2 line of ^{87}Rb is used for trapping, as the less complicated D_1 is more suitable for memory operations. Figure 4.3 shows the level scheme associated with the D_2 line and the associated transitions of interest.




Figure 4.2: Schematic of the ^{87}Rb MOT vacuum chamber and associated glass cell. The cell is coated on the external faces with anti-reflection coatings for 780 – 795nm. Continuous operation of the ion pump ensures the vacuum pressure remains on the order of $\sim 10^{-9}$ Torr.

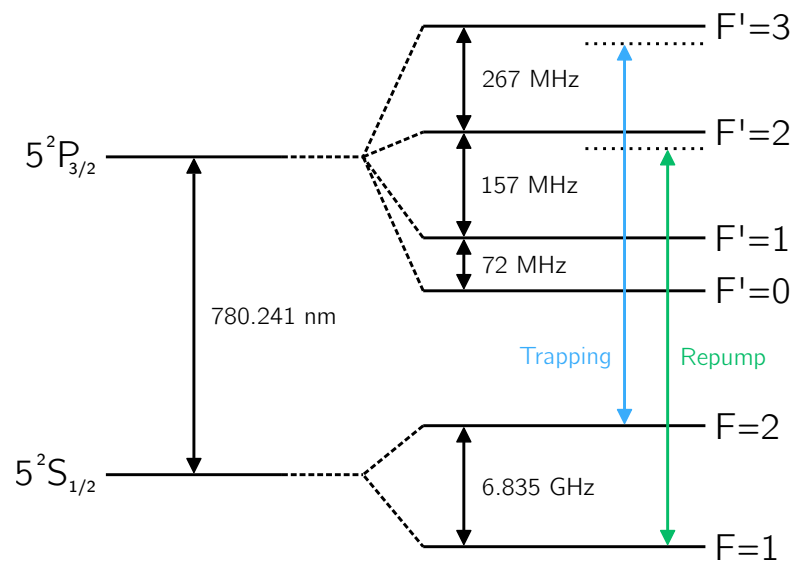


Figure 4.3: D_2 transition of ^{87}Rb showing the detuned trapping and repump transitions used when loading the MOT. The hyperfine splitting due to coupling between the nuclear spin ($I=3/2$ for Rb) and the total angular momentum of the electron.




Figure 4.4: Cooling beam geometry relative to the glass cell. The trapping and repump beams co-propagate and are retro-reflected to create the standing waves in all 3-dimensions. The trapping region is defined by the intersection of the three beams.

The trapping transition used is $F = 2 \rightarrow F' = 3$ while the repump transition is $F = 1 \rightarrow F' = 2$. The trapping transition is chosen as it is a cycling transition providing the best efficiency for trapping large numbers of atoms effectively. The beams used for cooling and trapping are derived from MOGLabs tunable Littrow diode lasers running at $\approx 780\text{nm}$ coupled with tapered amplifiers. Using the diode laser and tapered amplifier combination, $\sim 400\text{mW}$ of trapping power and $\sim 80\text{mW}$ of repump power can be delivered to the MOT. The trapping and repump beams are combined on a beam splitter before being split into 3 separate beams to create the 3D trapping volume. Fig. 4.4 shows the geometry of the combined cooling beams which are configured for $\sigma_+ - \sigma_-$ operation via $\lambda/4$ wave plates. To create the standing wave in each dimension, the cooling beams are retro-reflected with the assumption that the majority of the beam intensity will be transmitted without significant attenuation by the atomic cloud.

4.1.3 Saturated absorption locking

While there are many ways to lock a laser to a given frequency, one of the easiest ways when working with Rb is to use some form of saturated absorption spectroscopy (otherwise known as Doppler-free spectroscopy). To successfully lock our diode lasers to a given D_2 hyperfine transition of ^{87}Rb , we are required to control our laser frequency to within $\sim\text{MHz}$. Fortunately, we can simply measure and lock to a given atomic line using spectroscopy. However if one simply measures a probe transmitted through a vapour cell containing Rb, one will quickly find that the hyperfine transitions are completely washed out by the Doppler broadened linewidth ($\sim 500\text{MHz}$). To alleviate this problem, saturated absorption may be used with a common setup during the experiments in the present work as shown in Fig. 4.6. Here a retro-reflected strong probe beam is sent through a ^{87}Rb cell while the laser frequency is scanned over some

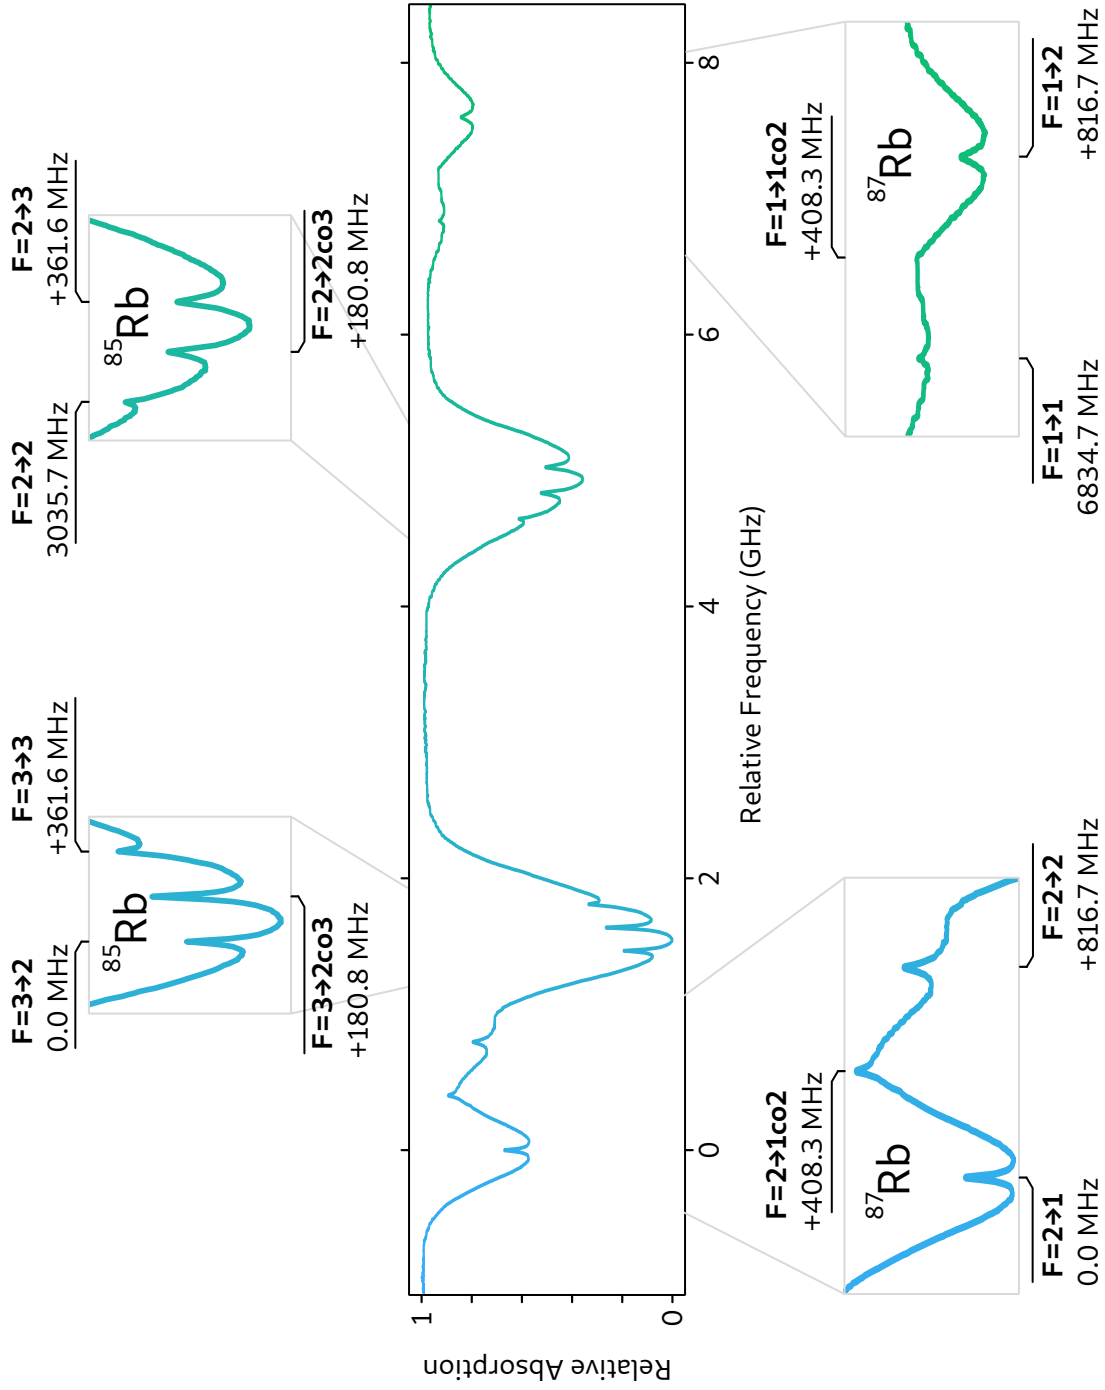


Figure 4.5: Saturated absorption measurement of the D₁ line of Rb. Both ^{87}Rb and ^{85}Rb are present. Each set of hyperfine transitions are highlighted with the transition frequencies relative to the leftmost transition for ^{87}Rb and ^{85}Rb respectively. Note that the ground state splittings are different for ^{87}Rb and ^{85}Rb .

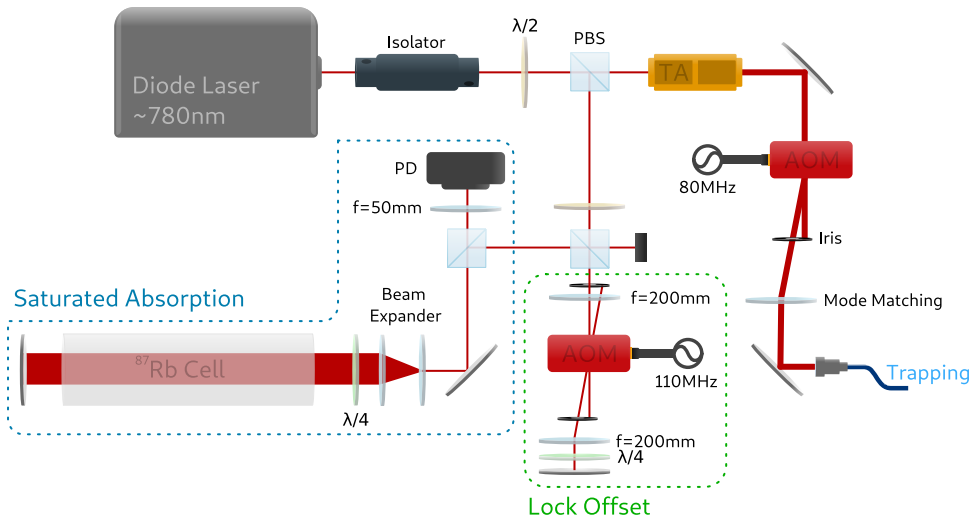


Figure 4.6: Optical setup used to lock the diode laser onto the cooling transition. The saturated absorption measurement provides a locking signal which can be locked to a particular hyperfine transition. This setup is cloned with the exception of the lock offset block for locking a second diode laser onto the repump transition.

range of interest. While the frequency is detuned from the transition, there is no effect and we see the standard Doppler broadened line. However, as we approach resonance for a given transition, the probe beam will initially saturate atoms with a velocity close to that of 0. On the second pass through the atoms, we will find an absorption dip around this velocity class due to this saturation, corresponding to the Doppler free linewidth of the transition. Fig. 4.5 shows the measured absorption profiles for the D_1 transition. From these absorption profiles, a locking signal can be generated by either modulating the laser diode current or a magnetic field (via driving coils surrounding the Rb cell) and locked to using the standard Pound-Drever-Hall (PDH) technique [25].

4.1.4 Magnetic fields

To facilitate coalescence of cooled atoms towards the centre of the MOT, we require 3D coils to produce the desired magnetic field. For quantum memory experiments, it is advantageous to elongate the MOT in the direction of probe propagation. This helps to maximise the usable optical depth of the ensemble rather than wasting atoms in a spherical distribution around the centre of the trap. This is achieved using race-track coils in anti-Helmholtz configuration, which are elongated along the propagation axis. Fig. 4.7 shows the coil geometry relative to the glass vacuum cell and the resultant field. During steady state loading, the magnetic field gradient in the trapping volume is generally around 6 Gs/cm, spiking up to 28 Gs/cm at the peak of the transient compression sequence.

An additional challenge to quantum memory experiments is that any extraneous fields will substantially decrease the memory life time and the effectiveness of optical pumping and cooling. To this end, we are also required to implement 3 sets of

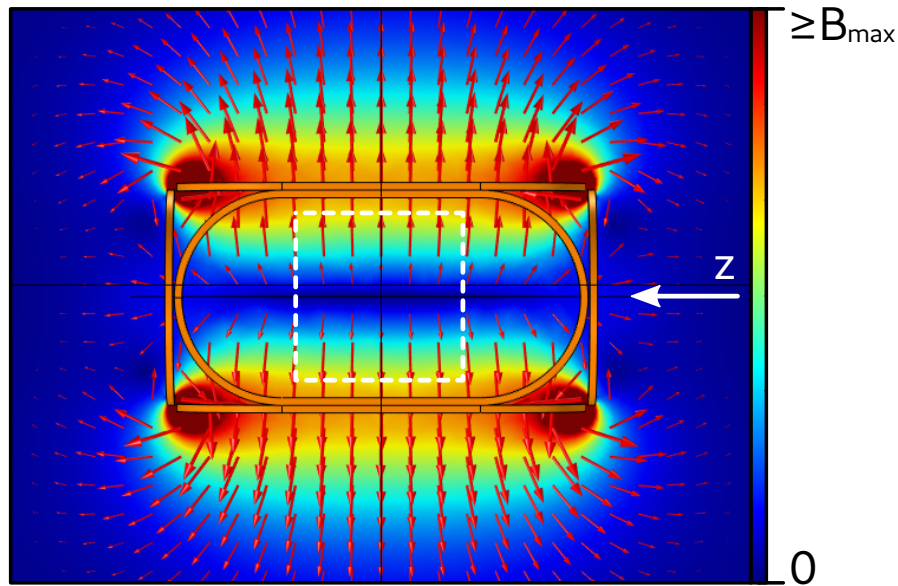


Figure 4.7: Geometry of the race track coils used to generate the elongated MOT. The dashed box in the centre represents the location of the vacuum cell with the storage axis for memory operations labelled as the z -axis. The field profiles are simulated using a **FDTD** simulation in COMSOL. The field magnitude is given at a slice through the centre axis which possesses radial symmetry about the z -axis. The vector field arrows are logarithmically scaled to the magnitude of the field. For clarity the field from the capping coils is not shown which would add additional confinement along the z -axis.

Helmholtz coils to cancel any background magnetic fields such as that from the Earth. To ensure field uniformity and optical access, the diameter of the coils is much greater than that of the experimental region and covers the entire extent of the optical table (2.4×1.2 m). During the transient compression phase, there will be rapidly changing magnetic fields which necessarily will induce Eddy currents into surrounding magnetic materials. To mitigate this, where possible non-magnetic metals or plastics are used. Additionally the vacuum chamber and cell are elevated off the table via a Delrin mount, further reducing the effect of Eddy currents induced in the optical table and mounts.

4.1.5 Optical pumping

It is often advantageous or necessary to work out of a particular magnetic sublevel (m_f). In the case of quantum memory experiments, often the choice of transitions to work with is dictated by their relative strength, with the goal being to choose transitions which maximise memory efficiency. In any case, without optical pumping, the population of the atomic vapour will be spread amongst the various m_f states. Optical pumping can be achieved by applying a bias field along the axis of propagation and using a suitably polarised pumping beam to shuffle the population to the desired m_f state. The bias field is implemented using a set of Helmholtz coils to ensure the field across the atomic ensemble is uniform.

4.2 EXPERIMENTAL CONTROL

Simply having a MOT running in steady state can be achieved with minimal electronics as all that is required is PDH locking to keep the diode lasers on the correct transition. However for more complicated experiments such as those involving quantum memories and atom light interactions, a larger degree of precision and control is required.

4.2.1 Generalised experimental control

The type of experiments with which we are concerned can be characterised by individual experimental runs which contain procedural subsets that are performed sequentially. This experimental run is then repeated at an effective duty cycle. For example, in a typical quantum memory experiment, the different procedures in an experimental run may be MOT loading, MOT compression, polarisation gradient cooling (PGC)/optical pumping, probe propagation and measurement. The general control schematics of the experiment are shown in Fig. 4.8. Here there are three asynchronous control loops that can run on FPGAs or any device that is capable of real time I/O operations. The first loop is concerned with the locking of various lasers systems. This will always run in the background and is independent of the main experimental loop. The second loop controls acquisition operations, using a National Instruments acquisition card (NI 5761). Timing for this loop is triggered by experimental timing. The third loop is the experimental loop which simply serves to implement control decisions at a specified time. Control decisions in this context simply correspond to either digital or analogue outputs, which may be connected to any relevant piece of lab equipment. Timing instructions are received by a National Instruments control chassis running a TCP client. The job of the TCP client is to connect to a designated address and establish a TCP socket connection over which the timing information will be sent.

4.2.2 Python *RunBuilder*

As the demanding timing requirements are segmented and delegated to FPGAs, construction of experimental runs can occur outside of the real-time operating system and LabVIEW oriented system of the National Instruments chassis. This has a few key advantages, namely that of a more convenient and developed language for constructing graphical user interfaces for ease of use and data manipulation in an interpretive language. The structure of the RunBuilder client treats each channel as an object which may correspond to a physical, digital or analogue channel. Each of these channel objects will contain a list of time-value pairs corresponding to the physical value that an output should reflect at the relevant time. A complete record of the behaviour of all connected channels can then be constructed and exported to the FPGAs via the TCP server. The FPGAs will cycle through these values according to the precise timing of the single cycle timed loops, which guarantee timing precision. In this way arbitrary experimental runs

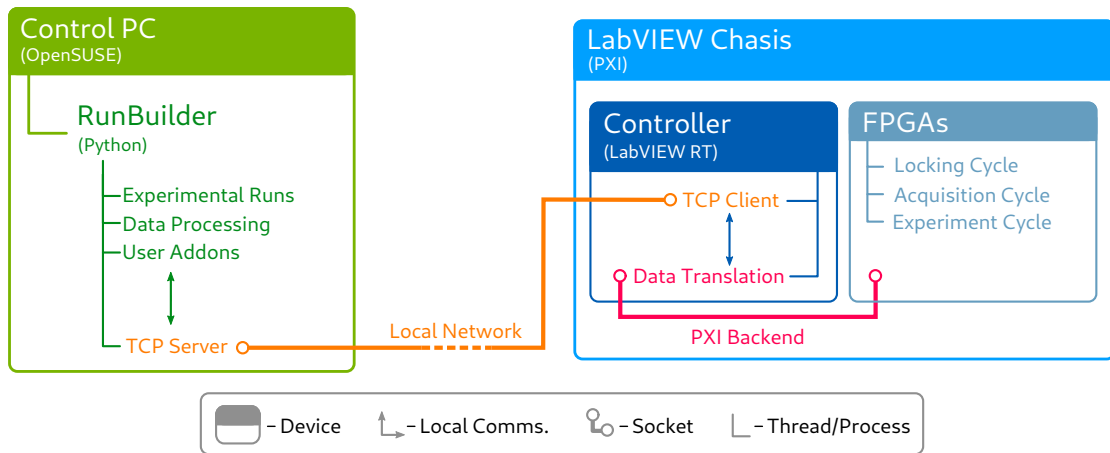


Figure 4.8: Schematics of the control scheme used for experimental control. The modular setup allows for automated experimental control which is useful for optimisation purposes. Extensions can be included as addons to each system. *FPGA* code is contributed by Geoff Campbell.

can be constructed for any conceivable experiment. Due to the modular construction of the experimental timings, automatic sweeps and measurements can be preprogrammed and automatically executed. This will become important for optimisation purposes described in Pt. [iii](#).

4.2.3 Atom loading

Using the experimental timing control, a generalised procedure for loading atoms into the trap can be constructed. This can be segmented into three main sections, loading, compression and *PGC*. During these stages, control is exerted over the trapping detuning frequency, repump detuning frequency and magnetic field strength of the trapping coils. Initially, we assume that there are only thermal atoms in the trapping volume. From this, an initial loading stage is used which aims at capturing a subset of these thermal atoms and loading them into the trap. While the duration for this can vary, typically this stage will last for about 400ms. The second stage is known as the compression stage. This stage is aimed at increasing the optical density in the interaction region (the region through which our probe beam will propagate). Pt. [iii](#) will cover in more detail the specifics behind finding an optimal sequence for future experiments using machine learning techniques. However for the quantum memory experiments covered in Part [ii](#), variations of the temporal dark SPOT method described in Sec. [3.3.5](#) are used. This period will last on the order of 10-20ms. The final part of the sequence is the *PGC* stage immediately after the compression. This stage combines optical pumping and *PGC* to cool the atomic vapour down further and facilitate pumping into the correct hyperfine and magnetic sublevel (dependent on the experiment).

4.3 CHARACTERISATION AND MEASUREMENT

4.3.1 Optical depth

As mentioned previously, OD is a useful quantity for determining the interaction strength along the propagation axis. As the OD grows, this measurement becomes erroneous when measured close to resonance. However it is possible to measure the OD still by taking absorption measurements at multiple detunings and fitting Eq. 2.25 to the absorption data. It should be noted that the OD will not only characterise the atomic density but rather the atomic density subject to the coupling between the optical mode and the atomic vapour. This will be a function of the beam parameters and the particular transition that the OD is measured on. Such comparison between ODs in different setups is somewhat fraught with danger. For optimisation purposes we may compare relative ODs to determine whether we have improved the trapping efficiency as the beam parameters and transition are held constant.

4.3.2 Imaging

It is often useful to look at the transverse profile of the atomic vapour for a number of reasons. Firstly in the case of the MOT used in the current work which has a cigar profile, one can use side imaging to diagnose issues regarding the laser cooling and magnetic fields. More commonly though, it is useful to align the optical pumping and probe beams.

Imaging can be done under two different conditions, fluorescence and absorption. Fluorescence imaging is the easier of the two and involves capturing scattered light from the laser cooling process on a CCD camera. Measurements taken under emission can be used to calculate the total atom number using some simple assumptions [18]. Absorption measurements are similar in setup with the addition of an off resonant expanded beam which is used to probe the transverse OD of the ensemble. Unlike the fluorescence measurement, the absorption measurement is destructive as it will optically pump the atoms during measurement. However in this way, it can be a more informative tool being that it is transition selective by construction. A schematic of the experimental imaging setup is shown in Fig. 4.9 along with an example of both types of measurement.

4.3.3 Temperature

Measurement of the temperature of the atomic vapour is often of interest as the expansion rate of the ensemble will tie directly into the memory lifetime. The ensemble temperature can be estimated via a ballistic drop, where the atoms are released from the trap (laser cooling and magnetic fields turned off) and allowed to undergo acceleration due to gravity as well as ballistic expansion. The velocity of the atomic ensemble,

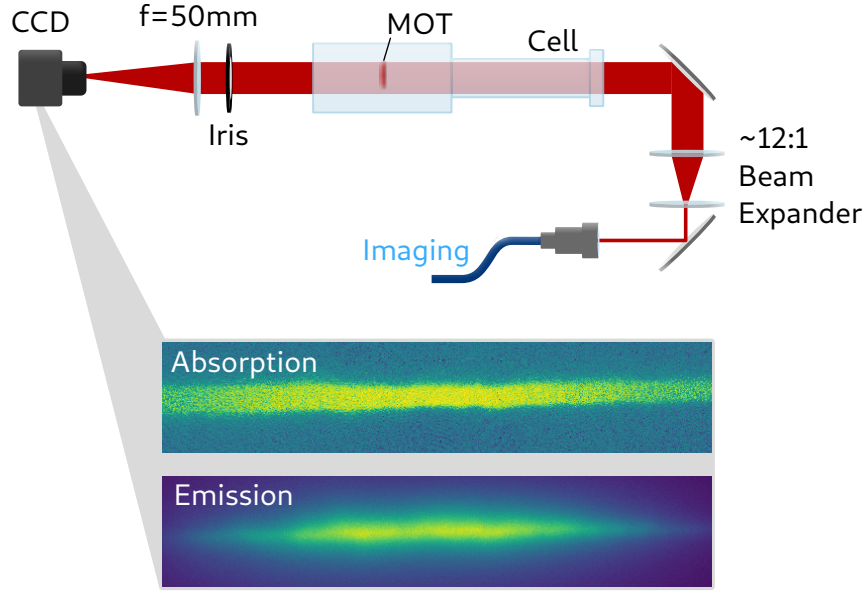


Figure 4.9: Experimental schematic for transverse imaging of the MOT. An example of both a fluorescence emission image and absorption image are shown. Note that the intensity of the absorption image has been inverted.

corrected for acceleration due to gravity, can then be used to determine the temperature of the ensemble. We assume an initial distribution [26]

$$N(x, y, z, t) = \frac{1}{\sqrt{8\pi\sigma_x^2\sigma_y^2\sigma_z^2}} \exp \left[-\frac{1}{2} \left(\frac{(x - gt^2/2)^2}{\sigma_x^2} + \frac{y^2}{\sigma_y^2} + \frac{z^2}{\sigma_z^2} \right) \right], \quad (4.1)$$

where N is the atomic density distribution, σ_i is the Gaussian radius of the expanded cloud in the i -axis at time t with the form $\sigma_i = \sqrt{\sigma_{i0}^2 + \sigma_v^2 t^2}$ where σ_{i0}^2 is the initial radius and σ_v^2 is the radial velocity. Note here that we have taken the x -direction to be the vertical direction, consistent with the experimental setup. The radial velocity will depend on the temperature of the cloud via the relation

$$T = \frac{M}{k_B} \sigma_v^2. \quad (4.2)$$

From this relation, we can determine the temperature of the ensemble using the absorption imaging techniques described in Sec. 4.3.2. Fig. 4.10 shows an experimental measurement where the atomic ensemble has undergone ballistic expansion. The depicted ensemble is undergoing expansion and widening, whilst also accelerating due to gravity.

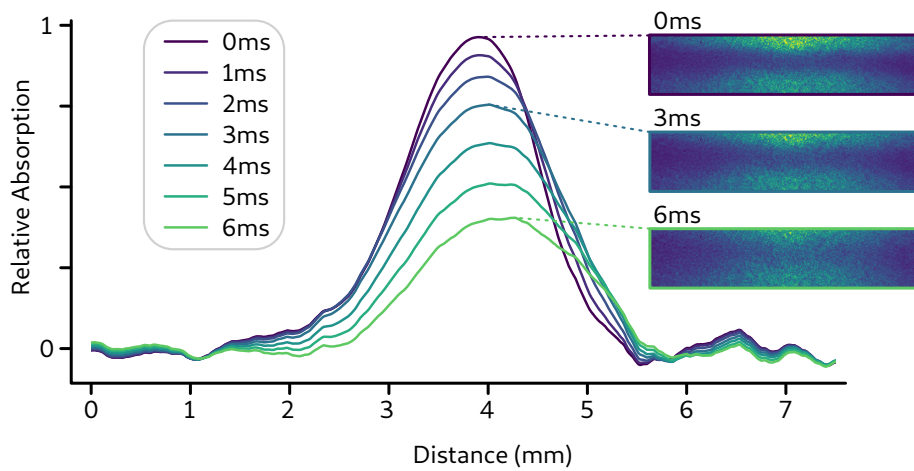


Figure 4.10: Ballistic drop measurement of the ensemble temperature with acceleration due to gravity. The images are captured under absorption measurement on the $F = 1 \rightarrow F = 2$ transition of the D_2 line. The measured temperature of the ensemble above is on the order of $\sim 240 \mu\text{K}$.

5

ATOMIC EMISSION

During the course of performing quantum memory experiments we observed a phenomena that we have simply come to call the atomic emission. This is characterised simply by a coherent axial emission along the direction of highest optical depth which can be sustained with only the trapping beams. We have performed experiments to characterise this emission and determine its source, the details of which are covered in this chapter.

5.1 LITERATURE

It has been known since Dicke first introduced the concept of superradiance to the atomic community [27] that ensembles of emitters can cooperate to create a kind of enhanced interaction, known as superradiance. Effectively, one can think of the ensemble of atoms acting like a single large atom with a collective decay, enhancing the rate of emission as N^2 , where N is the number of atoms [28]. Frequency shifts due to this collective behaviour can also be observed in the absorption and emission properties [29]. Since the initial conception there have been a number of demonstrations of superradiance in different systems such as ions [30, 31], quantum dots [32] and nanoplasmonics [33].

In quantum optics, it is often the case that the dimensions of the atomic ensemble are larger than the wavelength of the light, which complicates the situation somewhat. In this case, the collective emitters can still cooperatively enhance the stimulated emission via the use of entanglement between the N atoms in the ensemble [34, 35]. However, another effect similar to this behaviour can be observed which is sometimes referred to as superradiant light scattering [36]. Demonstrations of such included superradiant Rayleigh scattering in a BEC [37–39] where the scattering event provides a feedback mechanism for directional scattering.

The study into coherent scattering between emitting dipoles then was also examined from an optomechanical perspective with the assistance of cold atoms. Cold atoms, such as those in a MOT, can provide the densities required to exhibit the non-linearities associated with a high optical depth [40]. For the most part, however, these studies focused on cold atoms confined to some sort of resonator such as an optical cavity, with one of the earliest demonstrations showing preferential scattering into a cavity mode for a perpendicular pump beam [41]. The authors also observed that there was a frictional force present which decelerated the centre-of-mass motion of the atoms.

Additionally, BECs in cavities have also been shown to undergo the same process as well as exhibiting a quantum phase transition into a self organised supersolid [42].

After these initial results, the phenomenon now termed self-organisation was studied in cold atomic gases [43], especially within the context of collective atomic recoil lasing [44]. One of the defining characteristics of self organisation is pattern formation, which had been observed in non-linear optics before [45, 46]. The spatiotemporal structures emerge spontaneously when the system is driven from equilibrium, usually by an external pump beam of some description. These structures arise from dipole forces exerted on atoms by incident pump light [47]. It has also been shown that this can lead to atomic bunching, which can greatly enhance the non-linear response of a cold atomic gas [48, 49]. In the case of this self organisation into effective density gratings, light is preferentially scattered depending on the orientation of the pump with respect to the ensemble. In particular Ref. [49] shows that an optical cavity is not required to establish a feedback mechanism for self organisation.

Following this, there are few demonstrations in the literature of free-space self organisation and optical patterns. The most similar to the current work at hand is [50], which reported the observation of an optical emission from a ^{87}Rb MOT. In this work an elongated MOT similar to that of the present work, was illuminated with a pair of weak counter-propagating pump beams ($F = 2 \rightarrow F' = 3$ on the D_2 line) to create a standing wave while trapping of the MOT is halted. The result of this is the emission of two axial beams from the atoms, which decays due to atom loss. The authors demonstrated a steady state regime, which could be attained by turning down the intensity of the trapping beams and running the counter-propagating pumps. This emission also presented multimode properties such as different far field spatial patterns.

In follow up works [51, 52], the authors demonstrated a similar effect with counter-propagating lin- \perp -lin polarised light. The linearly polarised standing wave provides a mechanism for Sisyphus cooling, tightly confining the atoms to the lattice sites in the standing wave. Here, a fluctuation in the initial atomic distribution and optical fields can trigger a cascading wave-mixing effect to generate additional optical fields. These patterns however fluctuate shot-to-shot (on the order of atomic motion) indicating non-equilibrium phenomena.

Pattern formation in cold atomic ensembles is an ongoing research topic. For a more detailed look at the dynamics and theory associated with this phenomenon I direct the reader to the thesis of B. Schmittberger [53]. I will now outline our experimental observations and the key differences to the work described so far.

5.2 EXPERIMENTAL SETUP

As described in Ch. 4, we begin by creating an elongated ^{87}Rb MOT approximately 5 cm in length. With this setup, the OD typically achieved on the $F = 1 \rightarrow F' = 2$ transition




Figure 5.1: Experimental setup used when observing the atomic emission. Unlike the previous examples in the literature we do not require any additional beams other than those needed for laser cooling. The emission exits the atomic ensemble parallel to the axis with the highest optical depth. The inset shows a typical spatial pattern associated with the emitted mode.

of the D_1 line was on the order of $\sim 500 - 600$. As mentioned previously, examples in the literature have required the use of an additional pump beam to kick start and sustain the superradiant type phenomenon. Uniquely however, we find that over a certain OD threshold the emission will kick start and sustain itself. Generally for this to occur the trapping detuning must be kept in the range $-32 - -42$ MHz. The repump frequency was found in general to have little effect on the behaviour of the emission other than mediating the atom number, as such it was left at -9 MHz detuned. The geometry of the emission and trapping fields is indicated in Fig. 5.1. To characterise the emission, detection setups could be placed at either end of the MOT's major axis where the emission exits the cell and magnetic field regions.

5.3 CHARACTERISATION

Initially attempts were made to characterise the properties of the emission as it was an unexpected occurrence, given that the atoms were in a steady state trapping condition. This process highlighted a number of interesting features which are outlined below.

5.3.1 Steady state time dependence

The MOT can be run in a steady state regime, corresponding to the static loading phase depicted in Fig. 3.5. When running in this mode, the optical detunings and

magnetic fields are kept constant. Running the atoms in the steady state condition, the emission effectively runs continuously, as the main mechanism for loss is collisions with thermal atoms. As the MOT is running as per normal, the loss rate is already in equilibrium with the loading rate. Unlike [50] we do not require the trapping beams to be run at a lower power and can run the MOT continuously as would generally be done for achieving maximum atom number. A photodetector is placed at one side of the emission to measure the optical power as a function of time. Fig. 5.2 shows the acquired signals for two different trapping detunings where the magnetic field gradient is held constant at 6 Gs/cm. As is shown, the emission will generally oscillate within some characteristic frequency range dependent somewhat on the detuning of the trapping frequency. This frequency will change spontaneously over the course of a run, indicating sensitivity characteristic of the symmetry breaking observed in other works.

To determine the optical frequency of the light being emitted, the beat note between the emission and a reference beam was observed. Through this, it was determined that the emitted light is the same frequency as the trapping light to within ≤ 1 MHz, limited by the sampling rate of the photodetector. Additionally, by beating the emission directly with the trapping light we also find that the emission and trapping light are coherent.

5.3.2 Compression time dependence

While the majority of the MOT experimental cycle is static loading, the transient compression sequence varies the parameters in an attempt to reach a higher density. We observe, during this sequence, a dependence of the emission power on the atomic density. As mentioned previously, during an experimental run, it is typical to perform a transient compression sequence. Due to the symmetry of the trap the MOT is compressed towards the region of highest density. During this compression we observed the intensity of the emission increasing to a plateau, corresponding to the maximum density achievable. Fig. 5.3 shows the increase in emission intensity along with the time dependence of the magnetic field and repump frequency. The last 10 ms of the 20 ms compression run is shown as this is where the emission starts to appear past some critical density for this choice of trapping detuning. The plateau can be explained as the atom number saturates, due to competing effects of atom loss due to the decrease in repump efficiency and increased magnetic field potentials. At the end of the compression sequence, the atom loss rate is higher than that of the steady state regime. This is due to the fact that the final control values in the transient sequence, are suboptimal for trapping atoms. Thus, the MOT will tend towards the new equilibrium condition, imposed by these new control values. It can also be seen from Fig. 5.3 that the emission continues to oscillate as per the steady state regime. However in this case it is possible to observe the spontaneous jump from different characteristic frequencies.

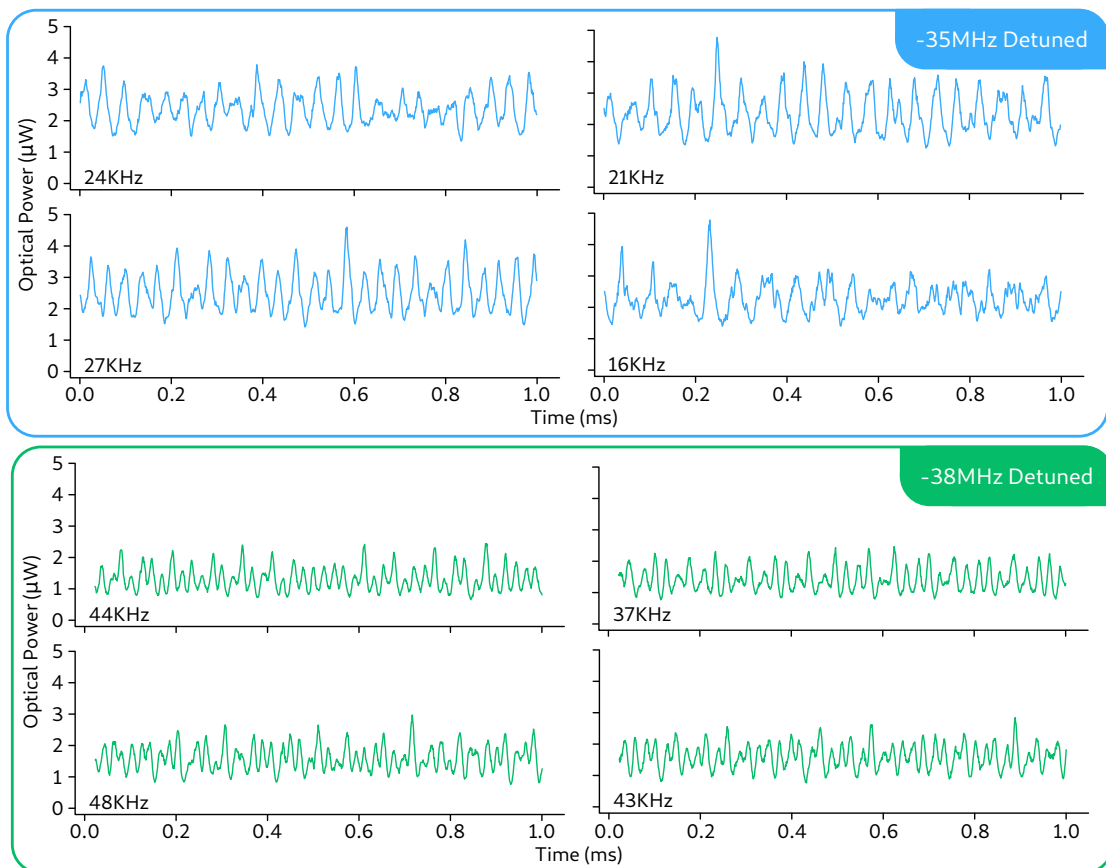


Figure 5.2: Observed time dependence of the atomic emission for two different detunings. The dominant oscillation frequency is given in the inset of each plot. A larger detuning also corresponds to a lower steady state atomic density.

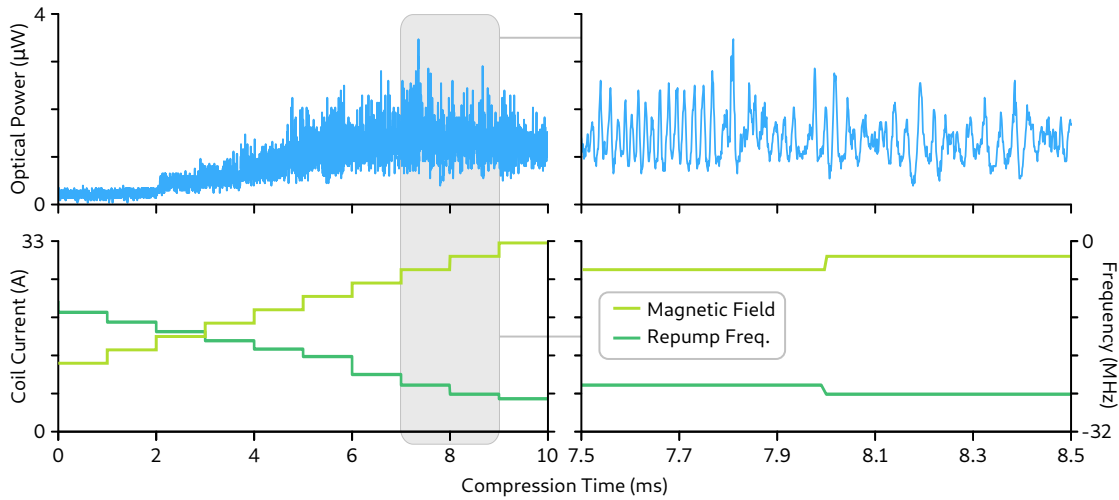


Figure 5.3: Measurement of the emission power showing the dependence on atom number during a transient compression sequence. The top row shows the evolution of the emission optical power over time. The bottom row shows the transient compression sequence parameters where the magnetic field is ramped from 6 Gs/cm to 28 Gs/cm and the repump frequency is detuned to induce a temporal dark SPOT. The highlighted area is expanded in the second column where the characteristic oscillation is once again observed. The compression sequence lasts for 20 ms however only the final 10ms is shown as the emission does not start until a critical density is reached for this lower trapping detuning of -32 MHz. The time dependence of the density for this sequence is depicted in Fig. 5.4.

Fig. 5.4 shows how the density and atom number changes during the compression sequence. We expect the relative density of the MOT to increase at the expense of total atom number. In the region where the emission originates the density is increasing. The critical density seen in Fig. 5.3, corresponding to the sudden appearance of the emission at 2 ms, occurs at 12 ms in Fig. 5.4.

5.3.3 Polarisation

The polarisation of the emission can be measured by construction of a basic polarimeter such as that shown in Fig. 5.5. By rotating the angle of the polariser, one can trace out a trajectory on the polarisation ellipse which can be fit to the equation adapted from [54] for when the fast axis of the quarter waveplate is kept at 0 degrees. Fig. 5.6 shows the measured intensity variation as a function of α , which can then be fit to determine the Stokes parameters. Due to the changing nature of the emission which jumps spontaneously between spatial patterns determining the polarisation is somewhat inaccurate. The large errors associated with each measurement are a direct result of the oscillations that occur during the steady state emission. Due to this, the polarisation is also changing faster than we are able to measure it. From this we can only draw the conclusion that the polarisation is indeed shifting as the emission drives its own instabilities. This conclusion is supported by observations using a CCD after a PBS, which show unstable areas of polarisation localised in the emission.

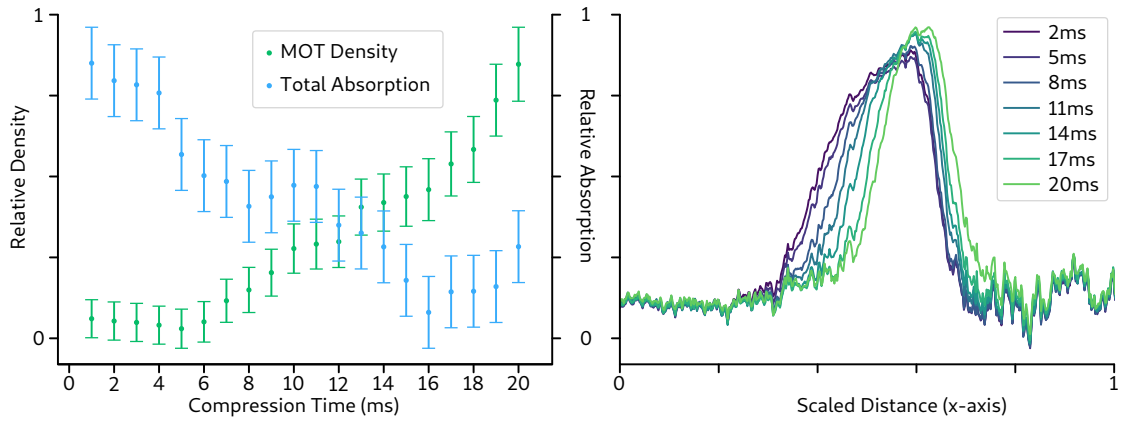


Figure 5.4: MOT density and absorption as a function of compression time. Here the total absorption is not a strict measure of atom number but is $\propto N$. The right plot shows the evolution of the spatial distribution of the MOT as the transient compression sequence continues. Here, absorption images have been integrated to show the distribution of the ensemble along the x-axis (see Fig. 5.1). The shift towards the right of the distribution is a result of imbalanced trapping beam powers due to the retro-reflecting geometry. As expected the density increases leading to a narrow spatial distribution, which is radially symmetric around the z-axis, similar to a cigar.




Figure 5.5: Polarimeter setup for determining the Stokes parameters. Either α or β may be swept and the intensity measured. This does necessitate the ability to measure faster than physical changes in the system that may manifest separate polarisations.

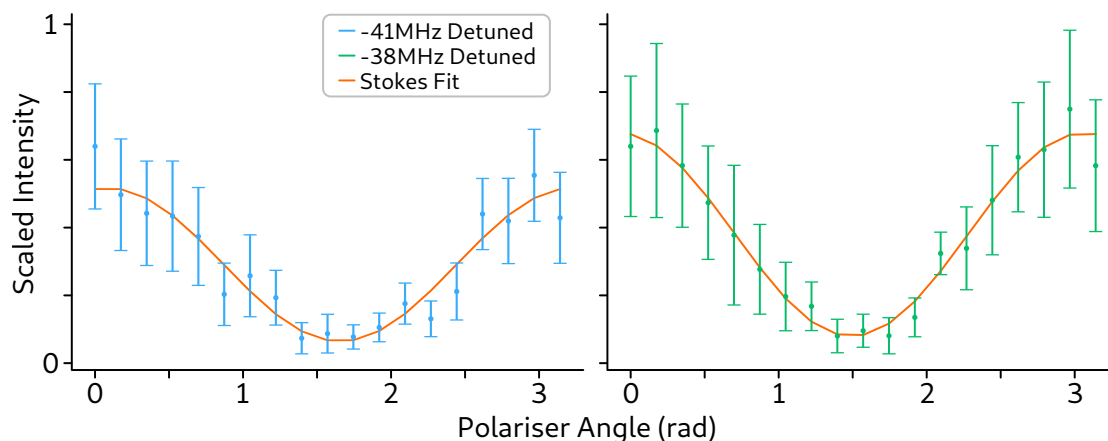


Figure 5.6: Change in measured intensity as a function of polariser angle. From these measurements an approximate polarisation can be obtained by fitting the Stokes parameters. However the dynamic nature of the emission precludes an accurate measurement as evidenced by the large errors associated with each measurement.

