Spin qubits in semiconductor devices: automation of measurement, optimisation, and control.



A THESIS PRESENTED
BY
SEBASTIAN B. ORBELL
TO
THE DEPARTMENT OF MATERIALS

FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

TRINITY COLLEGE,
UNIVERSITY OF OXFORD
MARCH 2024

©2024 – Sebastian B. Orbell
all rights reserved.

Thesis advisors: Natalia Ares, Andrew Briggs Sebastian B. Orbell

Spin qubits in semiconductor devices: automation of measurement, optimisation, and control.

Abstract

In the nascent field of semiconductor quantum computing, the current state-of-the-art systems are not far removed from the advanced physics experiments conducted in academic research labs. The automation of the tuneup and control of these systems therefore represents a critical step forward. This thesis aims to contribute to the transition from this academic-centric approach to a more integrated and automated platform, essential for the realisation of a scalable semiconductor quantum computing processor.

Central to this thesis is the development of three approaches which in turn automate: the discovery of spin qubit operational regimes; the characterisation and optimisation of qubit properties; and the design of optimal quantum gate operations. The first consists of a deep reinforcement learning decision agent and a convolutional neural network computer vision model, which enable the automatic and efficient discovery of bias triangles. These transport features, measured here in a GaAs double quantum dot, define the operational regime of a spin qubit, and their identification represents a crucial part of a fully automated tuning pipeline. The second approach is a Bayesian framework for the automated optimisation of spin qubit properties and discovery of operational "sweet-spots". The efficacy of this framework is demonstrated by optimising the Rabi frequency of a hole spin qubit in a Si finFET and the Q-factor of a singlet-triplet hole spin qubit in a Si/SiGe planar device. Lastly, this thesis introduces an algorithm for designing optimal quantum control schemes to achieve robust high-fidelity gate operations in spin qubits. This algorithm is capable of computing gradients through the time-evolution of quantum systems and incorporates realistic experimental constraints related to hardware and control electronics, such as non-linear transmission, cross-resonant driving, and finite bandwidths which can be tailored to each experiment. I also introduce a numerical sampling method which enables the design of pulse schemes which are robust to coloured noise spectra, in addition to quasi-static errors in Hamiltonian parameters, to mitigate incoherent errors and ensure applicability in real-world quantum computing scenarios. I demonstrate this method by simulating a driven hole-spin system, optimising pulse envelopes for single-, simultaneous single-, and two-qubit gates. I proceed to illustrate a method for the construction of a differentiable digital twin for a quantum processor using a Neural controlled differential equation model. This model provides a mechanism to bridge the gap between model uncertainty and experiments.

Contents

1 Introduction **1**

2 Background and related work **6**
- 2.1 Background . . . 7
  - 2.1.1 Quantum Computing . . . 7
  - 2.1.2 Semiconductor Qubits . . . 9
  - 2.1.3 Bias triangle charge transport . . . 14
  - 2.1.4 Two-state systems . . . 15
  - 2.1.5 Loss-DiVincenzo qubit . . . 18
- 2.2 Related work . . . 21
  - 2.2.1 Automated tuning of semiconductor qubits . . . 21
  - 2.2.2 Optimal quantum control . . . 25

3 Methods **29**
- 3.1 Charge-state measurements . . . 30
  - 3.1.1 DC measurements . . . 31
  - 3.1.2 rf-measurements . . . 32
  - 3.1.3 Charge sensing . . . 34
  - 3.1.4 Video-mode measurements . . . 35
- 3.2 Bayesian methods for automating and optimising quantum experiments 40
  - 3.2.1 Bayesian inference measurement analysis . . . 40
  - 3.2.2 Bayesian optimal experimental design . . . 42

4 Deep Reinforcement Learning for the Efficient Measurement of Quantum Devices **44**
- 4.1 Introduction . . . 45
- 4.2 Description of the algorithm . . . 50
- 4.3 The deep reinforcement learning agent . . . 54
- 4.4 Experimental results . . . 57
- 4.5 Discussion . . . 63

5 Bayesian methods for the automatic discovery of "sweet-spot" qubit operating regimes **66**
- 5.1 Introduction . . . 67

5.2 Devices and measurement setup . . . 69
5.2.1 Rabi chevron measurements . . . 69
5.2.2 Single Rabi trace measurements . . . 72
5.3 Quality metric analysis . . . 74
5.3.1 Bayesian inference . . . 74
5.3.2 Fourier analysis . . . 79
5.4 Bayesian optimisation of quality metrics . . . 82
5.5 Experimental results . . . 84
5.5.1 A Loss DiVincenzo qubit in a Si finFET device . . . 84
5.5.2 A Singlet-Triplet qubit in a planar SiGe device . . . 86
5.6 Discussion . . . 88

6 Robust optimal quantum control for high-fidelity gates with spin qubits **90**
6.1 Introduction . . . 91
6.2 Robust optimal quantum control . . . 94
6.3 Algorithm . . . 95
6.3.1 Pulse distortions . . . 98
6.4 Applications, optimal control of hole spin qubits . . . 101
6.4.1 Hole spin qubit Hamiltonian . . . 101
6.4.2 Applications . . . 105
6.5 Discussion . . . 112

7 A differentiable digital twin: Using Neural CDEs to simulate quantum systems **115**
7.1 Introduction . . . 116
7.2 Neural Controlled Differential Equations . . . 119
7.3 Network architecture . . . 121
7.4 Simulating quantum systems . . . 121
7.4.1 Singlet-Triplet qubits . . . 121
7.4.2 Model training . . . 122
7.4.3 Results . . . 123
7.5 Optimal quantum control . . . 124
7.6 Discussion . . . 127

8 Conclusion **129**
8.1 Reflections on automated scientific control and discovery . . . 130
8.2 Reflections on spin qubits . . . 131
8.3 Summary of results and applications . . . 133
8.4 Limitations and future research . . . 134

References **157**

Appendix A Singlet-Triplet qubit Hamiltonian **159**

Appendix B The Schrieffer-Wolff transformation **161**

Appendix C The Haar Measure **164**

Glossary

2DEG . . . . . Two Dimensional Electron Gas.

2DHG . . . . . Two Dimensional Hole Gas.

AC . . . . . Alternating Current.

ADVI . . . . . Automatic Differentiation Variational Inference.

AI . . . . . Artificial Intelligence.

AWG . . . . . Arbitrary Waveform Generator.

CMOS . . . . . Complementary Metal-Qxide-Semiconductor.

CNN . . . . . Convolutional Neural Network.

CNOT . . . . . Controlled Not.

COCOA . . . . . Complex and Optional Constraints Optimization with Auto-differentiation).

CPU . . . . . Central Processing Unit.

DC . . . . . Direct Current.

DQN . . . . . Deep Q-Network.

DRL . . . . . Deep Reinforcement Learning.

EDSR . . . . . Electric Dipole Spin Resonance.

ELBO . . . . . Evidence Lower Bound.

ESR . . . . . Electron Spin Resonance.

finFET . . . . . fin Field Effect Transistor.

GP . . . . . Gaussian Process.

GRAPE . . . . . GRadient Ascent Pulse Engineering.

HH . . . . . Heavy Hole.

KL . . . . . Kullback-Leibler.

LH . . . . . Light Hole.

ML . . . . . Machine Learning.

MSE . . . . . Mean Squared Error.

MUB . . . . . Mutually Unbiased Bases.

MW . . . . . Micro-Wave.

NCDE . . . . . Neural Controlled Differential Equation.

ORBIT . . . . . Optimized Randomized Benchmarking for Immediate Tune-up.

PSB . . . . . Pauli Spin Blockade.

PWC . . . . . Piece Wise Constant.

QD . . . . . Quantum Dot.

RF . . . . . Radio-Frequency.

RL . . . . . Reinforcement Learning.

SOI . . . . . Spin Orbit Interaction.

"I was within and without, simultaneously enchanted and repelled by the inexhaustible variety of life."
F. Scott Fitzgerald

Acknowledgments

To my friends, and to my family, To Joseph Hickie, William Wombell, Barnaby van Straaten for their intellectual discussions, bike maintenance, endurance sport competitions, and friendships. To Elodie Freymann, Maša Mićunović and Emma Madden for their stimulating philosophies, aesthetic insights, and enduring love.

Contributions

Statement

The content of this thesis is my original work with the input of named collaborators.

Papers

Deep reinforcement learning for efficient measurement of quantum devices
S. B. Orbell, V. Nguyen, D. T. Lennon, H. Moon, F. Vigneau, L. C. Camenzind, L. Yu, D. M. Zumbühl, G. A. D. Briggs, M. A. Osborne, D. Sejdinovic N. Ares
npj Quantum Inf 7, 100 (2021)

Reducing leakage of single-qubit gates for superconducting quantum processors using analytical control pulse envelopes
Eric Hyyppä, Antti Vepsäläinen, Miha Papič, Chun Fai Chan, Sinan Inel, Alessandro Landra, Wei Liu, Jürgen Luus, Fabian Marxer, Caspar Ockeloen-Korppi, Sebastian Orbell, Brian Tarasinski, Johannes Heinsoo
arXiv:2402.17757 (2024)

Looping in the Human: Collaborative and Explainable Bayesian Optimization
Masaki Adachi, Brady Planden, David A. Howey, Michael A. Osborne, Sebastian Orbell, Natalia Ares, Krikamol Muandet, Siu Lun Chau
AISTATS (2024)

Talks

Deep Reinforcement Learning for Efficient Measurement of Quantum Devices
S. B. Orbell
APS March Meeting 2021 talk.

Bayesian methods for the efficient measurement of Rabi oscillations
S. B. Orbell
Machine Learning for Quantum 2021 video poster.

Bayesian methods for optimising qubit quality factors
S. B. Orbell
APS March Meeting 2023 talk.

ML for optimising spin qubit experiments
S. B. Orbell
Infleqtion quantum AI workshop 2023 contributed talk.

A note on points of view

In the introductory sections, the prefaces to each chapter, and the concluding remarks, I will adopt a first-person narrative. This choice reflects a personal perspective on the field and the research undertaken within this thesis. Conversely, for the methodology and results chapters, I will assume a third-person narrative. This shift aims to enhance clarity and maintain adherence to the conventional standards of scientific writing. To facilitate a clear distinction between these narrative styles, I will explicitly delineate my contributions and the contributions of others in the subsequent note.

Description of contributions

My contributions and the inputs of my collaborators are as follows:

In the paper **"Deep reinforcement learning for efficient measurement of quantum"**, the results of which appear in Chapter 4, I wrote the code to run the experiments, ideated and developed the pre-classifier and classifier, trained the DRL agent, and CNN classifier, ran the experiments, and wrote the manuscript. V. Nguyen had the idea of using deep reinforcement learning to discover bias triangles, designed the DRL network architecture, and wrote the code for this network. D. T. Lennon, and F. Vigneau assisted in the lab with experimental setup. H. Moon, M. A. Osborne, and D. Sejdinovic aided with the conceptualisation of the DRL agent. L. C. Camenzind, L. Yu, and D. M. Zumbühl fabricated and provided the GaAs device. G. A. D. Briggs, and N. Ares assisted with the entire process, particularly in the provision of the resources, conceptualisation of the project, and assistance with the manuscript.

In Chapter 5, **"Bayesian methods for the automatic discovery of "sweet-spot" qubit operating regimes"**, I conceptualised the method, wrote the code for the analysis, optimisation, and experiment integration, performed the experiments, analysed the results, and wrote the text. J. D. Hickie, and B. van Straaten aided with the experimental setup, device tuning, and code integration. D. Jirovec, and G. Katsaros fabricated the Si/SiGe device, and aided with the experiment. A. V. Kuhlmann, S. Geyer, L.C. Camenzind, and D. M. Zumbühl fabricated the Si finFET device, tuned up the device and helped with the experiment setup. N. Ares assisted with the entire process, particularly in the provision of the resources, conceptualisation of the project, and assistance with the manuscript.

In Chapter 6, **"Robust optimal quantum control for high-fidelity gates in hole spin qubits"**, and Chapter 7, **"A differentiable digital twin: Using Neural CDEs to simulate quantum systems"**, I conceptualised the methods, wrote the code, performed the experiments, analysed the results, and wrote the manuscript. V. Michal assisted with the development of the Hamiltonian and noise model, and N. Ares assisted

with the entire process, particularly in the provision of the resources, conceptualisation of the project, and assistance with the manuscript.

Chapter 1

Introduction

"*It is entropy, not energy, that keeps stones on the ground and the world turning*".

Carlo Rovelli

Quantum computing is on the verge of a groundbreaking shift, poised to redefine the boundaries of computational power and efficiency. At the heart of this burgeoning field lies the promise of semiconductor quantum dot-based quantum computing, a domain that has witnessed substantial strides in recent years. Despite these advancements, the transition from experimental setups in academic labs to scalable, integrated quantum computing platforms remains a significant challenge. This thesis contributes to ongoing efforts to increase the scalability of semiconductor quantum dot-based quantum processors, with implications for the wider quantum computing community.