5.3.4 Feedback

As it had been observed that the emission was coherent and exhibited the characteristics of a non-linear unstable system, a natural question to ask may be what affects the strength of this emission. A particularly interesting phenomena that can be observed is a gain and feedback type response when retro reflecting one side of the emission back through the MOT. By placing a CCD at one end and a mirror at the other, the maximum amount of overlap between the emitted mode and retro reflected mode can be achieved. However the emitted modes are rapidly divergent, meaning that to collect all of the reflected mode into the detector would require a lens of a substantial diameter. This comes with a technical issue however, as a larger lens would collect more of the scattered light from the MOT as a result of the laser cooling process. This would in turn decrease the signal to noise ratio. The emitted modes themselves are poorly defined, which also precludes the use of a pinhole setup and, as seen in Sec. 5.3.3, have a non-static or poorly-defined polarisation.

The observed pulsing behaviour of the emission, with a retro reflected mode, is shown in Fig. 5.7 for a trapping detuning of -35 MHz. All of the traces displayed are a result of running the MOT in the steady state regime and letting the emission freely evolve. We observe a kind of superradiant effect where the emission will self reinforce to produce a steady pulsing behaviour. The total optical power of these traces can be as high as 1.8 ± 0.4 times that of the emission without retro-reflection. However, as can be observed from the trace of row 3 of Fig. 5.7, the maximum instantaneous optical power observed can be $\sim 3 - 4$ times that of the non-retro reflected case. We expect that the total optical power of the emission should be higher, given we are not collecting the full retro reflected mode due the divergence of the emission. An additional feature is that during periods of stable oscillation, not only is the maximum optical power




Figure 5.7: Measured oscillation patterns observed with a photodetector upon retro reflection of the emission from one side of the MOT. It can be seen that the retro reflected mode helps to reinforce the oscillations seen previously. The trapping detuning used is -35MHz which can be directly compared to Fig. 5.2 without retro reflection. The respective periodogram is given for each trace in the second column, often multiple spectral components are present which are distinct from harmonics.




Figure 5.8: (a)-(d) Far field profiles of the emitted modes from the atomic ensemble. The intensity is measured relative to the highest emitted intensity. The mode structure shown in (a) corresponds to the high intensity pulsing from Fig. 5.7. The trap detuning for (a)-(d) is constant at -35 MHz, with each image representing a different oscillatory behaviour. We have observed similar patterns are present for the full range of detunings at which this phenomena occurs.

higher, but the troughs in-between the pulses is also lower, suggesting that the atoms are collectively emitting and suppressing.

5.3.5 Far field profile

By measuring the far field profile it is possible to characterise the optical modes that make up the emission. Fig. 5.8 shows some typical spatial modes that have been observed during the free evolution of the emission with retro reflection. Each pattern also exhibits different oscillatory behaviour and different intensities. As the OD of the MOT for these experiments is high ($\sim 500 - 600$) distortion of the modes due to lensing by the ensemble is necessarily observed. The spatial patterns exhibit a distinct resemblance to the observed spatial patterns in the previous works which involved a pump beam creating a spatial organised structure [51–53]. The authors propose that the atoms self organise into bunched density gratings (see Sec. 5.3.6). Similar to previous works we also observe the spontaneous switching between modes that one would expect to occur as a result of a spatial reorganisation of the atomic ensemble. While previous works noted that this can occur on the order of $50 \mu\text{s}$, we observe a much faster time scale of $\leq 2 \mu\text{s}$ as limited by the camera shutter speed. As our temperature is ~ 2 orders of magnitude higher than previous work, we expect our velocity to be 1 order of magnitude higher. This would correspondingly lower the characteristic time for an atom to travel between adjacent lattice sites of bunching sites, to on the order of what we observed to be the switching time for these spatial mode jumps.

5.3.6 Physical description

Given the nature of the spatial pattern evolution and the collective radiative effects observed in the previous sections we can conclude that the observed effect is a result of a spontaneous self organisation of the atoms into bunched density gratings. The

preferential scattering and formation can be thought of as being a result of non-linear atom-light effects as a result of the bunching, which introduces additional non-linear terms into the susceptibility [53]. Thus we end up with two main effects to describe the observed phenomena, preferential scattering into an axial mode and wave mixing effects to generate patterns from the bunching induced non-linearity.

The main difference between this effect and previous results is that the system will spontaneously begin and moderate this behaviour in the absence of an external pump beam. In our case the pump beam is a result of collective preferential scattering into the axial mode. Further evidence of this is that for a fixed detuning we observe that by compressing the MOT along the axial direction using the capping coils to reduce the aspect ratio of the atomic ensemble we can effectively turn-off this phenomena. This preferential scattering can then be coherently enhanced to produce the amplification and suppression effects seen in Fig. 5.7. Currently the conjecture for describing the frequency of the intensity oscillations is the motion of the atoms within the trapping potential modified by the preferentially scattered light. The complex nature of the emitted mode, which is ill-defined, precludes a simple calculation of this frequency. However we would expect this to be on the order of \sim kHz which occupies the correct order of magnitude for the observed oscillation frequencies.

The far field patterns observed in Fig. 5.8 are effectively the same as seen in previous work [51–53], which have been shown to be a result of wave mixing effects within the ensemble induced via the density grating. The reduction in fidelity between our results and previous work is likely due to the effective pump in our system having an ill-defined spatial mode and polarisation which will introduce additional perturbations and instabilities into the imposed density grating. Additionally lensing due to the high OD will perturb the modes.

5.4 CONCLUSIONS AND OUTLOOK

The unexpected appearance of a coherent beam emitted from the cold atomic ensemble presented an opportunity to study a unique phenomenon. While similar effects have been observed in the literature, the present work demonstrates the first observation of this effect without the use of an additional pump beam, instead finding this effect to be self driven from preferential scattering in the atomic ensemble. This preferential scattering again leads to the spatial self organisation of the atoms into an effective density grating. Following this we also observe the far field profiles that one would expect from wave mixing processes.

There are however still a number of questions that remain unanswered to be addressed in future work. Due to the ill-defined nature of the pump beam a better option than analytical treatment is an empirical measurement to determine the behaviour and structure of the atomic bunching. This could be achieved by probing vibrational resonances [55]. Such a measurement and characterisation will help to confirm the

present conjecture regarding the observed intensity oscillations. Additionally it may be prudent to perform more accurate polarisation measurements by post-selecting the mode shape observed on the CCDs with the polarisation measurements. This capability is unattainable with the current CCDs due to the fast switching of the modes. Further characterisations could also be undertaken to determine the frequency and polarisation of different spatial elements within the emission, such as the wave mixing components. Additional investigation into this phenomena may help to further the understanding of cold atom dynamics and present new opportunities for coherent atom-light interactions which may be useful for quantum information processes.

6

QUANTUM MEMORIES

This chapter will introduce the foundations required for understanding quantum memories and the research conducted in relation to quantum memories in the present work.

6.1 QUANTUM COMPUTING AND MEMORIES

The world has recently seen the rise of an information economy, where services and wealth creation can be based entirely on the processing of information. The advent and technological rise of computing technology and commercial electronics was of course the main contributor to this in the last four decades, along with distributed communication and information networks such as the internet. Now the backbone of our society is heavily dependent on these technologies and services, following what could be described as the path of least resistance towards scalability.

A natural question that should arise is what is the follow up to this technology, what possibly could compete with this in an information age? With the success of Von Neumann type architectures [56] it is hard to imagine another type of computing technology that might follow. While additional architectures have been proposed such as that of biomimetic technologies [57–59], computing architectures that mimic biological processes, the next step that has received the most attention is that of quantum computing. Proposed quantum computing architectures themselves are in general no different to the Von Neumann model of computing, generally containing a processor and some form of storage and interconnects. The main advantage touted by proponents of this technology is of course the promise of superior computing power [60], leveraging the unique properties of quantum mechanical systems such as entanglement and superposition to efficiently compute the answers to common problems [61, 62]. While the work on quantum algorithms¹ continues, one of the most useful examples that is often flagged as the “killer-app” is Shor’s algorithm [63] for efficiently finding the prime factors of a given integer. Of course advances in such a problem have the potential to reshape the encryption and secure communication landscape, however for the most part this can be easily mitigated with so called quantum hardened encryption [64], which adds additional transforms and hashes that cannot be efficiently computed by a quantum computer. Nevertheless there remains a good case for quantum computing such as complex calculations in quantum chemistry [65] that can enable drug discovery or determine the solution to the nitrogen fixing

¹ quantumalgorithmzoo.org has a large showcase of useful algorithms for the interested reader.

problem that could enable efficient and scalable means of food production. For a more detailed review on quantum computing, I direct the reader towards the review by Steane [61] or Nielsen and Chuang's book [62] for a good introduction.

Quantum computation may refer to information processing alone, but often this term is taken to mean not only the information processing aspect, but instead includes by assumption the entire architecture involved. Consider for example a standard personal computer [56], be it desktop, laptop or hand held device. Whatever the device may be it will necessarily spend a large amount of its time doing nothing. The nature of these devices is that they are not always required to be doing something, so a large percentage of their runtime will be spent with the CPU at idle, which is to say waiting for instructions. However even when idle information still needs to be stored. At any given time the CPU will make calls to the on-board RAM and potentially the permanent storage such as HDDs or SSDs. Memory is indeed a large part of the complex operations undertaken in information processing. Take for example the case of a simulation, where the result of one calculation depends on another, or perhaps several. It is often the case that these results may not be available at the same time for a variety of reasons, thus they must somehow be stored. For quantum computing, one might naively assume that a storage operation could happen as it does now: with some copy operation into RAM or storage. However the naivety of this view comes from neglecting the no-cloning theorem present in quantum mechanics [66]. Effectively the no-cloning theorem prohibits the cloning of an unknown quantum state. Since the whole point of quantum computing is to encode our information on the quantum states of our chosen architecture, this will effectively prohibit us from copying our state into some form of memory, without either destroying much of the information we have encoded or already knowing the outcome of the computation. Another way of rephrasing this argument is simply, quantum information cannot be transmitted over a classical channel efficiently, i.e. any classical memory. Thus it follows that to perform quantum computation in the general sense we will of course need quantum memories as well.

An additional application of quantum memories presents itself in the realm of communications. In recent history there has been much controversy regarding the right to privacy among the members of nation states, with unsolicited spying having been uncovered in several nations, both domestic and international. Regardless of an individuals philosophical views on privacy rights, there will always remain a case for the secure transmission of data between two points, be it financial, medical or military. The search for a provably secure method of transmission was first answered in the seminal paper describing the BB84 protocol [67, 68]. This is the first example of a quantum key distribution protocol for provably secure communications based on non-commuting variables, to create a one-time pad encryption scheme. This is of course all well and good but the more practical question remains, how does one send quantum information between two (or more) parties? Often light is used as the transmission medium [69] however if the parties are separated by continents or even

line of sight, this process will require some form of repeater. This is due to the fact that optical fibres over long distances will reduce the fidelity of transmitted states such that distances greater than 100km are no longer feasible. As noted previously one cannot use a classical repeater and instead must resort to a quantum repeater. Each repeater in some sense requires a quantum memory to facilitate the required operations of a quantum repeater such as entanglement swapping [70]. Thus to construct a full quantum network, quantum memories are a must.

6.2 QUANTUM MEMORY THEORY

In Ch. 2 we introduced the atom-light theory relevant for understanding laser cooling and coherent interactions. I will now address the relevant theory for understanding atom-light interactions in the context of atomic ensembles which is presented in the present work.

6.2.1 Collective operators

Generally up until this point we have considered single emitters or atoms interacting with some optical field. In the case of a MOT we require a description of a collective excitation as our ensemble of two level atoms interacts an optical field. Firstly we define our atomic operators $\hat{\sigma}_{ij}$ to represent the atomic coherence between the $|i\rangle$ and $|j\rangle$ states. We wish to assume the ensemble has a high enough density such that we can treat it as a continuous distribution. In doing so we can consider slicing our ensemble into thin slices along our propagation direction (z-direction) where each slice contains $N_z \gg 1$ atoms. This allows us to define the collective atomic operators as

$$\hat{\sigma}_{\mu\mu}(z, t) = \frac{1}{N_z} \sum_{i=1}^{N_z} \hat{\sigma}_{\mu\mu}^i(t) \quad (6.1)$$

$$\hat{\sigma}_{eg}(z, t) = \frac{1}{N_z} \sum_{i=1}^{N_z} \hat{\sigma}_{eg}^i(t) e^{-i\omega_{eg}(t-z_i/c)}. \quad (6.2)$$

We can now describe the behaviour of the ensemble over some length $[0, L]$ in the z-direction.

6.2.2 Extension to Λ -system

For quantum memory operations it will be advantageous to work with a three level system. The reasons for this are simply a matter of memory performance. If one tries to store a coherent excitation on the excited state of a given atom then the maximum memory lifetime one can expect to achieve is dictated by the excited state lifetime Γ .

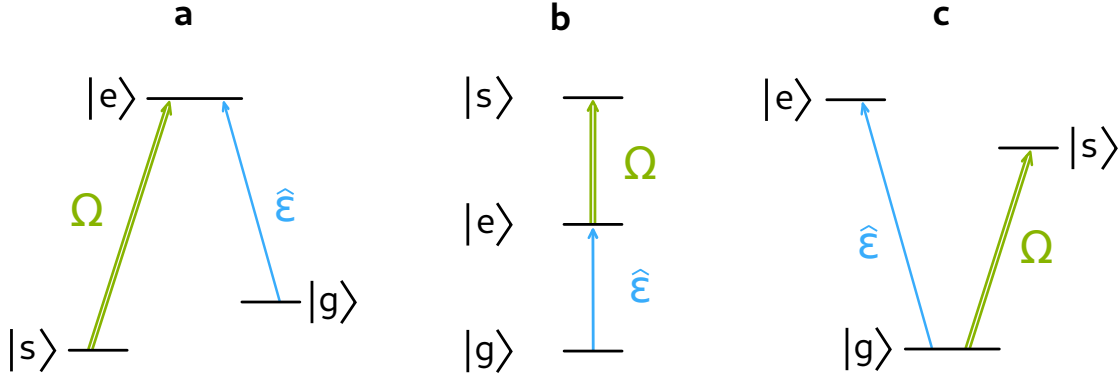


Figure 6.1: Depictions of various three level schemes. For the present work we are most concerned with (a) the Λ -scheme, however there also exists the (b) Ξ and (c) V configurations.

Thus for our purposes it is useful to consider how one might go about storing on a long lived ground state. A number of 3 level schemes exist as shown in Fig. 6.1 however we will primarily deal with the Λ scheme shown in Fig. 6.1a.

Suppose we have a 3-level atom as we did in Sec. 2.4.4 with states $|g\rangle$, $|e\rangle$ and $|s\rangle$ constituting a Λ -type scheme. As shown in [71] one can derive the collective operators for a three level system as

$$\hat{\sigma}_{\mu\mu}(z, t) = \frac{1}{N_z} \sum_{i=1}^{N_z} \hat{\sigma}_{\mu\mu}^i(t) \quad (6.3)$$

$$\hat{\sigma}_{eg}(z, t) = \frac{1}{N_z} \sum_{i=1}^{N_z} \hat{\sigma}_{eg}^i(t) e^{-i\omega_{eg}(t-z_i/c)} \quad (6.4)$$

$$\hat{\sigma}_{es}(z, t) = \frac{1}{N_z} \sum_{i=1}^{N_z} \hat{\sigma}_{es}^i(t) e^{-i\omega_{es}(t-z_i/c)} \quad (6.5)$$

$$\hat{\sigma}_{sg}(z, t) = \frac{1}{N_z} \sum_{i=1}^{N_z} \hat{\sigma}_{sg}^i(t) e^{-i(\omega_{eg}-\omega_{es})(t-z_i/c)}. \quad (6.6)$$

In general we consider a weak probe $\hat{\mathcal{E}}$ with frequency ω_p that couples the states $|e\rangle$ and $|g\rangle$, and a strong classical control field Ω_c with frequency ω_c that couples the states $|e\rangle$ and $|s\rangle$. We will assume that the control field is large compared to the probe and thus spatially uniform across the ensemble, allowing us to approximate the problem as 1 dimensional.

A full treatment of the derivation for the equation of motion of each of these operators can be found in [72–74], however for brevity we will simply present them:

$$\partial_t \hat{\sigma}_{gg} = -ig\hat{\mathcal{E}}\hat{\sigma}_{eg} + ig\hat{\mathcal{E}}^\dagger\hat{\sigma}_{ge} + \frac{\Gamma}{2}\hat{\sigma}_{ee} + \Gamma_{coll}(\hat{\sigma}_{ss} - \hat{\sigma}_{gg}) \quad (6.7)$$

$$\partial_t \hat{\sigma}_{ss} = -ig\Omega_c\hat{\sigma}_{es} + ig\Omega_c^*\hat{\sigma}_{se} + \frac{\Gamma}{2}\hat{\sigma}_{ee} + \Gamma_{coll}(\hat{\sigma}_{gg} - \hat{\sigma}_{ss}) \quad (6.8)$$

$$\partial_t \hat{\sigma}_{ee} = ig\hat{\mathcal{E}}\hat{\sigma}_{eg} + ig\Omega_c\hat{\sigma}_{es} - ig\hat{\mathcal{E}}^\dagger\hat{\sigma}_{ge} - ig\Omega_c^*\hat{\sigma}_{se} - \Gamma\hat{\sigma}_{ee} \quad (6.9)$$

$$\partial_t \hat{\sigma}_{es} = i\Omega_c^*(\hat{\sigma}_{ee} - \hat{\sigma}_{ss}) + ig\hat{\mathcal{E}}^\dagger\hat{\sigma}_{gs} - (\gamma_{es} - i(\Delta - \delta))\hat{\sigma}_{es} \quad (6.10)$$

$$\partial_t \hat{\sigma}_{ge} = ig\hat{\mathcal{E}}(\hat{\sigma}_{gg} - \hat{\sigma}_{ee}) + i\Omega_c\hat{\sigma}_{gs} - (\gamma_{ge} + i\Delta)\hat{\sigma}_{ge} \quad (6.11)$$

$$\partial_t \hat{\sigma}_{gs} = i\Omega_c^*\hat{\sigma}_{ge} - ig\hat{\mathcal{E}}\hat{\sigma}_{es} - (\gamma_{gs} + i\delta)\hat{\sigma}_{gs}, \quad (6.12)$$

where Γ_{coll} is a decay term corresponding to transitions induced by collisions, g is the atom-light coupling strength, γ_{ij} are relaxation rates for the ij coherences due to dephasing and collisions, $\Delta = \omega_{es} - \omega_c$ is the one photon detuning and $\delta = \omega_{eg} - \omega_p - \Delta$ is the two photon detuning. The convenience we have employed here is to work in the weak coupling regime. This allows us to take a similar approach as in Sec. 2.4 where we worked with a bright classical field, while making the substitution

$$\Omega_c \rightarrow g\hat{\mathcal{E}}. \quad (6.13)$$

In this way $\hat{\mathcal{E}}$ still may be a quantised field. Additionally we can make some further assumptions to simplify these equations of motion. In the case where the probe field is weak and the control field is strong we assume that the steady state population is effectively all in the ground state $|g\rangle$, which is to say

$$\hat{\sigma}_{gg} \approx 1 \quad \hat{\sigma}_{ee} \approx 0 \quad \hat{\sigma}_{ss} \approx 0 \quad \hat{\sigma}_{es} \approx 0. \quad (6.14)$$

In this case we may simplify Eqs. 6.7-6.12 to only

$$\partial_t \hat{\sigma}_{ge} = ig\hat{\mathcal{E}} + i\Omega_c\hat{\sigma}_{gs} - (\gamma_{ge} + i\Delta)\hat{\sigma}_{ge} \quad (6.15)$$

$$\partial_t \hat{\sigma}_{gs} = i\Omega_c^*\hat{\sigma}_{ge} - (\gamma_{gs} + i\delta)\hat{\sigma}_{gs}, \quad (6.16)$$

which correspond to the atomic polarisation and atomic coherence respectively. A common notation within the literature is to refer to the atomic coherence as a spin-wave. We now have a set of equations describing the evolution of the atomic state, however we must also include a term for the propagating probe field as the interaction is spatially distributed amongst the extent of the ensemble. Taking the equation for the propagation of light in an atomic ensemble with N atoms (assuming a uniform distribution of atoms)

$$(\partial_t + c\partial_z)\hat{\mathcal{E}}(z, t) = igN\hat{\sigma}_{ge}(z, t), \quad (6.17)$$

and combining with Eqs. 6.15-6.16 we now have a series of coupled differential equations that describe our interaction between the two fields and the atom. The control field is assumed to uniformly illuminate the ensemble such that it is not treated with a propagation equation. It is convenient to define two new operators [71]

$$\hat{P}(z, t) = \sqrt{N}\hat{\sigma}_{ge} \quad \hat{S}(z, t) = \sqrt{N}\hat{\sigma}_{gs}, \quad (6.18)$$

referred to as the polarisation operator and spin-wave operator respectively. We can reformulate our set of coupled equations using these operators as

$$(\partial_t + c\partial_z)\hat{\mathcal{E}} = ig\sqrt{N}\hat{P} \quad (6.19)$$

$$\partial_t\hat{P} = ig\sqrt{N}\hat{\mathcal{E}} + i\Omega_c\hat{S} - (\gamma_{ge} + i\Delta)\hat{P} \quad (6.20)$$

$$\partial_t\hat{S} = i\Omega_c^*\hat{P} - (\gamma_{gs} + i\delta)\hat{S}, \quad (6.21)$$

where we see the the collective enhancement that gives rise to a coupling strength up to $g\sqrt{N}$. Additionally in contrast to [71] we are neglecting the noise terms for simplicity. It is also useful in this case to define our on resonance OD as $d = g^2NL/\Gamma c$. We will also rescale our optical field by $\sqrt{c/\Gamma}$ and transition to a moving frame $z' = z + ct$. Finally we replace the spatial coordinate with a normalised position ζ which spans $[0, 1]$. Written in terms of OD, d , our equations now become

$$\partial_\zeta\hat{\mathcal{E}} = i\sqrt{d}\hat{P} \quad (6.22)$$

$$\partial_t\hat{P} = i\sqrt{d}\Gamma\hat{\mathcal{E}} + i\Omega_c\hat{S} - (\gamma_{ge} + i\Delta)\hat{P} \quad (6.23)$$

$$\partial_t\hat{S} = i\Omega_c^*\hat{P} - (\gamma_{gs} + i\delta)\hat{S}. \quad (6.24)$$

6.2.3 Raman transitions

The name generally given to the type of interaction we are describing is a Raman transition or Raman line, where a coherent absorption process occurs between the two ground states via the excited state. This type of coherent interaction is useful as previously mentioned for storing light on a long lived ground state, but is also used as a spectroscopic technique in various fields. A common approximation to make to further simplify the equations of motion is known as the Raman limit. In this case we wish to assume the one photon detuning is much larger than the natural line width of the excited state which is to say, $\Delta \gg \Gamma$. If we assume that the excited state mostly remains

empty in this case and does not vary dramatically, we may adiabatically eliminate the excited state such that $\partial_t \hat{P} = 0$. In this case we can now find an expression for \hat{P} as

$$\hat{P} = \frac{\sqrt{d}\Gamma\hat{\mathcal{E}} + \Omega_c\hat{S}}{\Delta - i\gamma_{ge}}. \quad (6.25)$$

Substituting this into our equations of motion we now have the simplified equations of motion

$$\partial_t \hat{S} = i\sqrt{d}\Gamma \frac{\Omega_c^*}{\Delta'} \hat{\mathcal{E}} - (\gamma' + i\delta') \hat{S} \quad (6.26)$$

$$\partial_{\xi} \hat{\mathcal{E}} = \frac{id\Gamma}{\Delta'} \hat{\mathcal{E}} + \frac{i\sqrt{d}\Omega_c}{\Delta'} \hat{S}, \quad (6.27)$$

where the following substitutions have been made

$$\Delta' = \frac{\Delta^2 + \gamma_{ge}^2}{\Delta + i\gamma_{ge}} \quad \gamma' = \gamma_{gs} + \frac{\gamma_{ge}|\Omega_c|^2}{\Delta^2 + \gamma_{ge}^2} \quad \delta' = \delta - \frac{\Delta|\Omega_c|^2}{\Delta^2 + \gamma_{ge}^2}. \quad (6.28)$$

It is clear from Eq. 6.28 that there is now additional loss from control field scattering as a result of power broadening of the excited state linewidth. This will contribute directly to increasing spin relaxation, to the detriment of memory lifetime. Additionally the two photon detuning is modified by the AC-Stark shift which will shift the energy of the atomic levels. While in general this is accounted for automatically experimentally by tuning to the resonance of a given transition for a given control field power, it can also be useful in modifying energy levels to engineer useful interactions between incident optical fields and atoms in this state.

In the absence of dispersion, the equations of motion we have derived here behave as if the atom was a two level system. For the cases we are interested in this will be generally true, however it is not the case for all systems where excited states may be close enough to have a significant effect.

6.3 EFFICACY AND LOSS

For a quantum memory to be of any practical use there are a number of criteria that a candidate technology must satisfy in some fashion to obtain the coveted title of “useful quantum memory”, as after all a fibre loop is by some definitions a quantum memory. The following serves as an overview of quantities considered to be useful for a quantum memory to have.

6.3.1 Efficiency

Potentially the most simple measure of a memories usefulness is efficiency, which is defined as a ratio between the input and output energies. As a quantum memory cannot perform a copy operation due to the no-cloning theorem, the general operation of a quantum memory can be simplified as storage and retrieval of the input state. During this process, any losses (such as atomic motion or incoherent scattering in our case) will contribute to a decrease in memory efficiency. The lower baseline is 50% as this is what is required to beat the no-cloning theorem, i.e. a memory scheme could be replaced by attempting to clone the state. This can be framed in terms of security, which is to say that if the memory is performing below this threshold then it is impossible to tell whether the stored state is the true state or a cloned copy of the state which has been stored elsewhere. Formally the efficiency of a quantum memory can be stated

$$\eta = \frac{|E_R|}{|E_I|}, \quad (6.29)$$

where E_R and E_I correspond to the energy of the retrieved and input states respectively.

6.3.2 Lifetime

The lifetime of a quantum memory is another feature which has a classical analog, however the definitions may vary in certain cases. In all cases it is measured with respect to the decay of the output state with respect to the input state, be it 50%, $1/e$ or $1/e^2$. The relevant lifetime will of course depend on the desired application, for example repeaters may require anywhere from milliseconds to hundreds of seconds depending on the application [67]. In general, similar to efficiency, it is better to have this quantity be as large as possible.

6.3.3 Bandwidth

The bandwidth of a quantum memory is important when considering the practical aspects of storing pulses of light. Typically a quantum memory based on some ensemble will have a bandwidth limited by the linewidth of the atomic transitions, for our cold atom case this will be on the order of MHz. However single photons generated from quantum dots (QDs) or conventional spontaneous parametric down conversion (SPDC) will have bandwidths on the order of GHz to THz [75]. Photon generation in atomic ensembles can mediate this problem to some degree by providing bandwidth and frequency matched photons however the brightness and efficiency of these sources is generally far lower than that of SPDC and QDs. Additionally high bandwidth is useful for things like pulse multiplexing and dual mode storage.

6.3.4 Fidelity

Fidelity is the first more “quantum” quantity that we attribute to memory performance. Simply measuring the efficiency of an operation does not ensure that the quantum nature of the encoded information has survived the process. For example, if we imagine an operation where an optical pulse is incident on a detector and we measure this pulse and then create a pulse with the same temporal shape and energy then in effect our pseudo-operation could have unity efficiency. However our fidelity in this case should be small as we destroyed all the encoded quantum information during our readout process. One way of formulating fidelity is to consider the input and output state of our quantum memory. In effect we wish to measure the overlap between the two states, thus we could define fidelity as

$$\mathcal{F} = |\langle \psi_{out} | \psi_{in} \rangle|^2, \quad (6.30)$$

where ψ represents the respective output and input states. We have assumed in this definition that our states are both pure states. However, using the wave function is not ideal as it is not experimentally accessible. In the more general case where we may have a mixed state we instead can define our fidelity as

$$\mathcal{F} = \left(\text{Tr} \sqrt{\sqrt{\rho_{out}} \rho_{in} \sqrt{\rho_{out}}} \right)^2, \quad (6.31)$$

where the ρ 's denote the output and input states represented as density matrices [76]. As the density matrix may be reconstructed using some form of tomography process, this measure of fidelity is more experimentally accessible.

6.3.5 Multi-mode capacity

Storage of a single mode in a quantum memory makes any practical device limited in capability. To create a more general purpose device it is advantageous to have the ability to store multiple modes independently such that information may be retrieved and stored independently in each mode. If a device can store information in more than one mode it is said to be multi-mode. One of the most basic multi-mode structures is the storage of different spatial modes in a quantum memory [77]. However one may also devise schemes for splitting the available bandwidth into discrete spectral components, or temporally multiplexing incoming pulses.

6.3.6 Additional considerations

The quantities listed above are the most commonly reported within literature and arguably the most useful to consider. However, there are a number of other features that one may consider in the search for a useful quantum memory.

INTEGRATION To Integrate with existing technologies and infrastructures, such as optical fibre backbones that network the globe, there are a few practical considerations. One such consideration is the suitability of the operating wavelength. Telecom networks are generally designed to operate at an optical wavelength of 1550 nm. This is due to the inherently low loss ($\sim 0.2\text{dB/km}$ [78]) of optical fibres at this wavelength which simplifies long haul transmission infrastructure. It is conceivable that any practical memory will either need to operate at this wavelength, or alternatively at some wavelength that can be up/down-converted to this wavelength [79].

MINIATURISATION The commercial success of optical fibre networks is due in part to the “plug-and-play” nature of this technology as it has matured. Simply put, racks can be hot-swapped which contain amplifiers, multiplexers, lasers and other relevant technology to run a communication network. It is necessary that any competing memory architecture has at least the potential capability to be miniaturised at some point in the future. This limits the usefulness of techniques that will rely on large and energy intensive devices such as cryostats, dilution fridges or vacuum chambers. With miniaturisation generally comes a decrease in manufacturing cost as well which will at some point play a role. Additionally quantum repeater setups will inevitably make it onto space based communication systems. When housed in a satellite the cost of implanting a technology is directly related to the size and weight of a given component, thus miniaturisation is an important factor for future quantum network infrastructure.

IN-MEMORY OPERATIONS An interesting prospect which is not a blanket criteria for all quantum memories, is the ability to perform operations within a memory. This extends the usefulness of a quantum memory beyond that of simple storage operations, introducing an ability to perform light matter interactions in a controlled manner. It has been shown that linear optical quantum computing can be performed using only quantum memories [80] and other strategies such as stationary light provide a promising avenue for performing cross phase modulation in-memory [81, 82].

CONTROLLABLE While initially the qualifier “controllable” seems a blatantly obvious condition for any useful quantum memory, the actual intent is more subtle. For example while an atomic frequency comb is in essence controllable, in that the frequency comb is written into the ensemble and a state is stored, the time at which the state is retrieved is fixed by this initial operation. This means the time at which the state is returned is not necessarily reconfigurable unless using an additional spin shelving techniques [83]. In general this will introduce additional losses and reduce memory efficiency. The

debate regarding the exact definition of configurable is outside the scope of the present work, but it suffices to say that the more configurable and controllable a memory is, the more useful it will be in general. For example fully on-demand quantum memories can be used in efficient distributed quantum computing schemes [84].

6.4 QUANTUM MEMORY LITERATURE

The definition of a quantum memory can be quite broad depending on the nomenclature used. For the purposes of this review I am mainly interested in optical quantum memories that one might intend for use in quantum interconnects and repeaters. I will try to categorise this into similar groups based on the architecture and protocols involved.

6.4.1 Delay lines and cavities

As mentioned before, one of the simplest approaches to implementing a quantum memory is using a delay line, usually in the form of a fibre spool. Referring back to our measures of efficacy, it is not difficult to see that only some of these criteria are not met. Efficiency of delay lines will depend on whether it is a free space line such as a Harriott delay line or fibre spool. In general working at telecom wavelength 1550 nm, for a 15 km delay line we would expect 3 dB of loss corresponding to an efficiency of $\sim 50\%$ with a 70 μs delay.

The use of high-Q photonic-crystal cavities has also been explored as an avenue for compact tunable delays by modifying the quality factor, Q , via a transverse pumping beam [85]. A delay of up to 1.45 ns was demonstrated, however the efficiency of this method quickly drops off with higher delay times.

The lack of controllability limits the application of either process to quantum memories and repeater technologies. In general a more on demand approach is required to facilitate useful memory operations, however such processes have found applications in heralding single photon sources. These delays are also generally useful for on-demand entanglement creation as demonstrated in [86].

6.4.2 Slow light and EIT

Slow light is the name generally given to any effect that results in the slowing of an optical pulse propagating in some medium. An example is the controllable slow light effect that is observed to accompany EIT. The idea is to use the non-linear effect that is present during the two-photon absorption process, for example in a Λ -type system. When the two-photon detuning is close to resonance the probe absorption is greatly reduced which is also accompanied by an extreme change in dispersion in this transparency window [87]. This dispersion modification in turn reduces the group

velocity of the propagating pulse, effectively compressing the pulse such that a large pulse may fit entirely within an atomic ensemble. With the pulse contained within the medium (generally some atomic ensemble) the control field intensities can be adiabatically ramped to 0, effectively closing the transparency window and transferring the optical energy into atomic excitation. Later when one wishes to retrieve the stored pulse, the control field is adiabatically increased again, opening the transparency window and releasing the stored pulse [88].

The initial theoretical groundwork for EIT was laid by Harris et al. in 1990 [89] showing that the application of this strong coupling field one could obtain a transparency window. The first experimental demonstration was followed shortly in 1991 by Bollinger et al. in Strontium [90]. It was then realised that EIT had the potential to operate as a protocol for quantum memories in atomic ensembles [91, 92], which has been studied and demonstrated in both warm [93, 94] and cold atoms [95, 96]. Notably in a recent result Hsiao et al. demonstrated a retrieval efficiency >90% in cold Cs atoms [97]. EIT has also demonstrated high fidelity with polarisation degrees of freedom [98, 99], spatial light structures [100] and squeezed light [101]. Storage times for EIT can vary from μs and ms in atomic vapours (warm and cold) through to seconds in solid state memories [102, 103].

6.4.3 Atomic frequency comb

As noted in section 6.3.6 atomic frequency combs (AFCs) have been proposed as a potential scheme for implementing a quantum memory [104]. This scheme requires at least a two level system with states $|g\rangle$ and $|e\rangle$, however it is useful to have an additional metastable ground state $|s\rangle$ for shelving purposes. The AFC quantum memory scheme is depicted in Fig. 6.2. Regarding the $|g\rangle \rightarrow |e\rangle$ transition we wish to have a narrow linewidth, corresponding to long coherence times, but with large inhomogeneous broadening. That is to say that the collective absorption profile of the ensemble is larger than an individual atom. The idea is then to spectrally shape this broadened transition such that for some detuning, Δ , atoms that are integer multiples of Δ remain in $|g\rangle$, while all other atoms are shelved in some auxiliary state that will not take part in this protocol. This shelving operation may be done via some optical pumping process. If the width of the frequency comb is denoted Γ , a pulse that is spectrally larger than Δ but smaller than Γ may be completely absorbed, provided the spectral density of the ensemble is high enough. The collective state will de-phase, later re-phasing after a time $2\pi/\Delta$ which corresponds to a collective re-emission. By shelving this collective excitation on the state $|s\rangle$ with an additional coupling beam, this storage time can be extended to $T_s + 2\pi/\Delta$, where T_s is limited by the coherence time of the spin-wave.

While the initial observation of this photon echo effect occurred as early as the 1980s [105] it wasn't until much later that the use of AFC as a quantum memory protocol was proposed [104] and subsequently demonstrated [106, 107]. Currently the highest efficiency achieved with AFC to date is 56% [83] using an impedance matched cavity

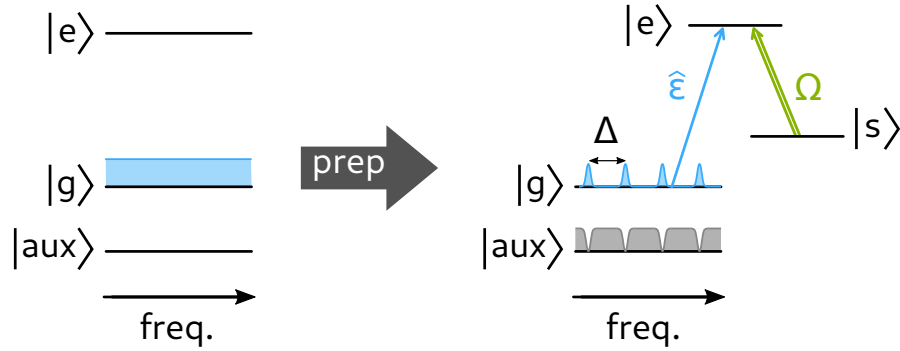


Figure 6.2: Preparation of an inhomogeneously broadened ensemble into an atomic frequency comb with comb spacing Δ . Once the atomic ensemble has been prepared a signal field may be stored and additionally shelved on a metastable state to increase the storage lifetime.

with a solid state crystal embedded inside. However the appeal of the AFC protocol is not so much the efficiency but rather the inherent multi-mode storage capabilities, with 64 single photon mode [69] and 1060 classical mode storage [108] having been demonstrated. Frequency multiplexing has also been demonstrated on AFC platforms [109, 110], as well as entanglement operations between telecom wavelength photons and AFC memories [111].

6.4.4 Raman memory

Raman scattering was first proposed as a quantum memory [112, 113] as a way to leverage the long lived Raman coherence with high efficiency. The situation is in some ways similar to EIT except that the desired outcome is that the probe light is fully absorbed and coherently stored on the metastable state via consecutive read/write pulses. An absorption window is opened at some far detuning as opposed to a transparency window. The bandwidth of the memory is effectively dynamically created by the control field by dressing the atomic states and providing a larger bandwidth for the probe field to couple to [114]. As shown in Sec. 6.2.3 the effective Raman dynamics operate at large detunings which is also advantageous to avoid inhomogeneous dynamics such as Doppler broadening. A disadvantage of this technique is the requirement for large control field intensities which introduces four wave mixing noise, however it has been demonstrated that this can be heavily suppressed with the addition of a surrounding cavity [115].

Raman memories have been extensively demonstrated in warm atomic vapours [114, 116] with a recent paper demonstrating an efficiency as high as 82% [117]. Additionally by using the vibrational states of Hydrogen molecules [118] and phonon modes in diamond [119], THz bandwidths have also been demonstrated. Raman memories have also been demonstrated in cold atom ensembles [120–122] as they provide a better platform for efficiency with less Doppler broadening.

A simple way to improve the efficiency of the Raman memory is to use a backward retrieval protocol [114, 123]. The retrieval of the stored signal field relies on the symmetric reversal of the storage dynamics, as such the optimal spin-wave has an initially small amplitude with the largest amplitude at the exit interface. This situation cannot be created with a forward retrieval geometry providing a straight forward argument for using a backward retrieval geometry. Generally this type of scheme is not used due to the experimental overhead required to implement such a setup, however it has demonstrated success in improving the efficiency of Raman processes [124].

6.4.5 Controlled reversible inhomogeneous broadening

In AFC we see that the scheme takes advantage of inhomogeneous broadening to produce a frequency bandwidth that may be used for the creation of a frequency comb. An alternate approach, that allows on-demand retrieval of stored photons, is to have some mechanism which renders the inhomogeneous broadening controlled and reversible. So called controlled reversible inhomogeneous broadening (CRIB) schemes were born from the idea of using photon echoes as a quantum memory. If we consider an ensemble absorbing a signal probe which has been inhomogeneously broadened, then we expect different parts of the ensemble to accumulate phase $e^{-i\delta_j t}$ where δ_j is the detuning of the j^{th} atom due to this broadening [125]. The aim is to rephase the atomic ensemble such that when all the phases of the atoms are equal, a re-emission of the stored signal is recovered. One way this is possible is to invert the detunings experienced by the ensemble, such that the ensemble now rephases symmetrically after dephasing for a time t_p with re-emission occurring at $2t_p$. Other methods rely on the use of π -pulses to induce the rephasing of the population, however this requires a complete population inversion which can be difficult to achieve efficiently.

6.4.5.1 Gradient echo memory

An example of a CRIB type memory, that uses a spatially dependent field to create the inhomogeneous broadening of the atomic ensemble, is gradient echo memory (GEM). This scheme is depicted in Fig. 6.3. As the broadening is spatial, different frequency components of the stored pulse will be stored at different spatial locations, thus the spin-wave observed is simply the Fourier transform of the input signal. Generally the gradient imparted will be monotonic as this has the additional bonus that when the gradient is flipped, re-emitted light is only resonant with the current spatial location, thus upon exiting the ensemble, the signal experience little to no re-absorption as the signal is off resonant with other spatial locations. This allows GEM schemes to obtain high efficiencies. Additionally the simplified control of the broadening (usually magnetic fields or AC-Stark shifts) and spatial storage of frequency components allows a GEM type scheme to perform spectral processing.

Initially GEM was demonstrated in a solid state memory (yttrium with praseodymium dopants) using linearly dependent electric fields, yielding an efficiency of 26% limited

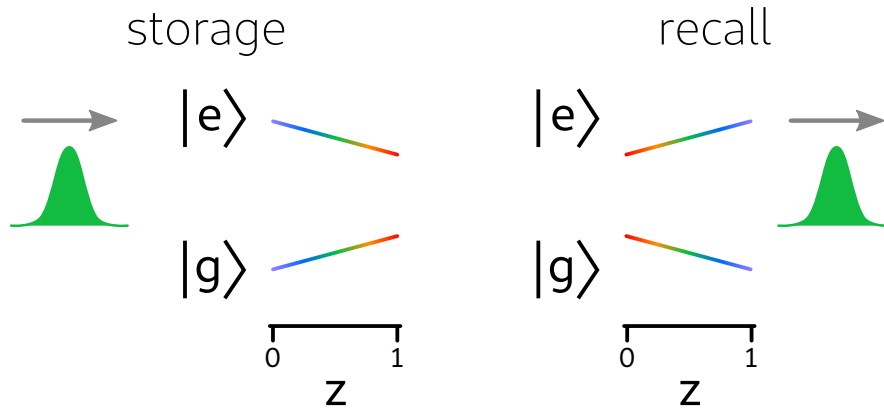


Figure 6.3: Storage of an input optical pulse in an ensemble with a spatially dependent inhomogeneous broadening. After some storage time the sign of the broadening is reversed and the stored pulse is re-emitted. The relative frequency of the transition for a given spatial location is depicted by the colouring.

by OD [126]. Following this GEM was demonstrated in warm ensembles [127, 128] and cold ensembles [16, 129] using Rb, both with efficiency up to 87% using magnetic field gradients as the CRIB mechanism. Using this high efficiency technique GEM has also demonstrated higher order spatial-mode storage [77] and dual-rail storage capabilities [130].

6.4.5.2 Revival of silenced echo

Another demonstrated CRIB scheme is revival of silenced echo (ROSE), which seeks to reduce the complexity of the ensemble preparation such as with GEM [131]. In this scheme after initial absorption, the ensemble is re-phased with a π -pulse, however the photon echo is suppressed. A second π -pulse then re-phases the ensemble a second time, yielding signal retrieval. While this scheme indeed simplifies the preparation step, it does increase the fidelity required in the π -pulse re-phasing as this occurs multiple times. Efficiencies up to 40% have been demonstrated with this protocol [132].

6.4.6 Summary

From the literature review above it is clear that quantum memories are still an active area of research and development with a rich array of potential candidates, from schemes to platforms, available at present. Fig. 6.4 shows an overview of the main schemes presented with the maximally achieved figure for each metric. There is still much remaining work to create a useful quantum memory in terms of the additional metrics discussed in Sec. 6.3.6, for example miniaturisation. However the last decade has seen much progress including the first realisation of a memory efficiency greater than 90%. In the present work we will be mainly concerned with improvements to the Raman memory scheme, using the Rb MOT presented in previous chapters.

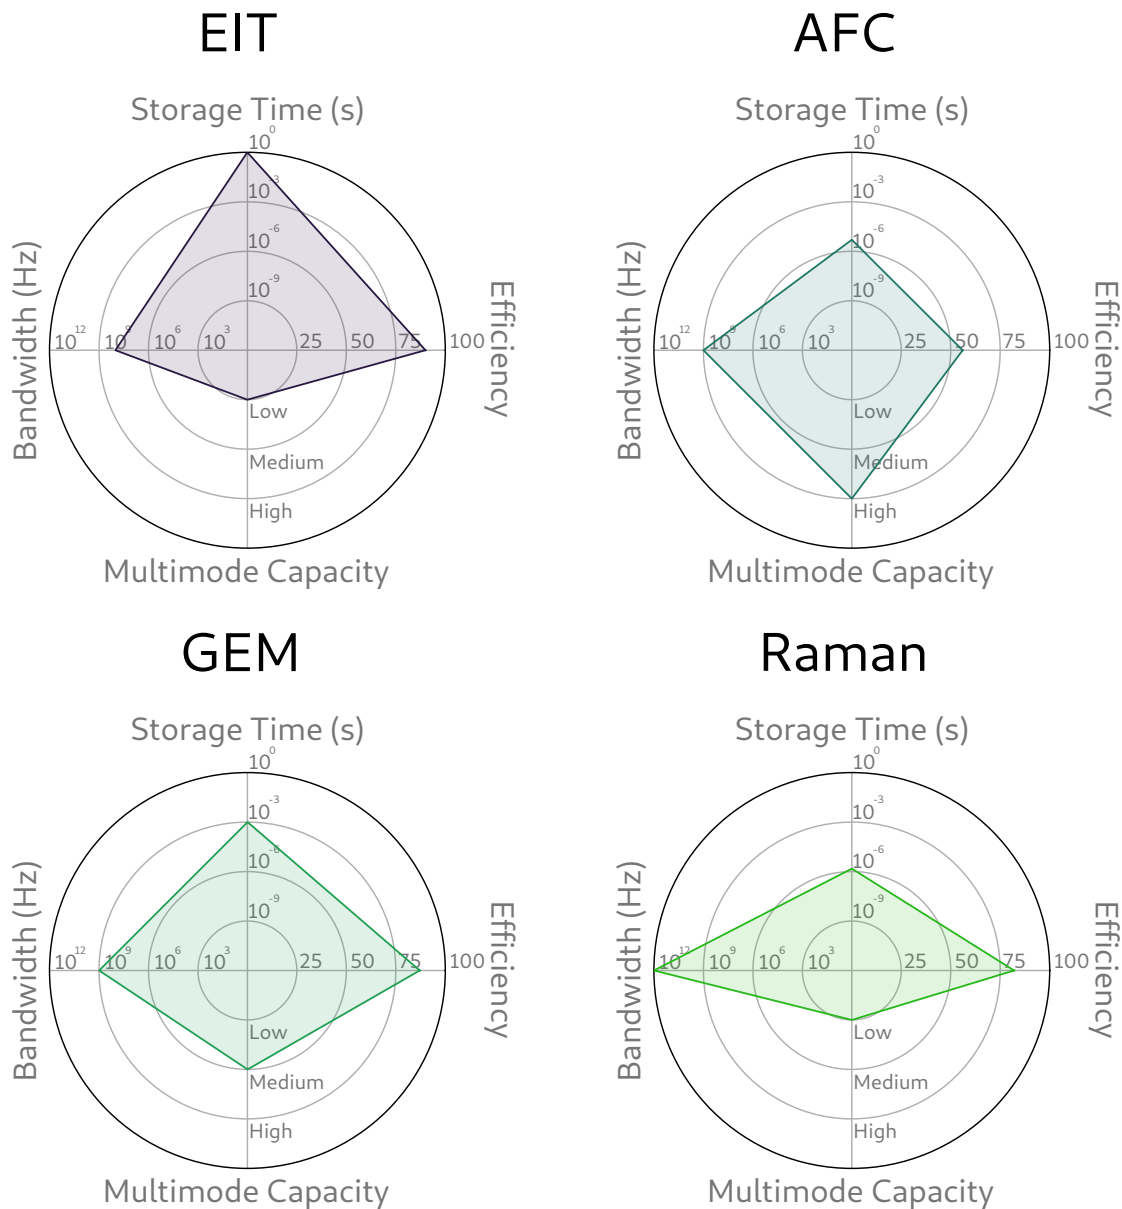


Figure 6.4: Summary of the maximum results for four parameters of interest for each of the main quantum memory schemes currently being pursued. Note that this is independent of platform and as such concurrent access to the maximum of each metric may not necessarily be possible. However this does serve as a useful guide to illustrate the successes of each scheme.

7

BACKWARD RAMAN MEMORY

In the previous chapters I have covered the basic theoretical framework for describing the operation of a quantum memory. This chapter will cover the theoretical basis for the operation of a backward Raman memory and report on the experimental demonstration of this memory.

7.1 FROM SCATTERING TO MEMORY

In Sec. 6.2.3 we briefly introduced the notion of a Raman transition or Raman scattering. For a 3 level Λ -scheme, this inelastic scattering occurs via a virtual energy level which is detuned from the excited state by a frequency Δ as shown in Fig. 7.1a. By applying a far detuned optical field from the ground state to this virtual energy state, the emission of a Stokes photon (in the case where $|g\rangle$ has a lower energy than $|s\rangle$), the converse situation will lead to an anti-Stokes photons) can be affected. This emitted photon will have greater/lower energy corresponding to the energy difference between the $|g\rangle$ and $|s\rangle$ states, as one of the atoms in the ensemble has been transferred to the $|s\rangle$ state.

The emission of the Stokes/anti-Stokes photons can be enhanced by stimulating Raman scattering where an additional strong coupling field is used to stimulate this emission. The population shuffling that occurs as atoms are transferred to the $|s\rangle$ creates the collective atomic excitation known as a spin-wave that was outlined in Sec. 6.2.2. This process is coherent, the quantum nature of the information stored on the light¹ will be preserved when mapped onto the atomic excitation, with the converse also being true during a recall operation. It is now that we have arrived at the Raman memory scheme illustrated in Fig. 7.1. Here a weak signal field is mapped onto the spin-wave via a strong classical coupling field, using the described two photon resonance.

7.2 FORWARD RAMAN MEMORY

As discussed in Sec. 6.4.4 it is more common for Raman type experiments to involve a forward recall geometry generally due to experimental simplicity. We previously derived Eqs. 6.22 - 6.24 that describe the interaction of a 3-level Λ -scheme with a bright classical control field and a weak probe.

¹ The exact reasoning behind this can be discussed in the context of the Jaynes-Cummings model of atom-light interactions and the linearity of the operators involved. This is however outside the scope of the work at hand.

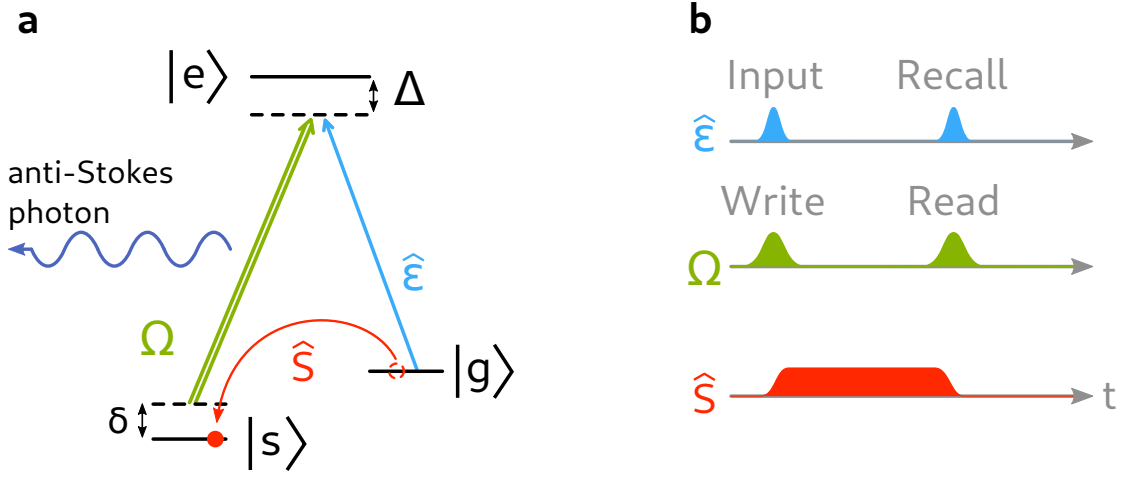


Figure 7.1: (a) Stimulated Raman scattering process with 2 optical fields applied to a 3-level system. The anti-Stokes photon has gained energy from the difference between the two ground state energies. The transition of an atom between the ground states can be used as a quantum memory protocol by considering the timings shown in (b). The \hat{S} operator is the spin-wave coherence that is generated as a result of this storage operation.

7.2.1 Memory protocol

The Raman memory protocol uses a read and write pulse applied to the atomic ensemble. A bright classical control field facilitates the use of the Raman absorption line, to store the incident probe light as an atomic spinwave. After absorption the control field is switched off until some time later at which point a read pulse is used to restore the stored probe. This approach is laid out diagrammatically in Fig. 7.1b.

It is useful to consider the dynamics of the absorption taking place as a result of the coherent interaction being driven by the presence of the control field. The susceptibility for a 3-level atom is given by

$$\chi(\omega) = \frac{2a_0}{k} \frac{i\Gamma}{\Gamma - i\Delta + |\Omega_c|^2(\gamma_{gs} - i\delta)^{-1}}, \quad (7.1)$$

where a_0 is the on resonance absorption coefficient and k is the wavevector [2]. If the one-photon detuning is swept we can map out the resonance and phase shift associated with the interaction of the fields with the 3-level atom. As expected, if $\Omega_c \rightarrow 0$ then we simply recover the Lorentzian absorption profile for a 2-level atom shown in Fig. 7.2a. The addition of the control field displays the EIT effect used for the slow light and two-photon resonant Raman memory schemes. As we increase the two-photon detuning we find the emergence of the narrow Raman absorption line that is used for memory purposes as shown in Fig. 7.2c. As noted before, the strong control field dresses the atomic states and controls the width of the absorption profile which is given by $\Gamma_R = \frac{\Omega_c^2 \Gamma}{4\Delta^2}$.

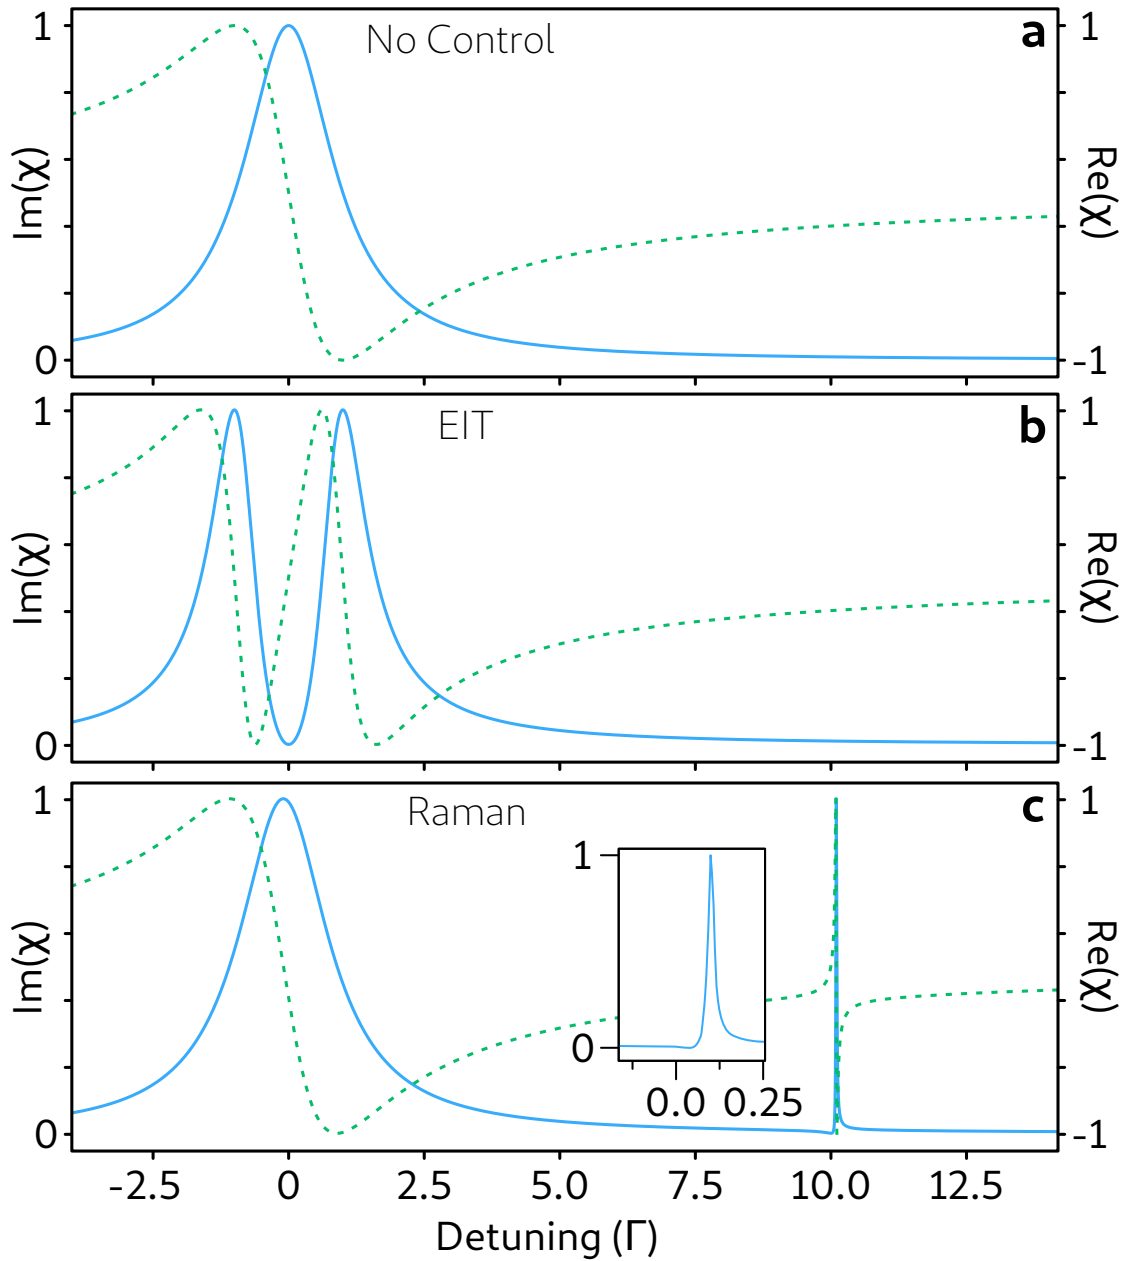


Figure 7.2: $\chi(\omega)$ for different probe detunings and control configurations. For each subplot the blue curve represents the imaginary (absorptive) part of χ and the green dashed curve gives the real (dispersive) part. (a) shows the case where $\Omega_c = 0$ which returns the Lorentzian profile for a 2-level atom. (b) shows the two photon resonance case where $\delta = 0$. Here the EIT window is observed. (c) shows the emergence of the narrow Raman absorption line where the two-photon resonance $\delta = 10\Gamma$. The inset shows the narrow absorption line starting at 10.0γ . For all plots the parameters used are $\Gamma = \Omega_c = 5$ and $\gamma_{gs} = 0$.

7.2.2 Numerical simulation

To study the operation of a Raman memory we can numerically solve these equations. While we have previously derived the Raman approximation, it is relevant to the current work to consider the full 3-level treatment as the experimental implementation in our system does not necessarily warrant the use of the far detuned approximation (for an ensemble this is $\Delta \gg d\Gamma$). The Raman memory can be demonstrated numerically by simulating the equations of motion (Eqs. 6.22 - 6.24). Further information regarding the simulation procedure and use of fruit fly algorithm can be found in Appendix B. I will first examine the case of operating on the narrow Raman line which results from a non-zero two-photon detuning ($\delta \neq 0$). The most simple case of turning off the control field is simply a “hard-off” or piecewise function. Fig. 7.3 shows the result of such a simulation where the input probe pulse is a Gaussian pulse of width Γ . The simulation is run on a system similar to the best performance we have observed in our MOT corresponding to an $OD = 1000$. The control field Rabi frequency is set to 29Γ , $\Delta = 100\Gamma$ and $\delta = 10\Gamma$. It can be seen from Fig. 7.3 during the write stage the probe pulse is mapped onto the atomic spin-wave, corresponding to the extinction of the control field. Here it remains during storage until the read pulse is applied which corresponds to the re-emission of light from the ensemble. It can be noted that rapid oscillations occur at the switching of the control field. This effect is due to the immediate extinction of the control field, which is non-physical. It effectively is a manifestation of the oscillating light-shift from the abrupt extinction of the bright field. Additionally in this simulation γ_{gs} is taken to be negligible ($\gamma_{gs} \approx 0$).

Clearly in this simulation there is much room for improvement as the efficiency is only $\approx 26\%$. Efficiency here can be gained from two improvements. Firstly the operation on the narrow Raman line at far detuning requires a large amount of control field power. While we could improve the recall by increasing the power of the control field, we can relax this requirement by operating in an intermediate scheme of two-photon resonance (i.e. $\delta = 0$), relaxing this power requirement. Fig. 7.4 shows a simulation of the two-photon resonant Raman memory with a hard control field, optimised for recall efficiency ($\approx 84\%$). The optimised Rabi frequency is found to be 10.5Γ , which is 87% less power than that of the far detuned memory. It can be noted however there still remains appreciable pulse distortion on the recalled signal.

The second improvement that can be made is to optimally shape the control field pulse to optimally map the optical excitation to a spin-wave coherence. Additionally this will also reduce scattering loss from applying the control field for longer than necessary. While such optimal shaping can be done in an analytical sense with some assumptions (see [133] for a good discussion) we can also solve this numerically similar to an online optimisation which will be thoroughly presented in Pt. iii. Details regarding the optimisation process can be found in Appendix B. Given the input probe field is a Gaussian generated from some laser source as this maps efficiently onto the atomic ensemble, the ideal control shape should also look like some Gaussian. The free parameters we wish to control are the Rabi frequency of the read and write pulses,

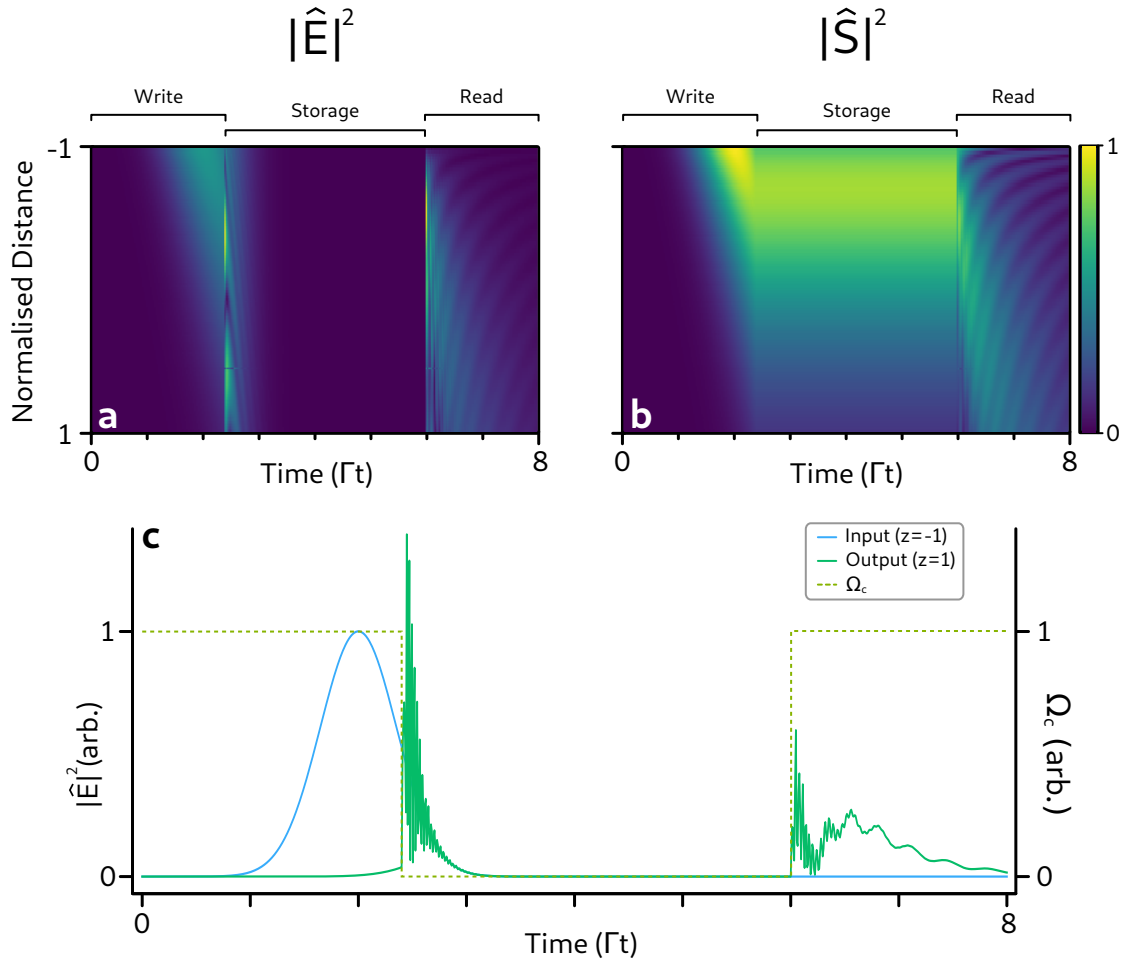


Figure 7.3: Numerical simulation of a Raman memory. Plots (a) and (b) show the amplitude of the electric field (probe) and spin-wave operators respectively. The atoms are assumed to extend uniformly over the normalised distance $-1 \rightarrow 1$. Intuitively it can be seen that the incident probe pulse is mapped onto the atomic spin-wave (write + storage) and converted back to a propagating field some time later (read). (c) shows the time trace that could be measured experimentally with a photo-detector. The recall efficiency is $\approx 26\%$.




Figure 7.4: Numerical simulation of a Raman memory operating on two-photon resonance. The control field has been optimally placed and the Rabi frequency optimally tuned to produce a maximum efficiency of $\approx 84\%$.

the initial input pulse timing and the width of the read and write pulses. As the read and write pulses should have the same width this provides a total of 4 parameters to optimise. Fig. 7.5 shows the results of this optimisation which yields a memory efficiency of $\approx 91\%$. Note that we have made a trade-off between efficiency and a minor amount of pulse distortion, corresponding to pulse compression. This however can be corrected by increasing the width of the read pulse. The Rabi frequency of the write and read pulses corresponds to 5.8Γ and 8.5Γ respectively. Some pulse distortion will always be present due to the fact that the re-phasing of the ensemble is driven by control field inducing free induction decay. Additionally efficiency will always be lost to incoherent absorption as the field propagates through the ensemble upon recall.

7.3 BACKWARD CONFIGURATION

It is possible to increase the efficiency of the Raman memory by adding an additional optical field and applying the read control pulse in the counter propagating direction. Due to the phase matching condition, it is possible to ensure the recalled signal is mapped onto a counter propagating coherence. As shown by [134] for the re-phasing operation to produce a complete time reversal process, it is also required that the detunings of the read and write pulses are anti-correlated. To include this additional field in the simulation, we need to re-write the equations of motion we had previously derived to include these terms. To simplify the numerical procedure it is advantageous to move to the adiabatic limit, such that \hat{P} is adiabatically eliminated. This is a valid approximation as it was observed in the previous simulations that \hat{P} remains negligibly populated, despite the closer detuning. Starting from Eqs. 6.26 - 6.28 we add the counter propagating fields to obtain

$$\partial_t \hat{S} = i\sqrt{d}\Gamma \frac{\Omega_+^*}{\Delta'_+} \hat{\mathcal{E}}_+ + i\sqrt{d}\Gamma \frac{\Omega_-^*}{\Delta'_-} \hat{\mathcal{E}}_- - (\gamma' + i\delta_{Stark}) \hat{S} \quad (7.2)$$

$$\partial_{\zeta} \hat{\mathcal{E}}_{\pm} = \pm \left(\frac{id\Gamma}{\Delta'_{\pm}} \hat{\mathcal{E}}_{\pm} + \frac{i\sqrt{d}\Omega_{\pm}}{\Delta'_{\pm}} \hat{S} \right), \quad (7.3)$$

where the + subscript indicates forward propagating fields and – similarly denotes the backward propagating fields. The third term in expression Eq. 7.2 has been simplified to include all the scattering terms into γ' . Working on two-photon resonance, δ' reduces to the total Stark shift from the control fields. Using this simplified model the backward and forward retrieval schemes can be compared. Fig. 7.6 shows a simulated direct comparison using this model which has been optimised for both schemes independently. The maximum efficiency for the forward retrieval is found to be $\approx 75\%$ which is lower than the previous simulation due to the inclusion of control field scattering and Stark shift terms. Additionally the efficiency is also reduced by incoherent absorption as the field propagates through the ensemble as a high OD of ≈ 1000 is used. It can be seen from Fig. 7.6a that the recalled pulse is distorted, as well as exhibiting a comb like

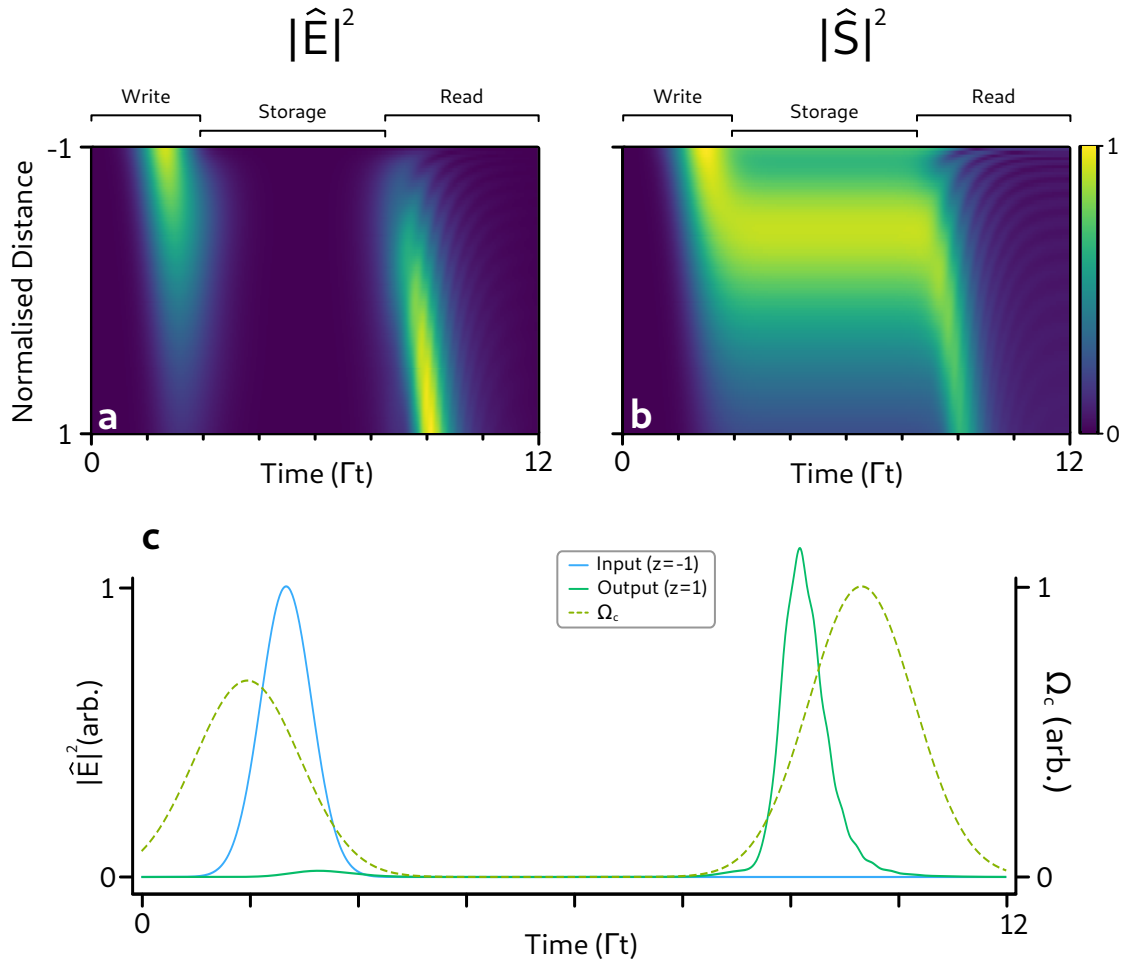


Figure 7.5: Numerical simulation of a Raman memory operating on two-photon resonance showing (a) the electric field operator amplitude, (b) the spin-wave operator magnitude and (c) the time trace of the input, recall and control fields. The control field has been optimally shaped as a series of two Gaussian pulses to give an efficiency of $\approx 91\%$. Note, the control fields have been normalised to the read pulse Rabi frequency.

structure from the free induction decay associated with the control field re-phasing the atomic coherence. These effects can be mitigated by employing the backward retrieval scheme shown in Fig. 7.6b. Using this scheme the absorption process is simply time reversed and the only pulse distortion that is observed is a compression due to the high intensity recall pulse. This can be mitigated as before by increasing the pulsewidth to account for this distortion. Using this scheme a maximum retrieval efficiency of 94% is observed. The relevant operators for each configuration are shown in Fig. 7.6c-h.

7.3.1 Phase matching

To maximise storage efficiency it is necessary to ensure that the overlap between the generated and optimal spin-waves is maximised. This is achieved for co-linear storage, however due to non-degeneracy of the ground states, there is necessarily a mismatch between the probe and control wave vectors. To maximise this overlap it is necessary to introduce an angle between the probe and control field pairs, θ , which will ensure the retrieved probe is co-linear as shown in Fig. 7.7. This angle will be modified by the presence of the large OD ensemble which introduces dispersion, however this can be corrected experimentally with careful alignment. Additionally atomic motion will reduce the retrieval efficiency if $2\pi/|\Delta k| < D$, where D is the distance over which the atoms travel during the storage time and Δk is the wave vector of the spin-wave. For a ^{87}Rb MOT with a temperature of $\approx 200 \mu\text{K}$ this limits the storage time to on the order of 280 ms.² However the more relevant time scale is the few milliseconds that it takes for the atoms to accelerate due to gravity out of the interaction volume. To ensure the stored and retrieved probe are co-linear, the atomic states must be chosen such that the storage state $|s\rangle$ is energetically lower than that of the initial ground state $|g\rangle$, otherwise only the control fields may be made co-linear.

7.4 EXPERIMENTAL IMPLEMENTATION

As shown in the prior sections, the backwards retrieval configuration has a number of advantages over that of a forward Raman memory. While most studies of Raman memories have been conducted in warm atomic vapours, which are advantageous for their experimental simplicity, there are a number of useful qualities cold atom systems exhibit, justifying the increased experimental complexity. Firstly the low temperatures reduce losses due to diffusion associated with phase matching, atom-atom collisions and broadening. Usually the OD associated with cold systems is comparatively lower than that of warm systems, however the cold atom system in the present work is capable of generating an OD on the order of 1000 on particular transitions. As such the cold atom experiment outlined in Ch. 4 is uniquely suited to perform such an

² Here I have assumed a ground state splitting of 6.8 GHz for ^{87}Rb .

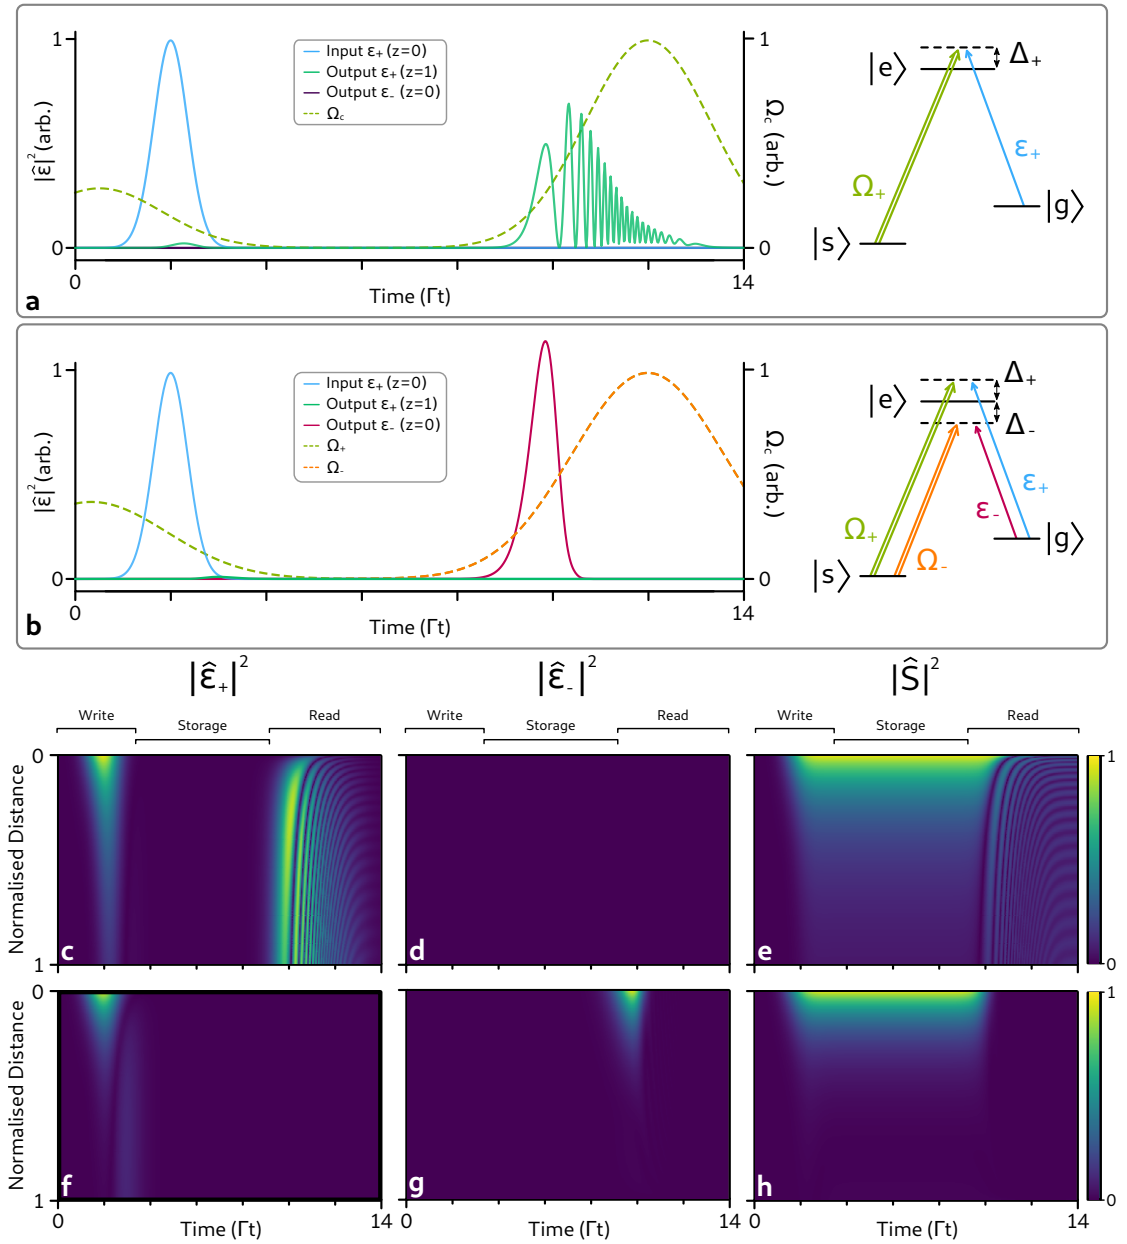


Figure 7.6: Numerical simulation of a Raman memory operating on two-photon resonance using the simplified two level model. The forward retrieval time trace is shown in (a) with a efficiency of 75%. (c) - (e) shows the magnitude of the operators. The backward retrieval time trace in (b) yields an efficiency of 94% with (f) - (h) showing the magnitude of the operators. The simulation parameters used are $OD = 1000$, $\Delta_{\pm} = \pm 33\Gamma$, $\delta = 0$ and $[\Omega_{write}, \Omega_{read}] = [3.31, 11.54]$ and $[4.0, 10.7]$ for the forward and backward retrieval respectively. In both cases the Gaussian read/write pulses were optimised for Rabi frequency, write input time and pulse bandwidth, with the bandwidth being kept symmetric for simplicity. The respective bandwidths were Γ for the input pulse, 1.8Γ and 2.2Γ for the forward and backward retrieval respectively.




Figure 7.7: Phase matching conditions required to satisfy the co-linear storage and retrieval of the probe fields. The plotted vectors correspond to the wave vectors of each respective field. The difference in their wave vectors determine the wave vector of the atomic spin-wave.

experiment. Here we will detail the experimental implementation of the backward retrieval configuration on the ^{87}Rb MOT.

7.4.1 Memory scheme

The memory is chosen to operate on the D_1 line due to the reduced complexity of the level structure. As the storage state must be energetically lower than that of the initial ground state we choose $|g\rangle = |5^2S_{1/2}, F = 2, m_f = 2\rangle$, $|s\rangle = |5^2S_{1/2}, F = 1, m_f = 0\rangle$ and $|e\rangle = |5^2P_{1/2}, F' = 1\rangle$ as this leverages the strongest Clebsch-Gordan coefficients. This also allows the polarisation filtering to be applied as the control and probe beams can be orthogonal circularly polarised. Additionally the choice of an edge m_f state allows for a high optical pumping efficiency.

7.4.2 Atom Preparation

The preparation of the atomic ensemble is performed according to Sec. 4.2.3, on the D_2 transition of ^{87}Rb . The trapping fields are red detuned from the $F = 2 \rightarrow F' = 3$ transition while the repump runs red detuned on the $F = 1 \rightarrow F' = 2$ resonance. The ensemble must be optically pumped into the $m_f = 2$ Zeeman sub-level required for the memory scheme. A 0.5 Gs bias field is applied along the storage (longitudinal) axis of the ensemble, which enforces the quantisation axis. An optical pumping beam, with intensity $\sim 0.7 \text{ mW/cm}^2$, is derived from the trapping field line which is σ_+ -polarised to pump atoms into the $|g\rangle$ state, while the repump field is applied to clear out the $F = 1$ level.

7.4.3 Optical Fields

A schematic of the optical paths is shown in Fig. 7.8. The laser light for the probe and control is generated by a MSquared SolsTiS titanium-sapphire laser which is locked to $|s\rangle \rightarrow |e\rangle$ transition on the D_1 line of ^{87}Rb using a saturated absorption lock. Two separate optical paths are used to produce the control and probe light.

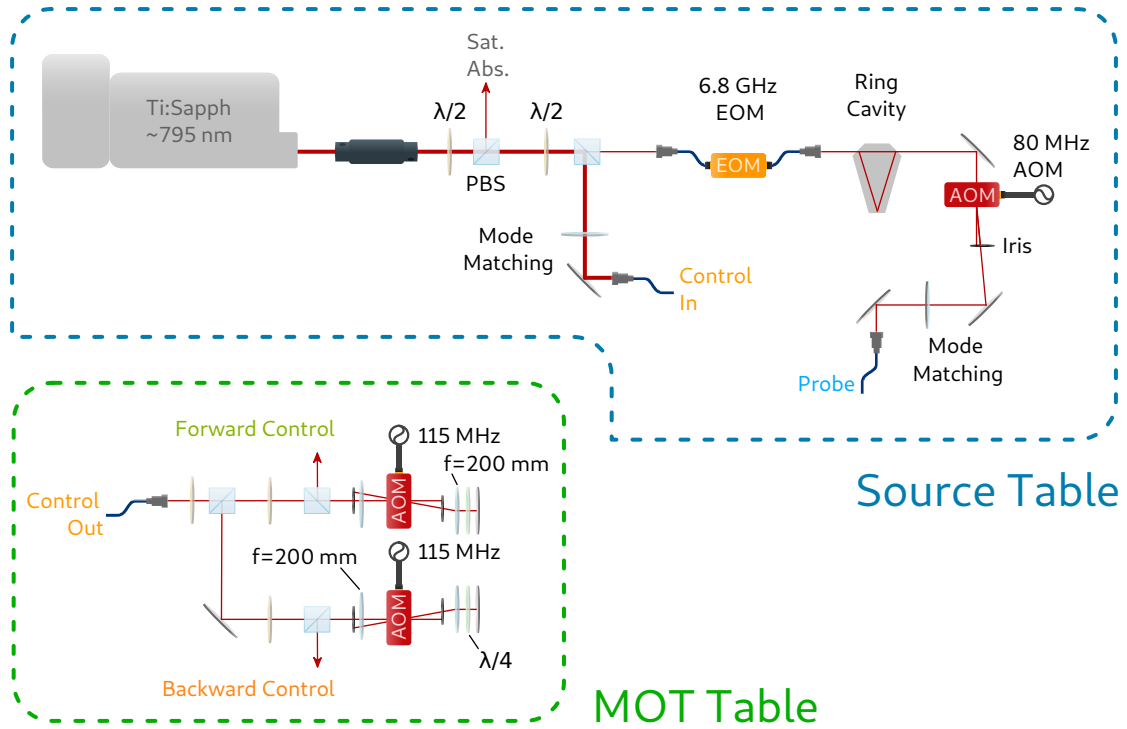


Figure 7.8: Optical setup for generating the correct frequencies used in the experimental implementation of the backward Raman memory. Beams are fibre coupled to the **MOT** table where the main vacuum chamber assembly is located.

The control fields are symmetrically detuned 230 MHz about the $|s\rangle \rightarrow |e\rangle$ transition, using two separate double passed acousto-optic modulators (AOMs) which select the anti-symmetric diffraction orders. The probe light needs to be resonant on the $|g\rangle \rightarrow |e\rangle$ transition. This is achieved on the second optical line where a fast electro-optic modulator (EOM) is used to generate 6.8 GHz side-bands. This modulated light is sent to a ring cavity which is used to filter to particular sidebands. The sidebands can be locked to using PDH locking, providing the ground state splitting shift in the probe frequency that is required. The difference in one-photon detuning is imparted via an additional AOM after the cavity. This is also used to temporally shape the probe field from continuous wave to a Gaussian temporal profile. Both the probe and control beams are fibre coupled into polarisation maintaining fibres, such that they can be placed on the main experimental optical table. Separation of the laser sources and main experimental table is crucial to remove major sources of magnetic and electrical interference such as GHz sources and Faraday isolators that may affect the operation of the memory. Due to constraints of the fibre EOM, the probe is generally on the order of 1 mW of optical power at the fibre output. The control beam is limited by the damage threshold of the optical fibres and is generally on the order of 200 mW at the fibre output, which is then divided into the forward and backward controls fields. By performing the frequency shifting of the control fields on the **MOT** table, phase noise between the control fields is minimised compared to separate fibres and any fluctuations in intensity due to temperature drift is common to both control fields.

7.4.4 Alignment and detection

Fig. 7.9 shows a schematic of the beam path and detection scheme involved. The ground state splitting of 6.8 GHz enforces an angle of $\theta = 6$ mrad between the control and probe beam to satisfy phase matching conditions for counter propagating fields. Experimentally this provides a challenge when it comes to both alignment and filtering of the control beam from the probe beam that we wish to measure. A folded beam path of 6 m was used to add an additional level of filtering (in concert with polarisation filtering). An unfortunate consequence of this was the introduction of interferometric instability from vibrational modes of the large posts used to elevate the beams to the MOT height. While this bears no consequence for the detection of a single run, it however increases the complexity of the alignment process, often requiring a large number of iterative steps to achieve sufficient alignment. This setup will drift over the course of an acquisition session and requires daily realignment. The optical pumping beam is injected with a small pick-off mirror in a similar way to the control fields and propagates along the axis of the MOT with a small angle. This angle is optimised to achieve maximum pumping efficiency. Beams are aligned to the atomic emission (see Ch. 5) using CCD cameras at distinct distances. Alignment to the emission enforces that the beams are incident on the most dense region of the MOT, post compression sequence, providing a higher fidelity alignment than side imaging. The final fine alignment is performed using Raman lines and absorption measurements.

As shown in Fig. 7.9, two calibrated continuous mode avalanche photo-detectors (APDs) were used on either end of the path length to measure the transmitted leakage and recall efficiencies. The dual detection setup allows for both the input and recall pulses to be measured within a single experimental run such that all losses involved are common to both modes. Incoherent absorption due to the large OD can be accounted for by measuring the difference between the forward probe with and without atoms present, subject to no coupling beams. Using this technique at a one-photon detuning of 230 MHz, loss due to incoherent absorption is $\sim 5\%$.

We carry out spatial, temporal and polarisation filtering to remove leakage of the control field, which is many orders of magnitude brighter than the probe. Polarisation filtering is achieved using a Glan-Taylor prism which provides ~ 30 dB of isolation. Additionally the probe beams are focused through a $150 \mu\text{m}$ pin-hole which gives approximately 90% transmission, with 15 dB of further isolation. A significant amount of control field leakage is visible on the APDs even after all this filtering. This can be finally removed temporally by placing a pick-off on 0th order of the normally blocked second-pass, of the control field AOMs. This provides a relevant timing for the control field which can be used to subtract the control field from the APD signals.

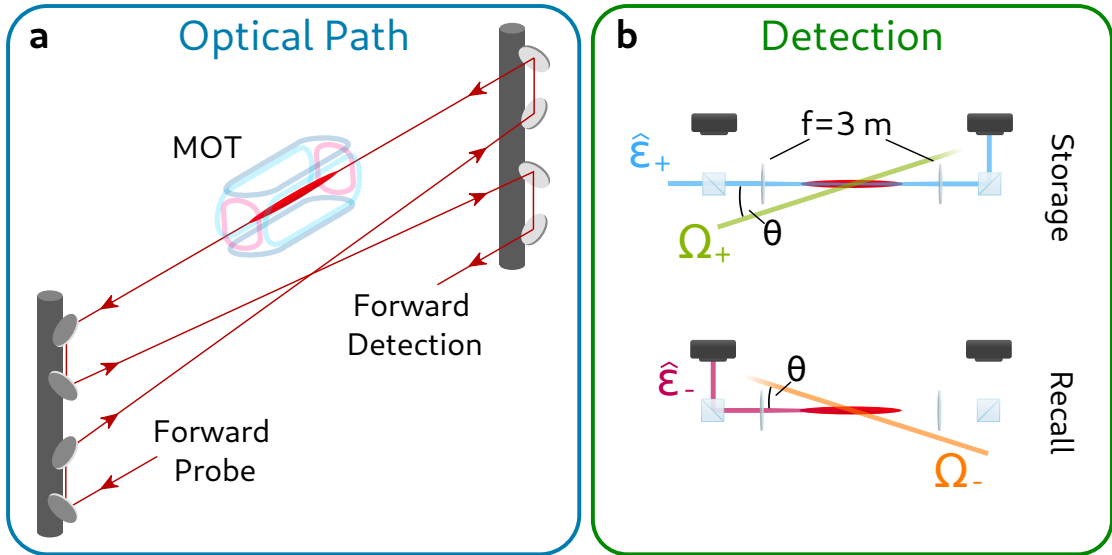


Figure 7.9: Simplified schematics of the optical path length and detection setup used for detecting the backward recalled probe. (a) shows the folded beam path used to elevate the optical fields to the MOT height. Phase matching and filtering requires that the total optical path length is 6m from when the probe and control are combined. (b) shows the detection scheme used during the quantum memory experiment. Note that for simplicity the angle is exaggerated and that normally the entire probe interaction region is illuminated by the control fields.

7.4.5 Machine learned compression and PGC sequence

The final efficiency of the memory is highly dependent on the final density of the atoms within the interaction volume. This is because we rely on the collective interaction of the atoms with the incident probe. To this end, the compression sequence was optimised via a Gaussian process (GP) optimisation method implemented using the M-LOOP package [135]. The compression sequence used is an optimised version of the temporal dark SPOT outlined in Sec. 3.3.5. The magnetic field and repump frequency ramps are parameterised by two parameters each, corresponding to the gradient and offset of a linear ramp of the form $x_0t + x_1$ which is sampled at 1 ms intervals for the 20 ms of ramp duration. By clipping the values of the ramp to within the physical limits and safe boundaries of the experimental setup, the GP learner can find the optimum ramp values. An additional 6 parameters were devoted to set the trapping detuning during a 1 ms PGC phase. Meanwhile the trapping and repump fields are applied without the magnetic trapping coils. Upon the completion of this stage the optical pumping sequence is carried out. The GP learner is used in an online optimisation context where feedback is received via an off resonant probe (~ -50 MHz) that is sent through the atomic ensemble. The absorption of this probe is measured and provided to the GP learner as a proxy for OD. This optimisation can be run daily to combat experimental fluctuations and can provide OD improvements of up to 20%. A full diagram of the loading and memory timings are shown in Fig. 7.10.