The quantum computing systems of today, even those which are well-funded and state-of-the-art, are still not far removed from the physics experiments installed for academic research. To unlock the immense potential of quantum computing, a paradigm shift is essential – from manual, labour-intensive processes in research laboratories to an automated, integrated platform. Spanning four distinct yet interconnected chapters, this thesis first develops an understanding of the physics and inherent limitations of semiconductor quantum dot-based devices, and proceeds to contribute methodologies that enable the automation of the tune-up, measurement, and control of spin qubits. This automation is crucial for the advancement and practical application of quantum computing technologies. Specifically, I present chapters which in turn demonstrate; the efficient discovery of spin qubit operational regimes, bias triangles; the automatic characterisation and optimisation of the spin qubit operational regime; and the development of a framework which enables the design and analysis of optimal quantum control schemes, with applications to spin qubits.

The first results chapter, titled "**Deep Reinforcement Learning for Efficient Measurement of Quantum Devices**" (Chapter 4) demonstrates the use of deep reinforce-

ment learning in the automated measurement of a GaAs double quantum dot device. This approach, which mirrors human learning by adopting a strategy based on actions and their resultant rewards, is applied to automate decision-making processes in quantum devices. Specifically, the paper focuses on double quantum dots, demonstrating the algorithm's capability to autonomously identify transport features within large parameter spaces. These features, crucial yet challenging to locate due to their presence in featureless regions, are identified rapidly by the algorithm, often in less than 30 minutes. This significant finding, achieved using a duelling deep Q-network architecture, showcases deep reinforcement learning's potential in quantum device applications. Its success emphasises this approach's applicability for decision-making in quantum measurement and operation.

The second results chapter, "**Bayesian methods for the automatic discovery of 'sweet-spot' qubit operating regimes**" (Chapter 5), highlights the importance of engineering qubits with optimal properties for the development of error-resistant quantum systems. I present an algorithm that automates the optimisation of qubit metrics, such as the Rabi frequency, and the discovery of operational "sweet-spots" using Bayesian inference, Fourier analysis, and Bayesian optimisation. The versatility of this method is demonstrated through its application in optimising quality metrics across different qubit parameterisations and distinct devices, including a Loss DiVincenzo qubit in a Si finFET, and a Singlet-Triplet qubit in a planar Si/SiGe device.

In the third results chapter, "**Robust optimal quantum control for high-fidelity gates with spin qubits**" (Chapter 6), I present a method tailored for developing optimal control schemes for quantum systems, with a focus on hole spin qubits. This approach is designed to mitigate coherent errors and enhance robustness against inco-

herent errors in order to implement high-fidelity quantum gate operations. A unique feature is its in-situ handling of pulse distortions, such as filter effects and non-linear transmission functions, crucial to minimising coherent errors in quantum gates. Central to its efficacy is a numerical sampling method which enables the design of control schemes resilient to decoherence and quasi-static errors. The chapter demonstrates the method's effectiveness through case studies involving single-, simultaneous single-, and two-qubit gates in hole spin qubit systems. The control schemes are particular to the device, control electronics stack, and noise environment under investigation, but the implications and methods are general.

Lastly, in Chapter 7, **"A differentiable digital twin: Using Neural CDEs to simulate quantum systems"**, I present a novel methodology which leverages Neural Controlled Differential Equations (Neural CDEs) to create a digital twin of a quantum system. The framework provides a solution to bridge the gap between experiments and incomplete theoretical models. The resultant model is fully differentiable and can be utilised to design optimal control schemes via stochastic gradient descent. This method incorporates both the advantages of differentiable models and the adaptability of Neural CDEs to handle complex noise models and yield an accurate differentiable model of the quantum system.

Collectively, these papers contribute significantly to the field of quantum computing, offering solutions for the automated measurement, optimisation, and control of spin qubit quantum devices. This thesis not only aims to deepen the understanding of semiconductor quantum dot-based quantum computing but to also bridge the gap between academic research and industrial applications. This research provides an insight into the potential of automated systems to change the field of quantum computing, potentially

paving the way for more advanced, efficient, and scalable quantum processors.

Chapter 2

Background and related work

"*What could I say to you that would be of value, except that perhaps you seek too much, that as a result of your seeking you cannot find*".

Hermann Hesse

The context of this work is within the field of semiconductor quantum dot based quantum computing, with an intent to also have an impact on the wider quantum computing community. In order to realise the vast potential of this new computing paradigm, we must make a significant shift away from academic research labs powered by PhD students, towards an integrated and automated platform. The aim of this thesis is to contribute approaches which will enable the automation of the tune-up, measurement, and control of spin qubits.

In the subsequent section, I delve into the theoretical and experimental background underpinning this thesis and discuss some related work. By examining the historical developments, key theoretical constructs, and recent technological breakthroughs, we gain insights into the challenges and opportunities that have shaped spin qubit quantum processors. Additionally, a review of related work is crucial for identifying gaps in existing methodologies and technologies, thereby motivating the need for the novel approaches introduced in the following chapters. This section sets the stage for understanding the significance of the proposed tools and methods in advancing the automation and optimisation of spin qubits, highlighting how this research contributes to and extends the current body of knowledge.

2.1 Background

2.1.1 Quantum Computing

Richard Feynman infamously posed the question as to whether we could simulate an arbitrary physical system, including quantum systems, on a universal classical computer. He concluded "Nature isn't classical, dammit, and if you want to make a simulation of nature, you'd better make it quantum mechanical". Due to the properties of en-

tanglement and the superposition of states, simulating a quantum system on a classical computer would quickly become intractable as the size of the system is scaled. Instead, a new quantum mechanical paradigm of computing must be considered, where computers use a quantum bit or a qubit.

In the 1970s, David Deutsch developed the concept and outlined the machinations of a universal quantum processor. He originally proposed the machine, now known as a quantum computer, as an experimental test for Everett's many worlds interpretation of quantum mechanics[1,2]. Subsequently, Peter Shor developed an algorithm for integer factorisation in polynomial time[3]. Shor's algorithm was shown to be almost exponentially faster than the most efficient classical algorithm. This was the first proof of principle result that demonstrated the huge computational advantages which quantum computers offer. It is postulated that quantum computers will have profound influences in many fields from chemistry[4] to artificial intelligence[5] and thus a large international effort to mature the technology has ensued[6].

Universal Quantum Computation

The physical properties of a system needed for universal quantum computation are well known. Loss and DiVincenzo outlined these criteria in 1997[7]. Principally, the system must exhibit an isolated two-state quantum object known as the qubit. Secondly, the process of preparing the quantum state in the qubit system must be reliable. The information encoded in the quantum state must be retained for a length of time which is at the least on the order of magnitude of the computational time. In practice, this means that the decoherence of the qubits must be minimised. Any quantum algorithm can be constructed from the XOR two-qubit gate in combination with single qubit gates. There-

fore, the system must have accurate one and two qubit gates which natively constitute a Universal gate set. Finally, the system must enable strong quantum measurements which denote a complete collapse of the wavefunction to produce the eigenvalue of some operator.

Many physical systems have been proposed as potentially scalable architectures for quantum computing including superconducting circuits[8,9], trapped ions[10,11], neutral atoms[12], defects in diamonds[13], and many more[6]. Each of these systems has its respective advantages and disadvantages for the prospect of developing a scalable universal quantum computer. Semiconductor quantum dots have many unique properties that lend themselves to creating scalable quantum architectures including; a long relaxation time, the potential for all electrical control, and a high degree of tunability[14]. It is expected that the community building these systems will be able to draw on the huge wealth of knowledge developed by the standard semiconductor industry. Furthermore, the industrial processes used to develop classical computing chips could be modified to mass-produce quantum computing chips once the technology has matured.

2.1.2 Semiconductor Qubits

The experimental realisation of the proposal by Loss and DiVincenzo requires an array of coupled quantum dots. Depending upon the chosen qubit system, discussed later in section 2.1.4, each dot must be occupied by a finite number of electrons and have tunable interdot tunnel couplings.

In early quantum dot architectures, an electron was confined to zero dimensions by engineering different materials in all three dimensions. For example, a semiconductor with a small bandgap can be enclosed by a semiconductor with a large bandgap to form

potential barriers[15]. A second approach is to engineer a heterostructure which creates a lower dimensional confinement potential and to use electrostatic gates to localise the electron in the remaining dimensions. This method allows for tunability of the inderdot tunnel couplings as well as control over the electron occupancy of the quantum dot. There has also been significant progress in the development of donor spin qubits[16]. Here, for example, the $I = 1/2$ nuclear spin of a ^{31}P donor in Si is used as the qubit basis and the resonant frequency is controlled by using an electrostatic gate to modulate the proximity of the donor electron[17]. Two-qubit gates are implemented using electrostatic gate electrodes by controlling an electron mediated nuclear exchange interaction[16].

This thesis will focus on quantum dots composed of a structure made of two semiconductors with different bandgaps (for example GaAs and the larger band gap semiconductor $(Ga_{1-x}Al_x)As$[18]) in a lateral quantum dot architecture. The larger band-gap material is doped and the carriers (electrons or holes) concentrate at the surface of the two semiconductors, typically 90 nm below the sample surface[17]. A thin sheet of highly mobile carriers called a two-dimensional electron gas (2DEG) or hole gas (2DHG), is formed at this interface. On the top of the surface, a pattern of metallic gates is fabricated using electron beam lithography. Conventionally, these devices have been made using GaAs based heterostructures[18], however, such devices are significantly limited by decoherence due to the hyperfine interaction caused by nuclear spins. Instead, there has been a growing effort by the community to develop lateral quantum dot devices based on isotropically purified Si or Ge heterostructures[17].

Commonly, in GaAs heterostructure systems, negative voltages are applied such that the electrons in the 2DEG are repelled locally just underneath the gates. These are referred to as depletion mode gates and, although they are popular in the community

for GaAs devices, such devices suffer from disorder-driven decoherence and tunability complexity when fabricated on silicon heterostructures[19]. Alternatively, accumulation mode devices use overlapping gates with positive potentials to confine electrons and have demonstrated superior stability and tunability in tunnel couplings in silicon devices[20,14]. The superior control achieved with accumulation mode devices is largely due to the scale of the spatial extent of the confinement potential. Accumulating the charge under a single gate also reduces the relative capacitive couplings of the charge to adjacent gates[20]. In either mode, the gate architecture and choice of gate voltages are imprinted onto the 2DEG. The gate architecture and electrostatic potentials are designed such that one or more small islands of electrons are formed, isolated by tunnel barriers from the rest of the 2DEG (quantum dots are typically $10 - 100$ nm in diameter)[18].

Hole spin qubits in Si and Ge quantum dots

In this thesis, I focus on the operation of hole spin qubits in Si and Ge based devices, and so introduce them in more detail here.

In the field of quantum computing, particularly in the development of hole spin qubits, several key properties have been identified as critical to their operation and coherence[17,21,22]. One of the fundamental characteristics of hole spins is that they inhabit valence-band states formed by atomic p-orbitals. This unique property significantly influences the operation of hole spin qubits. Hole spins exhibit a reduced susceptibility to interactions with nuclear spins compared to conduction-band electrons. This is primarily due to the vanishing Fermi contact interaction in holes, resulting in an anisotropic hyperfine interaction. This interaction, influenced by dipolar and orbital terms, is weaker by approximately an order of magnitude compared to that experienced

by conduction-band electrons. Another notable aspect is the absence of a valley degree of freedom in hole spins. In the context of Si quantum dots, the valley splitting of electrons is highly dependent on the atomistic features of the device. These features are often challenging to control uniformly across large arrays of quantum dots[23], presenting a significant advantage for hole spins.

The feature of hole spins with the most profound implications is their strong spin–orbit interaction. This interaction facilitates in situ electric control of spin couplings and allows for effective spin manipulation through local electric drives[24]. In contrast, the spin–orbit coupling in Si conduction electrons is relatively weak, necessitating auxiliary elements, such as micromagnets[25], for single-spin manipulation. However, the strong spin–orbit coupling in hole spins is also a potentially limiting property in the development of a quantum computing platform. While it enables enhanced control, it also increases the susceptibility of spin degrees of freedom to charge noise[26].

In experimental applications, hole spin qubits in planar Ge devices have shown significant potential[27]. High fidelity single- and two-qubit gates have been achieved in four-qubit processor based upon this device design[28]. Moreover, both Ge/Si core/shell nanowires and Ge hut wires exhibit strong spin–orbit interactions, a common feature that can be exploited for ultrafast single-spin manipulation[29,30]. The realisation of hole-spin qubits in CMOS devices is particularly promising[31,32]. These devices are not only advantageous due to compatibility with existing foundry fabrication processes but also due to tuneable sweet spots[33] and strong spin-photon coupling[34].