Figure 7.10: Relative timings of the experimental run for loading the atomic ensemble, compression, PGC, optical pumping and the experimental memory run. The dead time is included to allow the dissipation of eddy currents induced from fast switching of the magnetic field coils. Figure is adapted from [136].

7.5 EXPERIMENTAL RESULTS

By implementing the methods described above, a backwards retrieval Raman memory was successfully demonstrated. The OD was calibrated by fitting the Raman line shift to different control field powers. The OD was measured to be ≈ 500 during the time at which the Raman memory was operated.

7.5.1 Storage efficiency

Initially a Gaussian pulse, of pulse-width³ $5 \mu\text{s}$, is stored in the memory with the optimal control pulse-width corresponding to a Gaussian preceding the probe by $5 \mu\text{s}$ with a width of $9 \mu\text{s}$. This pulse width is consistent with the simulated optimum pulse results shown in Fig. 7.6. The efficiency of the memory operation fluctuates on a daily basis due to experimental conditions such as fluctuations in the Rubidium dispenser rate, magnetic environment and laser stability. The highest reproducible efficiency for a storage time of one pulse width was found to be $65 \pm 6\%$. The error associated with this measurement owes to the optical losses in the complicated optical path, which can fluctuate on the order of minutes. An example of the memory operation with a similar efficiency is shown in Fig. 7.11. Here the recalled pulse has been optimised experimentally to preserve the pulse-width of the input pulse and thus a high temporal overlap. At the time of the publication [136], these results represented

³ pulse-width here is defined as $1/e^2$.




Figure 7.11: Experimentally observed storage and backward recall of a $5 \mu\text{s}$ width pulse using the backward Raman memory protocol. The efficiency of the storage operation shown here is $63 \pm 6\%$ for a storage time of $5 \mu\text{s}$. The time axis is zeroed to the centre of the forward control pulse as this is the mechanism that controls the timing of the memory operation. The backward control field is centred on $15 \mu\text{s}$.

the highest efficiency optical Raman memory observed, compared with previous efficiencies residing around 30%. Additionally this figure also represented the first Raman memory to operate above the 50% threshold set by the no-cloning theorem. However this has notably been improved upon by Guo et al. who recently achieved a high performance memory with an observed efficiency up to 82% [117].

A useful metric for comparison of various quantum memory protocols is the time-bandwidth product which characterises the bandwidth and storage time via the relation $T\Delta\omega$, where T is the storage time and $\Delta\omega$ is the bandwidth. This is useful as often the characteristics of different platforms and protocols can be vastly different, whereas the time-bandwidth product allows one to quantify the number of pulses one could typically expect to store during the decay time of the memory. Raman type memories which often have a high bandwidth sometimes are not capable of storing efficiently, for example in comparison to GEM memories ($\sim 87\%$ efficiency). In the present implementation we were able to demonstrate the efficient storage of a 360 ns width signal for $1.5 \mu\text{s}$ as shown in Fig. 7.12. This operation has a time-bandwidth product of 160, corresponding to the maximum we were able to observe during this experiment. This puts the operation of our memory in a high fidelity regime where efficiency has not sacrificed the bandwidth and vice-versa. For comparison the results by Guo et al. [117] report a time-bandwidth product of 86.

7.5.2 Storage decay

To characterise the dynamics of the memory operation it is important to examine the decay properties of the memory. This can be performed by increasing the storage time and recalling the stored pulse to examine the decay in efficiency as shown in Fig. 7.13. There are effectively two main contributors to the decay in efficiency, inhomogeneities




Figure 7.12: High bandwidth storage operation yielding a time-bandwidth product of 160. The pulse stored has a $1/e^2$ -width of 360ns and is stored for one pulse width for a retrieval efficiency of 40%. The characteristic decay time is found to be $\tau = 60 \mu\text{s}$. Figure is adapted from [136].

in the magnetic field environment causing de-phasing and atomic motion. The atomic motion is characterised by a Gaussian distribution which accounts for atoms moving out of the interaction beam. Here we have taken diffusion to be effectively negligible as the distance travelled during the storage time is $\ll \lambda_S$, the wavelength of the spin-wave. Any additional magnetic field inhomogeneities will also add an exponential decay term as they modify the coherence of the stored state, which is crucial to the time reversal symmetry leveraged in the backwards Raman scheme. Thus to characterise the performance of the memory we fit the equation

$$\eta(t) = \eta_0 e^{-t^2/2\tau_v^2} e^{-t/2\tau_B}, \quad (7.4)$$

where η_0 is the memory efficiency at a recall time of 0, τ_v is the decay time due to atomic motion and τ_B is the decay time due to magnetic interference. The fit shown in Fig. 7.13 yields $\eta_0 = 0.63 \pm 0.01$, $\tau_v = 160 \pm 8 \mu\text{s}$ and $\tau_B = 127 \pm 10 \mu\text{s}$, where the errors are determined from the standard error of the fit parameters.

While there is some ability to limit the atomic velocity contribution by cooling the sample, acceleration due to gravity is always present. During the memory protocol the optical and magnetic trapping fields cannot be active, meaning this acceleration is unavoidable with the current experimental configuration. This issue can in some manner be avoided by aligning the storage axis with gravity, such that the only meaningful velocity is due to ballistic expansion, however this would require a complete experimental overhaul.

As discussed previously in Sec. 4.1.4, cancellation of the Earth's magnetic field and eddy currents is achieved by compensation coils and the isolation from magnetic materials. While these are effective methods there still remains some amount of magnetic




Figure 7.13: Experimentally observed efficiency decay as a function of storage time. Atomic motion and extraneous magnetic fields are responsible for the decay in efficiency as described by Eq. 7.4 which determines the fit shown in green. Errors in the efficiencies are determined from the standard deviation of the experimental acquisitions.

inhomogeneity which causes de-phasing effects during memory operation. It may be possible to remove such effects by increasing the shielding however this would add considerable experimental overhead. Additionally, the applied bias field homogeneity could be further improved by redesigning the coils. The storage time achievable is sufficient however for most general applications.

7.5.3 Losses

There are still a number of experimental challenges that remain to be addressed if the efficiency of this protocol is to be further increased. While further increasing the one photon-detuning could remove loss due to incoherent absorption, this would only account for a $\sim 5\%$ loss in efficiency. Additionally the control fields used to read/write the input states cannot be considered to be of uniform intensity. Inhomogeneities in the control field power will create spatially dependent light-shifts and dephasing mechanisms, contributing to losses. The field inhomogeneity is due to the practical challenge of creating a uniform beam while maintaining the optical power required for efficient storage. In future experimental redesigns it may be possible to design a telescope using cylindrical lenses which can better utilise available power while satisfying this condition.

7.6 CONCLUSIONS AND OUTLOOK

The possibility for the improvement of the Raman memory protocol via the use of the backward retrieval scheme has been predicted for some time without receiving experimental validation. I present here numerical simulations and the first experimental implementation, which demonstrate the strengths of the backward retrieval case in an optically dense ensemble of atoms. The result at hand, which is an important milestone, more than doubled previous efficiency benchmarks set by forward retrieval protocols with a modest time-bandwidth product of 160.

Improvements to this experiment mostly come in the form of attempts to increase efficiency. As mentioned, already control field inhomogeneity can be decreased as one of the main contributors to loss. Future research directions will include machine learned compression and loading cycles which will find a balance between high optical depth and low temperatures. This will limit losses to atomic motion within the ensemble. The ground work for such machine learning endeavours is laid out in Part [iii](#). While the current work relies on relatively bright input states, a large amount of work is continuing into the inclusion of the low bandwidth single photon source developed by the experimental quantum optics group at the University of Queensland [75]. This provides an opportunity for further study as it will allow high efficiency retrieval of truly quantum states of light (Fock states). However this will also introduce stringent filtering requirements as the control fields used in the current setup will more than likely preclude any ability to detect at the single photon level. A final research direction of some promise is the ability to construct arbitrarily shaped control fields using machine learning methods to optimise the storage and recall efficiency of input states. This could vastly reduce the power requirement while increasing operating fidelity.

Part III

MACHINE LEARNING

“The most exciting phrase to hear in science, the one that heralds the most discoveries, is not "Eureka!" (I found it!) but "That's funny..." ”

- *Issac Asimov*

8

INTRODUCTION TO MACHINE LEARNING

Machine learning has progressed rapidly in the last decade, finding more widespread applications than anyone initially predicted. In some sense this explosion in research and technology is similar to automation in the manufacturing sector. Tasks that were deemed intractable for software or machines were suddenly within reach. Certainly one of the great successes of machine and deep learning frameworks has been in the realm of image recognition. A naive perspective on the classification of images may be mislabelled as trivial, after all humans do it on a day to day basis. Clearly there exists some mapping between visual information and categorical labelling. Upon closer inspection however this mapping is quite complex, information used by an observer has multiple layers. For example there is visual data such as colour, material properties and shape which are all important. However there is also contextual information which can be obtained without direct measurement such as size, location and adjacent objects. Combining all of this information to produce meaningful results can be a daunting task.

An illustrative example is to simplify the task even further, let us for the moment consider task of determining if, given a particular visual scene, there exists a door. This is a task undertaken by humans daily, but remains a challenging concept for robotics researchers designing free-agents that interact with the real world. If presented with a door most people, past a certain age, can correctly identify with high efficiency that the object they are presented with is in fact a door, irrespective of cultural and social differences. The key here is "*past a certain age*". This implies that at some point we learnt what a door was, newborn children are of course unlikely to perform this task with meaningful efficacy. This information was assimilated from repeated interaction with such objects through many different sources, both physical (such as tactile feedback, operation, and watching others operate) and contextual (such as reading about doors in books, and seeing them in movies and TV shows). The exact mechanism of how this occurs is still up for debate, but it is clear that humans can learn new information through many different channels.

Now consider how one might programatically solve such a problem. Let us say we have access to visual information, such as through a camera. A basic approach may be to scan the scene for rectangular shaped objects and compare their similarity to a database of door images that we have collected. This similarity measure may include many metrics such as material and size. This approach quickly falls down on a few fronts though. Firstly if I were to show a circular door, this approach will fail to identify it. So we modify this first stage and look for circles and rectangles. This new occurrence will also substantially increase the number of entries to the door database. In fact

each new occurrence of a door will increase the dataset required to effectively identify doors, this will also decrease the speed of my algorithm which relies on lookups and comparisons. A model such as this requires massive overhead in terms of information to function accurately. It also has high efficiency within the training set, i.e. my known list of doors, but extremely poor generalisability outside of the training set, i.e. doors I have never seen before.

So how may we improve this algorithm further? Well, looking at this from a human perspective I might say, doors are human sized or larger openings that allow the transition to different domains. This abstract definition covers the majority of doors, but requires me to have some contextual information such as the size of objects I'm looking at, and the distinction between two areas that may be divided by a door. Size information can be obtained contextually, or physically by measurement. This is to say if given a photo of a scene I can identify a door if there are things around it which have a generally standardised set of sizes, i.e. chairs, bookcases and other household items. If it's not a photo I can just measure it of course. In this case I have increased the complexity of my algorithm. I now need to identify objects in a scene to get a sense of size. I am also required to determine domains, such that I may identify whether a door is involved. An implementation of such an algorithm would not require me to catalogue every door, or even every household object to compare sizes to. The cost of this is that I may misclassify things more often. For example a relatively large window may look like a door, it even functions similar to a door, but I should not classify it as a *door*. However while I loose some efficacy in the training set, my algorithm can much better generalise to different contextual situations and new doors.

A detail that we have glossed over above is exactly how one may gain contextual information programatically from a scene as this is not, in general, a trivial task. In the following sections I will lay the foundation for general machine learning techniques and deep learning as a platform for efficient computation of these seemingly hard problems.

8.1 WHAT IS MACHINE LEARNING ANYWAY?

In general machine learning refers to any process that attempts to distil an algorithmic approach to solve a particular problem using computational means. This distinguishes this process from simply assigning a skilled operator/human to the task. In general this is achieved by starting with some complex model which represents a mapping from the inputs to outputs and incrementally adjusting the free parameters until the desired outcome is achieved. This will not necessarily be a unique model, but it is a model with some measure of accuracy. This process of incrementally adjusting the model parameters is referred to as training or learning [137]. Ideally the trained model will be able to generalise to performing this task outside the set of training data, at which point the model can be said to have learned the task. This model can then be deployed

as the algorithmic approach to solve this particular problem, with this particular set of inputs and outputs.

8.1.1 Types of machine learning

Categorising the types of machine learning is a complex task given the overlap in methodology. Broadly machine learning (ML) can be divided into three major categories: supervised learning, unsupervised learning and reinforcement learning. Supervised learning is the most common comprising of classification and regression type problems. The term supervised implies that labelled training data is used to train the models on the expected mapping. For example in image classification a set of training data may consist of $k, m \times n$ -dimensional images, which have a corresponding list of k labels. The output in this case may be a vector representing the probability of the input image being classified as a particular label. State of the art algorithms have demonstrated accuracies of up to $\sim 88\%$ on the ImageNet database, which contains more than 14 million natural images [138, 139].

Unsupervised learning in contrast works with data that has no labels. Unsupervised learning aims to find hidden patterns or structures in unclassified data. In this scheme, the models are not so much attempting to provide a "right" answer, but instead exploring structures within the data which can provide inference [140]. These techniques can perform complex tasks such as anomaly detection [141] or clustering, where often it is the case that the figure of merit for detecting such an event is unknown or difficult to formulate.

Finally there is reinforcement learning (RL). This style of learning is generally aimed at learning to interact with an environment, or system, in some efficient manner. RL algorithms inform agents to maximise some reward or performance parameter, making them ideal for control systems and games. Most recently an algorithm named AlphaZero has demonstrated a mastery of the classic games, Chess, Shogi and Go, outperforming the respective classical algorithms considered to be state of the art [142].

Fig. 8.1 depicts the difference between each of these approaches. While in general these distinctions between the models hold, in practice a complete model will often employ multiple approaches, for optimal performance.

8.2 DEEP LEARNING

I have described a number of machine learning methods that are aimed at solving practical problems. However, the detail that is mostly glossed over is the models (and training). Construction of a model that can represent data in some meaningful way and potentially provide a useful mapping is a problem that has existed in machine learning since its conception. Until recently, construction of these models and representations required a careful, engineered approach to ensure stability and efficacy. The advent of

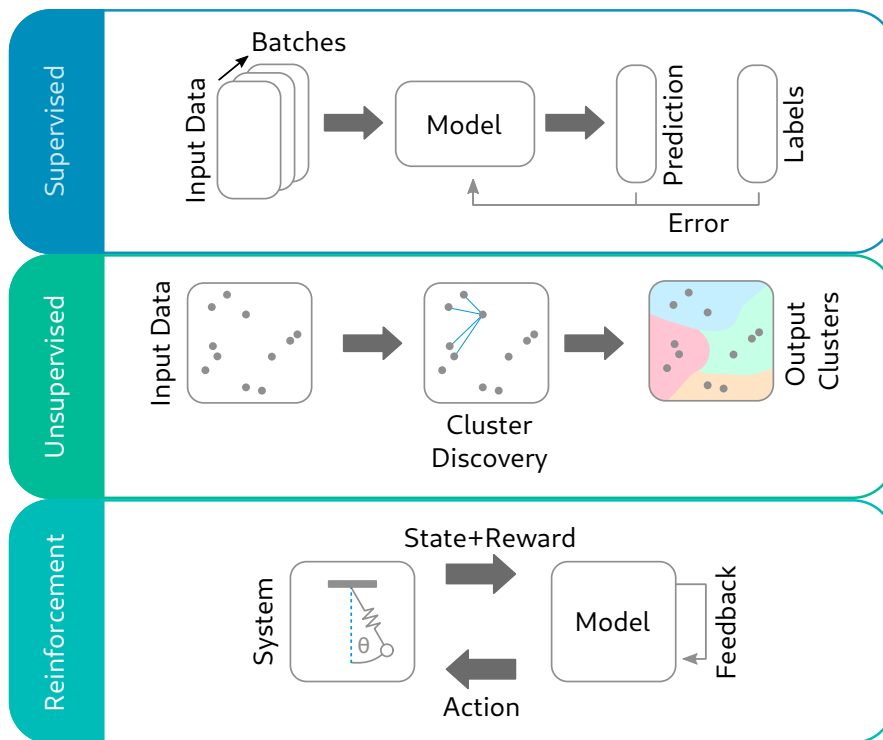


Figure 8.1: Comparison between the different model types during the training stage. Once each model is considered trained, the prediction stage would simply consist of passing the input for each model to receive an output.

deep learning techniques provided a convenient and computationally efficient means to perform such a mapping. Deep learning in essence refers to the representation of a mapping from input to output data using multiple layers of representation. This is done by constructing multiple layers of simple non-linear elements, that can approximate a desired mapping to some accuracy [143]. The important aspect of this approach is that these intermediate layers are not designed in the sense of early machine learning algorithms, but instead are learned.

8.2.1 Artificial neural networks

The building blocks of deep learning approaches are known as ANNs, layers of connected elements which can be used to approximate functions. ANNs derive their name from neural networks such as that found in the brain, after which ANNs are loosely modelled¹. A neuron is the base structure in these networks and performs the mapping

$$f(x) = A(x \cdot w + b), \quad (8.1)$$

where x is the input data, f is the output of the neuron, A is some non-linear activation function where the input is scaled by some weight matrix w , and b is a bias term. Note

¹ As we learn more about the human brain through efforts in neuroscience, this becomes even more loose.

we have not assumed the dimension of x or f , and in-fact they need not be the same. The activation function here is designed to mimic a phenomenon observed in the brain. Neurons in the brain will "fire" once a threshold electrical potential builds up across the cell wall. This is in essence how neurons in the brain process information [144] and is the basis for physical approaches looking to emulate this spiking behaviour for computing purposes [145].

8.2.2 Function approximation

We are interested in the ability to approximate some arbitrary function, as this is the core problem we need to address in our representation. If we assume our input resides in some space \mathcal{I} and our output resides in some space \mathcal{O} , then we are interested in determining a function, f , that creates a mapping, $f : \mathcal{I} \rightarrow \mathcal{O}$. In reality, access to this function is a difficult problem, so instead we wish to find some function f^* which is approximately equal to f , such that, $\epsilon = |f(x) - f^*(x)|$ is sufficiently small. The definition of sufficiently small will of course depend on the problem.

8.2.3 Multilayer perceptrons

One way to represent a function is to use the most basic form of neural network, known as a multilayer perceptron (MLP). Fig. 8.2a shows the optimal configurations for one and two neuron networks aimed at approximating a Gaussian function. Here we have arbitrarily chosen the tanh activation function. MLPs consist of layers of neurons (also called perceptrons), with weighted connections between some, or all, of the neurons in adjacent layers. Returning to our definition for a neuron, it is not hard to see that the ability to represent an arbitrary function will be extremely limited. Effectively this will only be a rescaling of the activation function. However with the addition of a second neuron we can begin to construct a larger class of functions. With the addition of one extra neuron we can already begin to represent a highly non-linear function. As the complexity of the function we wish to approximate increases so will the number of layers and neurons needed to represent this function. Fig. 8.2b shows the extension to two connected layers. Here as the input propagates through the graph it will be rescaled and offset via the weights and biases. For a given neuron n in layer m , we denote the input to the neuron's activation function

$$z_m^n = \sum_i A(z_{m-1}^i) \cdot w_{m-1,m}^i + b_m^n, \quad (8.2)$$

where the sum is performed over each connected neuron, $w_{m-1,m}^i$ are the weights between layer m and the previous layer, b_m^n is the bias for neuron n , and A is the activation function for the previous neurons. The output of this neuron is thus $A(z_m^n)$.

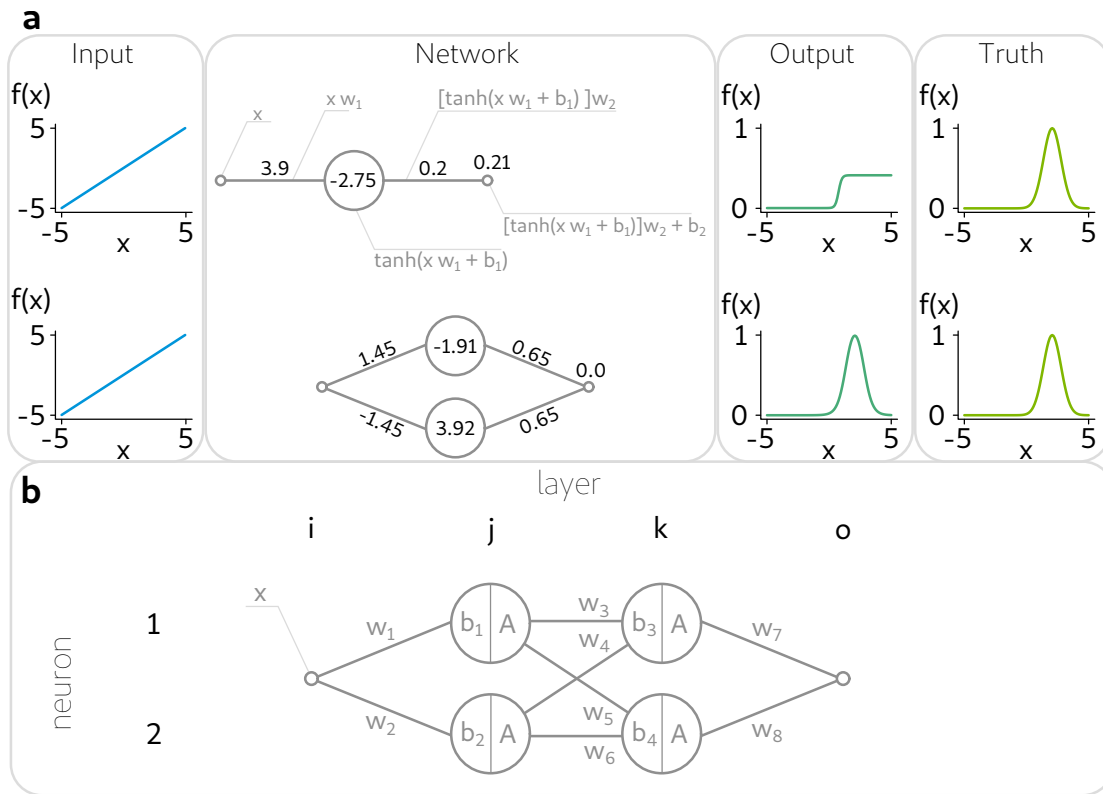


Figure 8.2: (a) Approximation of a Gaussian function using two network structures with 1 layer (depth 1). The first structure consists of a single neuron and poorly approximates the function. With the addition of a second neuron we can already begin to approximate a Gaussian with high efficiency. The labelled parts of the neuron show the construction of the operations from input to output. The weights and biases for optimal overlap are given for both networks. Weights are depicted as connecting lines, whereas biases are shown in or above the circles. Each circle represents either neurons or input/outputs. (b) depicts a network with two layers that are fully connected. As the input propagates through the graph, from left to right, the input is constantly rescaled and offset with the weights and biases. Additionally at each neuron the non-linear activation function A is applied.

Thus it is clear successive layers will apply the non-linear activation function, resulting in a highly non-linear mapping from input to output.

While a simple one dimensional mapping may require few neurons as above, a [MLP](#) for image processing may require millions of interconnected neurons. Provided our network dimensions are high enough we can always find some combination of functions that provide the approximate mapping we require [146]. The question becomes how to learn this mapping.

8.2.4 Loss functions

Let us suppose we have some model, $f(x; \theta)$. We wish to approximately perform the mapping, $f : x \rightarrow y$, subject to the model parameters θ . Our goal will be to determine θ , such that our mapping is accurate, to within some acceptable error. For any method that requires us to learn the model parameters θ , we will need to have some measure of how

close we are to the "truth" or actual values, \mathbf{y} . In general this is called a loss function², the choice of which will depend on the specific application. For example if the output of the model represents a probability distribution $p(\mathbf{y}|\mathbf{x};\theta)$, such as in classification where the output vector corresponds to a discrete probability density function, an appropriate choice of loss function is the cross-entropy between the training data and predictions [143]. For the work at hand we will almost exclusively deal with the mean squared error (MSE) given by

$$MSE = \frac{1}{N} \sum_{i=1}^m (y_i - \tilde{y}_i)^2, \quad (8.3)$$

where \tilde{y} are the predicted values and m is the number of samples in the training data. The loss function in this case is simply a measure of the average distance from the predicted values to the truth. While this is a common loss function, one may use any kind of loss function providing it is smooth and differentiable, as will become clear in the following sections.

8.2.5 Training neural networks

So far I have described the general topology of a MLP and the associated loss function that is required to determine the accuracy of the output. Now we must understand how to go about training these models, from some initialised state, to a useful model. A method for achieving this is known as backpropagation, but to implement this we require gradient descent. [147].

GRADIENT DESCENT If we consider each parameter $\theta_i \in \theta$, our aim is to find an update rule that we may apply iteratively, that takes us closer to a representation of our function. An update rule that may be applied simply is gradient descent

$$\theta_i^{t+1} = \theta_i^t - \alpha \frac{\partial L(\mathbf{x}, \theta_i^t)}{\partial \theta_i}, \quad (8.4)$$

where the t superscript denotes the time step, L is the loss function and α is a scaling factor known as the learning rate. Effectively at each time step the parameters are adjusted, descending the gradient of the loss function, thus reducing the loss. The size of this step is dictated by α , which in more advanced algorithms becomes adaptive (see [148]). The parameters that we are interested in updating in our model are the weights and biases of the MLP we described in Eq. 8.2.

² In the literature this is also called a cost function, but I will use loss to differentiate it from the optimisation cost function.




Figure 8.3: Structure and connected weights and biases to form a 3×3 MLP. The indicated neurons have been chosen to display how information flows through the network as a function of the weights and biases. The central neurons ($n=2$) have been highlighted to demonstrate which weights feed into the neuron from each respective layer. It is assumed that all neurons have the same activation function $a(x)$. Two neuron indices are required to distinguish the weights, however we have dropped the second index for clarity since we will only deal with the central neurons. Note that the superscript is not an exponent, but instead are neuron indices.

8.2.6 Backpropagation

From Eq. 8.4 it is clear that we will need access to the gradient of the loss function. Considering the nested structure of the neurons inside the MLP, it is not immediately straightforward how to access this gradient. This is where backpropagation comes in. For the remainder of the discussion we will refer to the notation used in Fig. 8.3, which gives the schematic for a fictional MLP structure that will inform this discussion. Our aim is to determine $\partial L / \partial \theta_i$ where θ_i is one of the weights in our model. Let us start with the weights, where we will break this into two separate cases.

FINAL LAYER Let's assume that our weight is in the final layer, we will denote this weight w_{lo}^2 . The partial derivative of the loss function can be expressed explicitly using the chain rule

$$\frac{\partial L}{\partial w_{lo}^2} = \frac{\partial z_o^2}{\partial w_{lo}^2} \frac{\partial N_{o,out}}{\partial z_o^2} \frac{\partial L}{\partial N_{o,out}}. \tag{8.5}$$

Each of these terms can be calculated with relative ease. The first term corresponds to the change in the input with respect to the weight. Since $z_o^2 = \sum_n a(z_l^n)w_{lo}^n + b_o$, this derivative is simply the output from the previous neuron $N_{l,out}^2 = a(z_l^2)$. The second term is the change in the output with respect to the input. The neuron takes the input and applies the activation function, thus this term is simply the derivative of the activation function. The final term is the change in the loss function with respect to the output and is a simple calculation. For the MSE with a single training sample, it simply reduces to $2(y_i - N_{o,out})$. We can now compute the derivative and apply the update rule given by Eq. 8.4.

HIDDEN LAYER Let us now suppose instead that the weight belongs to a hidden layer, i.e. any layer other than the final layer. For illustration let us suppose it is the weight w_{kl}^2 . As before we can use the chain rule to expand our derivative,

$$\frac{\partial L}{\partial w_{kl}^2} = \frac{\partial z_l^2}{\partial w_{kl}^2} \frac{\partial N_{l,out}^2}{\partial z_l^2} \frac{\partial L}{\partial N_{l,out}^2} \quad (8.6)$$

Similar to before the first and second terms simply become the output from the previous layer before scaling and the derivative of the activation function. The third term requires more care however. A change in the output of neuron N_l^2 will change every output that it is connected to in the final layer. In our simplified model we have only one output neuron, but for the sake of generality let's assume we had P neurons in the last layer. Then the third term can be found by adding the constituent parts,

$$\frac{\partial L}{\partial N_{l,out}^2} = \frac{\partial z_o^1}{\partial N_{l,out}^2} \frac{\partial L}{\partial z_o^1} + \frac{\partial z_o^2}{\partial N_{l,out}^2} \frac{\partial L}{\partial z_o^2} + \dots + \frac{\partial z_o^P}{\partial N_{l,out}^2} \frac{\partial L}{\partial z_o^P} \quad (8.7)$$

For our example we will only have the second term from the above equation, however we can treat all terms in the same fashion. The first factor is simply the change in the input of the last layer with respect to the output of our current layer. This yields w_{lo}^2 . The second factor is the change in the loss function with respect to the input of the last layer. Conveniently this is simply the same as Eq. 8.5 divided by a factor $\partial z_l^2 / \partial w_{kl}^2$. Since we have already computed these terms we can simply propagate this backwards and use the result we previously calculated. In fact this applies to all the terms in Eq. 8.7 and is how backpropagation works. If we start by calculating the gradients at the back of our graph (at the output), then we will have precomputed elements needed for the next layer. This entire process can be repeated for the biases with no appreciable difference.

8.3 HYPERPARAMETERS

The ability to train a neural network is only part of the job. There are a number of choices prior to training that will effect how well the model performs. The parameters of these are generally called hyperparameters.

8.3.1 Layers and neurons

The effectiveness of a model is in many cases determined by the choice of network structure. In the case of a [MLP](#) this relates simply to the number of hidden layers and neurons per layer. It is not necessary that the number of neurons remains constant throughout the network. There is no hard rule for how many layers or neurons a model should have when aiming to represent a particular data set. It is not true that the addition of more layers and neurons leads to increased accuracy, often the opposite is true overfitting may become an issue. In general one can empirically find a threshold in which the number of neurons per layer is suitable for the data at hand. However one may wish to impose certain restrictions such as in an autoencoder type setup [\[149\]](#). Autoencoders are an unsupervised learning technique that attempt to learn a signal whilst ignoring the noise in a data set. This is done with a network structure where the number of neurons towards the center of the network decreases, before increasing again towards to the output. The idea is that the reduced representation (less neurons) will prioritise the signal as this information is common to the dataset, whilst noise is suppressed being common only to the sample.

While it is not easy to exactly quantify in what manner the depth (layer number) and neuron number will affect the model, [Fig. 8.4](#) serves as an illustration to demonstrate the dependence in a 1D case. As can be seen increasing the depth allows the model to fit more complex components such as the data with increasing frequency even while undergoing exponential decay. However towards the upper end overfitting starts to become more prevalent. This could be improved by increasing the size of the training set. The introduction of more neurons increases the complexity of the representation of each data point within the model. As such the model begins to overfit the data on a shorter time scale. It can also be seen that the training becomes less stable due to the introduction of local minima into the loss landscape. Local minima can cause the training algorithm to become trapped, worsening or halting the training. This manifests as sharp spikes in the loss function over time.

8.3.2 Regularisation

One of the ways to mitigate overfitting in deep learning models is known as regularisation. Regularisation aims to incentivise the training process to learn the model, in a way that does not lead to overfitting. The most common of these are listed below in more detail. For a complete discussion on regularisation, I direct the reader to [\[143\]](#).

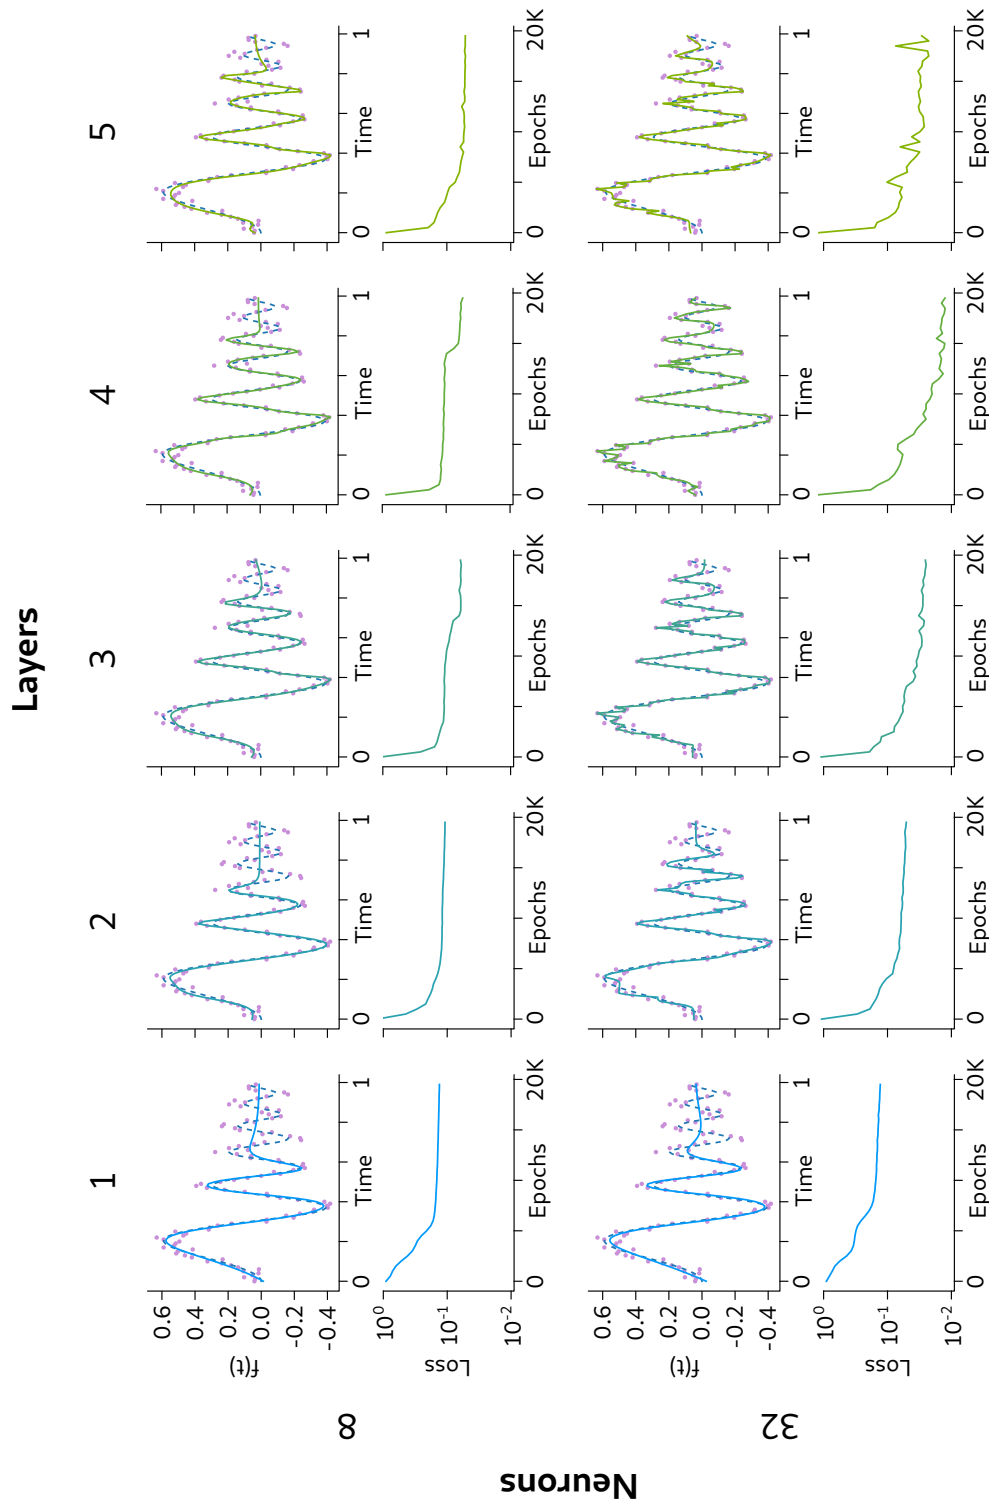


Figure 8.4: Learned fits for $f(t)$ given noisy sampled data, for different network topologies. Depths of 1 to 5 are shown with 8 and 32 neurons per layer. The dashed line indicates the true relationship while the purple points indicate the noisy samples on which the models are trained. The solid line is the predicted fit from the model. The training loss for each network as a function of epoch is also shown.

L1 AND L2 By modifying the cost function of the model to include penalties for undesirable training outcomes, the training process can be incentivised to learn more general fits. L2 penalties are the most commonplace modification, adding a cost that is proportional to the sum of the square of the weight magnitudes, $\sum |w_i|^2$. This has the effect of forcing the model to use all of the weights, dictating that no single weight should be large. L1 loss is less commonly used however is still important. The cost term is proportional to the sum of the weight magnitudes, $\sum |w_i|$. This has a sort of opposite effect to L2, instead returning sparse estimates. This is useful as it includes variable selection into the training process. For large models this is computationally efficient, as variables that have little or no effect in the model are disregarded.

DROPOUT This technique, as implied by the name, deactivates a percentage of neurons (and their subsequent inputs and outputs) during iterations of the training stage. The effect is that the model cannot rely on any one specific set of weights or neurons to fit the data, instead it must learn a more generalised model. The introduced randomness also helps induce training noise, which is useful to stop the gradient descent algorithms from getting caught in local minima.

EARLY STOPPING In some cases when using a validation set, kept aside from the training set, it is possible to avoid overfitting by monitoring the validation loss. When the model begins to overfit then it no longer represents the general trend or structure inherent in the data, instead tailoring specifically to the training data. This can be observed as an increase in the validation loss after a certain point, while the training loss continues to decrease. If the model training is halted when this begins to occur, then overfitting can be avoided.

DATA AUGMENTATION Data augmentation aims to reduce overfitting by increasing the size of the training data set. In many cases acquiring new data may be impossible or costly. In this case the existing data set may be extended by augmenting existing samples. For example, if one was designing a model to detect handwritten characters, then the data set may be extended by performing random transformations on the training images, such as scaling, shifting, shearing etc. Care must be taken when augmenting the dataset so as not to introduce biases or remove important information, which may in fact decrease the accuracy of the model.

BAGGING AND ENSEMBLES Bootstrap aggregating (bagging) is an ensemble method that aims to reduce problems relating to generalisation by combining the parameters of several trained models. The general idea is that from some random starting initialisation, the errors introduced during training will not be common to all models, thus they can be averaged out using different models. Bagging takes this one step further by constructing different datasets which are sampled from the initial data set with replacement. Each model is then trained on its own respective data set, with the final

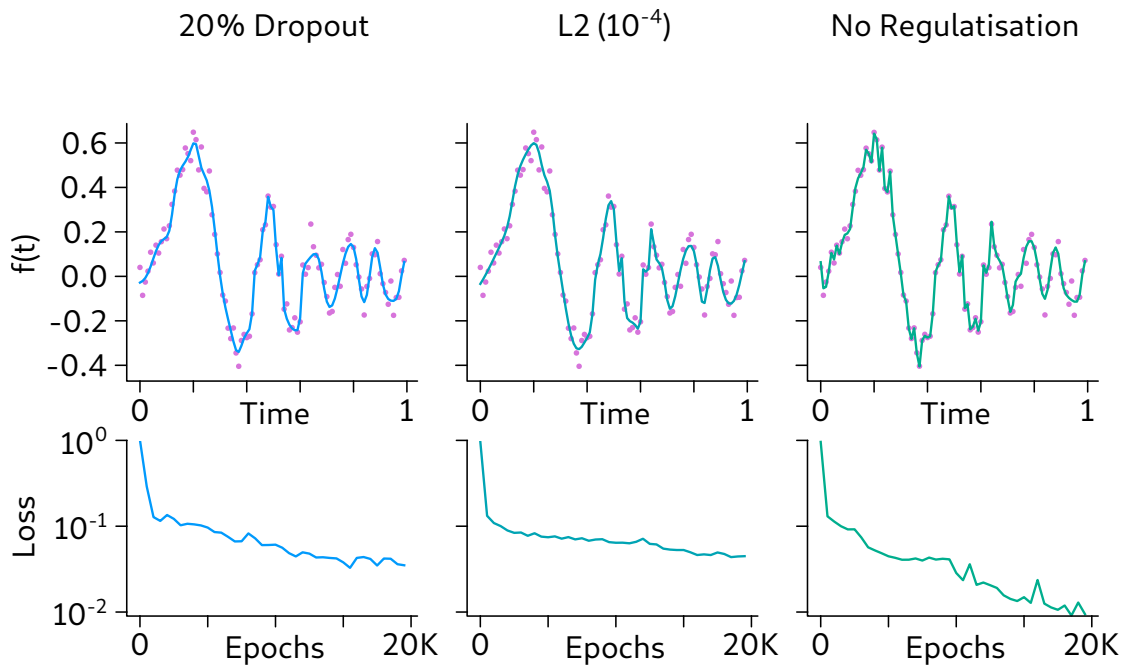


Figure 8.5: Different regularisation strategies applied to a MLP of depth 4 with 32 neurons per layer. The overfitting tendency is mitigated via the use of dropout and L2 regularisation. Note that the loss cannot be compared directly as the L2 regularisation includes a term into the loss function. In general the choice of activation function is decided empirically.

model being an averaged model of this ensemble.

In general some combination of these regularisation methods will be used to mitigate overfitting. Care must be taken to ensure that the model is not over regularised as this can lead to unintended underfitting. Fortunately this can be measured empirically by comparing the predictions of the model to the desired outcomes. Fig. 8.5 shows a direct comparison between three regularisation techniques, applied to a similar data set to that of Fig. 8.4. It can be seen that the overfitting can be mitigated using these strategies.

8.3.3 Activation functions

There are a number of activation function choices that will change the behaviour of the model. Any non-linear function can in practice serve as an activation function, however there are a few qualities that are convenient for an activation function to possess. In general they are monotonic functions that exhibit some sort of switch like behaviour, effectively turning on at some point. Monotonicity also helps with convergence during training operations. The activation function should also be continuously differentiable. Additionally it should be computationally efficient to compute this derivative, as this will be performed extensively during training. A few of the more common activation functions, along with the predictions from a trained model using these functions, is




Figure 8.6: Different activation functions applied to the neurons of a MLP of depth 4, with 32 neurons per layer. The activation function is plotted below the fit of each model to the random walk data.

shown in Fig. 8.6. The choice of activation function will determine the properties of the model fit. For example RELU is preferred in image processing for its ability to create hard boundaries in feature space. For regression an activation such as tanh or GELU [150] may be preferred to generate smooth features.

8.3.4 Initialisation

The initialisation of the weights and biases of a model can have a significant impact on the ability of the network to learn. This can be most easily intuited by considering an example where the output data sits in the range, $[-10, 10]$ for some set of inputs X . If initialised poorly, the model may output values on a range $[-100, 100]$ or greater, and possibly not symmetric around 0. In this case the distance between the desired representation and the starting representation is much bigger than it needs to be. The exact distance will depend on structures within the data itself. Training could become difficult or intractable in the worst case. This however can be mitigated in two ways. Firstly normalising the input and output data ensures that the order of magnitude of the weights is somewhat independent of the dataset, as it is always performing a mapping between two normalised distributions. Secondly a method developed by He et al. [151], known as He normal initialisation, provides a robust method for initialisation, which takes into account the size of the previous layer. The weights are still initialised randomly, however, their size is restricted, providing an efficient gradient descent.

8.3.5 Training algorithms

So far when talking about training models I have only referenced gradient descent (GD). However the truth of that matter is that if we were only to use GD then we would have a hard time training our network. If we have chosen our activation functions correctly, the loss landscape should be smooth and differentiable. However this is not to say that there will not be local minima, as neural nets are non-convex (see Ch. 9). Additionally, for efficient training we wish to use batches of data rather than singular training points at a time.

An extension to GD is stochastic gradient descent (SGD), which seeks to estimate the gradient using a batch of training data. This estimate introduces a source of noise, as such this method is less likely to get stuck in local minima, as the gradients never truly go to 0. A trade off is that the learning rate must be decreased over time as even at the optimal point the gradient estimate will be non-zero. Pure SGD can be slow to converge so it is common place to add momentum to speed up this process. Momentum effectively keeps a running average for past gradients (subject to some decay) that accumulates, similar to physical momentum [152]. This can help to escape local minima and traverse steep gradients quicker.

A popular training approach now is to use algorithms with adaptive learning rates/momentum such as AdaGrad [153] or Adam [148]. Indeed Adam is possibly the most popular approach to neural net training at present. Both methods present approaches to adaptive learning rates/momentum, which attempt to estimate the step size that should be taken based on the previous gradients and momenta.

The correct choice of training algorithm will depend on the problem at hand. For example in certain circumstances Adam may be faster than SGD but less stable around the minima. An approach that is often taken is to get close to a trained model using Adam, then switch to using a suitably tuned SGD algorithm to finish the process. There are also a number of hyperparameters that relate to the training process that may be tweaked.

LEARNING RATE Whether the algorithm in question has an adaptive learning rate or not, the learning rate will need to be chosen. The implementation of learning rate will differ for different algorithms, however, it generally will relate to the step size taken towards the minimum. Choosing the learning rate is done empirically, by looking at the training loss as a function of epoch. From this trend it can be determined whether the learning rate requires decay (in the case of SGD) or whether other decays should be increased such as momentum decay in the case of Adam.

BATCH SIZE The batch size controls how many samples will be used to perform a parameter update at a time. Choosing the batch size will depend on the data and the response of the model. Smaller batch sizes will tend to learn short term characteristics of the dataset more readily, while large batch size will tend to average over short term

fluctuations. The level of compromise between these two effects will depend on the problem.

EPOCHS This parameter sets how many training steps will be performed. In a single training step the parameters will be updated once per batch. This can be performed multiple times until a desired accuracy is acquired or, more likely, the validation loss starts to increase.

8.3.6 Choosing hyperparameters

As it can be seen there are a large number of hyperparameters that must be chosen in order to create a suitably functioning model. The difficulty in getting these parameters right will depend on the type of model and the structure of the data. Hyperparameters may be chosen and optimised by hand, or by using the model output as an empirical measure of performance. They can also be selected using a grid search or optimisation method, however, this can be costly for larger models. Normalisation of the input and output data reduces this complexity by enforcing certain hyperparameters that depend on the data to be independent.

8.4 EXTENDED NETWORKS

So far I have only talked about the most basic model, MLPs, however there are a number of extensions to deep learning models that are useful for different contexts [154]. I will discuss a few of the main models in the following³.

8.4.1 Convolutional neural networks

A simple extension are models known as CNNs. CNNs are models that include layers that enact convolution kernels on multidimensional input data, or data that has some relationship with nearby data points. First proposed in 1989 [155], CNNs have demonstrable success with applications in image processing and audio processing. The idea that CNNs attempt to exploit is that higher dimensional data structures contain information as series/arrays of data points, rather than individually. Take for example an image of a house. Given any one pixel, it is difficult to tell what this might belong to, but looking at the general shapes in a picture (peaked roof, flat vertical walls, exterior openings such as windows and doors) one can quickly identify a house. Kernel operations have existed in image processing for a long time, however, they have always been engineered to perform a particular task. CNNs seek to learn these kernels to produce accurate representations of higher dimensional structures. Some examples of image kernels are shown in Fig. 8.7, along with a general CNN structure used for image classification.

³ For a full discussion I refer the interested reader to [143] for a number of excellent overviews and insights.

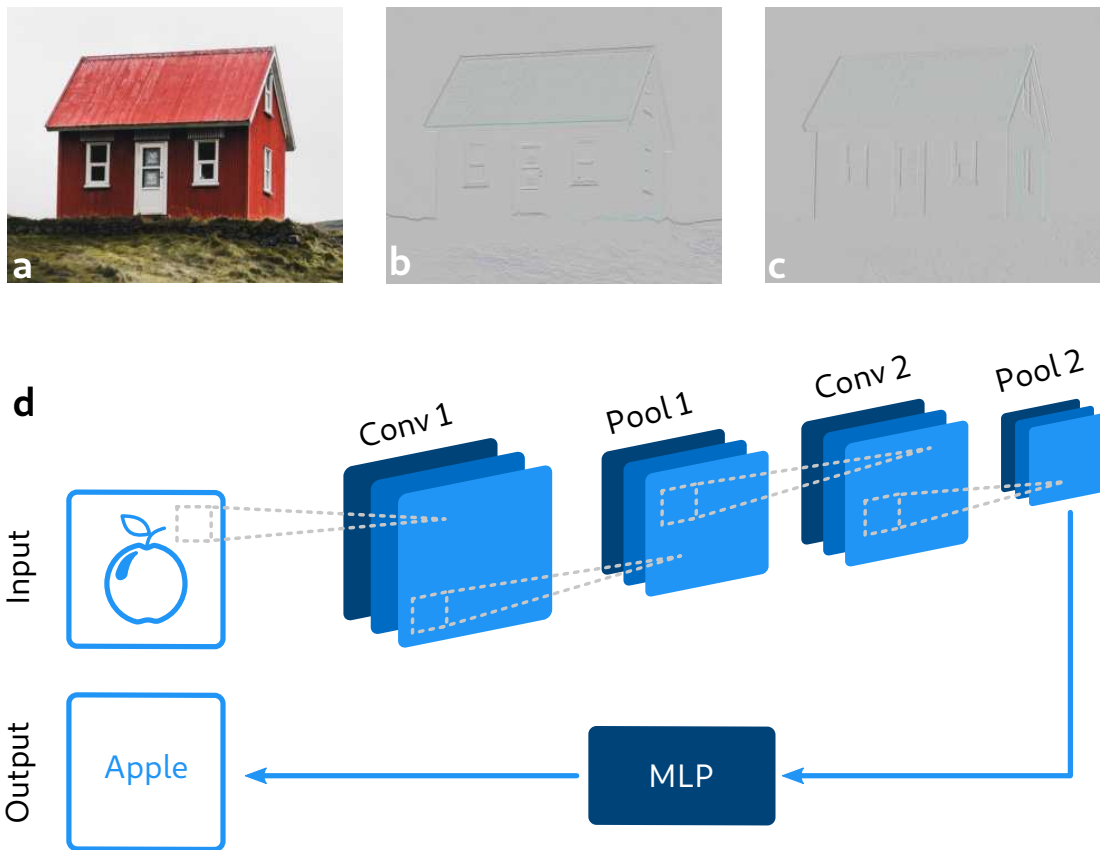


Figure 8.7: Demonstration of different convolutional filters that can be applied to images. (a) shows the original image of a house, (b) a horizontal edge detect kernel and (c) a vertical edge detect kernel. (d) Shows a standard CNN architecture that may be used for image classification. The pooling layers reduce the dimensionality of each layer, while the 'conv' layers perform a convolution. Original image in (a) by Luke Stackpoole.

8.4.2 Recurrent neural networks

Models that are designed to handle sequential data, such as text or time series traces, are known as recurrent neural networks (RNNs) [156]. These models are particularly useful when information needs to be shared across different parts of the model. This can be particularly important when dealing with data such as languages. For example consider the sentence, "I just ate an apple." There are many ways that I may rephrase this sentence and a language model would need to learn each of those examples. However, if a model is capable of sharing information between different parts of the model, this allows processing to occur irrespective of the structure.

8.4.3 Actor-critic and Deep Q networks

Deep Q networks and AC networks are used in reinforcement learning to approximate optimal policies for control. The idea is that policies (decisions based on a set of inputs) may be learned via an ANN or similar structure by repeatedly interacting with a system.

From this an optimal policy for exploiting a desired behaviour or outcome may be learned. [AC](#) networks will be further discussed in [Ch. 11](#).

8.4.4 Generative adversarial networks

A relatively recent addition to the deep learning models is the introduction of generative adversarial networks ([GANs](#)) [[157](#)]. As implied by the name, this architecture involves two competing models in order to learn to generate new data. The first network is a generator, which simply generates new data given some random seed input. The second network is the discriminator, which aims to determine whether the input data was sampled from the generator or the training set, given as a probability. The basic incentive structure is zero-sum (although this can be modified) such that the discriminator is rewarded for correctly classifying the data, while the generator is rewarded for fooling the discriminator. When the model has converged then the generators output is indistinguishable from the training set and the discriminator outputs a 50% probability in all cases. While [GANs](#) can be notoriously difficult to train, they have demonstrated success in photo realistic human faces [[158](#)] and recovering data from images in astronomy [[159](#)].

9

DEEP LEARNING OPTIMISATION

Optimisation problems are everywhere. It is an empirical fact, that the natural world is continually performing optimisation tasks through different means. Thus it is no wonder that naturally optimisation problems permeate our society. Most people will have come across even a basic problem at some point in their life: optimal route for picking up multiple passengers or optimal spending pattern for achieving certain goals. For simple tasks often the solution can be calculated analytically or with a brute force approach; consider all the options and select the best. However in the real world with problems that have many complex, interacting variables, this becomes quickly intractable. As with many questions in life, the answer to the question "how do I find an optimal solution" is, *it depends*. In this chapter we will attempt to elucidate this problem and present a new approach to solve complex optimisation tasks. The field of numerical optimisation is diverse and we will not seek to cover it in its entirety, but instead focus on the area that is most pertinent to the current work.

9.1 CONVEX VS NON-CONVEX

The distinction between convex and non-convex problems is an important one. Let us start with the general form of an optimisation problem [160]

$$\min_{x \in \mathbb{R}^p} f(x) \mid x \in \mathcal{C}, \quad (9.1)$$

where x is the controllable variable(s), f is some objective function that defines the performance, and \mathcal{C} is the constraint set that x is subject to. Our aim here is to minimise the objective function, subject to our constraints. Whether the problem is convex is determined by the objective function and its respective properties. The constraint set can also determine whether a problem is convex or not by segmenting regions of the objective function. A problem is deemed to be convex if a line may be drawn between any two points on the function with the line remaining above the function at all times. If this is not true then the problem is said to be non-convex. A diagrammatic representation of this is shown in Fig. 9.1.

Convex problems are substantially easier to optimise due to the inherent lack of steep local minima. Many conventional methodologies are extremely efficient at this task. However there are unfortunately a great number of problems that are non-convex, in fact most problems can be coded as non-convex optimisation problems.




Figure 9.1: Two different objective function topologies demonstrating the difference between convex and non-convex functions. The red section of the non-convex landscape represents a violation of the convex constraint. The presence of local minima in the landscape breaks prevents the use of efficient convex optimisers to consistently find the global minimum.

Non-convex topologies range from trivial to extremely difficult to optimise over, for a few reasons. A given landscape may contain many, steep local minima, saddle points, flat regions and widely varying curvature. Generally speaking non-convex optimisation is at least NP-hard, additionally lacking the tool kits generally available for convex optimisation [160]. This is generally due to the fact that theoretical guarantees or considerations are either weak or missing entirely. Convergence on a non-convex space is also difficult or impossible to measure without a priori knowledge of the landscape itself.

9.2 SOLVING NON-CONVEX PROBLEMS: A REVIEW

There do exist methods for solving non-convex problems aside from simply heavily sampling the entire landscape (an approach which quickly becomes intractable). One approach is to attempt to approximate the problem such that it becomes convex [161]. Once the problem is convex, appropriate methodologies may be applied to find an approximate solution which provides a starting point for a local-solver to take over.

9.2.1 Global optimisation

Global optimization strategies such as Bayesian optimisation/Gaussian process also provide a method for solving optimisation problems, however this can be slower than a grid search if evaluation of the objective function is fast. These processes represent the parameters as multivariate distributions and update their expectation values according to the observed data. Gaussian process optimisers have demonstrated success in many problems such as optimising the formation process of BECs [135, 162] and improving muon shielding [163]. Often, global optimisation strategies in general aim to solve the problem approximately, with local optimizers providing local convergence [164].

A common approach is to combine local optimisation with grid or random searches, such that a series of local optima are identified, from which the best optimum can be selected. This method has been used to optimise the operation of fast gates for ion trap quantum computing [165]. Generally global optimisation is applied to problems with low dimensionality, where finding the true global optimum is of high value. This is due to the fact that these methods are computationally and time intensive.

9.2.2 Evolutionary algorithms

Biologically inspired, evolutionary algorithms (EAs) are a class of algorithms starting with a population of candidate models that are mixed, selected and controlled [166]. The idea is to mimic Darwinian evolution in order to select well-performing models, thus many of these methods can be derivative free, decreasing the number of function evaluations. EAs perform well when the number of optimisation parameters is large [167] and have been applied successfully to design problems [168, 169], event selection in high energy physics [170], chip design [171] and quantum gate operation [172]. A downside to EAs is that they require tuning to perform well, with tuning of the hyperparameters often being problem specific. Improper tuning can lead to issues regarding convergence and local minima traps.

9.2.3 Stochastic methods

Stochastic optimisation (SO) is the name given to algorithms which perform probabilistic operations as a means of accelerating the optimisation process. The simplest form of SO is to randomly pick points until you find a good one, however such a strategy is bound to perform poorly on average. While EAs could be included in here, I will take this to refer to more inherently probabilistic methods such as Monte-Carlo methods and simulated annealing [173]. In the case of Monte Carlo methods, parameters receive random shifts which may be accepted as the new parameter set if they are below the minimum observed. Simulated annealing takes a similar approach however the probability of accepting the parameter change follows a "cooling schedule", analogous to an annealing substance. Both methods suffer from convergence issues, which is probabilistic and generally slow. However as they are not greedy algorithms, they are less likely to be susceptible to local minima traps [174]. Such SO methods have demonstrated success in optimising neural network models [175], solving the travelling salesman problem [176] and finding ground states of unknown potentials [177].

9.2.4 Metaheuristics

Metaheuristics are a class of iterative processes that exist outside of optimisation algorithms. Metaheuristics do not provide any guarantee regarding convergence or theoretical rigour, but instead are observed empirically to provide approximate solu-

tions to difficult problems [178]. They are based on the idea that an efficient search for solutions may be performed by balancing the exploration-exploitation trade off in an empirically valid manner [179]. As such many of these algorithms are organically inspired such as the fruit fly optimization algorithm [180] used in Ch. 7. EAs denote a category of metaheuristics, however there are a number of categories that are also naturally inspired such as particle swarm algorithms [181] and water wave optimization [182]. The number of different metaheuristics is vast, with many being variants on the same process, I direct the reader to [178] for a comprehensive list. Despite their abundance, particle swarm algorithms in particular have been applied to a number of problems including antenna control [183], cancer classification [184] and speech coding [185] to name a few. A comprehensive review of the applications can be found in [186].

9.2.5 Neural net surrogates

An approach using ANNs is differential evolution or other EAs, accelerated using surrogate models [187–189]. A surrogate model is a model which is evaluated instead of the actual objective function. The surrogate model provides an estimate of the cost landscape. Some approaches use ANNs as a surrogate approximation to the true objective function, which can be optimised over. In this case evaluation of the surrogate model can be parallel and extremely well optimised, often quicker than the more complicated model or physical system. This has demonstrated success in many technical applications such as microwave circuitry [190] and aerofoil design [191].

9.2.6 Summary and issues

The inherent structure of non-convex landscapes creates a complex optimisation problem that has attracted a lot of attention due to its real world value. A number of approaches to tackle this problem are outlined above with varying degrees of speed and success. An existing problem with all the algorithms above is that approaches seeking an exact solution necessarily scale badly with the problem dimensionality. As such for large problems, on which the current work focuses, it is necessary to take a more approximate approach similar to that of surrogates or metaheuristics. Additionally, there remains the problem of generality. There exists a no-free-lunch theorem for optimisation algorithms [192], which is to say that averaged over all possible landscapes all algorithms perform equally well. Unless something is known about the landscape then it is difficult to choose the correct efficient approach. With all this in mind I will now present a new approach that aims to alleviate some of these problems.

9.3 DEEP LEARNING APPROACH

Deep learning coincidentally turned out to be extremely good at distinguishing structures in high dimensional data sets [154]. This realisation coupled with efficient methods for training and evaluation makes it an ideal candidate for representation of a complex cost landscape, as is done in surrogate methods. The method we present is the first instance of ANNs in charge of both model representation and decision making processes in the context of optimisation. The pseudocode for this algorithm is given in Alg. 1. The idea behind the algorithm is that the ANNs will not only be the surrogate model but also be in charge of generating the next points for testing based on the inherent instability of the training and independent initialisations.

Algorithm 1 Deep learning based optimisation algorithm

```

nets  $\leftarrow$  3 {number of ANNs  $\geq$  1}
locals  $\leftarrow$  100 {number of local learners}
N  $\leftarrow$  number of parameters
samples  $\leftarrow$  2N
memory  $\leftarrow$  []
f  $\leftarrow$  objective function
DE  $\leftarrow$  differential evolution learner (pseudo-random)
Initialise ANNs

{Begin the initial sampling}
for i = 0 to samples do
    next  $\leftarrow$  DE
    cost  $\leftarrow$  f(next)
    append (cost, next)  $\rightarrow$  memory
end for

{Begin main optimisation loop}
while halt condition not met do
    for k = 0 to nets do
        train ANNk on memory
        preds  $\leftarrow$  []
        {Find best predicted parameters}
        for l = 0 to locals do
            next  $\leftarrow$  LBFGS on ANNk output, started at random point
            append (cost, next)  $\rightarrow$  preds
        end for
        best_pred  $\leftarrow$  argmin(preds[:,0])
        new_params  $\leftarrow$  preds[best_pred, 1]
        cost  $\leftarrow$  f(next)
        append (cost, next)  $\rightarrow$  memory
    end for
    {Inject some random point to avoid bias}
    next  $\leftarrow$  DE
    cost  $\leftarrow$  f(next)
    append (cost, next)  $\rightarrow$  memory
end while

```




Figure 9.2: Different predicted landscapes from the ANNs trained on the sampled (grey) data. With a sparse sampling of the landscape the neural networks will disagree as to the location of the best minima, facilitating exploration.

9.3.1 Multiple networks

The use of multiple networks in this algorithm has an advantage similar to that of bagging and ensemble methods in general [193]. One of the main problems with optimisation metaheuristics is determining the ratio between exploration and exploitation. Often this requires tuning of a particular parameter. In the present work this is handled automatically through the stochastic nature of the ANN training process. Consider the situation where a function is sparsely sampled, the samples of which we wish to use to train our ANN model. After performing the training, the representation of the function will be accurate around the sampled data points, but won't have any reference elsewhere. Thus an ensemble of networks trained on the same sparse data set will often predict different global minima, an example of this is shown in Fig. 9.2. Each predicted minimum can then be tested in succession, verified or not, and finally added to the memory. In this way the exploration of the landscape is facilitated by the lack of knowledge in a particular region. In the limit of dense sampling, the models will agree on the function structure which can be used as a rudimentary measure of convergence.

For this method to be successful we require that the models overfit the data, if only somewhat slightly. This ensures the networks are more exploratory in the limit of sparse sampling. Intuitively the amount of over fitting will determine the level of exploration, which can then be tuned with regularisation. In this way, the models will not necessarily represent the landscape completely and accurately, but their representation will be useful. The discovery of a steep minimum will bias the models towards exploring this area, in which case exploitation is effected. We term this method a stochastic artificial neural network (SANN).

9.3.2 Generating predictions

The advantage of the SANN surrogate models is that generating a new point is as simple as optimising over the predicted landscape (see Fig. 9.2). This operation is far quicker than optimising over the actual objective function for a number of reasons. First the evaluation of the SANN models is generally much quicker than the objective function,

especially if the objective function is tied to the operation of a physical system. Secondly, as a result of the network graph approach to tensor models taken in Tensorflow [194] and similar libraries, the gradients required for efficient search methods such as LBFGS [195] are efficiently calculated. A new point is generated by starting multiple local LBFGS searches at randomised points in parameter space. The minimum point found by these searches is then used as the next predicted point. These searches can be performed in parallel as the SANN models can be cloned, whereas the true objective function may instead only be operated serially, such as in the case of a physical system.

9.3.3 Pseudo-random learner

The inclusion of the DE learner aims to provide an unbiased stream of data into the memory. DE is an EA that has demonstrated success on many optimisation tasks [196]. The inclusion of DE ensures that the algorithm never truly stops exploring, even after the models have converged (providing of course the convergence criteria of the DE is not reached). In practice one could use any algorithm, however it is better to use something that is initially exploratory in nature, as it is in charge of generating the initial sampling. To this end an EA with a relatively large population is suitable. Throughout the learning process the convergence of the DE learner is much slower than that of the SANN, thus while it is geared towards optimisation, the sampling will appear more exploratory, or pseudo random compared to the SANN predictions.

9.3.4 Initial sampling

The initial sampling is performed by the DE learner and forms the basis for initially training the SANN. Empirically it was found that the number of initial samples should be at least $2N$, where N is the number of controllable parameters to optimise. However this will depend on the landscape and can be altered depending on the initial sampling. In general it is better to have a variety of objective function values in the initial training set, as this will allow the SANN to better generalise. However this is not imperative, it may just mean the absence of such will produce a slower optimisation. An improvement over this sampling method is presented in Ch. 11.

9.3.5 In-loop optimisation

The main optimisation is performed as an in-loop optimisation, which is to say the minima identified by the SANN are tested, providing the next best guesses of where to go. In many cases, the objective function will directly represent the behaviour of a physical system. After the initial sampling, the SANN is given complete control of the system performance subject to a number of bounds. While the SANN models themselves are not bounded, the bounds are included in the LBFGS minima search. By having the SANN operate in loop, the search for the global minima can be more efficient, by

biasing the search with the generalised landscape represented by the **SANN** models. This can be extremely useful if evaluation of the objective function is costly. After each evaluation, the cost and parameter pairs are added to the memory, on which the **SANN** will be retrained to represent this updated information, similar to a Bayesian update process. In this way the algorithm continues improving its knowledge regarding the true objective function landscape.

9.3.6 Function evaluations

For the most part we will be concerned with how efficient the algorithm is in terms of function evaluations rather than computational time. The justifications for this is that in many cases, evaluation of the objective function is costly in either time, or resources. In either case, for a high dimensional system many of the methods previously discussed require many evaluations to perform effectively. This algorithm instead compromises on computational efficiency, as each network is re-trained after generating a new point, which is a costly process. We assume that this process is faster or comparable to the time at which it takes to evaluate the objective function, allowing us to leverage the power of the **SANN** to access better predictions about the next best point. Additionally with multiple networks training in parallel, one can choose the number of networks such that a new point is always ready to be tested. It will often be true that the evaluation of the objective function is much slower than the training process for example in physical experiments found in quantum optics (see Ch. 10) and industrial applications. This compromise in computational efficiency allows the **SANN** to generate more informed predictions, ultimately leading to the accelerated optimisation process (in terms of function evaluations). In general we can specify the network size such that the objective function is always evaluating and never waiting on a prediction.

9.3.7 Testing efficacy

During the design of Alg. 1, the efficacy was tested by benchmarking on a 10 dimensional problem with an objective function of the form

$$f = \sqrt{\sum_i (x_i - x_{min,i})^2}, \quad (9.2)$$

where x_i represents the i^{th} candidate parameter and $x_{min,i}$ represents the minimum coordinate for the i^{th} parameter which is randomised at the start of the problem. Effectively, each parameter represents an offset polynomial in that dimension, which should be a fairly easy optimisation. Fig 9.3 shows the results of such a benchmark in which multiple algorithms were applied to the problem of Eq. 9.2. The implementation of Gaussian process and differential evolution are through the M-LOOP package while the fruit fly implementation follows that of Appendix B. The **SANN** consisted of 3 **ANNs** with a network structure of 5 hidden layers, with 64 neurons per layer. Only one




Figure 9.3: Benchmark of different optimisations on the random polynomial problem. The best performing algorithm in terms of best minima found is the neural net algorithm (NN), followed by Gaussian process (GP), fruit fly (FF) and differential evolution (DE). Algorithms are restricted to 200 function evaluations (including training for the neural nets).

optimisation run is shown here, however this result is indicative of the results observed on repeated benchmarks. Further benchmarks and improvements are detailed in Ch. 11. As shown in Fig. 9.3, the *SANN* finds a better minimum than the other algorithms. It also approaches the minimum quicker. It should be noted that in general the *SANN* is a global optimiser and the final optimisation in the region surrounding the best known minimum should be performed with a local optimiser for best accuracy.

10

OPTIMISING A MAGNETO-OPTICAL TRAP

In Pt. [i](#) and [ii](#), the case was made for the use of a cold atomic ensemble for quantum information purposes, such as a quantum memory. One of the main limitations in many of these experiments is the scarcity of OD, due to the inherent difficulties in efficiently trapping cold atoms. This chapter will report on the application of the deep learning algorithm presented in the previous chapter to this problem and resulting experimental results.

10.1 THE PROBLEM

The dynamics of the steady state operation of MOTs are well described by various works [[13](#), [55](#), [197](#)]. The limit of the theory, in this case, is that approximations generally rely on the operating regime meeting certain assumptions such as steady state dynamics, low intensity, small OD or far detuning. Presently, it remains a challenging endeavour to construct a quantitative description that captures the complete atomic dynamics outside of these approximations. Unfortunately, for work such as that presented in previous chapters and quantum memory schemes in general, this regime is exactly where one is required to operate such a system. The complexity in completing such a description mostly owes to the complex dynamics involved which, when extended to 3 dimensions, become computationally intractable, involving many-body interactions, polarisation and intensity gradients, and complex scattering processes [[197](#), [198](#)]. The result of this is that many processes aimed at improving OD revolve around the use of intuition to create adiabatic ramps such as that presented in Sec. [3.3.5](#). Despite this there have been indications that operating outside of this regime can increase the efficiency of collection into cold ensembles [[199](#)], including an experimental demonstration of BEC formation without evaporative cooling [[200](#)]. From this it is evident that the area that could yield the biggest improvement in atom number, and hence OD, is the transient compression sequence of a typical experimental run.

10.1.1 Parametrisation

Parametrising the relevant experimental controls is a tricky endeavour. Choose the wrong parametrisation and the landscape can become extremely complicated, precluding efficient search via optimisation. While in general there is no rule for how to choose this, as with an ansatz in solving differential equations, it helps to know something about the problem. For example, if the problem exhibits periodic behaviour

it may be useful to parametrise in a way that exploits this (e.g. a periodic basis). During the compression sequence we wish to exert control over three main parameters over a period of 20 ms: trapping detuning frequency, repump detuning frequency and magnetic coil current. Intuitively we expect that the resulting time-value profiles of each should probably be monotonic, however we have already performed such an optimisation, as demonstrated in Sec. 7.4.5. For this optimisation we wish to assume nothing about the presumed nature of the ramps involved. Instead we wish to apply only the physical constraints of the system, thus allowing the optimisation to explore the full range of potential behaviours. The simplest parametrisation that allows such expressivity is allowing piece-wise arbitrary control over a set of values, at regularly sampled intervals in each control channel. To implement this the duration of the 20 ms compression sequence is subdivided into 21 time value pairs for each of the three channels, giving a total of 63 optimisation parameters. The final parameter in each sequence sets each respective channel to the value for the PGC sequence immediately proceeding the transient compression, thus each parameter represents a 1 ms duration in the transient ramp. This choice of parametrisation balances granular control during the sequence whilst keeping the dimensionality of the problem small enough to perform an effective optimisation. It is clear that a human operator could not hope to efficiently optimise these parameters by hand and would quickly succumb to biases based on intuition and perceived efficacy of the solution.

10.1.2 Online approach

Previous machine learning techniques with deep learning have demonstrated success with quantum systems, such as optimal control [201–203], design [204, 205] and prediction [206]. However these approaches are performed in an offline setting, relying on prior information. Online approaches have been demonstrated using Gaussian process models [135] and EAs [207–209]. The approach I present here is the first online approach to leverage the power of deep learning.

Online optimisation is well suited to this experiment, due to the nature of the duty cycle implementation and need for efficient exploration. The experiment is run continuously at a frequency of ≈ 2 Hz. The MOT loading (preparation, PGC and compression) accounts for 99% of the duty cycle, with $< 1\%$ reserved for measurement and experimentation. However the majority of atoms will not leave the trap between cycles as the atom drop during experimentation (which occurs when the atom trapping is ceased during quantum memory operations) does not last long enough to facilitate the full loss of atoms. Thus atoms are accumulated over successive runs until an equilibrium value is reached, at which point a measurement of the performance occurs. The performance of the trap will often vary on a scale of 1-2 hours due to various physical drift sources, such as the Rb dispenser, laser power and frequency drift, and alignment drift due to temperature fluctuations. This drift precludes offline optimisation which would require a relatively dense sampling to accurately represent

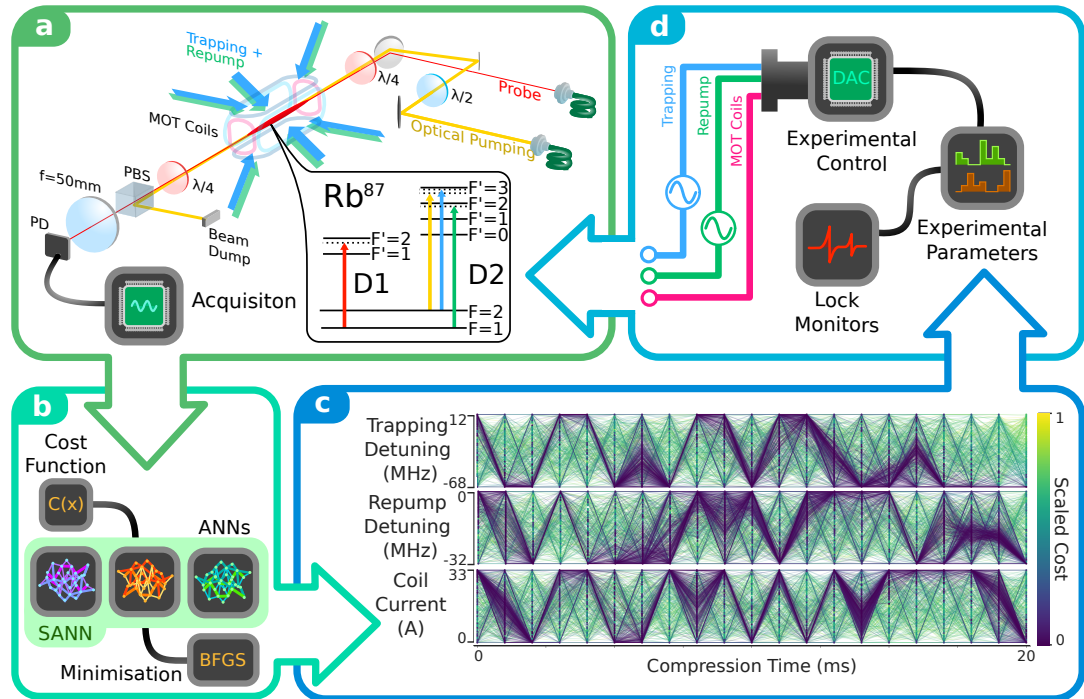


Figure 10.1: Experimental implementation of the online deep learning algorithm for optimisation of the OD. The feedback like loop starts by (a) sampling the response of the MOT via a probe beam for different parameter sets. These parameter, cost pairs are then used by (b) the SANN to train a representative model of the cost landscape. (c) shows the mapping between parameters and cost indicated by the line colour, observed over the experimental acquisition (lower cost is an improvement). The SANN predicts the next best set of parameters to examine. Parameters that are found by the SANN during a minima search are passed to the experimental controls (d) to be implemented experimentally.

the cost landscape, especially for a 63 dimensional optimisation. In this case it is advantageous to use an efficient online optimiser that can directly target the search without wasting evaluations in uninteresting parameter regions. Previously discussed methods require far too many function evaluations to efficiently optimise the landscape before drift becomes significant.

10.2 EXPERIMENTAL IMPLEMENTATION

The application of the deep learning algorithm to the experimental setup described in Ch. 4 is outlined in Fig. 10.1. The main barrier to entry is achieving automated experimental control, as this optimisation can (and should) be performed faster than a human operator. Fortunately as noted in Sec. 4.2, this is accessible by design and can simply be included as a module in the python RunBuilder. For the remainder of this section I will outline the different parts of the complete experimental process and its corresponding complexities.

10.2.1 (a) Cold atoms and measurement

For each new set of parameters, atoms currently in the trap must be relinquished. Without dropping the trap, atoms present in the trap can remain in the MOT, even with a new set of parameters with worse performance. These atoms would not otherwise have been trapped and while they will eventually leave the trap, the time taken to reach this new equilibrium is far greater than dropping the trap and reloading the MOT. An assumption of this method is that each run, or test of parameters, is independent from the last. Without the MOT drop this assumption is broken as current runs become correlated with the performance of the previous runs.

The atoms are prepared as outlined in Sec. 4.2.3, with the exception of the adiabatic ramps for our three control parameters, which are now replaced with the piecewise functions generated from our parameters. The final stage after the new compression sequences, is a 1 ms PGC stage and optical pumping, where final cooling takes place followed by optically pumping the atoms into the $m_f = 2$ magnetic sublevel with a σ^+ -polarised beam, detuned from the $F = 2 \rightarrow F' = 3$ D₂ transition. The exact detuning of the trapping beam during PGC and optical pumping beam is set by the final point of the machine learned compression sequence, as this frequency is controlled by an AOM common to both paths.

Up until this point we have not yet defined how to measure whether the system is performing well. As we are interested in improving the OD of the system, measuring this quantity would appear to be the correct technique. However as OD increases, gaining an accurate measurement becomes more difficult. For this method to work we do not require that we have an exact measure of OD, simply it is advantageous to have a fast measure of some quantity that is related to OD. For this we send an off-resonant probe pulse through the atomic ensemble which is measured on the other side by a bucket photo detector (PD). This probe provides us with a straightforward absorption measurement which will monotonically increase with OD. Our cost function in this case can be simply defined as

$$cost = \frac{1}{I_{ref}} \int_0^T p(t) dt, \quad (10.1)$$

where I_{ref} is the integral of a reference pulse sent through without atoms and $p(t)$ is the PD signal acquired for a time T . The cost function defined above returns a scalar value, where a larger value corresponds to less absorption and thus lower OD, and vice versa. An important addition is the scaling factor of the reference pulse. Before the new parameters are implemented and the atoms have been removed from the trap, a reference pulse is sent through which is used to appropriately scale the cost function. This removes fluctuations due to laser intensity, which were found to be a problem during experimentation. For the duration of the experiment the probe was detuned from the $F = 1 \rightarrow F' = 2$ D₁ transition. It was found that a detuning of -90 MHz

provided a good signal to noise ratio for the higher OD runs, while not precluding the optimiser from exploring lower OD regimes.

It has been demonstrated and catalogued before, that artificial agents will often cheat the task presented to them, if provided an opportunity to do so [210]. This system is no different, and oftentimes one must be careful to design the cost or measurement in such a way that removes such an opportunity. In the case of the present work, the agent has little ability to cheat at the task at hand, but over the course of experimentation we were able to identify one such case. In the high OD regime probe light which is close to resonance, can be lensed by the ensemble. While the optimiser does not have control over the relative position of the MOT, it can control the relative expansion to some degree with the magnetic fields. When approaching high OD, in certain cases the AI was able to cheat, by controlling the expansion such that less light (sometimes no light) was incident on the detector not via absorption, but by lensing. This problem was easily rectified by increasing the detuning and placing a large 2 inch optic ($f=50$ mm lens) to capture and focus the light onto the PD. However such problems are common and demonstrate the need to verify solutions generated by the optimisation process.

The measurement of the cost occurs after approximately 10 – 15 experimental cycles from the time of the MOT drop, although for poor parameter sets this can be substantially quicker as there is no loading time. The cost value sent to the SANN is monitored by repeated measurement of the cost at each cycle. Once the cost has reached an equilibrium value, within experimental noise, the acquired cost is sent to the SANN. In this regime 1 data point is gathered every 5 – 10 seconds on average.

10.2.2 (b) SANN training and minimisation

Once the relevant quantities have been measured to construct the cost function, the cost is sent to the SANN to be evaluated for training. As described in Sec. 9.3, once each of the predictions from the SANN has been tested, the returned costs are added to the running memory and the SANN is trained. We require that the SANN produce new predictions to test experimentally, which are found by an ensemble of L-BFGS local optimisers. For this optimisation task the SANN is comprised of 3 neural networks. The structure of each network is a depth of 5 with 64 neurons per layer. This network structure affords the required expressivity, while still allowing the networks to be trained in a sufficiently short time such that the experiment is not waiting for this computation. The networks are trained on modest hardware using only CPUs (Intel i7-920 2.67 GHz). Each respective network is trained using the Adam algorithm with early stopping and L2 regularisation. As we expect the response of the parameters to be generally smooth, GELU is used as the choice of activation function.

10.2.3 (c-d) Predictions and implementation

Continual application of the **SANNs** training and prediction loop leads to survey of the parameter landscape in a fashion that balances exploration and exploitation. Visualising a 63 dimensional space is not an easy task, graphically or otherwise, however due to the nature of the parameters constituting ramps it is possible to plot the generated ramps with respect to their relative cost as is done in Fig. 10.1c. Here we begin to see the emergence of particular ramp shapes that lead to a better performance (lower cost). Visualisation as such also allows the intuitive investigation into the importance of given parameter values. For example, in all cases the final point of the compression sequence must pass through a particular value to perform well. It is important to note that while each point is connected by a straight line, experimentally the switching time is $\ll 1$ ms, and thus the resultant sequence will instead look like piecewise steps.

Each of these predicted parameter sets is passed to the control systems which will implement the physical values. To ensure the stability of the optimisation a number of monitors were put in place which ensure that key systems remain locked. This corresponds to the trapping, repump and probe lasers, which require frequency stability to ensure a valid optimisation. In the case where one of these fell out of lock, the optimisation loop was halted and any measurements during that time automatically disregarded. As the networks train on the running memory, this is as simple as removing the offending cost and re-testing the parameter set. Following the implementation of the parameters on the experiment, the loop continues with part (a) and is repeated until some stopping criterium is reached. The criteria for halting the optimisation are: a maximum number of runs reached, convergence of the neural nets to a satisfactory regime, experimental drift tolerance exceeded, or human intervention.

10.3 EXPERIMENTAL RESULTS

While a number of experimental optimisation runs were performed on the cold atom system, I will examine the best achieved result for the initial part of this discussion. Sec. 10.3.6 will cover the identified local minima. Fig. 10.2 gives a complete overview of the best experimental optimisation result achieved. As shown in Fig. 10.2b the **AI** learner initially samples the landscape for the first 126 runs. After this the **SANN** begins the initial training on these training examples, with a batch size of 32. Once the initial training is complete, subsequent training on new examples will be much shorter. The **SANN** will begin to generate predictions after this initial training. The initial predictions fall within the same relative range as the **DE** algorithm owing to the early stopping techniques used to try and mitigate overfitting. However after only 100 runs the **SANN** is clearly and substantially outperforming the **DE** learner. Spikes in observed cost as a result of the **SANN** occurring later on correspond to an increased rate of exploration, as a result of overfitting. This is due to the fact that the **SANN** is now in a parameter region that is not changing much, with repeated observations, and thus it becomes

more likely to overfit the data. The optimisation was halted after a total of 703 runs on the grounds of convergence, with the optimal solution found at run 583.

10.3.1 Human and machine ramps

The human optimised ramp shown in Fig. 10.2a was optimised via manual adjustment and with the aid of a Gaussian process learner. The only difference between the constraints applied to this problem and the machine learned solution, are the biases and intuition applied to the problem, from the perspective of the skilled operator. In this case this solution builds on knowledge of the adiabatic ramps proposed as temporal dark SPOTs, and the intuitive understanding of the theory describing the steady state operation. The AI learner is however not burdened by such intuition and so called rationalities, placing merit on each possible combination until informed otherwise by the cost. As such the machine learner produces a completely unintuitive ramp, shown in Fig. 10.2c. Immediately there are a number of striking differences between this and the human solution. Firstly the solution is not monotonic, instead varying rapidly between the control rails, as determined by the physical limits. Secondly the solution is not continuous, showing little to no correlation between successive parameters. The stark contrast between the two solutions, highlights the fact that it is extremely unlikely that a human operator would seek to construct a solution with such a profile.

10.3.2 Quantifying success

Following the successful optimisation of the MOT loading, it is important to verify the physical behaviour of the system to determine that the intended behaviour is adhered to. As the absorption measurement is not a strict measure of OD, this should be the first port of call. A proper OD measurement can be obtained by varying the probe detuning symmetrically around resonance and fitting to Eq. 2.25. The observed absorption profiles are shown in Fig. 10.2d which correspond to ODs of 560 ± 8 and 960 ± 20 , for the human and machine runs respectively. Thus the machine learned run has improved the optical depth by $81 \pm 3\%$.

The observed increase in OD is also supported by the side absorption images shown in Fig. 10.2e. These integrated traces show the increase in optical depth around the centre of the distribution. It can also be seen that not only is the learned sequence increasing optical depth in the central region, but additionally shows a distribution of atoms above the main ensemble which does not vary substantially from run to run. While this feature does not contribute to usable OD, it is a feature that is only present in the learned sequences. This is further seen in the integrated traces shown below the absorption images, where the cloud can be seen as a tail on the distribution.

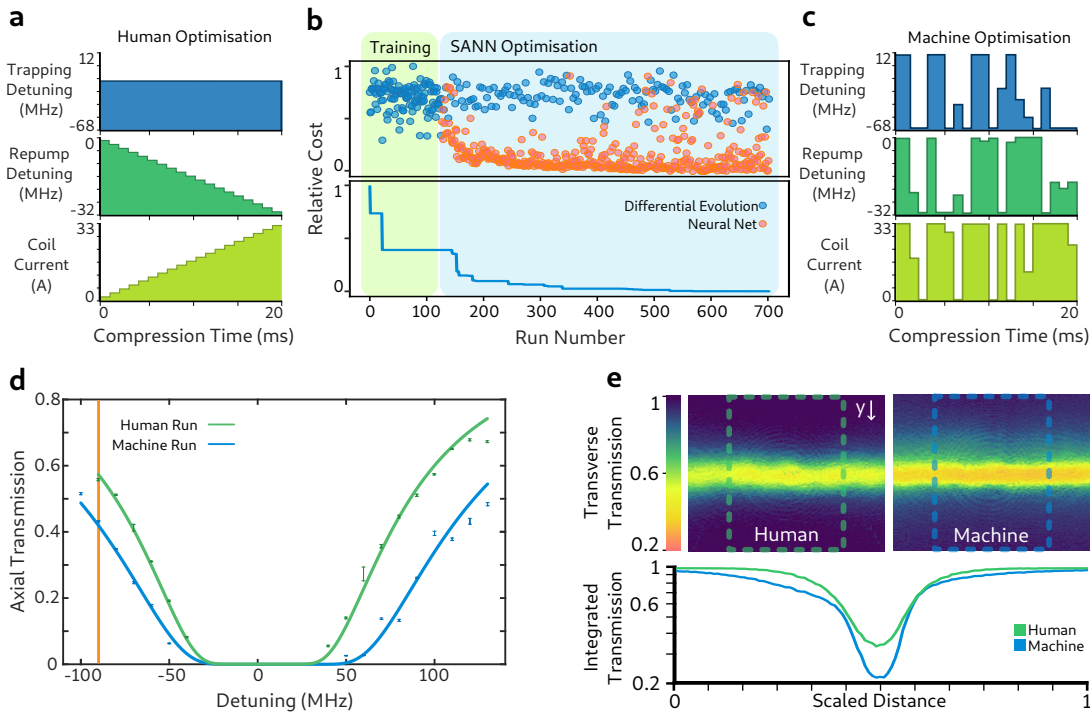


Figure 10.2: Results of the 63 dimensional optimisation using the in-loop AI algorithm. (a) Shows the original human ramp which was optimised using a Gaussian process learner. (b) the progression of the measure cost as a function of run number. The neural nets significantly outperform the DE learner with the optimal solution identified at run 583. The bottom plot of (b) shows the historical minimum. The resulting optimal ramp is shown in (c) displaying features which are neither continuous nor monotonic. The corresponding OD measurements for each run are shown in (d) with an OD of 580 ± 8 and 960 ± 20 for the human and machine runs respectively. The orange line represents the probe detuning during the optimisation. (e) shows a quantitative comparison between the human and machine learned ramps effect on the atomic ensemble. Side profile absorption imaging shows a distinct difference in density in the centre region with a halo of atoms surrounding the machine learned ramp. The integration of these absorption plots is shown below the side images on a logarithmic scale, with the dashed boxes indicating the integration region that is chosen to be free of image distortions.

10.3.3 Time dependence

As the compression sequence is time dependent, it is useful to look at the side absorption imaging over different parts of the run. Due to the variation of the repump and trapping frequencies, the atoms will often occupy different states throughout the sequence. To perform the imaging one must first pump the atoms into the state resonant with the imaging beam so as to image the entire population. This can be achieved by simply applying the trapping beam briefly (imaging is on the repump transition) to optically pump into the correct state. Fig. 10.3 shows a collection of selected side absorption images taken for both sequences using this technique. It can be seen that as expected the human sequence compresses the atoms monotonically. This is not the case for the machine learned run. Instead the machine learned run applies a multi-modal technique and releases the compression partially during the sequence. This is especially evident when comparing the images taken at 6 and 13 ms. These dynamics are similar to the release-and-capture techniques applied in optical lattices to increase lattice occupancy [211]. It should be noted that the images in Fig. 10.3 are not calibrated for OD and are instead modified for maximum visibility. Additionally the integrated trace at 1 ms intervals for each sequence is shown in Fig. 10.4. Here the evolution of the atomic spatial distribution is shown, with the monotonic nature of the human scheme clearly visible, compared to the machine learned scheme.

10.3.4 Temperature

One of the consequences of performing a transient compression sequence is heating of the ensemble. There are a few mechanisms for this, the first is that increasing the density will increase the number of re-scattering events that contribute to heating an atom. The second mechanism relates to the increased mean laser intensity experienced by a trapped atom, which can add additional heating rates [212]. The temperature of the ensemble can be monitored at different points throughout the compression sequence by performing a ballistic drop measurement at each point. Fig. 10.5 displays the evolution of the ensemble temperature over the course of the two respective sequences. The temperature of the human scheme increases over the course of the sequence nearly doubling from the start of the sequence. The temperature of the machine learned run is much more volatile, however the final 4 ms is spent cooling the ensemble, with the final temperature arriving at approximately 40% less than that of the human run. Additionally this late cooling feature is consistent across other minima observed in other optimisation runs.

10.3.5 Monitoring convergence

For global optimisation tasks, the idea of convergence has little meaning outside of certain classes of problem. A fortunate consequence of the approach taken with the

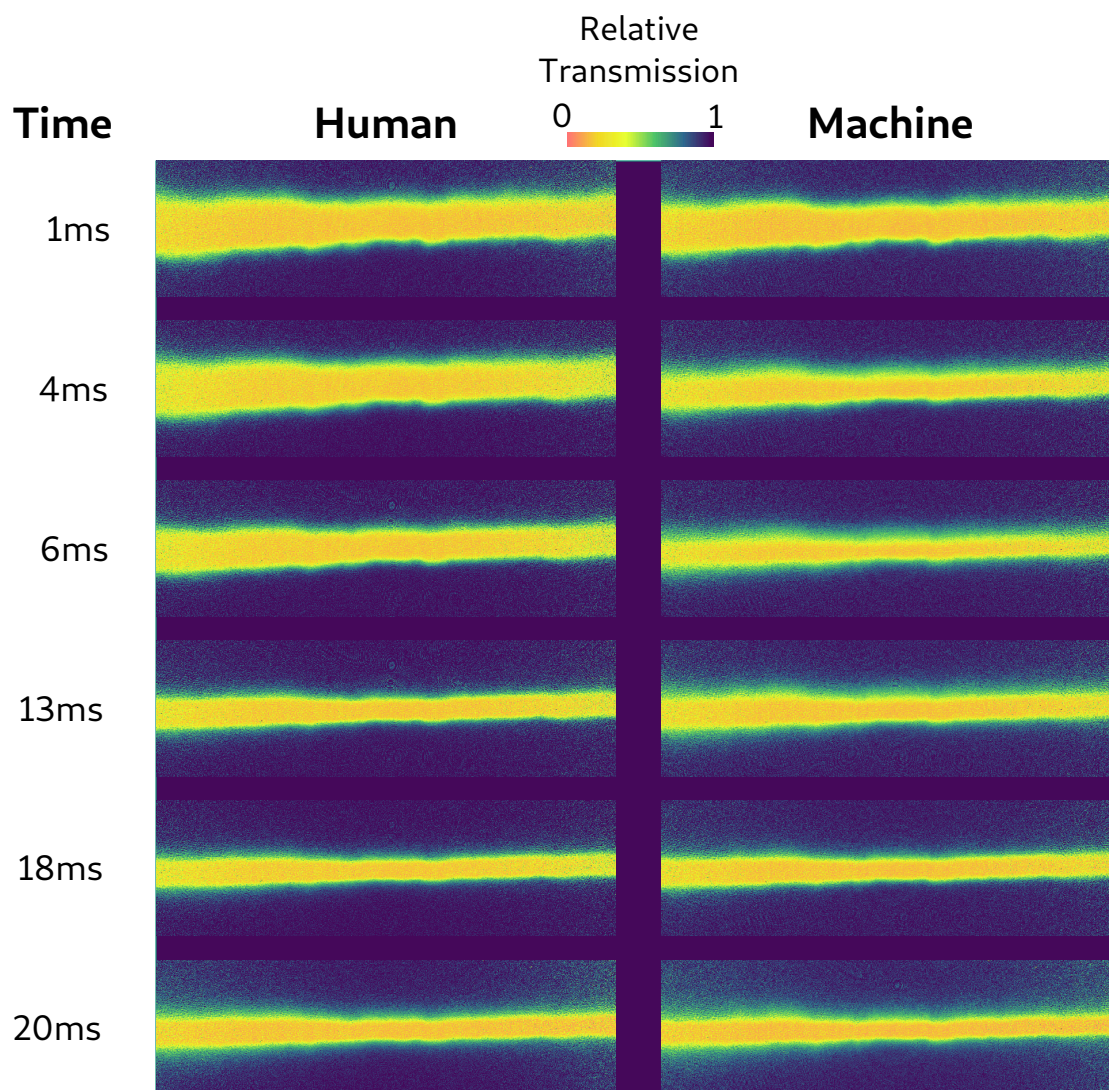


Figure 10.3: Side absorption images of the MOT at chosen times, to illustrate the difference in compression. While the human optimised sequence compresses the ensemble monotonically, the machine learned sequence opts for a multimodal technique, releasing the compression partially during the sequence. Note these images are not calibrated for OD.