Electronic Transport Properties

The principle interactions which dictate the electronic transport properties of semiconductor quantum dots are, at low temperatures, Coulomb blockade[35] and the quantum effects which arise due to the three-dimensional confinement potential.

Classically an array of quantum dots can be described by a network whereby each dot is capacitively coupled to each gate. In this classical description, the dot is coupled via a tunnel barrier, analogous to a capacitor, to the source, drain, and other dots[35]. At low temperatures and low bias, the charging energy of the dot is larger than thermal fluctuations, and Coulomb blockade restricts conductance across the dot. Electronic transport is energetically possible only when the electrochemical potentials of successive states (for any N) fall within the bias window so that $\mu_S \geq \mu(N) \geq \mu_D$ (here μ_S and μ_D denote the source and drain potentials respectively)[18]. This leads to a step-wise increase in the current as the bias is swept because additional transport channels enter the bias window. In each step, an electron can tunnel from the source onto the dot and then to the drain. This effect, known as single electron tunnelling, can occur to an intermediate state where the charge on the dot differs by e from the most thermodynamically stable state. This co-tunnelling proceeds, via either an elastic or an inelastic mechanism, as an activationless higher-order contribution to the tunnelling conductance[36].

Quantum dots, as a result of quantum effects, have discrete energy spectra and are therefore often described as artificial atoms[18]. Recently, this model has been shown to accurately describe the electron filling of s-, p-, d, and f- electrons in a silicon quantum dot and led to the proposal of potential new single valence qubits (with 1, 3, 5, 15, and 25 electrons occupying the dot)[37]. The discrete energy levels lead to sharp resonances in

electronic transport through the device. The energy differences between these resonant transitions are significantly smaller than the quantum dot's charging energy, the energy required to add an additional electron[35].

Semiconductor devices are often characterised by analysing their 2-D current-voltage maps. Double quantum dot regimes are recognised by the honeycomb pattern in these maps. Here, the edges of the honeycomb are visible due to off-resonance current, where the electrochemical potential of one of the dots aligns with its neighbouring lead. The co-tunnelling across these edges takes place via an intermediate virtual state and can be suppressed by increasing the width of the tunnel barriers[35].

2.1.3 Bias triangle charge transport

In semiconductor double quantum dots, bias triangle transport features arise due to the alignment of energy levels within the quantum dots and the application of an external bias voltage. A double quantum dot system can be modelled as two coupled quantum dots with discrete energy levels, denoted as E_L and E_R for the left and right dots, respectively. When a bias voltage V_{bias} is applied across the double quantum dot, it creates a potential difference that drives the transport of electrons from the source to the drain through the quantum dots.

The current through the double quantum dot system is governed by the sequential tunnelling of electrons, which occurs when the chemical potentials of the source (μ_S) and drain (μ_D) electrodes align with the energy levels of the quantum dots (see **Figure** 2.1).

The conditions for current flow can be described by the following equations:

$$\mu_S \geq E_L + n\hbar\omega_L \tag{2.1}$$

$$E_R + m\hbar\omega_R \leq \mu_D \tag{2.2}$$

$$E_L \leq E_R + eV_{bias} \tag{2.3}$$

where n and m are the quantum numbers corresponding to the discrete energy levels in the left and right quantum dots, and $\hbar\omega_L$ and $\hbar\omega_R$ are the energy spacings between the levels.

The alignment of these energy levels with the bias voltage results in regions of current known as bias triangles, which can be observed in stability diagrams (charge stability plots) of the double quantum dot system. Bias triangles with resonant tunnelling lines are formed from the charge degeneracy points in the current map. Resonant tunnelling in the bias triangles only occurs on the side of the triangles which align with the original triple points. However, inelastic and co-tunnelling processes lead to a finite current within the triangle. Pauli Spin Blockade (PSB), a phenomenon often exploited for spin-to-charge conversion, is exhibited across bias triangles in the forward (N+1, N+1) to (N, N+2) tunnelling direction[38].

2.1.4 Two-state systems

The quantum bit, or qubit, is a two-level quantum system that can be described as a linear combination of the two basis states $|0\rangle$ and $|1\rangle$. This superposition of states is characterised by the coordinates on the surface of a Bloch sphere. The operation of qubits for quantum information processing requires control in the form of arbitrary

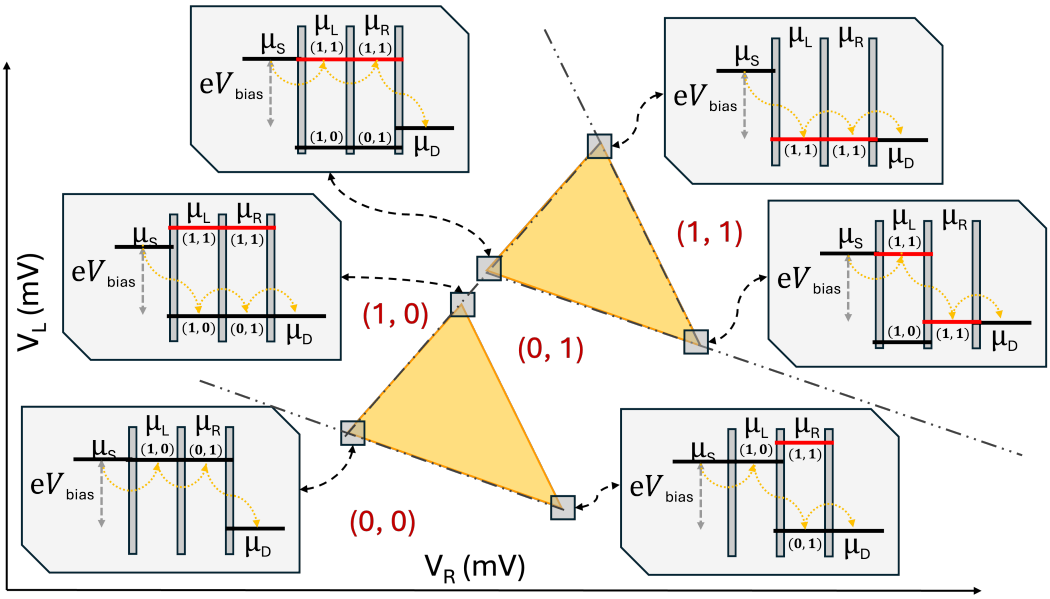


Figure 2.1: This figure illustrates the bias triangle transport features in a double quantum dot system under different gate voltage configurations applied to plunger gates above the left dot (V_L) and right dot (V_R) respectively. The central diagram shows the stability diagram with bias triangles corresponding to different charge states: (n, m), with n and m charges in the left and right dots respectively. The black dashed lines indicate the charge state transition points, and the yellow shaded region shows the region where charge transport is observable, termed bias-triangles. Inset diagrams detail the alignment of chemical potentials (μ_S, μ_L, μ_R, μ_D) of the source, left dot, right dot, and drain respectively and the application of the bias voltage (eV_{bias}) for electron transport. Each inset corresponds to a specific charge state within the bias triangle, showing the electron tunnelling pathways.

rotations on the x-, y-, and z-axes of the Bloch sphere. Numerous schemes have been developed by the semiconductor quantum dot community to define a qubit, some of which will be introduced below.

The spin qubit, first proposed by Loss and DiVincenzo[7], has basis states with opposing spin directions which are split by the Zeeman energy $E_Z = g\mu_B B_0$, where μ_B is the Bohr magneton and g is the gyromagnetic factor, under a static magnetic field B_0. Universal single spin qubit control has been demonstrated[39] using electron spin resonance (ESR) and achieving a gate fidelity of 99.6%. The gate operation time of ESR techniques is low because large oscillating magnetic fields are challenging to create locally and therefore electric dipole spin resonance (EDSR) manipulation using electrical fields is desirable[17]. For EDSR the electron/hole moves back and forth due to an oscillating electrical field. The electron/hole experiences a magnetic field gradient as it moves due to the interaction with an integrated micro-magnet, the hyperfine interaction[17], or the spin-orbit coupling of the material. Gate fidelities of 99.9% have been achieved using an integrated micro-magnet in a Si/SiGe quantum dot[40].

The singlet-triplet qubit has become one of the dominant qubit constructions within the literature due to the simple electrical control technique, high manipulation rate, and relatively high gate fidelities[41,42,43]. The singlet-triplet states of a pair of exchange coupled electrons/holes (in adjacent quantum dots) are $|S\rangle, |T_{0,\pm}\rangle$. An external magnetic field lifts the degeneracy of triplet states via the Zeeman interaction such that the triplet zero or the triplet minus and singlet states form the basis for qubit operation ($|0\rangle \equiv |S\rangle, |1\rangle \equiv |T_0\rangle$ or $|T_-\rangle$)[44]. Qubit state rotations are performed by detuning the exchange coupling, J, between the two dots for single qubit gates[45], and between pairs of dots to implement two-qubit gates[43].

Although control of singlet-triplet qubits can be achieved electrically, an inhomogeneous static magnetic field is still required for full single-qubit control and therefore universal quantum computation. For fully electrical universal control the exchange qubit was introduced. The exchange qubit states are comprised of three spins in a triplet quantum dot with coherent spin control in a subspace of two of the electrons[46]. Here, the need for an inhomogeneous static magnetic field is relaxed and all single-, and two-qubit operations can be implemented by modulating the exchange interactions between neighbouring dots. There are two varieties of exchange-only qubits, namely, exchange-only qubits[47], and resonant exchange qubits[48].

2.1.5 Loss-DiVincenzo qubit

We can describe the physics of a pair of coupled Loss-DiVincenzo qubits by constructing the Hamiltonian and using the Schrieffer-Wolf approximation (Appendix B) to simplify the dynamics. This model is central to the differentiable simulator of a two-hole spin qubit system used in Chapter 6. We can define the full Hamiltonian as

$$H = H_Z + H_t + H_\epsilon. \tag{2.4}$$

The Hamiltonian for electrons or holes tunnelling between sites 1 and 2, reads in operator notation:

$$\hat{H}_t = -a^\dagger_{1\sigma_1} t_{\sigma_1\sigma_2} a_{2\sigma_2} - a^\dagger_{2\sigma_2} t^*_{\sigma_1\sigma_2} a_{1\sigma_1}. \tag{2.5}$$

Here $\sigma_1, \sigma_2 = \uparrow, \downarrow$ are the pseudo-spin indices for sites (dots) 1 and 2[44], in the local basis that diagonalises the spin Zeeman Hamiltonian (computational basis), and $a^\dagger_{1\sigma_1}$ and $a_{2\sigma_2}$ are fermion creation and annihilation operators[49]. The tunnelling matrix takes the

general form[34]:

$$t = t_c \exp\left(i\eta \mathbf{n}\cdot\boldsymbol{\sigma}\right) = t_c(\cos\eta + i\mathbf{n}\cdot\boldsymbol{\sigma}\sin\eta), \tag{2.6}$$

with $\boldsymbol{\sigma} = (\sigma_x, \sigma_y, \sigma_z)$ the vector that contains the standard Pauli matrices. The elements of the matrix (2.6) are thus explicitly:

$$\begin{pmatrix} t_{\uparrow\uparrow} & t_{\uparrow\downarrow} \\ t_{\downarrow\uparrow} & t_{\downarrow\downarrow} \end{pmatrix} = t_c \begin{pmatrix} \cos\eta + in_z\sin\eta & (in_x + n_y)\sin\eta \\ (in_x - n_y)\sin\eta & \cos\eta - in_z\sin\eta \end{pmatrix}, \tag{2.7}$$

where $\mathbf{n} = (n_x, n_y, n_z)$ is a general unit vector, and the spin rotation angle $\eta = d/\lambda_{SO}$ is given by the ratio of the interdot distance and the spin-orbit length. Such spin-dependent tunnelling may arise from various causes, which include spin-orbit coupling and a mismatch between the local pseudo-spin bases[34,50].