Figure 10.4: Evolution of the atom spatial distribution, as a function of compression time, for both the human and machine learned sequences. The human sequence is monotonic whereas the machine learned sequence performs multiple compressions. The absorption traces are normalised to the imaging beam without atoms present, and the maximum observed absorption of each sequence.

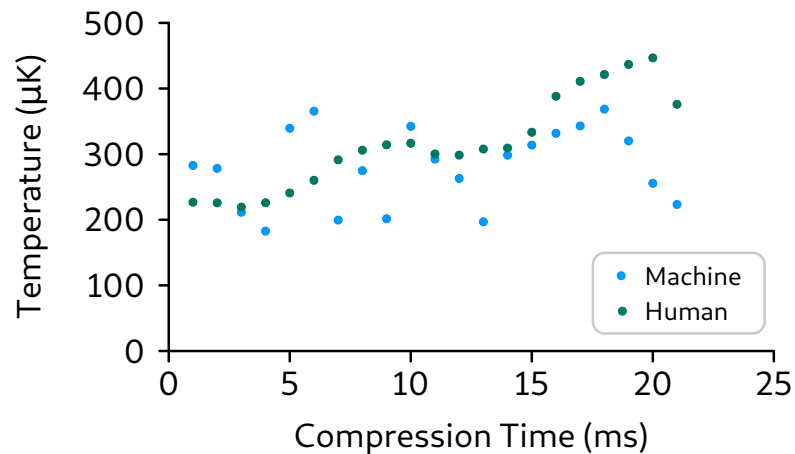


Figure 10.5: Evolution of ensemble temperature as a function of compression time. Each point is determined by performing a ballistic drop measurement after applying the sequence up until that time.

SANN model is that the model is continually making predictions on the efficacy of the proposed parameter sets. Thus we define the optimisation as having converged when the following set of criteria are met:

1. Each network in the **SANN** predicts the minima to be at the same point on average (ignoring exploration).
2. Each prediction generated by the **SANN** matches the observed cost value (ignoring exploration), i.e. the **SANN** well describes the local area.
3. Exploration is no longer yielding better minima after having explored for a time relevant to the scale of the problem.

Once the convergence criteria have been met the optimisation can be halted if it is deemed appropriate. Note that our convergence criteria above do not indicate convergence to a global optimum but instead are more a statement of the **SANN**'s capacity to continue efficient exploration. In the case of more complex landscapes these criteria become more difficult to satisfy and detect. In the case of the current optimisation the landscapes are relatively smooth. An advantage of the **SANN** network is that post optimisation one still has access to the representative landscape of the models. Using the model we can plot one dimensional slices for each parameter to observe the representative landscape as shown in Fig. 10.6. As seen in Fig. 10.6b the one dimensional slices are generally smooth corresponding to a minimum in the vicinity of these parameters. Fig. 10.6c shows the measures that can be monitored during an optimisation run to determine convergence. The first plot shows the convergence between the predicted and measured cost, scaled to the true cost. This effectively shows how optimistic the **SANN** is regarding a given set of parameters. The difference between these two values is shown in the second plot of Fig. 10.6c. The final plot shows the normalised distance from the current set of parameters to the best known set of parameters. This can be used to determine where in the parameter space the **SANN** is exploring. The overlap between the predicted and measured cost coupled with the increased exploration towards the end of the run satisfies criteria 2 and 3, while condition 1 is met via the networks exploring the same local area.

10.3.6 Solution efficacy and local minima

A natural question one might ask is for how long does the solution perform/remain valid. Many machine learning applications struggle with the problem of generality and will often find a solution unique to the system at hand at a given time. For our system, this is not the case. We find that the solution continued to function better than the human solution, with the same relative increase being demonstrably stable over a period of 6 months. However the solution itself is likely unique to this particular system in some fashion. That is to say the idiosyncrasies of the control system and experimental set-up that are unique to this system, are potentially being

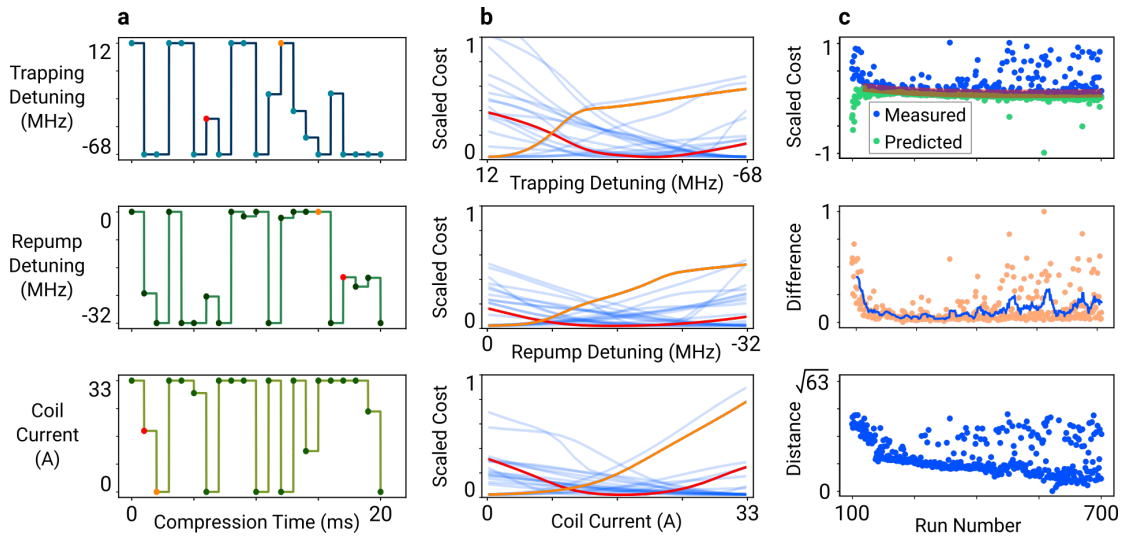


Figure 10.6: Outputs from the [SANN](#) network models. (a) shows the optimal run with the highlighted points being used to generate the similarly coloured slices in (b). The blue curves represent every other parameter slice not highlighted. (c) shows the convergence between the measured and predicted cost according to the [SANN](#) with the red region corresponding to experimental noise. The difference between these two values is shown in the middle plot of (c) with a moving average. The lower plot of (c) shows the distance of a given parameter set from the best observed parameter set which can be used as a measure of exploration within the space.

exploited by the [AI](#) to produce this increase. This is in some way supported by two observations. Firstly by running the optimisation with an experimental fault, the [AI](#) in certain cases produced entirely new solutions that did not perform as well, however still outperformed the human scheme. In this case the new solutions exploited faults, or simply compensated for them. Secondly upon redesign of the laser systems and modifications to the magnetic field driving circuit, the solutions needed to be re-optimised.

In addition to the best solution, through repeated application of the optimisation algorithm, a class of local minima were discovered that corresponded to [ODs](#) up to within a few percent of the best solution. Fig. [10.7](#) shows several of these local minima and their corresponding relative [OD](#) value. Due to the granular nature of the parametrisation it is difficult to find a correlation between these values. Through examination of these solutions, it would appear that the efficacy of the solution is less dependent on the absolute value of any given parameter, but instead relies more heavily on the value of a parameter relative to those around them with common blocks appearing in some of the solutions.

While the performance of any sequence will fluctuate with the daily experimental conditions, we find that the relative efficacy of each sequence is relatively constant. This finding allows us to conclude that the solutions found by the [AI](#) learner, represent a robust physical mechanism rather than only exploitation of the intricacies of the experimental setup.




Figure 10.7: Other identified local minima with their corresponding OD value. Multiple local minima can be identified via repeated application of the optimisation algorithm. Due to the large nature of the space there are classes of local minima with similar values.

10.3.7 Physical mechanisms

The stark contrast between the adiabatic ramps of the human sequence and the machine learned sequence, brings forth questions as to what physical mechanism could be at work here. An unfortunate consequence of working with cold atomic ensembles is that many of the dynamics of interest are either extremely difficult to measure or unattainable. With the current experimental setup we are in general limited to ensemble measurements, such as temperature, OD and state occupancy, due to the high atom number. That being said, there are a few key observations which provide hints as to the physics being exploited.

The dynamics of the MOT will change depending on the number of atoms present. This can be split into three main categories as shown by [197]: the temperature limited regime, multiple scattering regime and two component regime. For our static loading case we expect the MOT to be in the temperature limited regime, having relatively low density and with a spatial distribution close to that of Gaussian. Here the occupied volume is given by the temperature. As we approach the maximally loaded trap and/or begin compression, we enter the multiple scattering regime. In this regime the radiation trapping effect (re-scattering of trapping photons to cause heating) becomes relevant and, in this case, the density of the cloud is independent of the number of trapped atoms. In the centre of the trap the confining force is stronger due to the magnetic fields effect of the optical pumping. If this central region becomes filled then atoms will spill into the surrounding, weaker confinement region. In this case the atoms in the weaker confinement region will spread out with a radius much larger than that of the central dense region. This constitutes the two-component regime. Observing the absorption images of Fig. 10.2e we can see some evidence of this multimodal distribution.

The problem with the above treatment is that this only valid for the static regime, whereas transient regimes introduce complications. When operating transiently it may be that a given scheme can be considered quasi-static, moving from one static regime




Figure 10.8: Absorption images of the MOT captured during the compression sequence without pumping into the imaging state. The population will sometimes occupy the dark state for a period before finally ending up in that state towards the end of the run.

to the next, if the variation is small or slow enough. This is the case for the adiabatic sequence, but is however not true for the machine learned sequence which switches at the same rate or faster than the reconfiguration of the atomic distribution which is observed to take on the order of \simeq ms. Given the similar atomic distribution obtained by the learned sequence, it is likely that the transient path arrives at the two-component regime via a non-adiabatic path. It should be noted that this pathway may not be attainable for all systems e.g. nanofibres and different MOT setups. However there may be other pathways for which this deep learning approach may provide an avenue for discovery.

The temperature of the machine learned run is also consistently lower than that of the optimised adiabatic ramp. As the OD measurement is performed after the PGC, optical pumping and dead time (as would be the case for an quantum memory experiment), there is a non-zero amount of time for which the ensemble undergoes ballistic expansion, lowering OD. Atom loss will also occur if atoms in the trap are resonant to the trapping light. It can be observed that during the course of the learned sequence, the population will switch between the trapping and repump states as evidenced by Fig. 10.8. We theorise that this switching allows the machine learner to continue filling the trap (in concert with the release and capture style mechanism) while avoiding loss due to heating from radiation pressure and light assisted collisions.

These conclusions are supported in part by the observation that in the limit of low atom number the optimisation process cannot find runs that outperform the adiabatic run. In this case the system cannot leave the multiple scattering regime and enter the two-component regime due to lack of atom number, thus it is no surprise that this solution cannot outperform the adiabatic ramp.

10.4 SUMMARY AND FUTURE DIRECTIONS

Here I have presented the first experimental application of a deep learning directed optimisation algorithm, to a complex physical system, the magneto-optical trap. Using

the newly proposed algorithm in Ch. 9, a new transient compression sequence was identified which increased the OD of the ensemble by $81 \pm 3\%$. Additionally this learned scheme also lowered the temperature of the ensemble by $\sim 40\%$. In the preceding section, I propose a possible explanation for the increased effectiveness of the new scheme, linking back to the different regimes explored by early pioneers in the field of cold atoms.

There remains a number of unexplored avenues that can be addressed by future research. For starters, the proposed explanation regarding the different operating regimes can be further explored and quantified by a series of careful measurements regarding the loss rate and density of the ensemble, measurements that were not attainable during the current work. Unfortunately due to the complexity of the system involved simulation is intractable and development of such a theory is outside the scope of the current work.

Another direction is further exploration of the ability of the optimiser to find solutions. The optimal solution shows values which are hard against the parameter limits, corresponding to the physical limits of the system. To this end, a second MOT is currently being designed and constructed. With these limitations in mind, the physical constraints of the system have been expanded to accommodate larger detunings and higher magnetic field gradients. This will allow us to explore the possibility of whether the physical limits of the system are an artefact, or feature, of the optimisation.

In general this method has the potential to uncover rich physics, not bounded by human intuition. As quantum technology becomes more integrated and common place, in-loop optimisation will likely become a goto tool to actively improve the performance of such systems. To this end I believe that the application of this technique to other problems has much to offer.

In the previous two chapters, I presented a new kind of deep learning algorithm for optimisation purposes. While this algorithm is effective, the basic implementation does allow for some flaws regarding more complicated landscapes. In this chapter I will present improvements to this algorithm in an effort to improve relative strengths and address weaknesses. I will also present a number of benchmarks of this new algorithm. The new algorithm is a complete re-write in terms of the code base and additional features, it is named *Acuminata*.

11.1 PREVIOUS ISSUES

After running some benchmarks on the old algorithm it became apparent that while it worked well for certain landscapes, it was underperforming as the objective function became more complicated. Fig. 11.1a shows an example of such a problem, an arbitrarily constructed ‘eggshell’ like function with a global minimum centred on some point. The objective function for an N dimensional version of this problem is given by

$$\mathbf{X} = [x_1, \dots, x_N] \quad (11.1)$$

$$f(\mathbf{X}) = \sum_i (-2 \operatorname{sinc}((x_i - x_{min,i})) + 1) \cdot (1 - e^{-(x_i - x_{min,i})^2/3}), \quad \forall i \in [1 \dots N], \quad (11.2)$$

where $x_{min,i}$ gives the point at which a given parameter, x_i , will return the minimum value. For a 50 parameter problem, the neural net approach described previously is able to outperform the Gaussian process and DE algorithms, however as shown in Fig. 11.1b, the fruit fly algorithm is outperforming the current algorithm substantially. Complicated structures are observed to trap the SANN in regions of steep local minima, even with the DE learner providing additional information. In this case the SANN effectiveness is limited by the exploration ability of the DE learner which is far slower than the fruit fly algorithm. Additionally the effectiveness of the SANN is highly dependent on the initial sampling of the data. With more complicated landscapes this sampling becomes more crucial to the success of the algorithm and oftentimes the results are highly variable as a result.

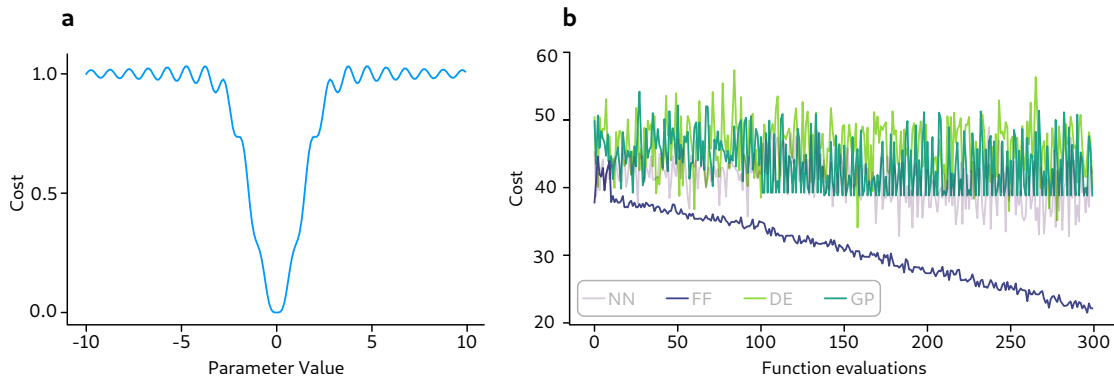


Figure 11.1: Optimisation over an eggshell benchmark function. (a) shows the general landscape shape for a 1 dimensional parameter slice, the total function for a multi-dimensional problem is simply the addition of each of these profiles. (b) shows the results of a 50 dimensional optimisation run using the previous deep learning algorithm, fruit fly, differential evolution and Gaussian process. As can be seen the fruit fly algorithm is now outperforming the neural nets.

11.2 METHOD IMPROVEMENTS

Fortunately many of the previously identified issues can be rectified by making some simple observations. There is an old adage that also applies in the machine learning field, ‘*garbage in, garbage out*’. This statement summarises the fact that machine learning, while powerful, is not magic. If one feeds any given model nonsense data, then one should expect to receive nonsense answers. This is true also in the context of optimisation. The power of our approach relies on the ability of deep learning to generalise the landscape in an efficient and meaningful manner. However, if the algorithm has never seen any good points, how can it possibly be expected to produce a good prediction? Clearly this will be limited by the effectiveness of our sampling. Furthermore, escaping local minima has thus far been left up to the stochastic elements of the algorithm, be it overfitting or the DE learner. While this approach will *eventually* yield results, the assumption that this process should be entirely stochastic is to throw away valuable, learned information about the landscape. In the following section I will outline various improvements to the algorithm and their respective motivations.

11.2.1 General cost functions

Until now we have assumed nothing about the physical landscape, be it structure, shape or otherwise. This is done to preserve generality in the approach to solving these problems. However many landscapes that are found within the physical world, that we may be interested in solving, often have certain characteristics. For example, there may be an overall coarse structure, the approximation of which is the starting point for many surrogate techniques. Consider the two cost landscapes shown in Fig. 11.2. Sub-figure (a) denotes a particularly hard problem to solve, if one has not sampled well, then potentially the algorithm will have no knowledge of one, or many, of these




Figure 11.2: Two separate objective function landscapes with different characteristics that pose a difficult optimisation problem. (a) presents a sampling problem, as knowledge of different parts of the landscape does not facilitate inference as to new local minima. (b) is a complex landscape that presents a more friendly approach, while there are local minima, an approximate solution to the problem can help to infer the location of better parameters (shown in green).

minima. Furthermore this information cannot be gained from knowledge of other parts of the landscape. (b) denotes another complex problem, and could be worsened by increasing the steepness of each local minimum, however a simplified model of the region can facilitate exploration of more valuable regions. While both problems can be solved with a stochastic approach, a more realistically valuable approach would be to use the information inherent in the landscape where possible. Thus we arrive at an assumption that we will use: the cost landscape has some useful coarse structure. In many real world applications this assumption is justified. Note that this assumption does not preclude us from solving either (a) or (b), it just allows us to accelerate the solutions for certain classes of problem.

11.2.2 The sampling problem

The effectiveness on the [SANN](#) network is heavily tied to the sampling for a number of reasons. Firstly, the initial sampling will create an initial bias in the network. As the training occurs continually on an ever increasing data set, the initial bias imparted by the sampling is difficult to remove, as it is continually reinforced. Secondly any introduced bias is reinforced via the fit-predict procedure, as the networks are more likely to predict good results in certain areas if they already return a good result. Due to this, it is important to impart as little bias as possible during the sampling process to produce better generalisations. The question of sampling is difficult, as is quickly discovered by taking the naive approach. One might assume taking a uniform distribution in each parameter will produce the desired outcome. Such an approach is shown in [Fig. 11.3](#). Taking the distance between any sampled point and the centre of the parameter space allows a measure of the region of the objective function that is being sampled. For a uniform distribution in parameter space we can see that the distribution is not uniform in distance (see the histogram of [Fig. 11.3](#)). The problem of course with this, is that the [SANN](#) will receive initially sampled data biased towards




Figure 11.3: Sampling method, where samples are determined by taking a uniform sampling of each parameter, to determine the coordinates in the objective function space. The scatter plot shows the distribution of 1000 samples. The histogram shows the frequency of the distance from the centre of the space $(0,0,0)$, which is not a uniform distribution.

a particular region of the objective landscape. As the training mechanism will bias towards more heavily sampled areas, already we have imparted a bias to the SANN. We wish instead to find a sampling technique that attempts to sample uniformly in terms of the distance, while still being stochastic in nature.

IMPROVED SAMPLING METHOD As shown in Fig. 11.3, sampling the parameters uniformly will create a distribution in distance space similar to a spherical surface with a characteristic width. The SANN will necessarily benefit from an unbiased sampling method which is difficult to attain, due to the unknown nature of the objective function. However we have empirically observed that for landscapes that exhibit coarse global features, modifications towards an unbiased sampling method provide the SANN with enough information to build an effective rudimentary model of the space.

11.2.3 Local minima methods

An observed issue with the previous algorithm is the proclivity to get stuck in local minima, especially if they are steep. Defining whether a minimum is global or local requires knowledge of the cost landscape, however for the most part we will always assume that any minimum encountered is local, facilitating exploration of the landscape. We wish to implement some method to escape local minima that takes advantage of knowledge of the landscape gained over time. Different methods for escaping local minima have been proposed in different contexts, such as restarts in derivative free optimisation [213], or momentum in gradient descent optimisers [148].

IDENTIFYING LOCAL MINIMA First one must establish criteria for determining if the algorithm is stuck in a local minimum. This is done by simply assessing how long the SANN has been in the region around the best known parameters. This is simply a

tolerance on the distance from the best known set of parameters. This tolerance can be modified to change the behaviour of the algorithm, forcing more or less exploration. For the most part the **SANN** will tend to get stuck in steep local minima, thus this technique of measuring the average localisation of the algorithm is an effective technique.

BUMPING TECHNIQUES The techniques used to escape local minima are termed bumping techniques, as they effectively try to 'bump' the **SANN** out of local minima. The techniques can be categorised into two main categories, stochastic and informed. The stochastic methods perform similar to that of Monte Carlo methods, choosing random parameters close to the local minimum. Informed methods attempt to use information gained from observed parameters up to this point, to infer the next best parameters. Without a priori knowledge of the objective function landscape, we are unable to determine which method will be the most effective for a given problem. As such, each method is cycled during the optimisation when it has been determined the **SANN** is stuck.

11.2.4 Exploration by overfitting

In general machine learning problems, overfitting is an undesirable behaviour that can decrease the generality of a given model. As mentioned previously, this algorithm relies on the willingness of the neural nets that comprise the **SANN** to overfit the observed data to some extent. This overfitting can be tuned via the regularisation coefficient applied to the weights, examples of which are given on the Ackley test problem in Fig. 11.4. As shown, this parameter can be crucial to the success of the algorithm, as too much or too little exploration will reduce the effectiveness of the algorithm. This choice of parameter will, in some manner, be mildly problem dependent. In general, this parameter can be chosen to bias towards exploration, in conjunction with exploitation methods.

11.2.5 Local exploitation

The effectiveness of a given global optimisation technique will always in some manner depend on the trade-off made between exploration and exploitation. The **SANN** is generally effective at regulating its exploration, however there are times when the networks can get stuck in certain areas via introduced biases or explore too much. Oftentimes it can be instructive to force the networks back to a particular area to reduce these biases and facilitate local exploration or exploitation. This also removes some of the dependence associated with the regularisation on the problem. The time spent in either exploration or exploitation can be determined dynamically for problems where one mode may be favoured over another. It was observed that the optimisation algorithm is stable with a 1:1 duty cycle between exploration and exploitation, subject to neither mode finding better minimum.




Figure 11.4: Different regularisations applied to the neural networks of the [SANN](#) to produce different exploration results. For the higher regularisation a lack of overfitting reduces exploration and thus the effectiveness of the algorithm. Conversely, too much overfitting also hampers to the effectiveness of the algorithm which avoids exploitation. To perform effectively, the algorithm is required to balance exploration and exploitation, as is the case for the middle plot which has identified a better minimum. The L2 regularisation coefficient for each optimisation run is noted above each plot.

11.2.6 New data generation

The previous version of the algorithm included a [DE](#) algorithm, to ensure a non-biased source of data and facilitate continued exploration. With the inclusion of the aforementioned bumping, exploration and exploitation techniques, the algorithm no longer requires the inclusion of this learner. The inclusion of these techniques is effective and stochastic enough that the evaluations otherwise reserved for the [DE](#) learner can be better utilised by the [SANN](#) predictions.

11.2.7 Improved performance

Applying the complete set of techniques outline above we can observe the performance of the new algorithm applied to the Ackley test problem. I have previously mentioned this problem, however we will now define it formally. The Ackley function is one of a number of optimisation benchmark functions used to assess the performance of a given procedure. The Ackley function is non-convex and defined as,

$$f(\mathbf{x}) = -a \exp \left(-b \sqrt{\frac{1}{N} \sum_{i=1}^N x_i^2} \right) - \exp \left(\frac{1}{N} \sum_{i=1}^N \cos(cx_i) \right) + a + e \quad (11.3)$$

$f(\mathbf{x}_{min}) = 0$, where $\mathbf{x}_{min} = (0, \dots, 0)$

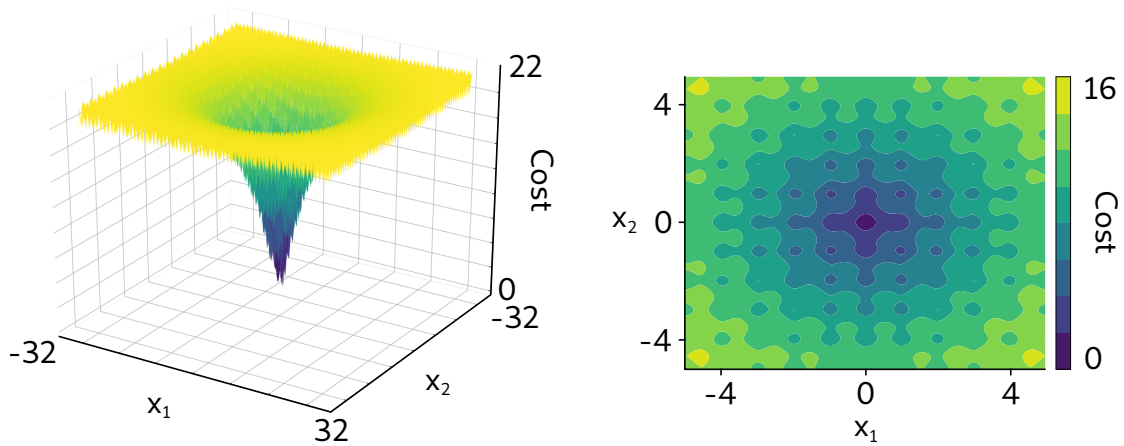


Figure 11.5: Ackley test problem landscape shown as both a 3-dimensional and contour plot, for the 2 parameter problem. The landscape exhibits a number of local minima with a global minimum at the centre of the landscape. The contour plot is shown for a smaller region of $[-5, 5]$ around the global minimum, illustrating the different length scales of the eggshell structure.

where the recommended settings are, $a = 20$, $b = 0.2$ and $c = 2\pi$. Here, e is Euler's number. This complicated function creates a difficult landscape with many local minima and a single steep global minimum at 0, with the problem generally solved within the bounds $[-32, 32]$. Fig. 11.5 depicts the 2-dimensional landscape which can be generalised to N dimensions.

To ascertain the performance of this algorithm with respect to the old algorithm and other popular methods of optimisation, we can attempt to optimise the Ackley function for a high number of dimensions, as this is what the deep learning algorithms do well, and, what this algorithm was designed for. Fig. 11.6 shows the results of a 100 dimensional test problem, where it is shown that the new algorithm is significantly outperforming the old algorithm and other popular techniques. The new algorithm (labelled **AC** after the naming of the code, *Acuminata*) initially only samples for 50 points before beginning the **SANN** prediction procedure. The initial discrete jumps in the first 50 runs correspond to the improved sampling technique, which has no a priori information. It can be seen that the improved sampling also provides a better representation of cost function in the initial 50 samples. This improvement does not vary substantially on repeated attempts. Additionally the switch between exploration and exploitation modes can be observed, with the exploration modes corresponding to the high cost evaluations interspersed throughout the optimisation.

11.3 OPTIMISATION BENCHMARKS

As noted briefly, there are a number of benchmark functions that are used to test the efficacy of a given technique. The appeal of using deep learning is the ability of the algorithm to generalise in high dimensional data sets. We are thus interested




Figure 11.6: Benchmark results of multiple optimisation algorithms on a 100 dimensional Ackley test problem. The costs observed by other optimisation algorithms indicate that they are not yet within the main central region of the global minimum and are still stuck on the local minima located on the initial edge. The benchmark algorithms are the old neural net algorithm (NN), fruit fly (FF), differential evolution (DE), Gaussian process (GP) and Acuminata (AC).

to determine how the algorithm scales with the problem dimensionality. For many global optimisation techniques, it is known that with increasing problem dimensions the efficacy of the solution found (for a fixed number of evaluations) often decreases dramatically. To test this, each algorithm is given 2000 evaluations with which to find the global optimum, for a problem of dimension N . For each benchmark function the number of dimensions is increased up to 300 dimensions, to identify the scaling of each algorithm's performance with problem dimensionality. A handful of benchmark problems common to the optimisation space were chosen as test problems, the specifications of which are given in Table. 11.1.

Fig. 11.7 shows the results of these benchmarks with increasing dimensions. As the fruit fly (FF) algorithm was found to outperform the other algorithms on the benchmark functions at hand, only the FF algorithm has been included for reference. FF is run 5 times (each with 2000 evaluations) on each problem to give an indication as to the variance of the method. Acuminata by contrast is approximately as stable in achieving its final value. With the interests of compute time in mind (FF is far cheaper in terms of compute time), Acuminata was tested once per dimension. However, one can still gather information on the methods variance via the portrayed linear relationships. As is shown in Fig. 11.7, Acuminata scales approximately linearly with increasing dimensions for each benchmark function, substantially outperforming FF for higher dimensions. From Fig. 11.7, it can be seen that Acuminata has the most trouble with the Hoop and Alpine functions. This is possibly due to the varying length scale of the features in each problem. However, despite this, Acuminata still scales better than FF for higher dimensions.

This series of benchmarks demonstrates Acuminata's effectiveness on a number of different benchmark functions. This is also evidence that the improvements to the

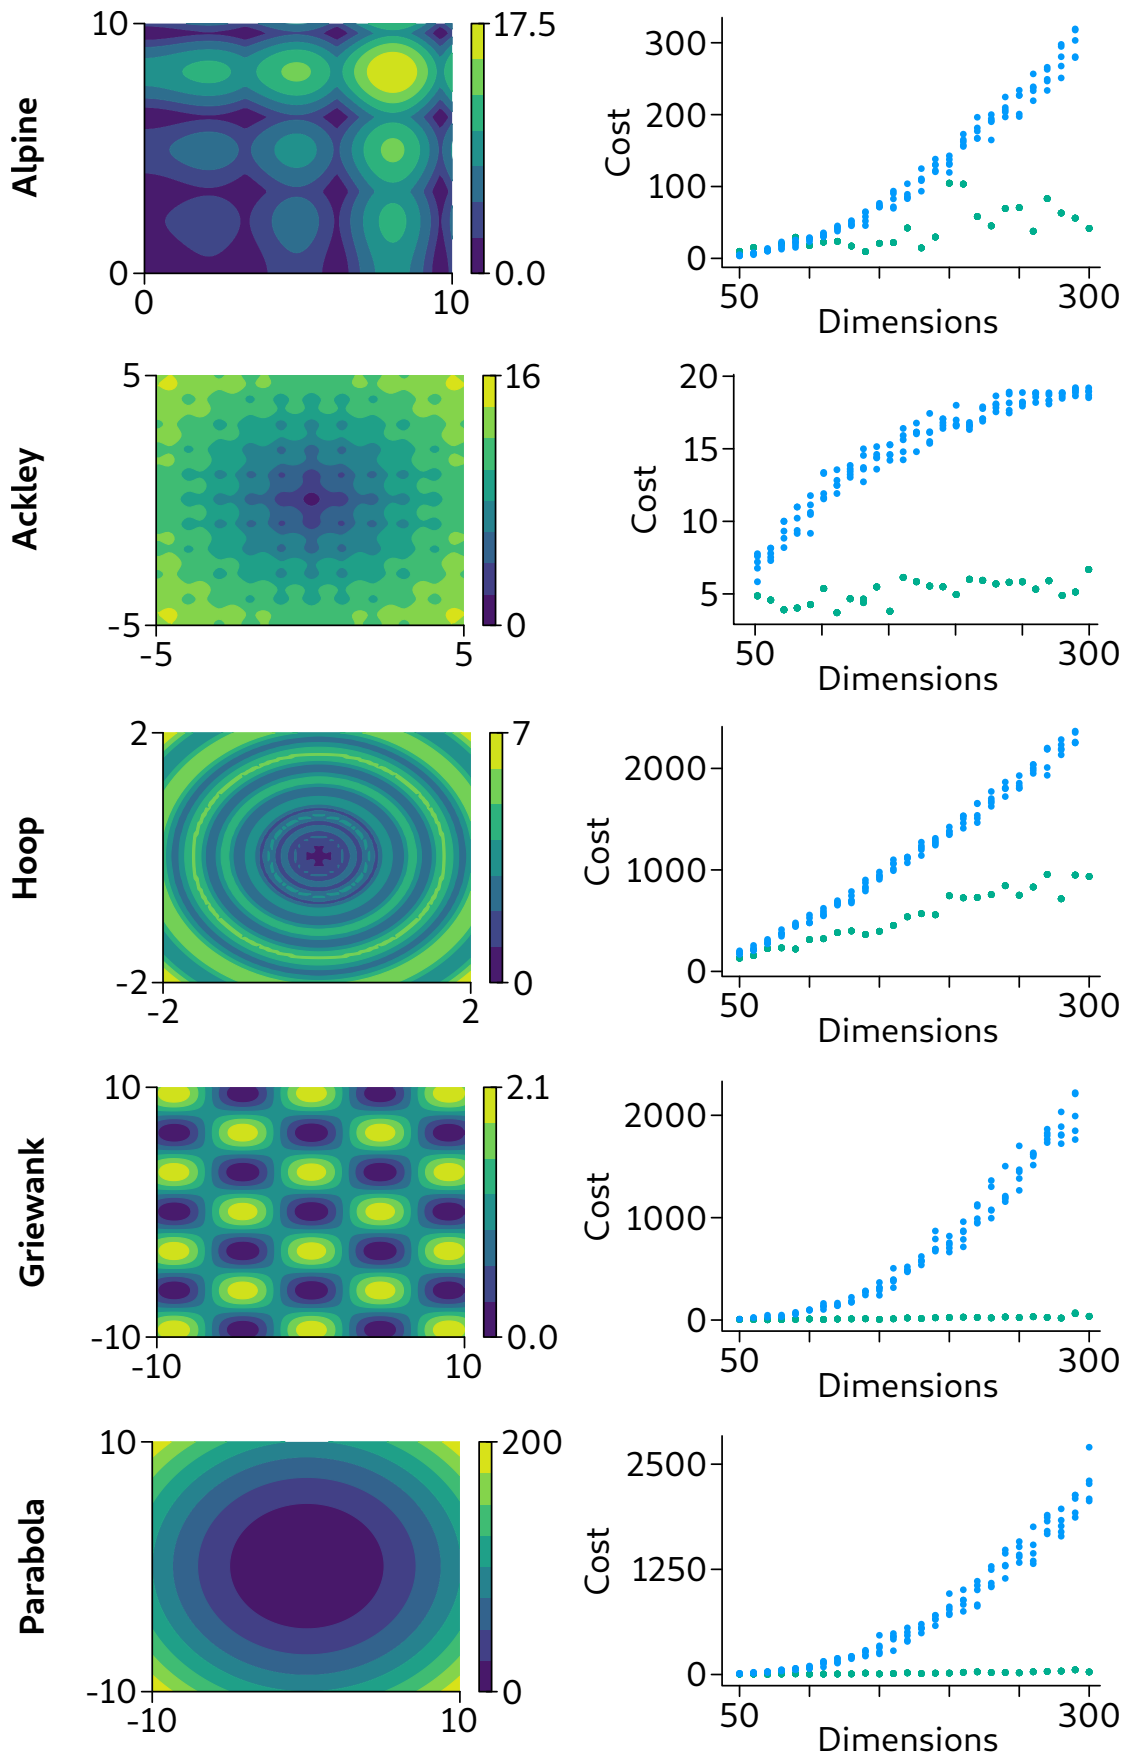


Figure 11.7: Results of applying the fruit fly (blue) and Acuminata (green) to different benchmark problems. The first plot shows the representative landscapes for a 2D version of each problem. The scatter plots represent the best cost found for each problem with increasing dimensions.

Name	Functional Form
Ackley	$f(\mathbf{x}) = -20 \exp \left(-0.2 \sqrt{\frac{1}{N} \sum_{i=1}^N x_i^2} \right) - \exp \left(\frac{1}{N} \sum_{i=1}^N \cos(2\pi x_i) \right) + 20 + e$ $\mathcal{D} = [-32, 32], f(\mathbf{x}) = 0 @ \mathbf{x} = [0 \dots 0]$
Alpine	$f(\mathbf{x}) = \sum_{i=1}^N x_i \cdot \sin(x_i) + 0.1x_i $ $\mathcal{D} = [0, 10], f(\mathbf{x}) = 0 @ \mathbf{x} = [0 \dots 0]$
Griewank	$f(\mathbf{x}) = 1 + \sum_{i=1}^N \frac{x_i^2}{4000} - \prod_{i=1}^N \cos\left(\frac{x_i}{\sqrt{i}}\right)$ $\mathcal{D} = [-600, 600], f(\mathbf{x}) = 0 @ \mathbf{x} = [0 \dots 0]$
Hoop	$f(\mathbf{x}) = \sum_{i=1}^{N-1} f_H(x_i, x_{i+1}) \text{ where,}$ $f_H(x, y) = (x^2 + y^2)^{0.25} \cdot (\sin(50(x^2 + y^2)^{0.1}))^2 + 1$ $\mathcal{D} = [-100, 100], f(\mathbf{x}) = 0 @ \mathbf{x} = [0 \dots 0]$
Parabola	$f(\mathbf{x}) = \sum_{i=1}^N x_i - x_{i,min}^2$ $\mathcal{D} = [-600, 600], f(\mathbf{x}) = 0 @ \mathbf{x} = [x_{1,min} \dots x_{N,min}], x_i \simeq \mathcal{U}(-100, 100)$

Table 11.1: Definitions, domians and minima of the various benchmark functions used. All benchmark functions are found within optimisation literature as standard test problems.

method have removed biases inherent in the previous methodology and increased the general efficacy of the approach.

11.4 REAL WORLD APPLICATIONS

Following the completion of Acuminata, it has been applied to a number of physical systems with demonstrable success. While Acuminata was designed with experimental optimisation in mind, the classes of problem that it can be applied to extend beyond this application. Put simply, there are a number of optimisation use cases within industry where the evaluation of an objective function are expensive, yet require optimisation through some method. For example within manufacturing there are machines which require manual optimisation by some experienced human operator. This can be an extremely costly process, with the knowledge of how to perform such an optimisation difficult to transfer in many cases.

To test the efficacy of Acuminata in a real world setting, a number of projects have been undertaken which comprise difficult optimisation tasks within a physics setting. These are outlined in the following sections.

11.4.1 Single electron transistor

Silicon quantum computing has attracted much attention in recent years due to the promise of scalable architectures that can leverage the extensive technology employed by the semiconductor industry [214, 215]. While there are a number of approaches to




Figure 11.8: Results of the optimisation of the single electron transistor operating point using Acuminata. While only 4 dimensional, the problem has a complex structure with experimental drift. The blue line is the observed costs and the green line is the running best. The 4 parameters correspond to sensor top gate, sensor left barrier gate, sensor right barrier gate and source drain bias of the single electron transistor.

quantum computing with silicon, one of the more promising is the confinement of single electrons within a silicon metal-oxide-semiconductor quantum dot device [216, 217]. These devices have demonstrated 1 qubit operations with control fidelities above 99.9% [218, 219] and 2 qubit control fidelities above 94% [217]. In all of these cases however, the operation of these quantum devices requires a certain amount of tuning. In particular readout of these devices occurs by monitoring the tunnelling of a single electron from the quantum dot to an attached reservoir via a single electron transistor charge sensor [220, 221]. The charge sensitivity of this device can be dynamically adjusted via the voltage of four attached gates, which are patterned onto the silicon substrate [220]. For the optimisation we wish to control 4 gates: sensor top gate, sensor left barrier gate, sensor right barrier gate and source drain bias. This set of gates controls the electron occupancy. Choosing the operating parameters before initiating control is generally achieved by performing a grid search which identifies a good operating regime. Due to the complex nature of the parameters, this requires a fine grain grid and/or an experienced human operator to be effective. Additionally the system can drift in a time shorter than the grid search completion, adding to the complexity of the search.

This system provides an ideal test bed for Acuminata, although the dimensions are small the landscape is complex due to the nature of the physics involved, containing many local minima. Fig. 11.8 shows the results of a 4 parameter optimisation which successfully identified a number of operating points. The cost is measured by acquiring a charge readout, with a given set of parameters, and calculating the variance of a set of moving variances. The ideal operating point is where the noise on the signal is small, but the charge jump can be easily discriminated from a zero reading. As can be seen, the cost landscape is extremely complicated, with small changes in parameters corresponding to a large change in the cost function.

METASELECTION As there are many operating points that may be selected, we employ a kind of metaselection post optimisation. This is done by taking the best 10 points and measuring their efficacy. Certain points will remain stable over long




Figure 11.9: 10 dimensional Ackley problem with a drift that increases over time. Without the compensation the algorithm struggles to follow the minimum. With compensation Acuminata follows the drift in the landscape over time. The first 200 runs of each optimisation constitute training.

time periods (hours to days), while other operating regimes will only exhibit transient efficacy on a much shorter time scale (seconds to minutes). As such if the same 10 points are measured after some time, these quasi-stable points can be identified and the best can be selected as the operating parameters.

COMBATING DRIFT A feature of this system that makes conventional optimisation difficult is the inherent sensitivity to temperature and stray electrical and magnetic fields, producing experimental drift. In general, for effective optimisation the observed experimental drift must be longer than it takes to acquire a single experimental run. This is not a stringent requirement with a majority of experimental setups satisfying this criteria. Drift compensation can be implemented by including an attention like system into the training methods. An example of this inclusion's effectiveness is shown in Fig. 11.9. Here drift is added to the 10 dimensional Ackley problem (in all dimensions) which the algorithm is required to track. It can be seen without the compensation the algorithm has difficulty following the drifting parameters due to the inherent biases in the model. However with compensation enabled, Acuminata will follow the drift while continuing to optimise. This type of compensation is not a feature that can be easily implemented on many of the global search strategies, which rely on a static landscape. However, the training and predictions of Acuminata and the deep learning, provide a flexible and intuitive structure for the inclusion of many such mechanisms.

After a solution has been identified, the observed parameter and cost pairs can be examined to determine the relative relationships between the parameters. This also allows one to investigate what parameters are drifting, and which parameters have more effect on the systems performance.

11.4.2 Quantum dot frequency shaping

Many quantum information processes in quantum optics rely on the ability to experimentally create single photon states for use as qubits. This often has stringent requirements such as indistinguishability, on demand generation and purity (i.e. a vanishing probability of multiple photon states). A promising platform for generating single photons is quantum dots embedded in cavity structures, having demonstrated high brightness while maintaining high indistinguishability [222]. Single quantum dots can be positioned within InGaAs pillar microcavities, in situ, using lithographic techniques, providing a pathway to scalable fabrication [223]. A remaining challenge is the reduction of noise introduced during the pumping process. For the quantum dot to emit a single photon it must first be coherently excited from the ground state, generally by a pump laser of some description. However, this provides an obvious challenge: pump light used to excite the quantum dot may be scattered by the cavity, introducing noise photons. Filtering can be attempted with polarising optical elements, given that the pump can operate with a defined polarisation. However, this will also facilitate the filtering of single photons emitted with this polarisation, reducing the brightness and efficiency.

One possible avenue to alleviating this problem is to spectrally shape the pump to preserve excitation efficiency, whilst simultaneously maximising the spectral separation between the pump and single photon emissions. Recently it has been demonstrated that this coherent driving can occur with a dichromatic approach [224]. However, this can also introduce additional losses within the solid state structure, which need to be carefully accounted for. This approach provides another case for deep learning optimisation, as the exact dynamics of a QD system can be difficult to characterise. Online optimisation of such a structure only requires some ability to measure the performance of the source, with the complexities simply contained in the system response. Arbitrary spectral shaping can be obtained using a setup as depicted in Fig. 11.10 [225]. As depicted, a SLM can be used to spectrally shape an incident beam by applying different 2D pixel masks. The full 2D pixel array of the SLM provides a high dimensional space that can be optimised, with a number of applicable parametrisations. Such a space is an ideal use case for the deep learning optimisation approach.

Currently, the optimisations that have been trialled have focused on using second order correlation measurements as the cost function, thus improving the spectral filtering by reducing the noise. Preliminary results have also shown promise in potentially optimising other components of the system. The ability to improve the brightness of the source while maintaining characteristics, such as indistinguishability and singular photons, is a problem that plagues all single photon sources.




Figure 11.10: General experimental setup for performing spectral shaping using a spatial light modulator (SLM). A diffraction grating is used to spatially separate the spectral components of an incident optical pulse. The SLM is positioned at the Fourier plane and can be used to modify the phase and amplitude of the different spectral components, by applying masks to the SLM pixels. The filtered beam is then recombined on a second diffraction grating.

11.5 FURTHER MACHINE LEARNING APPLICATIONS

Acuminata demonstrates the potential success of deep learning methods when applied correctly to a problem. The ability to efficiently explore higher dimensional spaces opens up a number of experimental opportunities that were previously out of reach.

11.5.1 Actor-critic networks

One such example is the problem of reinforcement learning applied to experimental setups. While reinforcement learning covers a wide variety of methods, I am mainly concerned with the AC network approach. In short, reinforcement learning aims to learn the best way to map a situation to an action, maximising some numerical reward [226]. A simple, naive approach to this may be a trial-and-error approach for example. However the field of reinforcement learning presents many more sophisticated approaches. In essence AC networks attempt to solve this problem by simultaneously learning the policy and value functions, for a given environment. A policy function provides a mapping between the current state of the environment and an action to be taken. An effective policy is one that always maximises our predefined measure of success, a reward function. The value function, on the other hand, maps a given state to a potential value. This value represents the total cumulative reward that an agent can expect to gain starting from a given state, i.e. while many paths may lead you to your goal, some paths will be better than others. While the reward function offers immediate feedback, the value network offers information regarding the long term value of a given state.

A schematic diagram of an AC network is presented in Fig. 11.11. The actor and critic models constitute neural networks of some form, generally MLPs. The actor model takes as an input the current state of the environment and maps to an action to perform.

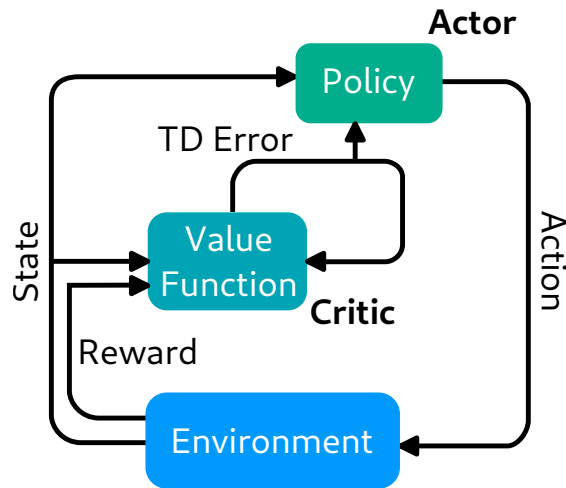


Figure 11.11: Actor critic schematic for reinforcement learning. The actor enacts an action on the environment and is provided feedback by the critic in terms of the TD error. The critic learns the value of each state which provides the basis for advice provided to the actor network.

The resultant reward and next state is fed to the critic which assesses the value of the state by the temporal difference TD error. The only feedback the actor receives is from the critic which informs the actor network how to improve its response. AC models are a fairly recent development in RL but they have demonstrated success in sequence prediction [227]. At present, the application of AC models to auto-alignment of an optical system is being pursued, the details of which are presented in Appendix A.

11.5.2 Learning transfer functions

A transfer function is a function which theoretically maps a system's inputs to its outputs. For a system with a given set of inputs u and outputs y , we are interested in determining the transfer function, $G(s)$, such that

$$Y(s) = G(s) \cdot U(s) \rightarrow G(s) = \frac{Y(s)}{U(s)} \quad (11.4)$$

where $Y(s)$ and $U(s)$ are the Laplace transforms of the output and inputs respectively. If one has access to the input and output functions then the transfer function is simply a deconvolution between y and u . However deconvolution operations are numerically difficult, often unstable and in certain cases unattainable by direct calculation. In some cases it is possible to infer a transfer function based on the behaviour of the system. This relies on the ability to infer the correct transfer function, based on a measure of correct system performance. This necessitates the ability to construct arbitrary functions, a feat that quickly becomes a high dimensional problem if the function exhibits even minor complexity. A scalable approach to this problem is presented in Appendix C, in the context of precise gravitational sensing using cold atom interferometry.

11.6 SUMMARY AND OUTLOOK

In this chapter I have presented a number of improvements to the deep learning optimisation approach, presented in Ch. 9, comprising a complete re-write into Acuminata V1.0. The improvements facilitate the ability of the algorithm to efficiently explore high dimensional spaces, with complicated and non-convex structures. The benchmarks presented in this chapter clearly demonstrate the ability of this method to scale to extremely high dimensions compared to that of conventional optimisation methods. The cost for this performance gain, is computational time. This method lends itself to physical situations, however, when the the objective function can be efficiently computed methods such as FF are to be favoured. Future improvements to the efficiency of training and minimisation procedures will help to improve this. The efficacy of Acuminata was also demonstrated on two additional physical systems, which are continuing work. Finally this chapter closes with additional applications in machine learning that are currently being concluded, with preliminary results found in Appendices A and C.

While the improvements in Acuminata are substantial, the application to additional systems will continue to yield improvements and additions, improving efficacy of the method. Potential future projects include, but are not limited to, optimisation of quantum memory pulses and storage operations, optimisation of 1 and 2 qubit gate pulses and readout in silicon based QDs, unitary operation design from sets of quantum gates and brightness optimisation in QD single photon sources. Aside from optimisation, there is also substantial opportunity for using the AC control type algorithms to integrate with quantum experiments which require online active control. The methods outlined for the transfer function construction allow a general purpose method for constructing arbitrary functions, which can prove useful as a basis for pulse construction in different bases, further facilitating optimisation techniques described above.

12 | CONCLUSION

The work covered in this thesis has been broadly aimed at both implementation of improved protocols for quantum memories and the subsequent optimisation of such systems in the general sense. Initially we covered the operation of a magneto-optical trap, a tool used by scientists for decades in fundamental studies of matter light interaction. Here we found a novel characteristic of the high OD, elongated ensemble. Namely, that this setup facilitates the coherent emission of scattered light spontaneously along the axis of highest OD. While this effect has been observed before, this is the first instance of such an effect occurring without prior stimulation from an additional beam. Additionally this coherent emission provides a method for aligning experimental beams to the highest density region of the ensemble.

Following this, the implementation of a Raman quantum memory with backward retrieval was detailed. We showed with simulation, that one can expect enhanced performance from backward retrieval dynamics, by sidestepping many of the pitfalls of forward configurations. However this configuration requires careful alignment to perform adequately. Fortunately the atomic emission provides a convenient mechanism for aligning these beams to the high OD region of the MOT. We demonstrated an efficiency of $65 \pm 6\%$ with a time-bandwidth product of 160. This marks the first experimental implementation of a backward Raman memory.

OD is the limit of many atomic quantum memory implementations, as it necessarily governs the interaction strength between incident optical fields and the atomic ensemble. With this in mind, we wished to increase the performance of the quantum memory schemes implemented in the cold atom MOT, by increasing the OD using machine learning. This push yielded the first application of a neural network controlled algorithm to a physical system, in an online optimisation context. The transient compression sequence was targeted as the point of optimisation, as this potentially could contain the most rich dynamics. Using this method we successfully improved the OD of the ensemble from $580 \pm 8 \rightarrow 960 \pm 20$. This method also uncovered a non-standard method for traversing different MOT regimes, and produces a much higher OD ensemble, with a lower final temperature.

Following the success of this approach to our physical system, we went on to explore the efficacy of this algorithm in different contexts. Unsurprisingly, for more complicated landscapes this algorithm was outperformed by some more mature approaches such as FF algorithm. A number of substantial improvements were made to the algorithm which avoided these pitfalls, improving the performance of this algorithm beyond conventional techniques. Through a series of standardised benchmarks we find that the performance of this new algorithm, known as Acuminata, scales linearly with the prob-

lem dimensions, up to at least 300 dimensions. Following this marked improvement, I presented some additional physical systems to which we are applying Acuminata and some preliminary results.

Finally, we touched on the future machine learning applications that we are presently pursuing, namely that of the AC control algorithm for auto alignment and the construction of a transfer function from experimental gravimeter and seismometer data. Both of these projects are detailed in the appendix and are currently undergoing preparation for publication.

Future research directions

There are a number of opportunities for extending the research outlined in this thesis. Firstly, there is the possibility to further study the atomic emission, as this may provide some insight into the collective behaviour of the ensemble and open further pathways to atom-light interactions. This can come in the form of more detailed measurements and characterisation, along with attempts to distil a simplified model.

Quantum memory investigations are continuing on many different platforms and there remains no shortage of future research directions for this system. To present these methods as viable quantum memory technologies we must continue to increase the efficiency of this system. By using the tools outlined in this thesis for machine learning it should be possible to improve the memory performance via optimisation. This can be extended to the single photon regime, with the possibility of learning efficient gate operations in memory, such as cross-phase modulation.

The deep learning approach to physical systems has opened up a myriad of opportunities for future research, some of which are being pursued already. On the optimisation side of things, we wish to continue applying this algorithm to novel physical systems, in both the quantum and classical domain. It has already demonstrated a propensity to elicit non-intuitive and physically interesting solutions to problems. We expect that the application of this algorithm to more systems will yield breakthroughs in research and help to facilitate quantum technology design and implementation.

In terms of general machine and deep learning techniques applied to physical systems, there is seemingly no end to the possible research directions. The actor-critic methods have the potential to be applied to a wide variety of control problems in the quantum domain, such as unitary gate design and device control. The ability to traverse high dimensional landscapes efficiently also opens up opportunities for construction of arbitrary pulse sequences using the methods described in Appendix C. Another potential avenue is the application of surrogate models to predicting chaotic and complex systems, to facilitate control and predictive methods.

To summarise, we have shown that even systems as studied and common as a magneto-optical trap, can still demonstrate surprising physics. Furthermore, that the application of machine learning to physics problems has the potential to yield novel and rich dynamics, unburdened by the limits of human intuition.

Part IV

APPENDIX

A

ACTOR-CRITIC MODEL FOR AUTO ALIGNMENT

An incalculable amount of time and effort is spent on aligning optical systems for quantum optics experiments and large scales experiments, such as LIGO. In general this is done by a human operator as the exact situation will determine the way in which this alignment should be performed. Once a given system is aligned, barring any major incident, the main source of mis-alignment is natural drift, which can be accounted for with various control and locking systems. One of the problems with designing a generalised auto-alignment system is that there is no standardised setup. That is to say each optical experiment or system will have its own characteristic lengths, lenses, optical devices and beam characteristics. Additionally aberrations and errors introduced by these optical elements will be unique to a given setup. This is generally suited to an experienced human operator, who can generally intuit the methods required to align a given system, having prior knowledge of how to work with a particular, or series, of optical elements.

A.1 REINFORCEMENT LEARNING PLATFORM

The generalised idea of [RL](#) is to learn an optimal policy for interacting with a system. A policy defines how an agent, interacting with the environment, will act when faced with a particular state of the system. In some ways this is analogous to the stimulus-response model in psychology [226], as the policy does not perform any evaluation or retrospection regarding the chosen behaviour. Indeed if a policy is said to be optimal, no evaluation is required. To determine an optimal policy, we include a reward function, which defines the criteria for determining the efficacy of a particular action. A reward function seeks to parametrise the world in some way. For example, if my aim is to keep a plant alive, watering the plant could return a positive reward. Here I may measure the 'healthiness' of the plant to determine my reward for the action of watering. Conversely putting the plant in a microwave oven should probably return a negative reward. The final key ingredient is the value function which seeks to provide information on the value of arriving at a particular state, by quantifying how much reward one can expect from a state in the future. It is not hard to see where calculating the value of a state would be useful. Many processes can be dealt with by determining a simple set of rules. For example in the game of chess, one rule might be, do not allow the queen to be captured. While this is a good rule for most situations, there are certain situations where it is advantageous to sacrifice the queen in order to achieve checkmate. In this case, calculating the value of the states would be very powerful, as it allows one to

determine an immediate win to the game. Determining the value function is often a non-trivial task. Considering a simple and deterministic game such as tic-tac-toe, it is clear to see that one can simply calculate all possible move combinations and have an exact measure of a given state's value. The same cannot be said for chess, with estimates placing the total number of possible games in excess of 10^{120} [228]. Clearly brute force is not the way to go for such problems. In practice value functions are learned through an agent's interaction with the environment. The combination of rewards and value estimation provides the mechanism with which the policy will be formed, the exact implementation of which depends on the approach.

A.2 GENERAL ACTOR-CRITIC

In Sec. 11.5.1 we introduced the concept of the AC network and its associated structure. To summarise, the AC structure contains two neural networks, an actor and a critic. These two networks work in tandem to learn the policy and value function simultaneously by interacting with the environment. The actor represents the policy, simply taking as an input the current state of the system, and performing an action on the environment. The critic seeks to approximate the value function, the idea being that an ANN should be good at representing the high dimensional structures of the value function. The question is of course how to inform the networks, to improve the on-policy performance.

A.2.1 Temporal difference learning

As shown in Fig. 11.11, both the actor and critic learn from feedback via the TD error. TD learning is considered to be a combination of ideas from both Monte Carlo and dynamic programming approaches [226]. TD aims to learn from experience by providing an immediate update to the value function, given by

$$V(s_t) \leftarrow V(s_t) + \alpha [R_{t+1} + \gamma V(s_{t+1}) - V(s_t)], \quad (\text{A.1})$$

where V is the value function evaluated on some state s at time t , α is a learning rate and R_{t+1} is the observed reward for the new state. The constant γ scales how much emphasis should be placed on the expected value of the state following the initial state. The total calculated value inside the brackets, scaled by α , is known as the TD error, often denoted by δ_t . The approach above is a special case of TD learning known as TD(0), as it provides an immediate update after only one step. TD methods are a bootstrapping method, as the estimate they provide is in some fashion based on an estimate. This can offer advantages to Monte Carlo algorithms which in contrast only

update after the completion of an episode. In the case of long episodes this can be tedious and costly. Currently it is an open question as to which method converges first.

A.2.2 Network losses

The AC networks are updated according to a set of losses, that adjust the response of the networks based on observed system response. Using the TD error, both the actor and critic networks are updated simultaneously. For each system interaction, we calculate the TD error, δ_t , as stated in Eq. A.1. The actor loss is based on the policy gradient loss and is defined by

$$L_{actor} = -\ln(\pi(A_t|s_t, \theta)) \cdot \delta_t, \quad (\text{A.2})$$

where A_t, s_t are the action and system state at time t , θ represents the networks weights and π is the policy. The appearance of the logarithm is due to the definition of the policy gradient [229]. Effectively, this loss function allows the actor to learn the correct policy via the TD error, by scaling the probability of performing that action.

The critic loss function simply needs to provide some mechanism for feedback, relating to the correct value of the state. The critic estimates a value for the state, based on the estimated value of the next state, with a discount factor, and the observed reward. The value function is supposed to reflect the value that an agent can expect to receive, when starting from a given state. Thus, the loss function of the critic is simply the mean squared error between the observed value and the previous value. This can be stated formally as

$$L_{critic} = \text{MSE}(R_{t+1} + \gamma V(s_{t+1}), V(s_t)). \quad (\text{A.3})$$

Using Eqs. A.2 and A.3, we can train the algorithm using our gradient descent algorithms, as per usual. In this case we use the Adam algorithm, which will adaptively change the learning as training progresses.

A.2.3 Network design

We design the network by testing on a basic problem. The test function consists of a continuous action space, with a goal position somewhere in the space. The reward is set to monotonically decrease with distance from this goal. When the agent has reached a position within 5% of the goal (± 0.1), the problem is termed complete. An additional reward is assigned for how close the agent is to the target position. For the purposes of testing we use the reward function

$$r(x) = -|x - 2| + 5 \quad (\text{A.4})$$

$$R_{t+1} = 10 \cdot [r(x_{t+1}) - r(x_t)] + \delta_{done} \cdot (10000 + e^{-(x_{t+1}-2)^2 \cdot 100}) \cdot e^{-N_{step}/100}, \quad (\text{A.5})$$

where $r(x)$ is a reward function with the maximum at $x = 2$, N_{steps} is the number of steps taken by the agent thus far and δ_{done} is a delta function = 1 when the problem is considered complete, and 0 otherwise. Eq. A.5 defines the reward for a given time step, arriving at the state x_{t+1} . The first term defines a differential reward, based on the movement of the agent with respect to the goal. The second term corresponds to an additional reward for solving the problem. Solving the problem yields an immediate spike in the reward of 10000, with a further reward given based on how close the agent was to the goal. In this problem the goal is located at $x = 2$. The additional reward is reduced by an exponential factor of how many steps it took to get there. This is necessary to ensure the agent reaches the goal in the shortest number of steps possible, by penalising behaviour that is not exclusively leading to the goal. The problem is solved on the interval $[-3, 3]$ with the action space restricted to lie within this same interval.

Clearly this problem is trivial and could be solved easily by a gradient ascent algorithm, however its value is in the ability to test the effect of different network structures and hyperparameters. To solve this problem, we construct an actor-critic agent with the parametrisation

$$x_t \rightarrow A(x) \rightarrow a_t \sim \mathcal{N}(\mu, \sigma) \quad (\text{A.6})$$

$$x_t \rightarrow V(x) \rightarrow v_x. \quad (\text{A.7})$$

Here the actor network maps a given state to an action, which is normally distributed. This is achieved by having the actor network map to a μ and σ for a given state. This distribution is sampled to determine the action that should be taken. The value network simply maps a given state to its approximate value.

Even with such a simple problem, the performance of the agent can be quite sensitive to the hyperparameters. Potentially the most important component of the agent is the critic (value) network. This is because estimates produced by the critic directly inform the actor as to how to interact with the system, as well as providing an estimate as to the long term value of a state. As such, a poor estimate of the value function will ultimately lead to poor agent performance, and in some cases divergence. We observe that for certain choices of network structure, the critic training can become inherently unstable. An example of this is shown in Fig. A.1, which contains the episode history for two separate agents. The agents have identical network structures for the critic, 2 hidden layers with 400 neurons per layer. The only difference between the two layers is the activation functions used. As shown in Fig. A.1, the network with RELU activations fails to converge on an optimal solution. 400 neurons per layer is a

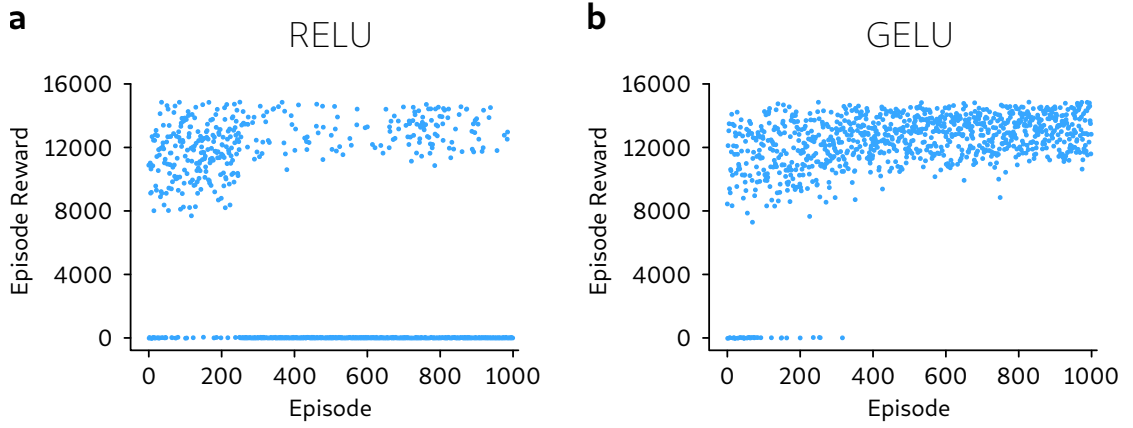


Figure A.1: AC agent episode history when attempting to learn the problem. The only difference between the two agents is the activation function used in the value network, RELU in (a) and GELU in (b).

far more complex network than is needed to solve this problem, but it is chosen to exacerbate an underlying problem. It is known that vanishing and exploding gradients are a problem in training neural networks [230, 231]. This can occur if careful thought is not given to the network structure and learning rates of the AC networks. In this case, GELU manages to converge on a solution as it is slightly more robust to these problems. However it can be demonstrated for high learning rates that this problem once again occurs.

By benchmarking on this problem, we find that one can optimise the hyperparameters to provide a fast and robust search of the solution space. In the present case, we find that a modest critic network provides robust performance, consisting of 2 hidden layers with, 5 GELU activated neurons per layer. For the actor network, 2 hidden layers with, 40 GELU activated neurons per layer is used. The actor network could be reduced in size also, however we find that it has less effect on the algorithms performance. At each episode the networks are trained using the TD(0) error, with which 3 training steps are taken. To demonstrate the general efficacy of these parameter choices, we also demonstrate this agent applied to the same problem as before, where Eq. A.4 is replaced with a modified 1D Ackely function, given by

$$r(x) = 20 - \text{Ackley}(x - 2). \quad (\text{A.8})$$

This produces a maximum reward at $x = 2$ with the structure shown in Fig. A.2b. As is demonstrated in Fig. A.2a, the AC agent can also solve this problem. The structure of the problem can be increased by multiplying the input to the Ackley function by some constant. This has the effect of adding more local minima and a steeper gradient around the maximum. Doing so, such that the entire Ackley domain is scaled into our reduced interval $[-3, 3]$, we find that the AC agent can still solve this problem, unencumbered by the increased structure.




Figure A.2: AC agent applied to the 1D Ackley function. (a) shows the episodic performance of the agent applied to this problem. (b) shows the structure inherent in the Ackley function that determines the reward structure. Clearly it is no longer monotonic.

A.3 PHYSICAL PROBLEM

Now that we have defined our AC agent, we are required to define the problem we wish to solve. As previously mentioned, we wish to create an automatic alignment system for optical experiments. The general experimental layout is shown in Fig. A.3. Six actuators control the alignment of the beam into a ring cavity. Four actuators are dedicated to 2 steering mirrors, which allow the agent to beam-walk. The final 2 actuators control the position of 2 lenses, which alter the beam waist and position. The present modes are monitored by scanning the cavity frequency with the back mirror mounted on a piezo. The cavity scan signal is monitored on the PD. A CCD is also used for mode selection, such that the correct mode on the PD scan can be determined via the mode's spatial profile. The system produces a number of challenges. Due to the mode structures, the optimisation is necessarily non-convex, unless a tight region around the optimum alignment is selected. Additionally, drift due to temperature and other external factors is present, modifying the alignment efficacy. Finally, the actuators also exhibit hysteresis, which can be corrected somewhat with proper control, however, it is still a significant challenge to the optimisation.

The reward function is determined by measuring the intensity of a particular peak. The reward can be formulated in many ways, for example it can simply be the intensity of the peak, the change in intensity or some measure of mode purity from the CCD. For the majority of testing, we use the intensity of the peak as the reward.

A.4 ALGORITHMIC APPROACHES

The AC agent constitutes the main algorithm of interest in this experiment. The AC model has a few distinct advantages that make it an advantageous approach. Firstly, drift is less of an issue, as the AC agent simply learns to control the system within this




Figure A.3: Optical setup for the auto alignment experiment into a cavity. 6 actuators are used to walk the beam, to maximise a particular mode. The cavity is scanned using a piezo to determine the present modes.

drift. This can also be countered somewhat, by assigning rewards based on the relative change in intensity. Secondly, while training may take longer than a standard deep learning network, once the network is trained it can also act as a control algorithm, correcting for drift and fluctuation after alignment is complete.

We wish to benchmark the [AC](#) agent against a number of algorithms. Acuminata is an obvious choice, given its ability to operate in high dimensional spaces. However as noted in previous chapters, this is an approximate global method, not an accurate local method. To improve the performance of the deep learning approach it will be necessary to determine a point at which the optimisation has become convex. At this point we can kick off a local optimiser, such as LBFGS, to refine the alignment. Given the drift and hysteresis present in the system, stochastic gradient descent may be a better choice. Additionally, as the parameter space is low, we can also test the [FF](#) algorithm, as we expect this to yield good results in the low parameter number regime.

A.4.1 Data acquisition and reward

The cavity frequency is scanned using a piezo at a frequency of 10Hz. The cavity spectrum is acquired and the peak of interest identified. During initial testing we are interested in the relative height of a given peak (usually the TEM00 mode). Thus we identify this peak and measure the intensity from the PD. From this intensity we can formulate a cost function or reward, depending on the algorithm. For Acuminata, we can simply define the cost as the negative of the intensity, thus a higher intensity is a lower cost.

Defining a reward function for the [AC](#) network is a little more complicated. Here our reward function must try to assign a reward to actions that are considered desirable.

Assigning the intensity as the reward is a simple approach, however, we also require that our algorithm reaches some final state that is considered aligned. In this case we have to assign a 'done' condition, which has a large reward associated with the completion of the goal. As we wish the intensity to continually increase, we expect that this goal should also tend to increase over time. One way to implement this is to assign rewards that are proportional to the relative improvement of the peak (see Sec. [A.5.2](#)).

A.5 PRELIMINARY RESULTS

Initially we start by testing the ability of the algorithms to align the cavity setup using only the steering mirrors. The lenses should have far less of an effect on the beam position. However, the mirrors will have the ability to remove all signal from the cost function, corresponding to no light going through the cavity. Although there are only 4 optimisation parameters, there will be large sections of parameter space which yield no signal.

A.5.1 Acuminata

We start by testing Acuminata, the results of which are shown in Fig. [A.4](#). We provide experimental bounds on the actuators of ± 5000 steps from a reference point at 0. These bounds allow the algorithm to align the beam such that it may exceed the cavity aperture. For this optimisation, the beam starts with a beam waist which is not matched to the cavity, but is sufficient for some light to be transmitted. The cost function we wish to minimise corresponds to the intensity of the TEM₀₀ peak. This peak is initially determined using a reference trace coupled with a CCD measurement. Acuminata performs 50 initial samples to determine the initial relationship between the parameters. As can be seen from Fig. [A.4a](#), during the initial sampling only one value is observed that corresponds to any meaningful transmission. By 100 runs Acuminata has determined that the optimum value lies outside of this position and begins moving the actuators away from this position. By run 250 the minimum, within measurement error, has been reached. Fig. [A.4b](#) shows the progression of the actuator position over the course of the optimisation.

A.5.2 Actor-Critic

Using the network structure outlined in Sec. [A.2.3](#), we apply the AC model to the alignment problem. Fig. [A.5](#) shows the results of the preliminary optimisation. As shown by the increasing frequency of high rewards, the AC model is learning how to align the system from random starting positions. It can be seen from Fig. [A.5b](#) that there still remains a large amount of variation on the actuator position, even for high rewards. This is due to the nature of the stopping condition we have implemented,




Figure A.4: Preliminary alignment using Acuminata. (a) shows the observed costs during the optimisation which lasted for 751 runs. (b) shows the position of the actuators as determined by Acuminata. The four actuators being controlled correspond to two separate steering mirrors with x and y control.

which forces the model to reach an ever increasing threshold before arriving at a terminal state. The state is characterised by the value of a designated peak on the PD signal. Thus the terminal state is simply a threshold value based on the measured intensity of this peak. The reward function is constructed such that the model receives an exponentially increasing reward the closer it is to this threshold. Upon exceeding this, a significantly larger additional reward is received, proportional to intensity of the peak. Upon exceeding the threshold 3 times, the terminal state intensity is increased.

The AC model requires a larger amount of tuning compared to other algorithms, such as Acuminata. This is due to the nature of the RL approach which uses a bootstrapping method for approximating the value function. The noise and hysteresis in the system necessarily perturbs the accuracy of this approximation. The model must therefore be tuned correctly to limit the effect of noise at the expense of slower training.

It should be noted the AC network should not be directly compared to the performance of Acuminata in this circumstance. While Acuminata aims at immediately optimising the system, the AC network is attempting to optimise and learn a generalised control scheme. The history of the reward function will necessarily appear to be more variable due to the fact that the AC network starts in a random position and works its

way to the best position. The advantage of this approach over Acuminata is that the AC network should be able to handle drift and hysteresis if parametrised properly. The AC model learns a generalised control method and could potentially be used as a coarse alignment tool in a daily routine.

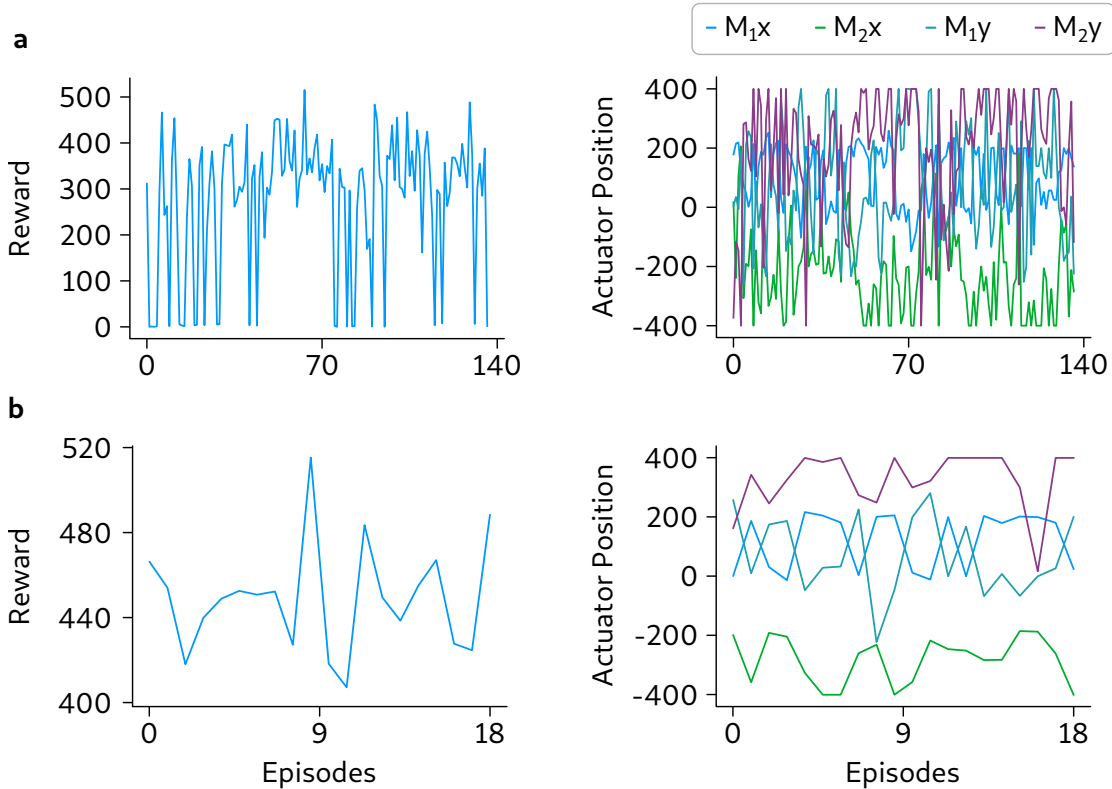


Figure A.5: Preliminary alignment using the Actor-Critic model. (a) shows the cumulative reward and final state for each episode. (b) shows the cumulative reward and final state for episodes which achieved a cumulative reward greater than 400.

A.6 CONTINUING WORK

Further work on this problem can be roughly divided into three sections: benchmarking, AC tuning and extensions. It is important to determine the best approach to solving this problem by benchmarking different algorithmic approaches such as fruit fly, Monte-Carlo and gradient descent. We expect that due to the inherent noise and non-convex nature of the problem Acuminata will outperform these methods, however, it is still important to characterise each method. Furthermore, Acuminata can be extended to include a local optimiser, such as LBFGS, that can be enabled when convergence is reached. This will ensure the identified minimum is optimised locally, possibly providing a better alignment than human operators.

Tuning is still required to get the AC network running efficiently. While we have demonstrated some success in learning to align the system we still have yet to demonstrate a stable approach to reaching this outcome. Improvements can also be made to

the reward function to increase this stability. Finally, the indication of the system state can be improved to better represent the experimental setup. Currently the actuator position is used as the system state. Drift and hysteresis mean that these parameters don't accurately represent the system over time. A better method may be to feed a down-sampled cavity trace to the AC model as a representation of the state. This would provide complete and accurate information as to the misalignment of the cavity. We have also yet to test the ability of the AC model to align the system after intentional misalignment once it has learnt to control the system.

There are further extensions which have the potential to yield interesting and useful results. In the preliminary results we have only used the 4 actuators attached to the steering mirrors. This can be extended to the full 6 actuators which will allow algorithms to adjust the beam size and divergence as well. Further coupling improvements could be gained by including an SLM to correct for aberrations. The introduction of a CCD camera to characterise the state could be used to allow the algorithm to align to specific higher order modes. In this case the reward/cost function could be altered to determine the purity of the transmitted mode.

B

FRUIT FLY OPTIMISATION

The fruit fly (FF) algorithm was originally introduced by [180], as a global search metaheuristic for solving optimisation problems. FF is inspired by the search patterns of fruit flies, when searching for a source of food. The procedure used in the current work, in pseudo code, is given in Alg. 2.

Algorithm 2 Fruit fly metaheuristic

```
tol  $\leftarrow 1 \cdot 10^{-5}$  {step tolerance}
max_iter  $\leftarrow$  {maximum iterations to solve}
x0  $\leftarrow [0 \dots 0]$ 
f0  $\leftarrow \infty$ 
bounds  $\leftarrow$  {list of (xmin, xmax)}
iter  $\leftarrow 0$  {current iterations}
pop  $\leftarrow []$  {list of FF agents objects}
while iter < max_iter do
  if iter  $\neq 0$  then
    best_f, best_x = max(pop.f), pop.x[argmax(pop.f)]
    x0  $\leftarrow$  best_x
    f0  $\leftarrow$  best_f
  end if
  for xarr in pop.x do
    index  $\leftarrow$  randint(0, x0.size)
    xnext  $\leftarrow$  x0
    xmin, xmax = bounds[index]
    radius  $\leftarrow$  (xmin - xmax)/2
    progress  $\leftarrow$  iter/max_iter
    step  $\leftarrow$  radius · exp(ln(tol) · progress) ·  $\mathcal{U}(-1, 1)$ 
    xnext[index]  $\leftarrow$  xnext[index] + step
    xarr  $\leftarrow$  xnext
  end for
  iter  $\leftarrow$  iter + 1
end while
```

This algorithm is used in various parts of this thesis to solve optimisation tasks, where speed of evaluation is crucial. The population, denoted as *pop*, contains a list of agents that evaluate an objective function during their search. If the task can be run in parallel, then each agent can also evaluate the function in parallel. A problem given to the population is required to have two things, an objective function and a list of bounds for each parameter. Using these two components the FF procedure facilitates a random walk around the space, similar to that of a Monte Carlo method. However as the problem progresses the step size becomes smaller, according to the tolerance and maximum iterations.

B.1 RAMAN MEMORY SIMULATION

Using this method, the parameters of the Raman memory simulation performed in XMDS [232], which was detailed in Ch. 7, can be optimised. The problem is defined by four controllable parameters, the write pulse timing, read/write pulse width, write pulse power and read pulse power. To evaluate the objective function, each of these parameters is written into the XML file used by XMDS to compile the C code simulation. The python script executes the XMDS simulation which outputs the results of the simulation to a HDF5 file. This file is opened by the python script, which uses the output to calculate the cost. Here the cost is simply defined as the ratio between the integral of the output and input pulses, providing a measure of the efficiency. The actual cost used by FF is $c(\epsilon) = 1 - \epsilon$, where ϵ is the efficiency. The code for this simulation can be found at [233].

C

LEARNING TRANSFER FUNCTIONS

Atom interferometry has attracted much interest for its applicability to precision sensing of inertial forces, such as acceleration and rotations. For a complete description of the theory and experimental complexities involved with atom interferometry, I direct the interested reader to [234]. Briefly, atoms are loaded into a MOT under UHV and cooled to temperatures on the order of μK . At these low temperatures, the cooled ensemble will begin to exhibit wavelike properties. These cooled atoms are then either dropped or launched vertically, where they undergo a splitting and recombination, interferometry sequence. The different phase, acquired from the unique interferometer paths taken by the matter-waves, can be used to precisely characterise local acceleration due to gravity. However, it is also possible to use these techniques for sensing other inertial forces [235, 236].

One of the problems that needs to be addressed experimentally is vibration isolation. Due to the nature of the interferometry involved, vibrations on optical elements, especially that of the retro-reflecting Raman beam mirror, introduce additional phase noise. Generally there are two approaches to dealing with this problem: active isolation systems and post-correction. Active isolation can be implemented using a passive vibration isolation stage, with active feedback from a low noise seismometer to voice-coil actuators [237, 238]. Post correction uses the error signal from the seismometer to record vibrations common to the mirror. These can then be filtered out in post. This additional method is useful for removing vibrations that are outside of the gain bandwidth of the active isolation feedback. As shown in [234], this can substantially improve the sensitivity of the gravity measurement.

We are interested in decreasing the noise of the interferometer measurements, by using post corrected values from the seismometer data. A simple correction that can be applied is

$$g_{vib} = \frac{1}{T^2} \int_0^{2T} f(t) \cdot S(t) dt, \quad (\text{C.1})$$

where $S(t)$ are the seismometer measurements for an experiment time $2T$. $f(t)$ is the sensitivity function of the interferometer and is given by

$$f(t) = \begin{cases} t & \text{if } t \leq T \\ 2T - t & \text{if } t > T. \end{cases} \quad (\text{C.2})$$

The implicit assumption is that the accelerations returned by the seismometer, is indeed the actual acceleration. However, in practice, the seismometer will have a response that is really a convolution between the true acceleration and the transfer function of the measurement apparatus. This includes the seismometers, filters and electronics. Thus instead we have

$$S(t) = h(t) * a(t), \quad (\text{C.3})$$

where $h(t)$ is the transfer function of the experimental apparatus, convolved with the true acceleration, $a(t)$. This convolution operation will affect the accuracy of the post correction as it necessarily filters the acceleration signal. While it is advantageous to correct for this, obtaining the transfer function directly can be difficult, or in certain circumstances, unattainable. To circumvent this issue we employ a novel machine learning strategy.

C.1 MACHINE LEARNING APPROACH

We require some method to learn the transfer function, or more generally, the form of an arbitrary function. Unlike regression, we do not have a candidate model to start from, as we assume the transfer function of intermediary devices in the signal processing and analogue-to-digital conversion are unknown. There exist methods for developing functional relationships within data, such as symbolic regression [239, 240]. While in practice symbolic regression can be powerful, solution detection can be slow and difficult to converge. As we do not require a free form expression to describe our transfer function, we present an approach that instead allows us to leverage the computational speed of global optimisation approaches, such as [FF](#).

Our approach is to represent the transfer function in frequency space, by generating filters from orthogonal basis functions. Once we have obtained the transfer function, we can perform a deconvolution operation (simple division) to obtain the true accelerometer values.

C.1.1 Function parametrisation

There are a number of possible ways to parametrise the space that we wish to construct. Unlike the previous approach in [Ch. 10](#), piecewise control of the individual frequencies is not the best approach. Firstly, this problem quickly become intractable with spectral resolution. Additionally, it is unlikely that, physically, the frequency response would have such sharp features. Quickly changing features can also introduce artefacts into the fast Fourier transforms, altering the true value. While there are a number of orthogonal basis that we may construct our function from, we opt for the Legendre polynomials.




Figure C.1: First five Legendre polynomials. The polynomials form an orthogonal basis from which arbitrary functions can be constructed.

The first five polynomials are shown in Fig. C.1. These polynomials allow for the construction of smooth transfer functions, in a domain that is closer to that of the likely transfer function. Other bases, such as the Fourier series, can also be used. However, as the function is not expected to be periodic, this basis requires many more basis functions, to accurately represent the transfer function.

The full parametrisation consist of N parameters, such that N is even. We require this, as we will construct both the real and imaginary parts of the transfer function in the following fashion

$$\text{Re} [H(\omega)] = \sum_{i=0}^{N/2} c_i \cdot P_i(\omega) \quad (\text{C.4})$$

$$\text{Im} [H(\omega)] = \sum_{i=0}^{N/2} c_{i+N/2} \cdot P_i(\omega), \quad (\text{C.5})$$

where P_i is the i^{th} Legendre polynomial, c_i is the i^{th} scaling coefficient and $H(\omega)$ is the transfer function. Here the Legendre polynomials are scaled so that they span the range of ω . The aim is to learn the coefficients, c_i , such that we have an approximate form for H . These coefficients become the parameters we will perform the global optimisation on. Fortunately, this task can also be performed in parallel and does not rely on serial operation, greatly increasing the speed of the computation. Since this function can be evaluated quickly, we will use [FF](#) algorithm to solve this optimisation task.

C.2 METHOD TESTING

To verify that we can infer a transfer function applied to a set of data, we begin by testing the method on some synthetic test cases. We start by generating some multi-




Figure C.2: Original generated spectrum centered around 50Hz, filtered with a narrow band filter. This test set is used to test recovery of the original spectrum via the proposed method.

modal test data, centred on 50Hz with a $1/e^2$ bandwidth of 12.7Hz. Additionally we add some noise with a bandwidth of 40Hz. As shown in Fig. C.2, we perturb this distribution by applying a convolution with some narrow band filter. We will attempt to recover the original spectrum, by minimising the MSE between the recovered and original signal. In reality we would not have access to this signal, however, we can formulate the cost function in other ways, as outlined in the proceeding section.

After performing the optimisation, we can see from Fig. C.3 that we have recovered the spectrum, after applying a deconvolution operation with our learned transfer function. We find that the relative deviation is highest at the edge of the spectrum. This is to be expected as these errors contribute the least to the overall signal, and hence to the MSE loss. This could be altered by defining a better loss metric.

c.2.1 Cost function

In the real problem, we do not have perfect information regarding the true signal. Instead we must construct a cost function relevant to the experimental set-up, and desired outcome. If we were simply to minimise the residuals, then we could unintentionally remove the slow drift in local gravity that we are trying to measure. The aim is to decrease environmental vibrations without removing the real gravity variation. A way around this problem is to minimise the two sample variance between the residuals of two adjacent runs. Doing so is equivalent to calculating the first point of the Allan variance curve [241]. This forces the optimisation to learn a transfer function that

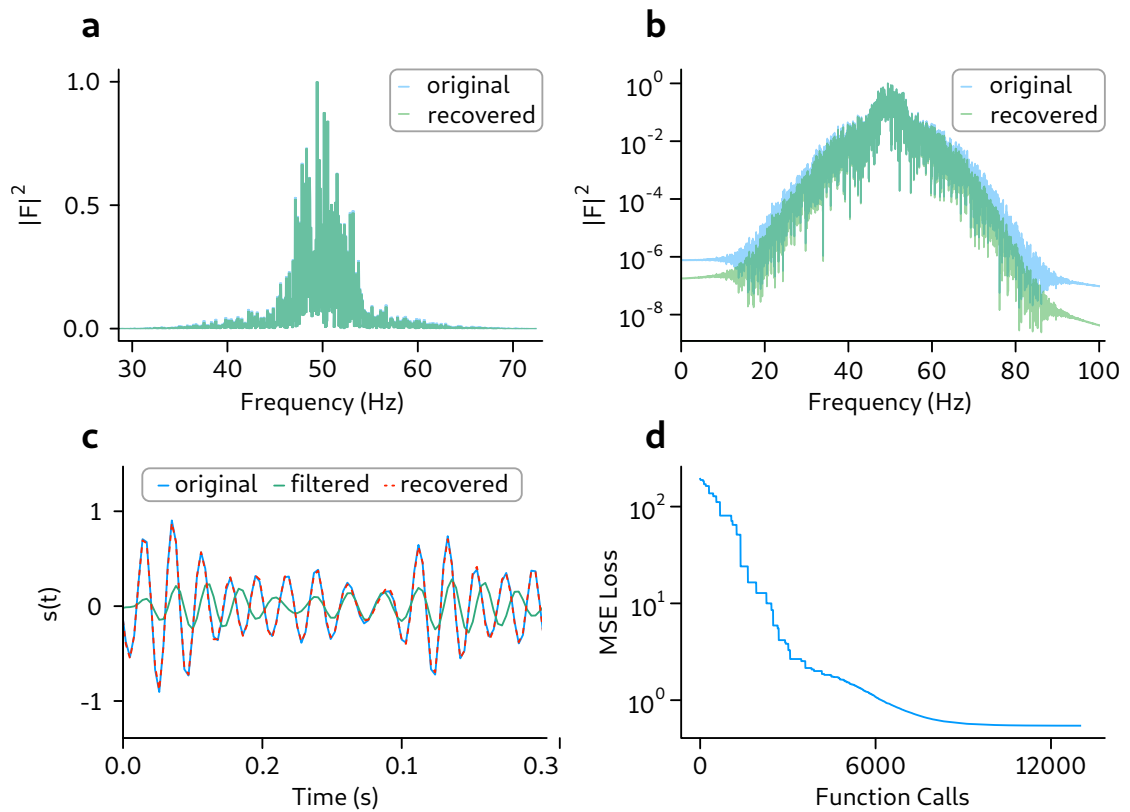


Figure C.3: Results of a test optimisation aimed at recovering the original signal. (a) shows the overlap between the original and recovered spectrum. (b) shows the same overlap but on a logarithmic scale. The edges of the spectrum show most relative divergence. (c) shows a comparison between the time series for the original, recovered and filtered signals. (d) shows the progression of the optimisation loss as a function of function evaluations.

returns the right variance in the short term, without removing the characteristic longer term drift in g . For the optimisation we use the cost function

$$C(\mathbf{X}) = \sqrt{\frac{1}{N-1} \sum_{i=1}^{N-1} ([g_{AI,t} - g_{Vib,t}] - [g_{AI,t+1} - g_{Vib,t+1}])^2}, \quad (\text{C.6})$$

where $g_{AI,t}$ and $g_{Vib,t}$ are the atom interferometer residuals and vibrational correction at time t respectively. The construction of the cost function also constrains the solutions to be physical, as non-physical phenomena such as violating causality would be selected against by the comparison of the signals.

c.2.2 Avoiding overfitting

As we are seeking to fit an arbitrary function, to minimize some average property over our samples, there is the possibility for the algorithm to overfit. This can be combated by monitoring a validation loss. Approximately 12.5% of the data is reserved as a validation set, which will not take part in calculation of the cost function during the optimisation run. This validation loss can be monitored to determine whether the learned transfer function generalises to, as of yet, unseen data.

C.3 CONTINUING WORK

We are currently in the process of applying this method to some gravity measurements, using a data set acquired from Humboldt University [234]. The aim is to provide a series of sensitivity benchmarks, with respect to various approaches, i.e. no correction, active isolation, and post correction.