In the presence of multiple fermions we use the occupation number representation of the quantum states[49], with basis $|\{n_{j\sigma_j}\}\rangle = |n_{1\uparrow}n_{1\downarrow}n_{2\uparrow}n_{2\downarrow}\rangle$. For two electrons or holes, we consider:

$$\begin{aligned} |1\rangle &= |1100\rangle \qquad (\uparrow\downarrow, 0), \\ |2\rangle &= |1010\rangle \qquad (\uparrow, \uparrow), \\ |3\rangle &= |1001\rangle \qquad (\uparrow, \downarrow), \\ |4\rangle &= |0110\rangle \qquad (\downarrow, \uparrow), \\ |5\rangle &= |0101\rangle \qquad (\downarrow, \downarrow). \end{aligned}$$

This includes the state of double occupancy on site 1 only because it is the most proximal one in energy for our range of dot detuning for qubit operation. In this basis we get the

tunnelling Hamiltonian matrix, in the explicit form

$$H_t = \begin{pmatrix} 0 & -t_{\downarrow\uparrow} & -t_{\downarrow\downarrow} & t_{\uparrow\uparrow} & t_{\uparrow\downarrow} \\ -t^*_{\downarrow\uparrow} & 0 & 0 & 0 & 0 \\ -t^*_{\downarrow\downarrow} & 0 & 0 & 0 & 0 \\ t^*_{\uparrow\uparrow} & 0 & 0 & 0 & 0 \\ t^*_{\uparrow\downarrow} & 0 & 0 & 0 & 0 \end{pmatrix}. \tag{2.8}$$

Taking into account that basis state $|1\rangle$ has extra energy $\Delta = U - V - \varepsilon \gg t_c$ compared with the states with single occupation of the two dots, with U the charging energy of the quantum dot, V the inter-dot Coulomb interaction energy, and ε the detuning energy of the double dot; we get at second order of perturbation theory (the Schrieffer-Wolff transform) the effective spin exchange Hamiltonian in the basis $\{|2\rangle, |3\rangle, |4\rangle, |5\rangle\}$:

$$H^{(2)}_{\text{eff}} = \frac{-1}{\Delta} \begin{pmatrix} |t_{\downarrow\uparrow}|^2 & t_{\downarrow\downarrow}t^*_{\downarrow\uparrow} & -t_{\uparrow\uparrow}t^*_{\downarrow\uparrow} & -t_{\uparrow\downarrow}t^*_{\downarrow\uparrow} \\ t^*_{\downarrow\downarrow}t_{\downarrow\uparrow} & |t_{\downarrow\downarrow}|^2 & -t_{\uparrow\uparrow}t^*_{\downarrow\downarrow} & -t_{\uparrow\downarrow}t^*_{\downarrow\downarrow} \\ -t^*_{\uparrow\uparrow}t_{\downarrow\uparrow} & -t^*_{\uparrow\uparrow}t_{\downarrow\downarrow} & |t_{\uparrow\uparrow}|^2 & t_{\uparrow\downarrow}t^*_{\uparrow\uparrow} \\ -t^*_{\uparrow\downarrow}t_{\downarrow\uparrow} & -t^*_{\uparrow\downarrow}t_{\downarrow\downarrow} & t^*_{\uparrow\downarrow}t_{\uparrow\uparrow} & |t_{\uparrow\downarrow}|^2 \end{pmatrix}. \tag{2.9}$$

The Zeeman term has the form;

$$H_Z = \frac{\mu_B \mathbf{B}}{2} \Sigma_{d=1,2} g_d \boldsymbol{\sigma} \tag{2.10}$$

given the g-factor g_d for the dots 1 and 2 respectively. For a field along the z-direction, this becomes

$$H_Z = \beta\Big(\Delta_g\big(\,|\downarrow,\uparrow\rangle\langle\downarrow,\uparrow| - |\uparrow,\downarrow\rangle\langle\uparrow,\downarrow|\,\big) + \Sigma_g\big(\,|\uparrow,\uparrow\rangle\langle\uparrow,\uparrow| - |\downarrow,\downarrow\rangle\langle\downarrow,\downarrow|\,\big)\Big), \tag{2.11}$$

where $\beta = \frac{\mu_B \mathbf{B}}{2}$, $\Delta_g = g_1 - g_2$, and $\Sigma_g = g_1 + g_2$.

The Hamiltonian term corresponding to the detuning is given by:

$$H_\epsilon = \epsilon \left| S_{(2,0)} \right\rangle \left\langle S_{(2,0)} \right|, \tag{2.12}$$

where the detuning $\epsilon = U - \epsilon'$ represents the additional energy contribution of the $(2,0)$ singlet state when compared to the $(1,1)$ singlet state, given the charging energy U and the one-particle picture energy difference between the dots.

2.2 Related work

2.2.1 Automated tuning of semiconductor qubits

As discussed (section 2.1.2), the ability to tune inter-dot tunnel couplings in laterally defined semiconductor quantum dots is a crucial advantage in the control and read-out of qubit states. However, the high dimensional parameter space (defined by the electrostatic potentials across the gates), which is an inevitable property of this tunability, presents the community with a number of complex engineering challenges. The parameter space requires thorough characterisation and exploration in order to locate regimes with the desired transport properties needed to define qubits. This parameter space, as in other types of qubit implementations, can consist of many dimensions that scale with the number of qubits. Charge traps and material impurities create unique disorder potentials which lead to variability between devices even those with identical architectures[51].

Due to the complexity of the tuning task, many experimenters habitually spend hours each day manually tuning devices using expert domain knowledge. This type of ap-

proach, while possible for devices comprising of a few quantum dots, quickly becomes intractable as the number of qubits in the array increases. Hence, if quantum dots are to be considered a feasible architecture for large-scale quantum information processing, the requirement to efficiently and automatically tune them must be achieved.

Automated tuning pipelines are generally divided into coarse-tuning and fine-tuning tasks. Here, coarse tuning is defined by the process of finding gate voltages for which single or double quantum dot (single electron regime) transport features are measured[52]. Fine-tuning algorithms, operate within the few-electron regime and aim to characterise tunnel couplings or load the device with a desired number of electrons. A core principle used in these approaches is the estimation of linear combinations of gates which allows for the control of each electrochemical potential and tunnel coupling individually. These linear combinations, termed virtual gates, are determined by the inverse of the cross capacitance matrix[53].

Van-Diepen et al. developed a procedure for fine-tuning the inter-dot tunnel coupling using virtual gates in combination with an automated feedback loop which converges to a pre-defined target value[54]. The cross-capacitance matrix was calculated using a difference metric between the processed data and a classical model of the avoided crossings. They developed and successfully tested this algorithm on a linear triple quantum dot device with a GaAs/AlGaAs heterostructure. Mills et al. proposed a similar algorithm but also introduced an image analysis toolbox developed in Python which further automated the procedure of fitting virtual gates[55]. This protocol measures cross-capacitances from pairwise stability diagrams, for nearest neighbour gates only, which enables control of individual electrochemical potentials or tunnel couplings.

These linear combinations of gates only compensate for crosstalk in small regions around

which they were calculated. The compensation of crosstalk to the tunnel barriers over larger regions of the parameter space is therefore a slow iterative process that requires many measurements. Some progress has been made to efficiently gain orthogonal control over all inter-dot tunnel couplings in a linear quadruple quantum dot. This method, which required relatively few measurements, was achieved by including the exponential dependence of tunnel couplings on gate voltages[56]. For virtual gates to be useful in completely automating the control of large arrays of quantum dots, developments must be twofold; the number and speed of the measurements required to calculate them must be reduced, and secondly, the range of validity of the mapping to virtual gates must be extended.

Automated Tuning Augmented by Machine Learning

Recently there has been a trend in the relevant literature towards automated tuning using machine learning frameworks. Machine learning has enjoyed great success in the automation of tasks that had previously proved difficult for machines. The marriage between automating the control of quantum experiments and machine learning, therefore, seems to be a promising development. The first attempts to control the tuning procedure follow a general framework with three principal components;

1. A machine learning tool for state classification using measured data.
2. A tailored optimisation function which couples a difference metric, between the current state and the desired state, to the experimental control parameters.
3. A numerical optimisation routine to navigate the multi-dimensional parameter space.

Kalantre et al. developed a method which follows this general framework[57] and their method was subsequently implemented experimentally by Zwolak et al.[58]. They trained a convolutional neural network (CNN) using simulated data to classify the transport features in an arbitrary sub-region, of gate voltage space, as short circuit, barrier, single dot, or double dot, which was in turn used to construct a score function. The tuning procedure was implemented as an optimisation problem of this score function, solved using the Nelder-Mead method.

Optimisation via purely numerical methods necessarily fails in these flat sub-regions of the current map. A new paradigm must be adopted if tuning protocols are to be robust to measurement noise and operate effectively over the whole parameter space. Prior knowledge of the optimisation landscape and regularisation against over-fitting could aid an algorithm in meeting these conditions. Well-developed machine learning tools such as generative or physically informed models and reinforcement learning will be crucial in future progress.

A paper, published by Teske et al.[59], applies a similar framework to the tuning of the inter-dot tunnel couplings. They use a quasi-Newton optimisation routine to attain user-defined tunnel coupling values. The method uses a Bayesian statistical approach, called a Kalman filter, to approximate the Jacobian and thus eliminate the need for costly finite difference measurements. The restricted optimisation gate voltage space in this fine-tuning setting is less devoid of features so this approach proved to be relatively successful. Contributions by Darulova et al.[60] and Durrer et al.[61] also use Machine Learning classification tools alongside heuristic formulas for navigation of the parameter space for coarse-tuning and fine-tuning, respectively.

Lennon et al. developed an algorithm which employs a more sophisticated algorithmic

workflow to robustly solve the coarse tuning problem in a high dimensional parameter space[62]. The most novel development of their method was the identification of a boundary hypersurface, in the gate voltage hyperspace, upon which interesting transport features are present. The principle of the hypersurface is very general as it is agnostic to device architecture, material, and dimensionality, such that it represents a central point for future algorithmic development. The bottleneck of this protocol, as with all others discussed, remains the time-intensive 2-D current scans. In order to tune larger arrays of quantum devices, this bottleneck must be addressed by: (i) using physics-informed and generative ML tools to limit the number and resolution of 2-D measurements needed[63], and (ii) developing faster measurement techniques such as dispersive readout[64,65].

2.2.2 Optimal quantum control

Quantum Error Correction (QEC) codes are the cornerstone of fault-tolerant quantum computation[66]. They pave the way for the creation of logical qubits that exhibit error rates low enough to support the execution of complex quantum algorithms. By incorporating additional physical qubits into these logical constructs, we can bolster their resilience to errors. Nevertheless, this increase in physical qubits brings about a rise in the overall error rate. It is imperative, therefore, to maintain the physical qubit error rate or gate infidelity beneath a certain threshold to enhance the fault tolerance of the logical qubits as the code size expands[67]. Thus, central to the advancement of Quantum Computing is the execution of high-fidelity quantum gates. The fidelity is fundamentally limited at the hardware level by incoherent noise processes[68,69]. Here, interactions between the computational system and its environment can destroy the coherent quantum information. However, even in the absence of decoherence, gate operations are plagued

by errors. Such errors, termed coherent errors, emerge from various sources, including state leakage[70], miscalibration, cross-resonance[71,72], non-adiabatic transitions[73], the impact of filtered signals, and of non-linear transmission functions[74,75].

In the quantum computing community, the quest for optimal control of qubits and the implementation of high-fidelity gates is an ongoing engineering challenge. The focus of this section is to examine relevant literature in areas such as gradient-based optimisation frameworks and control mechanisms for quantum systems. These works illuminate several avenues for advancements, while also highlighting limitations that our proposed approach seeks to address.

Approaches to optimal quantum control include deep reinforcement learning[76,77], analytical approaches to spectral engineering[78], and adiabatic schemes[79,80]. However, a fundamental technique employed for the optimal control of quantum systems through gradient descent is the GRAPE formalism, introduced by Khaneja et al.[81] in 2005. This method utilises analytical formulas to calculate the gradients of system evolution in the Schrödinger picture. The introduction of GRAPE marked a considerable advancement in computational efficiency compared to earlier methods that computed gradients using finite differences[82]. Nonetheless, the integration of GRAPE with realistic constraints of physical systems continues to be a substantial challenge for the community[83,84].

Modern advances in Machine Learning have motivated the creation of programming paradigms with automatic differentiation enabled for all operations. The quantum control community have used this to their advantage and there have been several efforts to realise quantum control methods optimised using gradient descent through models which better represent realistic experimental conditions. A pioneering contribution on this front was COCOA (Complex and Optional Constraints Optimisation

with Auto-differentiation), which was demonstrated using superconducting qubit simulations[83]. COCOA outperformed the original GRAPE proposal by learning the appropriate Fourier components of pulse envelopes which enabled them to impose realistic constraints on the pulse amplitude bandwidth.

There are other motivating reasons for moving to a simulation and control framework which uses automatic differentiation, such as Hamiltonian parameter estimation[85]. Ni et al. propose using such a framework to optimise the fidelity of quantum operations by propagating gradients through parameters related to the superconducting circuit device design, as well as the pulse shape[86]. Abdelhafez et al. opt to use automatic differentiation to optimise gates using realistic pulses in a multi-objective optimisation setting for protected superconducting qubits[87]. A collaboration between IBM and academia yielded a comprehensive method, which they dubbed C3 (Control, Calibration, and Characterisation) which is reliant upon using stochastic gradient descent for model fitting in addition to designing optimal control schemes[88]. Jirari et. al. employ automatic differentiation to investigate the regimes in which a system described by a driven spin-boson model is quasi-controllable[89].