BIBLIOGRAPHY

1. Shankar, R. *Principles of quantum mechanics* (Springer Science & Business Media, 2012).
2. Lambropoulos, P. & Petrosyan, D. *Fundamentals of quantum optics and quantum information* (Springer, 2007).
3. Einstein, A. On a Heuristic Viewpoint Concerning the Production and Transformation of Light. *Annalen der Physik* **17**, 132–148.
4. Horodecki, R., Horodecki, P., Horodecki, M. & Horodecki, K. Quantum entanglement. *Reviews of Modern Physics* **81**, 865 (2009).
5. Bell, J. S. On the Einstein Podolsky Rosen paradox. *Physics Physique Fizika* **1**, 195 (1964).
6. Freedman, S. J. & Clauser, J. F. Experimental test of local hidden-variable theories. *Physical Review Letters* **28**, 938 (1972).
7. Aspect, A., Dalibard, J. & Roger, G. Experimental test of Bell's inequalities using time-varying analyzers. *Physical Review Letters* **49**, 1804 (1982).
8. Thearle, O., Janousek, J., Armstrong, S., Hosseini, S., Schünemann, M., Assad, S., Symul, T., James, M. R., Huntington, E., Ralph, T. C., *et al.* Violation of Bell's inequality using continuous variable measurements. *Physical Review Letters* **120**, 040406 (2018).
9. Steck, D. A. *Quantum and atom optics* 2007.
10. Hänsch, T. W. & Schawlow, A. L. Cooling of gases by laser radiation. *Optics Communications* **13**, 68–69 (1975).
11. Wineland, D. & Dehmelt, H. Proposed $10^{14}\delta\nu/\nu$ laser fluorescence spectroscopy on Ti^+ mono-ion oscillator III (side band cooling). *Bulletin of the American Physical Society* **20**, 637–637 (1975).
12. Lett, P. D., Watts, R. N., Westbrook, C. I., Phillips, W. D., Gould, P. L. & Metcalf, H. J. Observation of atoms laser cooled below the Doppler limit. *Physical Review Letters* **61**, 169 (1988).
13. Dalibard, J. & Cohen-Tannoudji, C. Laser cooling below the Doppler limit by polarization gradients: simple theoretical models. *Journal of the Optical Society of America B* **6**, 2023–2045 (1989).
14. Metcalf, H. J. & van der Straten, P. *Laser Cooling and Trapping* (Springer, 1999).
15. Wildermuth, S., Hofferberth, S., Lesanovsky, I., Groth, S., Krüger, P., Schmiedmayer, J. & Bar-Joseph, I. Sensing electric and magnetic fields with Bose-Einstein condensates. *Applied Physics Letters* **88**, 264103 (2006).

16. Cho, Y.-W., Campbell, G., Everett, J., Bernu, J, Higginbottom, D., Cao, M., Geng, J, Robins, N., Lam, P. & Buchler, B. Highly efficient optical quantum memory with long coherence time in cold atoms. *Optica* **3**, 100–107 (2016).
17. Raab, E., Prentiss, M, Cable, A., Chu, S. & Pritchard, D. E. Trapping of neutral sodium atoms with radiation pressure. *Physical Review Letters* **59**, 2631 (1987).
18. Steck, D. A. *Rubidium 87 D line data* 2001.
19. Kohns, P, Buch, P, Süptitz, W, Csambal, C & Ertmer, W. On-line measurement of sub-Doppler temperatures in a Rb magneto-optical trap-by-trap centre oscillations. *Europhysics Letters* **22**, 517 (1993).
20. Wallace, C. D., Dinneen, T. P., Tan, K.-Y. N., Kumarakrishnan, A, Gould, P. L. & Javanainen, J. Measurements of temperature and spring constant in a magneto-optical trap. *Journal of the Optical Society of America B* **11**, 703–711 (1994).
21. Lindquist, K, Stephens, M & Wieman, C. Experimental and theoretical study of the vapor-cell Zeeman optical trap. *Physical Review A* **46**, 4082 (1992).
22. Haw, M., Evetts, N., Gunton, W., Van Dongen, J., Booth, J. L. & Madison, K. W. Magneto-optical trap loading rate dependence on trap depth and vapor density. *Journal of the Optical Society of America B* **29**, 475–483 (2012).
23. Ketterle, W., Davis, K. B., Joffe, M. A., Martin, A. & Pritchard, D. E. High densities of cold atoms in a dark spontaneous-force optical trap. *Physical Review Letters* **70**, 2253 (1993).
24. Yan-Xu, H., Yong-Hong, L., Chun-Hong, Z., Shu-Jing, L. & Hai, W. Realization of high optical density Rb magneto-optical trap. *Chinese Physics Letters* **26**, 023201 (2009).
25. Drever, R., Hall, J. L., Kowalski, F., Hough, J., Ford, G., Munley, A. & Ward, H. Laser phase and frequency stabilization using an optical resonator. *Applied Physics B* **31**, 97–105 (1983).
26. Chu, S., Hollberg, L., Bjorkholm, J. E., Cable, A. & Ashkin, A. Three-dimensional viscous confinement and cooling of atoms by resonance radiation pressure. *Physical Review Letters* **55**, 48 (1985).
27. Dicke, R. H. Coherence in spontaneous radiation processes. *Physical review* **93**, 99 (1954).
28. Scully, M. O. & Svidzinsky, A. A. The super of superradiance. *Science* **325**, 1510–1511 (2009).
29. Friedberg, R., Hartmann, S. R. & Manassah, J. T. Frequency shifts in emission and absorption by resonant systems of two-level atoms. *Physics Reports* **7**, 101–179 (1973).
30. DeVoe, R. & Brewer, R. Observation of superradiant and subradiant spontaneous emission of two trapped ions. *Physical Review Letters* **76**, 2049 (1996).

31. Casabone, B, Friebe, K, Brandstätter, B, Schüppert, K, Blatt, R & Northup, T. Enhanced quantum interface with collective ion-cavity coupling. *Physical Review Letters* **114**, 023602 (2015).
32. Scheibner, M., Schmidt, T., Worschech, L., Forchel, A., Bacher, G., Passow, T. & Hommel, D. Superradiance of quantum dots. *Nature Physics* **3**, 106–110 (2007).
33. Luk'yanchuk, B., Zheludev, N. I., Maier, S. A., Halas, N. J., Nordlander, P., Giessen, H. & Chong, C. T. The Fano resonance in plasmonic nanostructures and metamaterials. *Nature Materials* **9**, 707–715 (2010).
34. Roof, S., Kemp, K., Havey, M. & Sokolov, I. Observation of single-photon superradiance and the cooperative Lamb shift in an extended sample of cold atoms. *Physical Review Letters* **117**, 073003 (2016).
35. Röhlberger, R., Schlage, K., Sahoo, B., Couet, S. & Ruffer, R. Collective Lamb shift in single-photon superradiance. *Science* **328**, 1248–1251 (2010).
36. Yoshikawa, Y., Torii, Y. & Kuga, T. Superradiant light scattering from thermal atomic vapors. *Physical Review Letters* **94**, 083602 (2005).
37. Inouye, S, Chikkatur, A., Stamper-Kurn, D., Stenger, J, Pritchard, D. & Ketterle, W. Superradiant Rayleigh scattering from a Bose-Einstein condensate. *Science* **285**, 571–574 (1999).
38. Schneble, D., Campbell, G. K., Streed, E. W., Boyd, M., Pritchard, D. E. & Ketterle, W. Raman amplification of matter waves. *Physical Review A* **69**, 041601 (2004).
39. Yoshikawa, Y., Sugiura, T., Torii, Y. & Kuga, T. Observation of superradiant Raman scattering in a Bose-Einstein condensate. *Physical Review A* **69**, 041603 (2004).
40. Asbóth, J., Domokos, P, Ritsch, H & Vukics, A. Self-organization of atoms in a cavity field: Threshold, bistability, and scaling laws. *Physical Review A* **72**, 053417 (2005).
41. Black, A. T., Chan, H. W. & Vuletić, V. Observation of collective friction forces due to spatial self-organization of atoms: from Rayleigh to Bragg scattering. *Physical Review Letters* **91**, 203001 (2003).
42. Baumann, K., Guerlin, C., Brennecke, F. & Esslinger, T. Dicke quantum phase transition with a superfluid gas in an optical cavity. *Nature* **464**, 1301–1306 (2010).
43. Tesio, E, Robb, G., Oppo, G.-L., Gomes, P., Ackemann, T, Labeyrie, G, Kaiser, R & Firth, W. Self-organization in cold atomic gases: a synchronization perspective. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **372**, 20140002 (2014).
44. Bonifacio, R & De Salvo, L. Collective atomic recoil laser (CARL) optical gain without inversion by collective atomic recoil and self-bunching of two-level atoms. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **341**, 360–362 (1994).

45. Ackemann, T. & Lange, W. Optical pattern formation in alkali metal vapors: Mechanisms, phenomena and use. *Applied Physics B* **72**, 21–34 (2001).
46. Giusfredi, G., Valley, J., Pon, R., Khitrova, G. & Gibbs, H. Optical instabilities in sodium vapor. *Journal of the Optical Society of America B* **5**, 1181–1192 (1988).
47. Tesio, E., Robb, G., Ackemann, T., Firth, W. & Oppo, G.-L. Spontaneous optomechanical pattern formation in cold atoms. *Physical Review A* **86**, 031801 (2012).
48. Schmittberger, B. L. & Gauthier, D. J. Enhancing light-atom interactions via atomic bunching. *Physical Review A* **90**, 013813 (2014).
49. Ayllon, R., Mendonça, J., Gisbert, A., Piovela, N & Robb, G. Multimode collective scattering of light in free space by a cold atomic gas. *Physical Review A* **100**, 023630 (2019).
50. Greenberg, J. A. & Gauthier, D. J. Steady-state, cavityless, multimode superradiance in a cold vapor. *Physical Review A* **86**, 013823 (2012).
51. Schmittberger, B. L. & Gauthier, D. J. Spontaneous emergence of free-space optical and atomic patterns. *New Journal of Physics* **18**, 103021 (2016).
52. Schmittberger, B. L. & Gauthier, D. J. Transverse optical and atomic pattern formation. *Journal of the Optical Society of America B* **33**, 1543–1551 (2016).
53. Schmittberger, B. L. *Multimode Atomic Pattern Formation via Enhanced Light-atom Interactions* PhD thesis (Duke University, 2016).
54. Berry, H. G., Gabrielse, G & Livingston, A. Measurement of the Stokes parameters of light. *Applied optics* **16**, 3200–3205 (1977).
55. Grynberg, G. & Robilliard, C. Cold atoms in dissipative optical lattices. *Physics Reports* **355**, 335–451 (2001).
56. Von Neumann, J. First Draft of a Report on the EDVAC. *IEEE Annals of the History of Computing* **15**, 27–75 (1993).
57. Furber, S. Large-scale neuromorphic computing systems. *Journal of Neural Engineering* **13**, 051001 (2016).
58. Li, Y., Wang, Z., Midya, R., Xia, Q. & Yang, J. J. Review of memristor devices in neuromorphic computing: materials Sciences and device challenges. *Journal of Physics D: Applied Physics* **51**, 503002 (2018).
59. Schuman, C. D., Potok, T. E., Patton, R. M., Birdwell, J. D., Dean, M. E., Rose, G. S. & Plank, J. S. A survey of neuromorphic computing and neural networks in hardware. *arXiv preprint arXiv:1705.06963* (2017).
60. O'Brien, J. L. Quantum computing over the rainbow. *Physics* **1**, 23 (2008).
61. Steane, A. Quantum computing. *Reports on Progress in Physics* **61**, 117 (1998).
62. Nielsen, M. A. & Chuang, I. *Quantum computation and quantum information* 2002.

63. Shor, P. W. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *Society for Industrial and Applied Mathematics Review* **41**, 303–332 (1999).
64. Bernstein, D. J. *Introduction to post-quantum cryptography* 1–14 (Springer, 2009).
65. Georgescu, I. M., Ashhab, S. & Nori, F. Quantum simulation. *Reviews of Modern Physics* **86**, 153 (2014).
66. Wootters, W. K. & Zurek, W. H. A single quantum cannot be cloned. *Nature* **299**, 802–803 (1982).
67. Gisin, N., Ribordy, G., Tittel, W. & Zbinden, H. Quantum cryptography. *Reviews of Modern Physics* **74**, 145 (2002).
68. Bennett, C. H. & Brassard, G. *Proceedings of the IEEE International Conference on Computers, Systems and Signal Processing* 1984.
69. Simon, C., De Riedmatten, H., Afzelius, M., Sangouard, N., Zbinden, H. & Gisin, N. Quantum repeaters with photon pair sources and multimode memories. *Physical Review Letters* **98**, 190503 (2007).
70. Briegel, H.-J., Dür, W., Cirac, J. I. & Zoller, P. Quantum repeaters: the role of imperfect local operations in quantum communication. *Physical Review Letters* **81**, 5932 (1998).
71. Gorshkov, A. V., André, A., Lukin, M. D. & Sørensen, A. S. Photon storage in Λ -type optically dense atomic media. II. Free-space model. *Physical Review A* **76**, 033805 (2007).
72. Higginbottom, D. B. *Atom-light couplers with one, two and ten billion atoms* PhD thesis (Department of Quantum Science, Centre for Quantum Computation et al., 2018).
73. Everett, J. L. *Atom-light interfaces for quantum information processing* PhD thesis (Centre for Quantum Computation, Communication Technology, Research School of Physics, and Engineering, College of Science, The Australian National University, 2018).
74. Campbell, G. T. *Quantum memories and monolithic resonators* PhD thesis (Australian National University. Dept. of Quantum Science., 2015).
75. Rambach, M., Nikolova, A., Weinhold, T. J. & White, A. G. Sub-megahertz linewidth single photon source. *Applied Physics Letters Photonics* **1**, 096101 (2016).
76. Jozsa, R. Fidelity for mixed quantum states. *Journal of Modern Optics* **41**, 2315–2323 (1994).
77. Higginbottom, D. B., Sparkes, B. M., Rancic, M., Pinel, O., Hosseini, M., Lam, P. K. & Buchler, B. C. Spatial-mode storage in a gradient-echo memory. *Physical Review A* **86**, 023801 (2012).
78. Li, M.-J. & Hayashi, T. in *Optical Fiber Telecommunications VII* 3–50 (Elsevier, 2020).

79. Albrecht, B., Farrera, P., Fernandez-Gonzalvo, X., Cristiani, M. & De Riedmaten, H. A waveguide frequency converter connecting rubidium-based quantum memories to the telecom C-band. *Nature Communications* **5**, 1–6 (2014).
80. Campbell, G., Pinel, O., Hosseini, M., Ralph, T. C., Buchler, B. & Lam, P. K. Configurable unitary transformations and linear logic gates using quantum memories. *Physical Review Letters* **113**, 063601 (2014).
81. Everett, J. L., Campbell, G. T., Cho, Y.-W., Vernaz-Gris, P., Higginbottom, D. B., Pinel, O., Robins, N. P., Lam, P. K. & Buchler, B. C. Dynamical observations of self-stabilizing stationary light. *Nature Physics* **13**, 68–73 (2017).
82. Everett, J. L., Higginbottom, D. B., Campbell, G. T., Lam, P. K. & Buchler, B. C. Stationary light in atomic media. *Advanced Quantum Technologies* **2**, 1800100 (2019).
83. Sabooni, M., Li, Q., Kröll, S. & Rippe, L. Efficient quantum memory using a weakly absorbing sample. *Physical Review Letters* **110**, 133604 (2013).
84. Beals, R., Brierley, S., Gray, O., Harrow, A. W., Kutin, S., Linden, N., Shepherd, D. & Stather, M. Efficient distributed quantum computing. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* **469**, 20120686 (2013).
85. Tanabe, T., Notomi, M., Kuramochi, E., Shinya, A. & Taniyama, H. Trapping and delaying photons for one nanosecond in an ultrasmall high-Q photonic-crystal nanocavity. *Nature Photonics* **1**, 49–52 (2007).
86. Takeda, S., Takase, K. & Furusawa, A. On-demand photonic entanglement synthesizer. *Science Advances* **5**, eaaw4530 (2019).
87. Kasapi, A, Jain, M., Yin, G. & Harris, S. E. Electromagnetically induced transparency: propagation dynamics. *Physical Review Letters* **74**, 2447 (1995).
88. Fleischhauer, M., Imamoglu, A. & Marangos, J. P. Electromagnetically induced transparency: Optics in coherent media. *Reviews of Modern Physics* **77**, 633 (2005).
89. Harris, S. E., Field, J. & Imamoglu, A. Nonlinear optical processes using electromagnetically induced transparency. *Physical Review Letters* **64**, 1107 (1990).
90. Boller, K.-J., Imamoglu, A & Harris, S. E. Observation of electromagnetically induced transparency. *Physical Review Letters* **66**, 2593 (1991).
91. Andre, A, Eisaman, M., Walsworth, R., Zibrov, A. & Lukin, M. Quantum control of light using electromagnetically induced transparency. *Journal of Physics B: Atomic, Molecular and Optical Physics* **38**, S589 (2005).
92. Fleischhauer, M. & Lukin, M. D. Dark-state polaritons in electromagnetically induced transparency. *Physical Review Letters* **84**, 5094 (2000).
93. Julsgaard, B., Sherson, J., Cirac, J. I., Fiurášek, J. & Polzik, E. S. Experimental demonstration of quantum memory for light. *Nature* **432**, 482–486 (2004).

94. Novikova, I., Walsworth, R. L. & Xiao, Y. Electromagnetically induced transparency-based slow and stored light in warm atoms. *Laser & Photonics Reviews* **6**, 333–353 (2012).
95. Choi, K. S., Deng, H., Laurat, J. & Kimble, H. Mapping photonic entanglement into and out of a quantum memory. *Nature* **452**, 67–71 (2008).
96. Geng, J., Campbell, G., Bernu, J., Higginbottom, D., Sparkes, B., Assad, S., Zhang, W., Robins, N., Lam, P. K. & Buchler, B. Electromagnetically induced transparency and four-wave mixing in a cold atomic ensemble with large optical depth. *New Journal of Physics* **16**, 113053 (2014).
97. Hsiao, Y.-F., Tsai, P.-J., Chen, H.-S., Lin, S.-X., Hung, C.-C., Lee, C.-H., Chen, Y.-H., Chen, Y.-F., Ite, A. Y. & Chen, Y.-C. Highly efficient coherent optical memory based on electromagnetically induced transparency. *Physical Review Letters* **120**, 183602 (2018).
98. Cho, Y.-W. & Kim, Y.-H. Atomic vapor quantum memory for a photonic polarization qubit. *Optics Express* **18**, 25786–25793 (2010).
99. Wang, Y., Li, J., Zhang, S., Su, K., Zhou, Y., Liao, K., Du, S., Yan, H. & Zhu, S.-L. Efficient quantum memory for single-photon polarization qubits. *Nature Photonics* **13**, 346–351 (2019).
100. Parigi, V., D’Ambrosio, V., Arnold, C., Marrucci, L., Sciarrino, F. & Laurat, J. Storage and retrieval of vector beams of light in a multiple-degree-of-freedom quantum memory. *Nature Communications* **6**, 1–7 (2015).
101. Akamatsu, D., Akiba, K. & Kozuma, M. Electromagnetically induced transparency with squeezed vacuum. *Physical Review Letters* **92**, 203602 (2004).
102. Longdell, J. J., Fraval, E., Sellars, M. J. & Manson, N. B. Stopped light with storage times greater than one second using electromagnetically induced transparency in a solid. *Physical Review Letters* **95**, 063601 (2005).
103. Heinze, G., Hubrich, C. & Halfmann, T. Stopped light and image storage by electromagnetically induced transparency up to the regime of one minute. *Physical Review Letters* **111**, 033601 (2013).
104. Afzelius, M., Simon, C., De Riedmatten, H. & Gisin, N. Multimode quantum memory based on atomic frequency combs. *Physical Review A* **79**, 052329 (2009).
105. Mossberg, T., Flusberg, A., Kachru, R. & Hartmann, S. Total scattering cross section for Na on He measured by stimulated photon echoes. *Physical Review Letters* **42**, 1665 (1979).
106. Afzelius, M., Usmani, I., Amari, A., Lauritzen, B., Walther, A., Simon, C., Sangouard, N., Minář, J., De Riedmatten, H., Gisin, N., *et al.* Demonstration of atomic frequency comb memory for light with spin-wave storage. *Physical Review Letters* **104**, 040503 (2010).
107. De Riedmatten, H., Afzelius, M., Staudt, M. U., Simon, C. & Gisin, N. A solid-state light–matter interface at the single-photon level. *Nature* **456**, 773–777 (2008).

108. Usmani, I., Afzelius, M., De Riedmatten, H. & Gisin, N. Mapping multiple photonic qubits into and out of one solid-state atomic ensemble. *Nature Communications* **1**, 1–7 (2010).
109. Sinclair, N., Saglamyurek, E., Mallahzadeh, H., Slater, J. A., George, M., Ricken, R., Hedges, M. P., Oblak, D., Simon, C., Sohler, W., *et al.* Spectral multiplexing for scalable quantum photonics using an atomic frequency comb quantum memory and feed-forward control. *Physical Review Letters* **113**, 053603 (2014).
110. Saglamyurek, E., Puigibert, M. G., Zhou, Q., Giner, L., Marsili, F., Verma, V. B., Nam, S. W., Oesterling, L., Nippa, D., Oblak, D., *et al.* A multiplexed light-matter interface for fibre-based quantum networks. *Nature Communications* **7**, 11202 (2016).
111. Seri, A., Lenhard, A., Rieländer, D., Gündoğan, M., Ledingham, P. M., Mazzera, M. & de Riedmatten, H. Quantum correlations between single telecom photons and a multimode on-demand solid-state quantum memory. *Physical Review X* **7**, 021028 (2017).
112. Kozhokin, A., Mølmer, K. & Polzik, E. Quantum memory for light. *Physical Review A* **62**, 033809 (2000).
113. Nunn, J, Walmsley, I., Raymer, M., Surmacz, K, Waldermann, F., Wang, Z & Jaksch, D. Mapping broadband single-photon wave packets into an atomic memory. *Physical Review A* **75**, 011401 (2007).
114. Reim, K., Nunn, J, Lorenz, V., Sussman, B., Lee, K., Langford, N., Jaksch, D & Walmsley, I. Towards high-speed optical quantum memories. *Nature Photonics* **4**, 218 (2010).
115. Saunders, D., Munns, J., Champion, T., Qiu, C, Kaczmarek, K., Poem, E, Ledingham, P., Walmsley, I. & Nunn, J. Cavity-enhanced room-temperature broadband Raman memory. *Physical Review Letters* **116**, 090501 (2016).
116. Michelberger, P., Champion, T., Sprague, M., Kaczmarek, K., Barbieri, M, Jin, X., England, D., Kolthammer, W., Saunders, D., Nunn, J, *et al.* Interfacing GHz-bandwidth heralded single photons with a warm vapour Raman memory. *New Journal of Physics* **17**, 043006 (2015).
117. Guo, J., Feng, X., Yang, P., Yu, Z., Chen, L., Yuan, C.-H. & Zhang, W. High-performance Raman quantum memory with optimal control in room temperature atoms. *Nature Communications* **10**, 1–6 (2019).
118. Bustard, P. J., Lausten, R., England, D. G. & Sussman, B. J. Toward quantum processing in molecules: A THz-bandwidth coherent memory for light. *Physical Review Letters* **111**, 083901 (2013).
119. England, D., Bustard, P., Nunn, J, Lausten, R & Sussman, B. From photons to phonons and back: A THz optical memory in diamond. *Physical Review Letters* **111**, 243601 (2013).

120. Ding, D.-S., Zhang, W., Zhou, Z.-Y., Shi, S., Xiang, G.-Y., Wang, X.-S., Jiang, Y.-K., Shi, B.-S. & Guo, G.-C. Quantum storage of orbital angular momentum entanglement in an atomic ensemble. *Physical Review Letters* **114**, 050502 (2015).
121. Ding, D.-S., Zhang, W., Zhou, Z.-Y., Shi, S., Bao-Sen & Guo, G.-C. Raman quantum memory of photonic polarized entanglement. *Nature Photonics* **9** (2015).
122. Ding, D.-S. in *Broad Bandwidth and High Dimensional Quantum Memory Based on Atomic Ensembles* 91–107 (Springer, 2018).
123. Surmacz, K., Nunn, J., Reim, K., Lee, K., Lorenz, V., Sussman, B, Walmsley, I. & Jaksch, D. Efficient spatially resolved multimode quantum memory. *Physical Review A* **78**, 033806 (2008).
124. Chen, B., Zhang, K., Bian, C., Qiu, C., Yuan, C.-H., Chen, L., Ou, Z. & Zhang, W. Efficient Raman frequency conversion by coherent feedback at low light intensity. *Optics Express* **21**, 10490–10495 (2013).
125. Lvovsky, A. I., Sanders, B. C. & Tittel, W. Optical quantum memory. *Nature Photonics* **3**, 706 (2009).
126. Hétet, G., Longdell, J., Alexander, A., Lam, P. K. & Sellars, M. Electro-optic quantum memory for light using two-level atoms. *Physical Review Letters* **100**, 023601 (2008).
127. Hétet, G., Hosseini, M., Sparkes, B. M., Oblak, D., Lam, P. K. & Buchler, B. C. Photon echoes generated by reversing magnetic field gradients in a rubidium vapor. *Optics Letters* **33**, 2323–2325 (2008).
128. Hosseini, M., Sparkes, B. M., Campbell, G., Lam, P. K. & Buchler, B. C. High efficiency coherent optical memory with warm rubidium vapour. *Nature Communications* **2**, 174 (2011).
129. Sparkes, B., Bernu, J., Hosseini, M., Geng, J., Glorieux, Q., Altin, P. A., Lam, P. K., Robins, N. & Buchler, B. Gradient echo memory in an ultra-high optical depth cold atomic ensemble. *New Journal of Physics* **15**, 085027 (2013).
130. Higginbottom, D. B., Geng, J., Campbell, G. T., Hosseini, M., Cao, M. T., Sparkes, B. M., Bernu, J., Robins, N. P., Lam, P. K. & Buchler, B. C. Dual-rail optical gradient echo memory. *Optics Express* **23**, 24937–24944 (2015).
131. Damon, V., Bonarota, M., Louchet-Chauvet, A., Chaneliere, T. & Le Gouët, J.-L. Revival of silenced echo and quantum memory for light. *New Journal of Physics* **13**, 093031 (2011).
132. Dajczgewand, J., Le Gouët, J.-L., Louchet-Chauvet, A. & Chanelière, T. Large efficiency at telecom wavelength for optical quantum memories. *Optics Letters* **39**, 2711–2714 (2014).
133. Vernaz-Gris, P. *Preparation of large cold atomic ensembles and applications in efficient light-matter interfacing* PhD thesis (Research School of Physics and Engineering, The Australian National University, 2018).

134. Moiseev, S. & Tittel, W. Optical quantum memory with generalized time-reversible atom–light interaction. *New Journal of Physics* **13**, 063035 (2011).
135. Wigley, P. B., Everitt, P. J., van den Hengel, A., Bastian, J. W., Sooriyabandara, M. A., McDonald, G. D., Hardman, K. S., Quinlivan, C. D., Manju, P., Kuhn, C. C., *et al.* Fast machine-learning online optimization of ultra-cold-atom experiments. *Scientific Reports* **6**, 1–6 (2016).
136. Vernaz-Gris, P., Tranter, A. D., Everett, J. L., Leung, A. C., Paul, K. V., Campbell, G. T., Lam, P. K. & Buchler, B. C. High-performance Raman memory with spatio-temporal reversal. *Optics Express* **26**, 12424–12431 (2018).
137. Alpaydin, E. & Bach, F. *Introduction to machine learning* (MIT press, 2020).
138. Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K. & Fei-Fei, L. *Imagenet: A large-scale hierarchical image database in 2009 IEEE conference on computer vision and pattern recognition* (2009), 248–255.
139. Xie, Q., Hovy, E., Luong, M.-T. & Le, Q. V. Self-training with Noisy Student improves ImageNet classification. *arXiv preprint arXiv:1911.04252* (2019).
140. Hastie, T., Tibshirani, R. & Friedman, J. *The elements of statistical learning: data mining, inference, and prediction* (Springer Science & Business Media, 2009).
141. Eskin, E., Arnold, A., Prerau, M., Portnoy, L. & Stolfo, S. in *Applications of data mining in computer security* 77–101 (Springer, 2002).
142. Silver, D., Hubert, T., Schrittwieser, J., Antonoglou, I., Lai, M., Guez, A., Lanctot, M., Sifre, L., Kumaran, D., Graepel, T., *et al.* A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science* **362**, 1140–1144 (2018).
143. Goodfellow, I., Bengio, Y. & Courville, A. *Deep Learning* <http://www.deeplearningbook.org> (MIT Press, 2016).
144. Lodish, H., Berk, A., Zipursky, S. L., Matsudaira, P., Baltimore, D. & Darnell, J. *Molecular cell biology* 4th edition. *National Center for Biotechnology Information, Bookshelf* (2000).
145. Maass, W. Networks of spiking neurons: the third generation of neural network models. *Neural Networks* **10**, 1659–1671 (1997).
146. Hornik, K. Approximation capabilities of multilayer feedforward networks. *Neural Networks* **4**, 251–257 (1991).
147. Rumelhart, D. E., Hinton, G. E. & Williams, R. J. Learning representations by back-propagating errors. *Nature* **323**, 533–536 (1986).
148. Kingma, D. P. & Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
149. Bengio, Y., Lamblin, P., Popovici, D. & Larochelle, H. *Greedy layer-wise training of deep networks in Advances in neural information processing systems* (2007), 153–160.

150. Hendrycks, D. & Gimpel, K. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415* (2016).
151. He, K., Zhang, X., Ren, S. & Sun, J. *Delving deep into rectifiers: Surpassing human-level performance on imagenet classification in Proceedings of the IEEE international conference on computer vision* (2015), 1026–1034.
152. Polyak, B. T. Some methods of speeding up the convergence of iteration methods. *USSR Computational Mathematics and Mathematical Physics* **4**, 1–17 (1964).
153. Duchi, J., Hazan, E. & Singer, Y. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research* **12**, 2121–2159 (2011).
154. LeCun, Y., Bengio, Y. & Hinton, G. Deep learning. *Nature* **521**, 436–444 (2015).
155. LeCun, Y., Boser, B., Denker, J. S., Henderson, D., Howard, R. E., Hubbard, W. & Jackel, L. D. Backpropagation applied to handwritten zip code recognition. *Neural Computation* **1**, 541–551 (1989).
156. Rumelhart, D. E., Smolensky, P., McClelland, J. L. & Hinton, G. Sequential thought processes in PDP models. *Parallel Distributed Processing: Explorations in the Microstructures of Cognition* **2**, 3–57 (1986).
157. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A. & Bengio, Y. *Generative adversarial nets in Advances in Neural Information Processing Systems* (2014), 2672–2680.
158. Karras, T., Laine, S. & Aila, T. *A style-based generator architecture for generative adversarial networks in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2019), 4401–4410.
159. Schawinski, K., Zhang, C., Zhang, H., Fowler, L. & Santhanam, G. K. Generative adversarial networks recover features in astrophysical images of galaxies beyond the deconvolution limit. *Monthly Notices of the Royal Astronomical Society: Letters* **467**, L110–L114 (2017).
160. Jain, P. & Kar, P. Non-convex optimization for machine learning. *arXiv preprint arXiv:1712.07897* (2017).
161. Boyd, S., Boyd, S. P. & Vandenberghe, L. *Convex optimization* (Cambridge university press, 2004).
162. Davletov, E., Tsyganok, V., Khlebnikov, V., Pershin, D., Shaykin, D. & Akimov, A. Machine Learning for Achieving Bose-Einstein Condensation of Thulium Atoms. *arXiv preprint arXiv:2003.00346* (2020).
163. Baranov, A, Burnaev, E, Derkach, D, Filatov, A, Klyuchnikov, N, Lantwin, O, Ratnikov, F, Ustyuzhanin, A & Zaitsev, A. *Optimising the active muon shield for the SHiP experiment at CERN in Journal of Physics: Conference Series* **934** (2017), 012050.

164. Pintér, J. D. *Global optimization in action: continuous and Lipschitz optimization: algorithms, implementations and applications* (Springer Science & Business Media, 2013).
165. Ratcliffe, A. K., Oberg, L. M. & Hope, J. J. Micromotion-Enhanced Fast Entangling Gates For Trapped Ion Quantum Computing. *arXiv preprint arXiv:1902.06364* (2019).
166. Sloss, A. N. & Gustafson, S. in *Genetic Programming Theory and Practice XVII* 307–344 (Springer, 2020).
167. Tamaki, H., Kita, H. & Kobayashi, S. *Multi-objective optimization by genetic algorithms: A review in Proceedings of IEEE international conference on evolutionary computation* (1996), 517–522.
168. Hornby, G., Globus, A., Linden, D. & Lohn, J. in *Space 2006* 7242 (2006).
169. Koziel, S., Bekasiewicz, A., Kurgan, P. & Bandler, J. W. Rapid multi-objective design optimisation of compact microwave couplers by means of physics-based surrogates. *IET Microwaves, Antennas & Propagation* **10**, 479–486 (2016).
170. Teodorescu, L. & Sherwood, D. High energy physics event selection with gene expression programming. *Computer Physics Communications* **178**, 409–419 (2008).
171. Thompson, A. *An evolved circuit, intrinsic in silicon, entwined with physics in International Conference on Evolvable Systems* (1996), 390–405.
172. Zahedinejad, E., Ghosh, J. & Sanders, B. C. High-fidelity single-shot Toffoli gate via quantum control. *Physical Review Letters* **114**, 200502 (2015).
173. Kirkpatrick, S., Gelatt, C. D. & Vecchi, M. P. Optimization by simulated annealing. *Science* **220**, 671–680 (1983).
174. Press, W. H., Teukolsky, S. A., Vetterling, W. T. & Flannery, B. P. *Numerical recipes in C* 1988.
175. De Freitas, J., Niranjana, M., Gee, A. & Doucet, A. Sequential Monte Carlo methods for optimisation of neural network models. *Cambridge University Engineering Department, Cambridge, England, Technical Report TR-328* (1998).
176. Aarts, E. & Van Laarhoven, P. Simulated annealing: an introduction. *Statistica Neerlandica* **43**, 31–52 (1989).
177. Pavlyukevich, I. Lévy flights, non-local search and simulated annealing. *Journal of Computational Physics* **226**, 1830–1844 (2007).
178. Abdel-Basset, M., Abdel-Fatah, L. & Sangaiah, A. K. in *Computational intelligence for multimedia big data on the cloud with engineering applications* 185–231 (Elsevier, 2018).
179. Blum, C. & Roli, A. Metaheuristics in combinatorial optimization: Overview and conceptual comparison. *Association for Computing Machinery Computing Surveys* **35**, 268–308 (2003).

180. Pan, W.-T. *A new evolutionary computation approach: fruit fly optimization algorithm in 2011 Conference of Digital Technology and Innovation Management* (2011), 382–391.
181. Zhang, Y., Wang, S. & Ji, G. A comprehensive survey on particle swarm optimization algorithm and its applications. *Mathematical Problems in Engineering* **2015** (2015).
182. Zheng, Y.-J. Water wave optimization: a new nature-inspired metaheuristic. *Computers & Operations Research* **55**, 1–11 (2015).
183. Donelli, M., Azaro, R., De Natale, F. G. & Massa, A. An innovative computational approach based on a particle swarm strategy for adaptive phased-arrays control. *IEEE Transactions on Antennas and Propagation* **54**, 888–898 (2006).
184. Selvan, S. E., Xavier, C. C., Karssemeijer, N., Sequeira, J., Cherian, R. A. & Dhala, B. Y. Parameter estimation in stochastic mammogram model by heuristic optimization techniques. *IEEE Transactions on Information Technology in Biomedicine* **10**, 685–695 (2006).
185. Venayagamoorthy, G. K. & Zha, W. Comparison of nonuniform optimal quantizer designs for speech coding with adaptive critics and particle swarm. *IEEE Transactions on Industry Applications* **43**, 238–244 (2007).
186. Poli, R. Analysis of the publications on the applications of particle swarm optimisation. *Journal of Artificial Evolution and Applications* **2008** (2008).
187. Jin, Y. A comprehensive survey of fitness approximation in evolutionary computation. *Soft Computing* **9**, 3–12 (2005).
188. Jin, Y. Surrogate-assisted evolutionary computation: Recent advances and future challenges. *Swarm and Evolutionary Computation* **1**, 61–70 (2011).
189. Won, K. S. & Ray, T. A framework for design optimization using surrogates. *Engineering Optimization* **37**, 685–703 (2005).
190. Rayas-Sánchez, J. E. EM-based optimization of microwave circuits using artificial neural networks: The state-of-the-art. *IEEE Transactions on Microwave Theory and Techniques* **52**, 420–435 (2004).
191. Hacioglu, A. Fast evolutionary algorithm for airfoil design via neural network. *AIAA Journal* **45**, 2196–2203 (2007).
192. Wolpert, D. H. & Macready, W. G. No free lunch theorems for optimization. *IEEE transactions on evolutionary computation* **1**, 67–82 (1997).
193. Opitz, D. & Maclin, R. Popular ensemble methods: An empirical study. *Journal of Artificial Intelligence Research* **11**, 169–198 (1999).
194. Abadi, M. *et al. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems* Software available from tensorflow.org. 2015. <https://www.tensorflow.org/>.
195. Liu, D. C. & Nocedal, J. On the limited memory BFGS method for large scale optimization. *Mathematical Programming* **45**, 503–528 (1989).
196. Price, K. V. in *Handbook of Optimization* 187–214 (Springer, 2013).

197. Townsend, C., Edwards, N., Cooper, C., Zetie, K., Foot, C., Steane, A., Szriftgiser, P., Perrin, H & Dalibard, J. Phase-space density in the magneto-optical trap. *Physical Review A* **52**, 1423 (1995).
198. Hanley, R. K., Huillery, P., Keegan, N. C., Bounds, A. D., Boddy, D., Faoro, R. & Jones, M. P. Quantitative simulation of a magneto-optical trap operating near the photon recoil limit. *Journal of Modern Optics* **65**, 667–676 (2018).
199. Torrontegui, E., Ibáñez, S., Martínez-Garaot, S., Modugno, M., del Campo, A., Guéry-Odelin, D., Ruschhaupt, A., Chen, X. & Muga, J. G. in *Advances in atomic, molecular, and optical physics* 117–169 (Elsevier, 2013).
200. Hu, J., Urvoy, A., Vendeiro, Z., Crépel, V., Chen, W. & Vuletić, V. Creation of a Bose-condensed gas of 87Rb by laser cooling. *Science* **358**, 1078–1080 (2017).
201. Palittapongarnpim, P., Wittek, P., Zahedinejad, E., Vedaie, S. & Sanders, B. C. Learning in quantum control: High-dimensional global optimization for noisy quantum dynamics. *Neurocomputing* **268**, 116–126 (2017).
202. Palittapongarnpim, P., Wittek, P. & Sanders, B. C. *Controlling adaptive quantum phase estimation with scalable reinforcement learning in 24th European Symposium on Artificial Neural Networks, Bruges, April 27–29, 2016* (2016), 327–332.
203. August, M. & Ni, X. Using recurrent neural networks to optimize dynamical decoupling for quantum memory. *Physical Review A* **95**, 012335 (2017).
204. Ju, S., Shiga, T., Feng, L., Hou, Z., Tsuda, K. & Shiomi, J. Designing nanostructures for phonon transport via Bayesian optimization. *Physical Review X* **7**, 021024 (2017).
205. Li, C., de Celis Leal, D. R., Rana, S., Gupta, S., Sutti, A., Greenhill, S., Slezak, T., Height, M. & Venkatesh, S. Rapid Bayesian optimisation for synthesis of short polymer fiber materials. *Scientific Reports* **7**, 1–10 (2017).
206. Mavadia, S., Frey, V., Sastrawan, J., Dona, S. & Biercuk, M. J. Prediction and real-time compensation of qubit decoherence via machine learning. *Nature Communications* **8**, 1–6 (2017).
207. Amstrup, B., Toth, G. J., Szabo, G., Rabitz, H. & Loerincz, A. Genetic algorithm with migration on topology conserving maps for optimal control of quantum systems. *The Journal of Physical Chemistry* **99**, 5206–5213 (1995).
208. Geisel, I., Cordes, K., Mahnke, J., Jöllenbeck, S., Ostermann, J., Arlt, J., Ertmer, W. & Klempt, C. Evolutionary optimization of an experimental apparatus. *Applied Physics Letters* **102**, 214105 (2013).
209. Warren, W. S., Rabitz, H. & Dahleh, M. Coherent control of quantum dynamics: the dream is alive. *Science* **259**, 1581–1589 (1993).
210. Lehman, J., Clune, J. & Misevic, D. *The surprising creativity of digital evolution in Artificial Life Conference Proceedings* (2018), 55–56.

211. DePue, M. T., McCormick, C., Winoto, S. L., Oliver, S. & Weiss, D. S. Unity occupation of sites in a 3D optical lattice. *Physical Review Letters* **82**, 2262 (1999).
212. Gordon, J. & Ashkin, A. Motion of atoms in a radiation trap. *Physical Review A* **21**, 1606 (1980).
213. Cartis, C., Roberts, L. & Sheridan-Methven, O. Escaping local minima with derivative-free methods: a numerical investigation. *arXiv preprint arXiv:1812.11343* (2018).
214. Kane, B. E. A silicon-based nuclear spin quantum computer. *Nature* **393**, 133–137 (1998).
215. Loss, D. & DiVincenzo, D. P. Quantum computation with quantum dots. *Physical Review A* **57**, 120 (1998).
216. Leon, R., Yang, C. H., Hwang, J., Lemyre, J. C., Tanttu, T., Huang, W., Chan, K. W., Tan, K., Hudson, F., Itoh, K., *et al.* Coherent spin control of s-, p-, d- and f-electrons in a silicon quantum dot. *Nature Communications* **11**, 1–7 (2020).
217. Huang, W., Yang, C., Chan, K., Tanttu, T., Hensen, B., Leon, R., Fogarty, M., Hwang, J., Hudson, F., Itoh, K. M., *et al.* Fidelity benchmarks for two-qubit gates in silicon. *Nature* **569**, 532–536 (2019).
218. Yang, C., Chan, K., Harper, R., Huang, W., Evans, T., Hwang, J., Hensen, B., Laucht, A., Tanttu, T., Hudson, F., *et al.* Silicon qubit fidelities approaching incoherent noise limits via pulse engineering. *Nature Electronics* **2**, 151–158 (2019).
219. Yoneda, J., Takeda, K., Otsuka, T., Nakajima, T., Delbecq, M. R., Allison, G., Honda, T., Kodera, T., Oda, S., Hoshi, Y., *et al.* A quantum-dot spin qubit with coherence limited by charge noise and fidelity higher than 99.9%. *Nature Nanotechnology* **13**, 102–106 (2018).
220. Yang, C., Lim, W., Zwanenburg, F. & Dzurak, A. Dynamically controlled charge sensing of a few-electron silicon quantum dot. *American Institute of Physics Advances* **1**, 042111 (2011).
221. Morello, A., Pla, J. J., Zwanenburg, F. A., Chan, K. W., Tan, K. Y., Huebl, H., Möttönen, M., Nugroho, C. D., Yang, C., van Donkelaar, J. A., *et al.* Single-shot readout of an electron spin in silicon. *Nature* **467**, 687–691 (2010).
222. Somaschi, N., Giesz, V., De Santis, L., Loredano, J., Almeida, M. P., Hornecker, G., Portalupi, S. L., Grange, T., Antón, C., Demory, J., *et al.* Near-optimal single-photon sources in the solid state. *Nature Photonics* **10**, 340 (2016).
223. Dousse, A., Lanco, L., Suffczyński, J., Semenova, E., Miard, A., Lemaître, A., Sagnes, I., Roblin, C., Bloch, J. & Senellart, P. Controlled light-matter coupling for a single quantum dot embedded in a pillar microcavity using far-field optical lithography. *Physical Review Letters* **101**, 267404 (2008).
224. He, Y.-M., Wang, H., Wang, C., Chen, M.-C., Ding, X., Qin, J., Duan, Z.-C., Chen, S., Li, J.-P., Liu, R.-Z., *et al.* Coherently driving a single quantum two-level system with dichromatic laser pulses. *Nature Physics* **15**, 941–946 (2019).

225. Frumker, E. & Silberberg, Y. Phase and amplitude pulse shaping with two-dimensional phase-only spatial light modulators. *Journal of the Optical Society of America B* **24**, 2940–2947 (2007).
226. Sutton, R. S. & Barto, A. G. *Reinforcement learning: An introduction* (MIT press, 2018).
227. Bahdanau, D., Brakel, P., Xu, K., Goyal, A., Lowe, R., Pineau, J., Courville, A. & Bengio, Y. An actor-critic algorithm for sequence prediction. *arXiv preprint arXiv:1607.07086* (2016).
228. Shannon, C. E. XXII. Programming a computer for playing chess. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* **41**, 256–275 (1950).
229. Schulman, J., Moritz, P., Levine, S., Jordan, M. & Abbeel, P. High-dimensional continuous control using generalized advantage estimation. *arXiv preprint arXiv:1506.02438* (2015).
230. Pascanu, R., Mikolov, T. & Bengio, Y. *On the difficulty of training recurrent neural networks* in *International conference on machine learning* (2013), 1310–1318.
231. Bengio, Y., Simard, P. & Frasconi, P. Learning long-term dependencies with gradient descent is difficult. *IEEE transactions on Neural Networks* **5**, 157–166 (1994).
232. Dennis, G. R., Hope, J. J. & Johnsson, M. T. XMDS2: Fast, scalable simulation of coupled stochastic partial differential equations. *Computer Physics Communications* **184**, 201–208 (2013).
233. Tranter, A. D. *Raman Optimisation Code* https://github.com/tranterfied/raman_optimisation.git.
234. Freier, C. Atom interferometry at geodetic observatories (2017).
235. Riehle, F, Kisters, T., Witte, A, Helmcke, J & Bordé, C. J. Optical Ramsey spectroscopy in a rotating frame: Sagnac effect in a matter-wave interferometer. *Physical Review Letters* **67**, 177 (1991).
236. Mcguirk, J. M., Foster, G., Fixler, J., Snadden, M. & Kasevich, M. Sensitive absolute-gravity gradiometry using atom interferometry. *Physical Review A* **65**, 033608 (2002).
237. Senger, A. A mobile atom interferometer for high-precision measurements of local gravity (2012).
238. Freier, C. *Measurement of local gravity using atom interferometry* PhD thesis (Diploma thesis, Humboldt-Universität zu Berlin, 2010).
239. Schmidt, M. & Lipson, H. Distilling free-form natural laws from experimental data. *Science* **324**, 81–85 (2009).
240. Quade, M., Abel, M., Shafi, K., Niven, R. K. & Noack, B. R. Prediction of dynamical systems by symbolic regression. *Physical Review E* **94**, 012214 (2016).

241. Allan, D. W. Should the classical variance be used as a basic measure in standards metrology? *IEEE Transactions on Instrumentation and Measurement*, 646–654 (1987).