Spin qubits are prone to coherent errors originating from the microwave drive, such as pulse distortions and cross-resonance[74]. While the spin qubit community has yet to rigorously examine the design of robust optimal quantum control schemes using automatic differentiation, efforts have been made using alternative methodologies. For instance, fast quasi-adiabatic[90], and adiabatic protocols[73] have been proposed to optimise gate operations for spin qubits. Other methods achieve high fidelity pulses with realistic constraints, using numerical optimisers[91,92,93]. Optimal control methods, derived using GRAPE[69], and geometrical considerations[94] have been tested experimentally, however

so far they have failed to surpass the error correction threshold.

Chapter 3

Methods

"*The aim of life is to live, and to live means to be aware, joyously, drunkenly, serenely, divinely aware*".

Henry Miller

This chapter serves as a guide to the methodologies and techniques employed in the research undertaken during the course of my PhD. It is designed to offer both a theoretical and practical understanding of the tools and processes used in the tuning and analysis of spin qubits. This section delineates methods developed or used during this work, highlighting a few examples which are mostly absent from the literature and cover the scope of the thesis, including; experimental methods such as high-bandwidth "video-mode" spin state manipulation and readout; Bayesian optimal experimental design for characterising quantum systems; and a review of methods for approximating the Haar-measure in order to calculate the fidelity of two-qubit gates.

3.1 Charge-state measurements

Accurate measurement of charge states serves as a cornerstone for the initialisation, control, and readout of qubit systems. The need for high-fidelity charge state measurements is underscored by its direct impact on various quantum operations, including state preparation, readout, and consequently quantum gate fidelity. More precisely, inaccuracies in charge state determination can lead to erroneous initialisation, misclassified quantum state readout, and miscalibration of quantum gate operations.

Navigating to the qubit regime within a quantum dot system necessitates information about both the charge occupation and the tunnel rates. A map of the charge stability landscape is, therefore, instrumental for configuring the quantum dots into the specific physical regimes required for qubit operation. Charge transitions with particular parities between the relative charge occupations are required to enforce a spin-to-charge conversion mechanism. This mechanism is essential for reading out the qubit state, using Pauli spin blockade to translate the spin information into a charge configuration

that can be more readily detected[95,96]. Given these considerations, there is a pressing need for experimental methods that provide high-fidelity, high-speed, and minimally perturbative charge state measurements.

3.1.1 DC MEASUREMENTS

One of the primary techniques for probing the electronic properties of Quantum Dots (QDs) is through transport measurements, specifically when a QD is coupled to source and drain electrodes. These measurements leverage a quantum phenomenon termed the Coulomb blockade. In the context of this blockade, the QD typically impedes the flow of current. The obstruction arises due to the misalignment potential of the quantum dot (QD) energy level with the energy window, which is defined by the source-drain voltage, V_{SD}. This voltage measures the energy difference between the Fermi levels of the source, μ_s, and drain, μ_d.

To surmount this blockade and allow electron transport, the electrochemical potential of a specific QD level must be positioned between μ_s and μ_d. Such adjustments can be realised either by modulating the gate voltage V_G or by expanding the source-drain bias window V_{SD}. Systematically varying V_G with respect to V_{SD}, and measuring the source to drain current, enables the elucidation of the QD stability landscape. Sweeping both V_G and V_{SD} reveals a series of diamond-shaped regions, known as the Coulomb diamonds. Within these regions the electron count on the QD remains invariant, such that the discontinuities show changes in the charge occupations (i.e. transport). Similarly, charge occupations of quantum dots can be measured with respect to multiple gate voltages by applying a bias voltage (V_{SD}), sweeping the gates, and recording the DC signal using a current sensor. Such signals can be measured with high sensitivities using carefully

designed electronic systems with high-bandwidth current amplification.

However, DC measurements are significantly limited by low-frequency $1/f$ noise, reducing the signal-to-noise ratio. In addition to this, aggressive low-pass filtering is necessary to achieve a low effective electron temperature, reducing the measurement bandwidth. These effects mean that DC measurements have a low signal-to-noise ratio in regimes with fast transport or very low charge occupations, and therefore require long integration times to recover a signal. DC current measurements also require quantum dots to be coupled to ohmic contacts, which restricts the possible architectures of scalable devices.

3.1.2 RF-MEASUREMENTS

Radio frequency (rf) reflectometry, specifically in the dispersive regime, has emerged as the principle technique for the measurement of the charge and spin states of semiconductor quantum dot devices[97]. The resonant frequency and phase of an LC tank circuit are affected by a change in the impedance of a coupled device, such as a quantum dot. The total impedance Z_{DUT} of the Device Under Test (DUT) is given by:

$$Z_{\mathrm{DUT}} = R_{\mathrm{sys}} + \frac{1}{j\omega C_T} \tag{3.1}$$

where; Z_{DUT} is the total impedance of the DUT; R_{sys} is the Sisyphus resistance, representing energy dissipation due to dynamic processes in the quantum system; C_T is the sum of the quantum and tunnelling capacitances; and ω is the angular frequency[98]. The Sisyphus resistance and tunneling capacitance arise from irreversible charge relaxation processes that occur when the system is driven at frequencies at or above its relaxation rate. Sisyphus resistance, linked to quantum friction, involves phonon pumping mech-

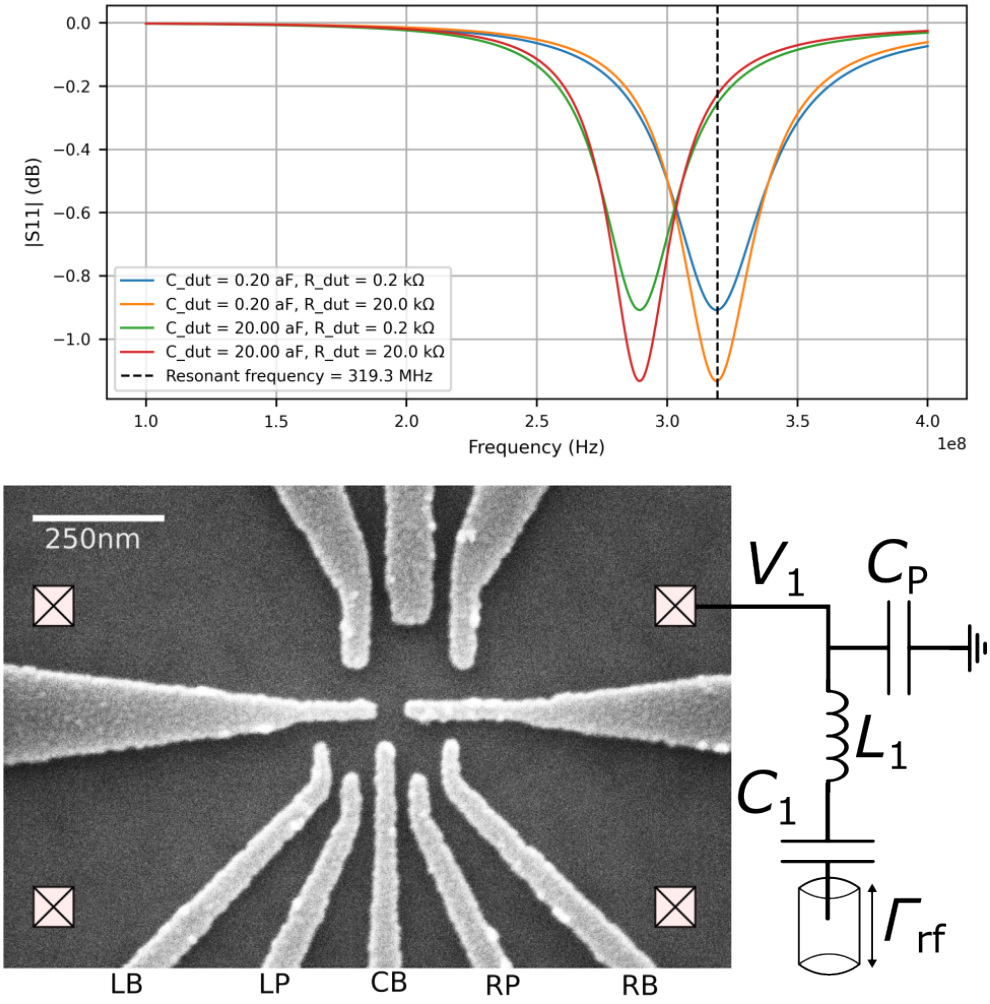


Figure 3.1: Upper: Magnitude of the reflection coefficient $|S_{11}|$ in dB as a function of frequency for an RF tank circuit with different DUT (Device Under Test) configurations. The DUT is modelled as having both a capacitive (C_{dut}) and resistive (R_{dut}) component, with values indicated in the legend. The resonant frequency of the tank circuit without the DUT is marked with a dashed vertical line at approximately 319.3 MHz. Here the circuit parameters are; $L_1 = 2.7$ nH; $C_1 = 92$ fF; parasitic capacitance C_P; and the characteristic impedance $Z0 = 766\ \Omega$, and resistance $R = 50\ \Omega$. The values used in this simulation are illustrative and intend to show the qualitative affects of shifts in the resistance and capacitance of the DUT on the amplitude of the reflected signal. **Lower:** Scanning electron microscope (SEM) image of a nanostructure, showing the layout of the electrodes used in the experiment. The circuit diagram on the right illustrates the RF tank circuit with the inductor (L_1), capacitor (C_1), and the RF source (Γ_{rf}).

anisms resulting in net power dissipation. In contrast, quantum capacitance emerges from reversible charge polarization, influenced by the discrete density of states and the non-linear characteristics of the energy levels. By precisely matching the impedance of the device to a high-frequency transmission line, it becomes possible to conduct measurements with a bandwidth reaching up to 200 MHz[99]. This allows for an operational frequency extending into the GHz range, thereby effectively surpassing the limitations imposed by 1/f noise. A change in the amplitude or phase of the reflected signal in parallel with the rf-tank circuit can be measured via demodulation with a reference signal (see **Figure** 3.1). As a result, rf-reflectometry enables faster, more sensitive measurements than DC current sensing.

3.1.3 Charge sensing

One advanced method to realise high sensitivity measurements of QD charge states is by employing a coupled charge sensing QD. A coupled charge sensing QD refers to an auxiliary QD capacitively coupled to the primary QD of interest. The purpose of this secondary dot is to detect variations in the electrostatic environment, particularly originating from the charge dynamics in the primary QD. The tunnelling of an electron into or out of the primary QD induces a change in the local electrostatic potential. Owing to the capacitive coupling between the two QDs, this modification also perturbs the energy states of the sensing dot. Measuring the change in impedance of the sensing QD via rf-reflectometry thus provides indirect knowledge about the charge dynamics of the main QD. The methodology offers extremely sensitive measurement capabilities, often with an accuracy at the single-electron level[97].

3.1.4 Video-mode measurements

An rf-tank circuit can be bonded to the ohmic contact of a charge sensing quantum dot, enabling the measurement of changes in the charge occupation of a capacitively coupled double quantum dot system. The high sensitivity and bandwidth of the dispersive measurement enable the measurement of the charge state, via resistive and capacitive changes, with an integration time on the order of a microsecond. This allows the construction of a measurement scheme whereby gate voltages are swept rapidly using AC signals. The increased speed of acquiring these stability diagrams opens up new avenues for efficient automated tuning and characterisation methods[64,65].

Tuning the rf-frequency and power

In tuning the frequency of an RF tank circuit, we employ a method which analyses the phase and amplitude differences in a dispersive signal both within and outside the Coulomb blockade regime. By sweeping the frequency of the rf-drive and recording the resultant signal differences, the frequency yielding the maximum absolute deviation is chosen as the optimal operating point. Similarly, the power is increased until the the maximum absolute deviation of this measurement converges or some threshold is met.

Tuning the charge sensor

We can investigate the charge occupation in a double quantum-dot device through dispersive readout, utilising a capacitively coupled charge sensor[65]. In this operational regime, a variation in the charge occupancy of the double-dot system induces a shift in the resonance frequency of the tank circuit bonded to the ohmics of the charge sensor. Conventional methodologies often employ a two-dimensional raster scan of the sensor

barrier gate voltages to discover the region of maximum gradient of the measured signal with respect to these voltages. This acts as a proxy which typically assures maximum sensitivity to capacitive shifts in the dispersive response of the charge sensor due to changes in the charge occupation of the device being measured.

We introduce an approach involving the measurement of the dispersive signal with respect to the device gate voltages by applying AC signals to these gates[65]. As an alternative to gradient-based methods described above, we calculate the standard deviation of these AC measurements as our objective function (**Figure** 3.2). This objective serves as a direct metric for our signal-to-noise ratio. We can map out this metric and find an optimum by changing the DC voltages across the sensor gates. This method does not require us to make dense raster scans in all gates to approximate gradients via finite differences, and therefore the optimisation of this objective function is not limited to the conventional pair of gates used in raster scans; it extends to other gates that influence the measurement sensitivity.

Spin selective measurements

In the experimental setup used to measure the Si/SiGe double quantum dot device reported in Chapter 5, we employed a two-dimensional fast spiral scanning technique to investigate the quantum states of our semiconductor device. This method involves using an arbitrary waveform generator (AWG) to rapidly vary the gate voltages applied to the device and measuring the corresponding change in the impedance of the sensor dot via rf-reflectometry (**Figure** 3.3). A rectangular window of the gate voltage space of the device was mapped out using a spiral scan pattern applied to the barrier gates of the double quantum dot using the AWG. For each discrete point (pixel) on this spiral, the

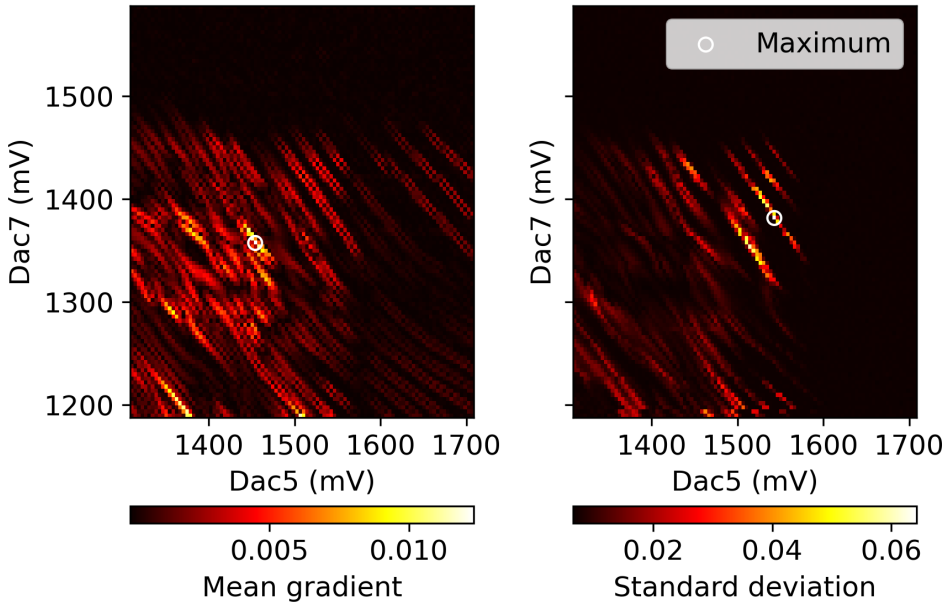


Figure 3.2: Optimisation of the sensitivity of the rf-reflectometry measurement of a charge sensor. Here, we measure the dispersive signal of an rf-tank circuit (**Figure** 3.1). This tank circuit is bonded to the ohmics of a charge sensor which is capacitively coupled to a double quantum dot in a Si/SiGe device measured at $\approx 50\,\mathrm{mK}$ in a dilution refrigerator. The measurements are reported by sweeping the barrier gate voltages (Dac5 and Dac7) of the sensor dot. **Left**: The conventional approach, where the gradient of a dispersive signal, with respect to the barrier gate voltages, is computed using finite differences during a sweep of the sensor barrier gates; the point of the maximum gradient, indicated by the white circle, denotes the location of optimal sensitivity. **Right**: The new method, involves the rapid acquisition of a stability diagram in the device gate voltage space using AC high-frequency signals from an Arbitrary Wave Generator (AWG). Here, the standard deviation of the dispersive signal is plotted while sweeping the sensor gates. The white circle represents the optimal voltage configuration, corresponding to the highest standard deviation, signifying the point of maximum sensitivity.

applied gate voltages correspond to specific coordinates in the two-dimensional voltage space defined by Channel 1 and Channel 2.

At each pixel of the spiral scan, we executed a pulse sequence designed to manipulate the charge and spin states of the holes. Assuming the pixel coordinate is such that the pulse begins in the (2, 0) charge state next to the transition to the (1, 1) charge state, the sequence begins by pulsing deep into the (2, 0) charge regime, which is expected to reinitialise the qubit and load a spin singlet state due to the Pauli exclusion principle. Following the initial loading, the pulse sequence rapidly jumps across the transition into the (1, 1) charge state, where the singlet and triplet spin states mix. The sequence is then reversed back into the configuration of the initial point, where the measurement is executed. If the system transitioned into a triplet spin state during this process, Pauli spin blockade would prevent further tunnelling, as the triplet state is forbidden in the (2,0) charge state due to the Pauli exclusion principle. This spin-selective tunnelling is measured by triggering the acquisition card, using the AWG, to demodulate and integrate the dispersive signal before the AWG moves the gate voltages to the next pixel in the measurement sequence.

This combination of fast spiral voltage scanning, a spin-manipulation and readout pulse sequence, and high-bandwidth measurements provides a powerful method for mapping the quantum state transitions, particularly for identifying the mixing of singlet and triplet spin states (**Figure** 3.3). By changing the DC voltages applied to the device gates continuously running this pulse sequence video-mode measurement we can investigate the charge stability diagram and rapidly locate regions within it where singlet-triplet mixing can be induced and read out. This provides a useful spin-qubit tuning mechanism which can be employed manually or automatically.

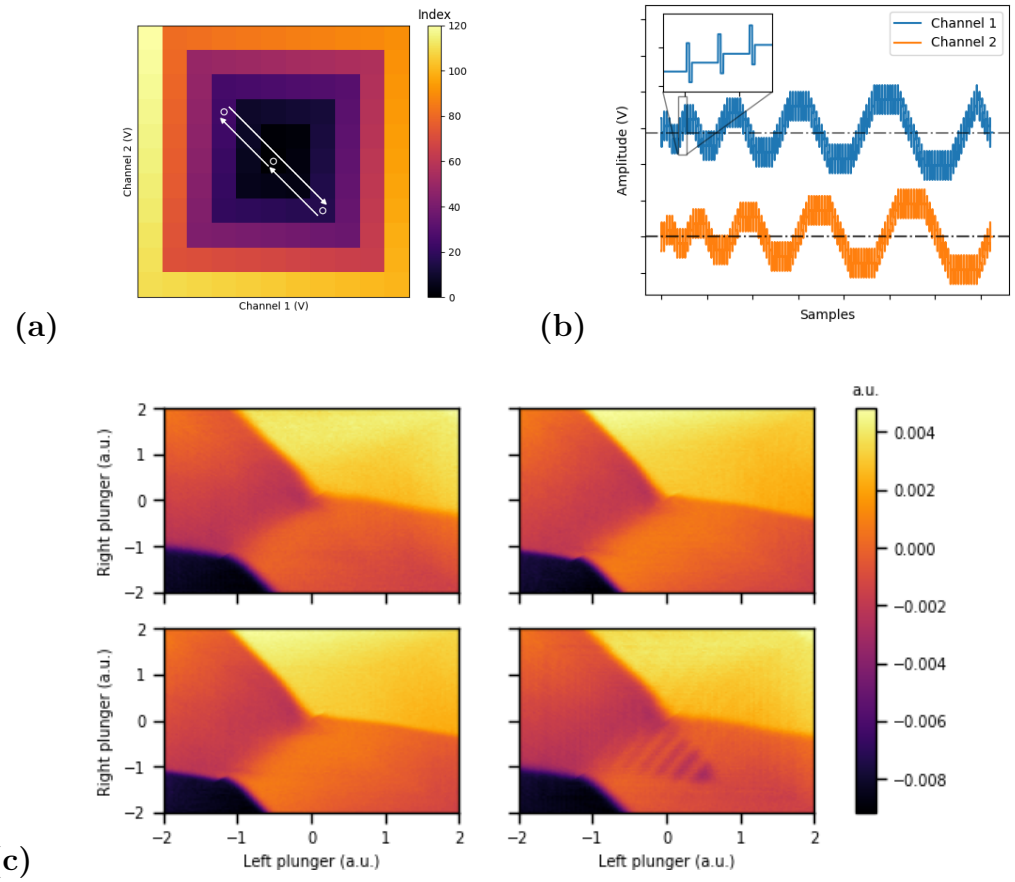


Figure 3.3: Fast 2D Scan and Spin-State Detection. (a): 2D fast spiral measurement of a 10x10 scan in device gate voltage space (Channel 1 vs. Channel 2), with pulse sequence (initialisation, operation, measurement) indicated by a white line for each pixel. The colour bar shows the pixel measurement index. **(b)**: Time-resolved AWG pulse signals for the full window measurement with pulsed spin manipulations applied at each pixel. **(c)** Reflectometry measurements of a charge stability diagram, using the tank circuit described in **Figure** 3.1, of Si/SiGe double dot with charge sensor at $\approx 50\,\mathrm{mK}$, showing interdot transitions using this "video-mode" technique. Conditions: Top left, $50\,\mathrm{mT}$ magnetic field without spin-manipulation pulse; Top right, spin-manipulation pulse without magnetic field; Bottom left, no magnetic field or pulsing; Bottom right, both pulsing and magnetic field. Observation: Spin manipulation and readout indicated by readout triangle, with Zeeman interaction causing singlet-triplet splitting under magnetic field and Pauli spin blockade visible with pulsing.

3.2 Bayesian methods for automating and optimising quantum experiments

The complexity and volume of data generated in contemporary quantum computing experiments often surpass the capacity of traditional analysis methods, necessitating a shift towards automation. Automated systems bring the advantage of handling large datasets efficiently, reducing human error, and providing consistent and repeatable results[100]. Furthermore, the design of experiments, which is foundational to the scientific method, benefits significantly from automation. It enables researchers to optimise experimental parameters swiftly and accurately, ensuring that the experiments are not only resource-efficient but also maximise the information yielded by those resources. This transition to automated systems is not just a matter of convenience but a necessary step towards more sophisticated, reliable, and high-throughput research methodologies[101].

3.2.1 Bayesian inference measurement analysis

Using Bayes' theorem, we calculate properties of posterior distributions given experimental data and a physics model. Here, the posterior probability distributions $p(\theta|d)$ over the latent variables in our physics model describe our belief over the parameters θ conditioned upon the data d. Bayes' theorem, expressed as

$$p(\theta|d) = \frac{p(d|\theta)p(\theta)}{p(d)}, \tag{3.2}$$

uses our prior knowledge of the distribution $p(\theta)$, in addition to a likelihood distribution $p(d|\theta)$, and an evidence term $p(d)$. We define the likelihood distribution by our belief that data d could be generated by the parameters θ. The evidence term $p(d)$ is a

normalising factor which can be included by ensuring that the integral of the posterior distribution's probability density equates to one.

Automatic Differentiation Variational Inference

Automatic differentiation variational inference (ADVI)[102] is an efficient inference method which is highly optimised by algorithms based upon stochastic gradient descent. Variational methods map the inference problem to an optimisation problem by use of a proposal family of parameterised distributions Q to approximate the true posterior P, both defined over the same probability space χ. In order to calculate the optimal approximation to the true posterior distribution, we use a metric to quantify the dissimilarity between this and our proposal. This metric, the Kullback-Leibler divergence D_{KL}[103], or relative entropy,

$$D_{KL}(Q||P) = \sum_{\theta \in \chi} Q(\theta) \log\left(\frac{Q(\theta)}{P(\theta|d)}\right) \tag{3.3}$$

can be interpreted as the amount of information lost when Q is used to approximate P. Calculation of the KL divergence requires exact knowledge of the true posterior distribution. However, as this is unknown, we instead use the conditional probability relation $P(\theta|d) = P(\theta, d)P(d)$ to define the variational free energy, $\mathcal{L}(Q)$, which is also known as the Evidence Lower Bound (ELBO)[104].

$$\mathcal{L}(Q) = \log(P(d)) - D_{KL}(Q||P) \tag{3.4}$$

$$= \sum_{\theta} Q(\theta) \left[\log(Q(\theta) - \log(P(\theta, d)))\right] \tag{3.5}$$

The log-evidence $\log(P(d))$ is constant with respect to θ and consequently, the minimum in the KL divergence can be calculated by maximising the ELBO. Variational

methods, however, require the development of custom optimisation algorithms for each specific problem. Automatic differentiation variational inference solves this bottleneck by automatically converting probabilistic programmatic models into custom variational algorithms. In this thesis, Bayesian inference and specifically ADVI was used to estimate the physical parameters, and calculate optimal experimental designs, of a Si finFET device by performing Rabi chevron measurements.

3.2.2 Bayesian optimal experimental design

Optimal experimental designs define a set of experimental conditions which are optimal with respect to a statistical criterion. Optimal designs achieve the same precision of parameter estimation, with a reduced experimental cost, as compared to non-optimal designs.

An optimal design d maximises the expected information gain (EIG) on a set of target parameters θ

$$EIG(d) = \mathbf{E}_{p(y|d)}\Big[H\big[p(\theta)\big] - H\big[p(\theta|y,d)\big]\Big] \tag{3.6}$$

where $H[\cdot]$ is the entropy, $p(\theta)$ is the prior, and the posterior is proportional to the product of the likelihood and the prior, $p(\theta|y,d) \propto p(y|\theta,d)p(\theta)$, according to Bayes' theorem. In this form, the optimal design represents the design that yields the highest reduction in posterior entropy over the target latent variables, as estimated using the expectation over possible future observations. Provided the predictive model is correct, this constructs an experimental design that is optimal from an information-theoretic perspective. We use Bayesian optimal experimental design to efficiently infer the Rabi and Larmor frequencies for a spin qubit in a Si finFET device, as reported in chapter 5, by sequentially measuring pixels of a Rabi Chevron (**Figure** 3.4).

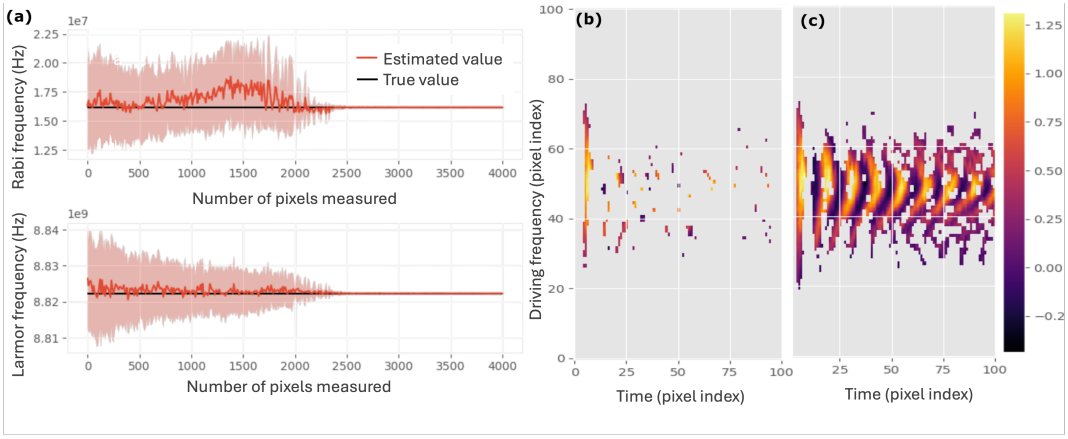


Figure 3.4: Convergence of parameter estimation during sequential measurements of a Rabi chevron. These experiments record the DC current trace, reported in rescaled units proportional to the ground state probability, performed by changing the driving frequency, and evolution time during a pulse sequence applied to a Loss DiVincenzo qubit in a Si finFET device[105] measured at $\approx 50\,\mathrm{mK}$ **(a)** The convergence of the inferred Rabi and Larmor frequencies (Hz), using Bayesian optimal experimental design to sequentially select pixels of a Rabi chevron, is illustrated by the mean and standard deviation across the number of measured pixels. **(b, c)** Measured pixels of a Rabi Chevron at two stages of the sequential measurement process demonstrating the data acquisition strategy's effectiveness in measuring maximally informative pixels, which are clustered around the resonant frequency. The x-axis indices represent the evolution time, while the y-axis indices represent the driving frequency.

Chapter 4

Deep Reinforcement Learning for the Efficient Measurement of Quantum Devices

"*Paths are made by walking*".

Franz Kafka

Automated discovery of spin qubit operating regimes.

Problem statement: The challenge is to efficiently and autonomously identify bias-triangle transport features in GaAs double quantum dot devices within the parameter space defined by the gate voltages of the device. This task is complicated by noisy transport measurements and barren plateau features of the parameter space.

Solution: This chapter demonstrates an approach for the efficient measurement of quantum devices based on deep reinforcement learning and convolutional neural networks[63]. Convolutional neural networks, a computer vision method, are used to autonomously identify the bias-triangle features, and the deep reinforcement learning agent is trained to make decisions to navigate through the barren device parameter space towards bias-triangles.

4.1 Introduction

Reinforcement learning (RL) is a neurobiologically-inspired machine learning paradigm where an RL agent will learn policies to successfully navigate or influence the environment. Neural network-based deep reinforcement learning (DRL) algorithms have proven to be very successful by surpassing human experts in domains such as the popular Atari 2600 games[106], chess[107], and Go[108]. RL algorithms are expected to advance the control of quantum devices[109,110,111,77,112,113,114,115,116,117,118,119,120,121,122,123,124,125], because the models can be robust against noise and stochastic elements present in many physical systems and they can be trained without labelled data. However, the potential of deep reinforcement learning for the efficient measurement of semiconductor quantum dot devices is still largely unexplored.

Singlet-triplet qubits encoded in double quantum dots[7] have demonstrably long coherence times[41,27], as well as high one-[42] and two-qubit[126,127,128] gate fidelities. Promising qubit performance was also demonstrated in single-spin qubits[40,39,129], and exchange-only qubits[47,130]. However, quantum dot devices are subject to variability, and many measurements are required to characterise each device and find the conditions for qubit operation. Machine learning has been used to automate the tuning of devices from scratch, known as super coarse tuning[52,131,62], the identification of single or double quantum dot regimes, known as coarse tuning[132,133], and the tuning of the inter-dot tunnel couplings and other device parameters, referred to as fine-tuning[134,59,135].

The efficient measurement and characterisation of quantum devices has been less explored so far. Previously, an efficient measurement algorithm for quantum dot devices combining a deep-generative model and an information-theoretic approach was developed[136]. Other approaches have developed classification tools which are used in conjunction with numerical optimisation routines to navigate quantum dot current maps[132,57,59]. These methods, however, fail when there are large areas in parameter space that do not exhibit transport features. To perform efficient measurements in these areas, which are often key for qubit operation, requires prior knowledge of the measurement landscape and a procedure to avoid over-fitting, i.e. a regularisation method.

In this chapter, we propose to use DRL for the efficient measurement of a double quantum dot device. Our algorithm is capable of finding specific transport features, in particular bias triangles, surrounded by featureless areas in a current map. The state-of-the-art DRL decision agent is embedded within an efficient algorithmic workflow, resulting in a significant reduction of the measurement time in comparison to existing methods. A convolutional neural network (CNN), a popular image classification

tool[137], is used to identify the bias triangles. This optimal decision process allows for the identification of promising areas of the parameter space without the need for human intervention. Fully automated approaches, such as the measurement algorithm presented here, could help to realise the full potential of spin qubits by addressing key difficulties in their scalability.

We focus on quantum dot devices that are electrostatically defined by Ti/Au gate electrodes fabricated on a GaAs/AlGaAs heterostructure (4.1**a**)[138,139]. All the experiments were performed using GaAs double quantum dot devices at dilution refrigerator temperatures of approximately 30 mK. The two-dimensional electron gas created at the interface of the two semiconductor materials is depleted by applying negative voltages to the gate electrodes. The confinement potential defines a double quantum dot which is controlled by these gate voltages and coupled to the electron reservoirs (the source and drain contacts). Depending on the combination of gate voltages, the double quantum dot can be in the 'open', the 'pinch-off' or the 'single-electron transport' regime. In the 'open' regime, an unimpeded current flows through the device. Conversely, when the current is completely blocked, the device is said to be in the 'pinch-off' regime. In the 'single-electron transport' regime, the current is maximal when the electrochemical potentials of each quantum dot are within the bias voltage window V_{bias} between source and drain contacts.

The algorithm interacts with a quantum dot environment within which our DRL decision agent operates to efficiently find the target transport features. The environment consists of states, defined by sets of measurements in gate voltage space, and a set of actions and rewards to navigate that space. This quantum dot environment has been developed based upon the OpenAI Gym interface[140]. Manual identification and characterisation

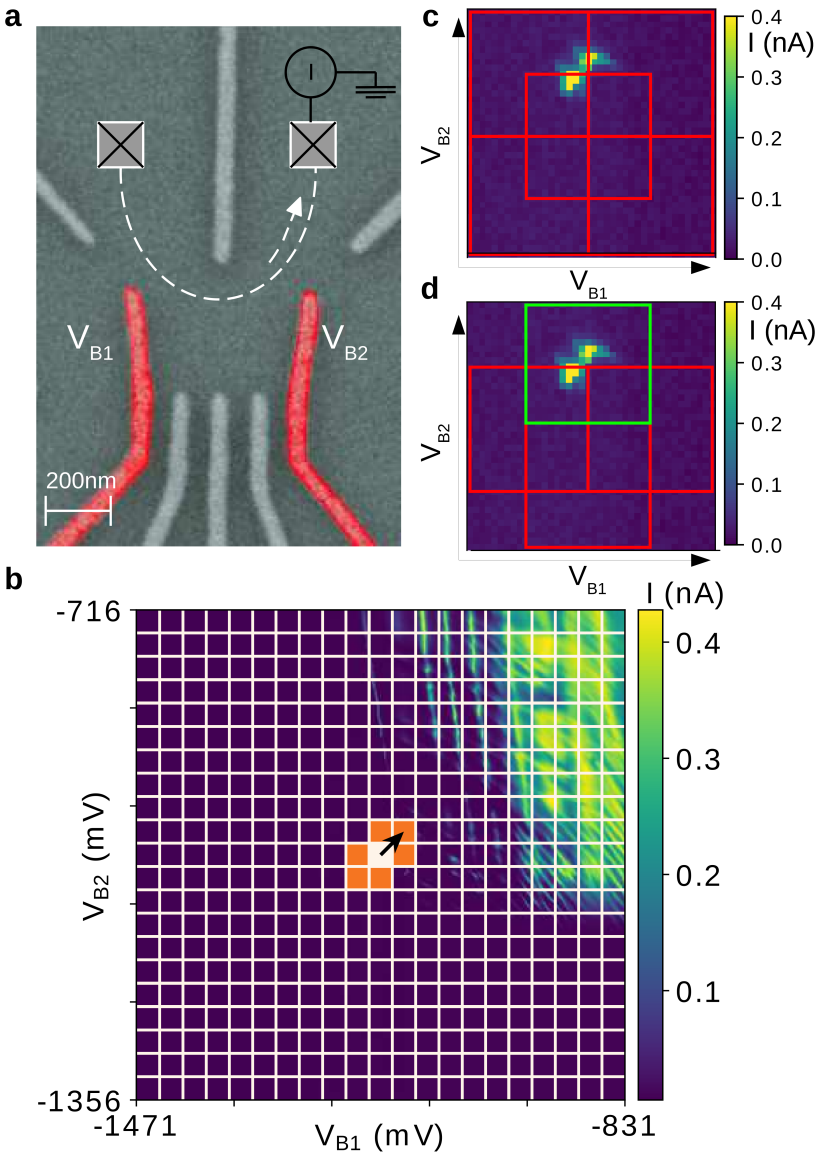


Figure 4.1: Overview of device architecture and quantum dot environment. **(a)** False-colour SEM image of a GaAs double quantum dot device. Barrier gates, labelled V_{B1} and V_{B2}, are highlighted in red. The arrow represents the flow of current through the device between source and drain contacts. **(b)** A current map. The white grid represents the blocks available for investigation by the DRL agent. The DRL agent is initiated in a random block (state) indicated by a filled white square. The filled orange blocks show the available action space for the DRL agent and the arrow shows a possible policy decision. **(c)** and **(d)** The nine sub-blocks defined within each block, a $32\,\mathrm{mV} \times 32\,\mathrm{mV}$ window in gate voltage space, to calculate a statistical state vector. These sub-blocks are equal in gate voltage size, five of them are shown in (**c**) and four in (**d**). The green sub-block in (**d**) contains bias triangles.

of transport features requires a high-resolution measurement of a current map defined by, for example, barrier gate voltages V_{B1} and V_{B2} while keeping other gate voltages fixed, an example of which is shown in **Figure** 4.1**b**. A super coarse tuning algorithm allows us to choose a set of promising gate voltages and focus on exploring the current map as a function of two gates, for example, the two barrier gates[62]. This is the gate voltage space our DRL agent will navigate.

The DRL algorithm takes the gate voltage coordinates found by the previous super coarse tuning algorithm[62], and divides the gate voltage space corresponding to the unmeasured current map into blocks. The size of the blocks is chosen such that they can fully contain bias triangles (blocks are shown as a white grid in **Figure** 4.1**b**). Devices with similar gate architectures often show bias triangles of similar sizes for a given V_{bias}. The DRL agent is initiated in a random block. The agent acquires a reduced number of current measurements from this block and makes a decision on whether a high-resolution measurement is required and on which block to explore next if bias triangles are not observed. The agent has a possible action space represented by a vector of length six; this means the agent can decide to acquire measurements in any of the four contiguous blocks ('up', 'down', 'left' or 'right') or in the two diagonal blocks that permit the agent to efficiently move between the 'open' and the 'pinch-off' transport regimes. These blocks correspond to an increase or decrease of both gate voltages, which strongly modulates the current through the device. The remaining two diagonal blocks, which correspond to a decrease of one gate voltage and an increase of the other, do not often lead to such significant changes in the transport regime and are thus not included in the agent's action space to maximise the efficiency of the algorithm. The DRL agent can be efficiently trained using current maps already recorded from many other devices.

This is because their transport features are sufficiently similar, even though the gate voltage values at which they are observed vary for different devices.

The decision of which block to explore next is based on the current measurements acquired by the DRL agent in a given block. The block is divided into nine sub-blocks (**Figure** 4.1 **c** and **d**) and the mean μ and standard deviation σ of the current measurements corresponding to each sub-block are calculated. These statistical values, constituting an 18-element vector, provide the agent with information of its probable location in the current map. The statistical state vector or state representation vector enables the DRL decision agent to abstract knowledge about the transport regime, distinguishing between 'open', 'pinch-off', and 'single-electron transport' regimes, with a reduced number of measurements. In this way, the state vector defines a state in the quantum dot environment.

This statistical approach, compared to the alternative of using CNNs to evaluate acquired measurements, makes the agent less prone to over-fitting during training and more robust to experimental noise. To decide whether the agent has found bias triangles in a given block, the algorithm uses a CNN as a binary classification tool. Combining a state representation based on measurement statistics and CNNs in a reinforcement learning framework which makes use of the experience of the agent navigating similar environments during training, our algorithm provides a decision process for efficient measurement without human intervention.

4.2 Description of the algorithm

The algorithm is comprised of different modules for classification and decision making (**Figure** 4.2). In the initialisation stage, two low-resolution current traces are acquired

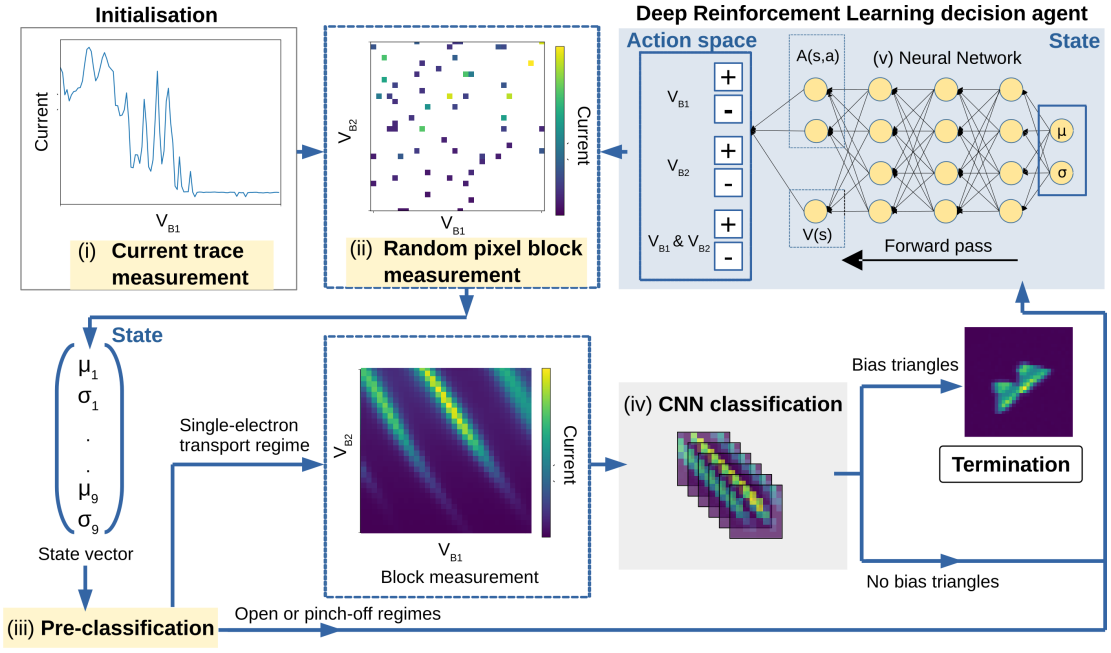


Figure 4.2: Schematic depicting the algorithmic workflow. (See main text for a full description) In the initialisation stage, starting from the gate voltage coordinates proposed by a coarse tuning algorithm, the algorithm measures low-resolution current traces as a function of $\mathrm{V_{B1}}$ ($\mathrm{V_{B2}}$) with $\mathrm{V_{B2}}$ ($\mathrm{V_{B1}}$) set to the maximum voltage given by the gate voltage window of interest (**i**). The algorithm then performs a random pixel measurement in the block corresponding to the proposed starting gate voltages (**ii**). In this measurement, mean current values and standard deviation are calculated for 9 sub-blocks within the block until convergence. The statistical state representation vector (state vector) obtained is then assessed by the pre-classification stage (**iii**). If the mean current value corresponding to any of the sub-blocks falls within threshold values given by the initialisation stage, then the block is pre-classified as corresponding to a possible single-electron transport regime. In this case, the block is explored further by performing a high-resolution scan. This block measurement is normalised and input into a CNN binary classification algorithm (**iv**). If the CNN identifies bias triangles, then the algorithm terminates. If either the pre-classifier or the CNN-classifier rejects a block, then the state vector is input into the DRL decision agent (**v**). The decision agent subsequently selects an action on the gate voltages which determines the next block to measure via the random pixel method.

by the algorithm as a function of V_{B1} (V_{B2}) with V_{B2} (V_{B1}) set to the maximum voltage given by the gate voltage window to be explored. The algorithm extracts from these measurements the maximum and minimum current values and their standard deviation, which will be used in a later stage by the classification modules. The gate voltage regions we explore are delimited by a 640 mV $\times$ 640 mV window centred in the gate voltage coordinates proposed by a super coarse tuning algorithm[62] and the current traces in this stage have a resolution of 6.4 mV.

The gate voltage region is divided in 32 mV $\times$ 32 mV blocks and the agent is initialised in a randomly selected block. The algorithm takes random pixel measurements of current within this block. Each pixel is 1 mV $\times$ 1 mV. As these measurements are performed, the algorithm estimates the 18-dimensional state vector given by μ and σ for each of the 9 sub-blocks in which the block is divided. Pixels are sampled randomly from the block until the statistics from the state representation have converged. Convergence is generally achieved after sampling fewer than 100 pixels, significantly less than the 1024 pixels in a block (**Figure** 4.3).

The state vector is first evaluated by a pre-classification module. A block is classified as corresponding to the single-electron transport regime if any of the nine μ values, corresponding to the nine sub-blocks, falls between two pre-defined values. These values are set to be 0.01 to 0.3 times the maximum current range detected in the initialisation stage. We have found that the choice of such hyperparameters does not have a significant impact on the performance of the algorithm. If the pre-classifier identifies the block as corresponding to the single-electron transport regime, a high-resolution current measurement (1024 pixels, 1 mV $\times$ 1 mV resolution) of the block is acquired. This block measurement is normalised and evaluated by a CNN binary classifier. For

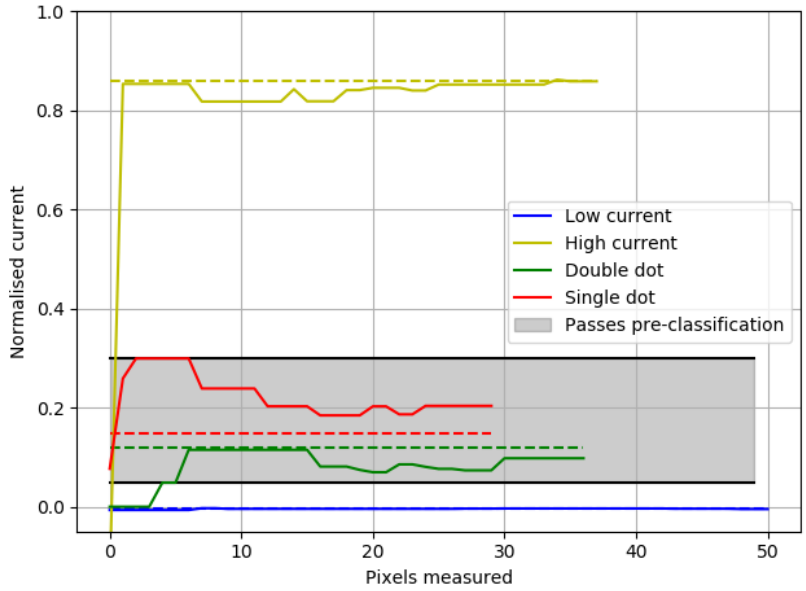


Figure 4.3: Pixel sampling convergence for the random pixel measurement in sub-blocks taken from pinch-off (low current), open (high current), single dot, and double dot regimes. The threshold values set for the pre-classification tools are indicated. Therefore, any sub-block within the grey shaded region will pass the pre-classification. The dashed lines represent the true means after measuring the sub-blocks using a $1\,\mathrm{mV} \times 1\,\mathrm{mV}$ resolution scan.

any output value greater than 0.5, the block is identified as containing bias triangles. If bias triangles are identified within the block, the algorithm is terminated. **Figure** 4.4**a** shows the blocks in a current map that would be identified by the pre-classifier as corresponding to the single-electron transport regime, while **Figure** 4.4**b** shows the blocks that would be evaluated by the CNN binary classification to determine if bias triangles are observed.

If the pre-classifier considers the block to correspond to the 'open' or 'pinch-off' regimes, or if the CNN does not identify bias triangles within the block, the DRL agent has to decide which block to explore next. With this objective, the state vector is normalised using the variance and mean current values obtained in the initialisation stage, and fed into a deep neural network which controls the DRL decision agent. The agent will then propose an action which it expects will lead to the highest long-term reward. This action a_t, given by $a_t = \arg\max_{a'} Q^{\pi}(s_t, a')$, is the action which maximises the Q-function for the agent's stochastic policy π in the state-action pair (s_t, a') at time t. The Q-function measures the value of choosing an action a' when in state s_t and therefore the action a_t represents the agent's prediction for the most efficient route to bias triangles. In our quantum dot environment setting, the action determines the next block to explore and the algorithm begins a new iteration.

4.3 The deep reinforcement learning agent

Our algorithm makes use of the deep Q-learning framework which uses deep neural networks to approximate the Q-function[106]. The Q-function is defined by $Q^{\pi}(s_t, a_t) = \mathbb{E}\left[R_t | s = s_t, a = a_t, \pi\right]$, which gives an expected reward R_t for a chosen action a_t taken by an agent with a policy π in the state s_t. This expected reward is defined as $R_t =$

$\sum_{\tau=t}^{\infty} \gamma^{\tau-t} r_\tau$, where $\gamma \in [0,1]$ is a discount factor that trades off the importance of immediate rewards r_t, and future rewards $r_{\tau>t}$. The agent aims to maximize R_t via the $Q^\pi(s_t, a_t)$ learnt by the neural network. In particular, we chose to implement the duelling deep Q-network (duelling DQN)[141] architecture for our DRL decision agent. This architecture factors the neural network into two entirely separate estimators for the state-value function and the state-dependent action advantage function[141]. The state-value function, $V^\pi(s_t) = \mathbb{E}_{a_t \sim \pi(s_t)} [Q^\pi(s_t, a_t)]$ gives a measure for how valuable it is, for an agent with a stochastic policy π in the search for a promising reward, to be in a given state s_t. The state-dependent action advantage function[141] gives a relative measure of the importance of each action, given by $A^\pi(s_t, a_t) = Q^\pi(s_t, a_t) - V^\pi(s_t)$. In duelling DQN, when combining the state-value function and the state-dependent action advantage function, it is crucial to ensure that given Q we can recover $V^\pi(s_t)$ and $A^\pi(s_t, a_t)$ uniquely. For this purpose, the advantage function estimator is forced to be zero at the chosen action a_t[141]. Our duelling DQN consists of three fully connected layers with 128, 64, and 32 units respectively. The duelling component is defined by a further fully connected layer with 64 units for the action advantage function estimator and one unit for the state-value function estimator.

This approach allows the agent, through the estimation of $V^\pi(s_t)$, to learn the value of certain states in terms of their potential to guide the agent to a promising reward. This is particularly beneficial in our case since different state vectors can correspond to the same transport regime and thus be equally valuable in the search for bias triangles. Consequently, the most beneficial actions in these states would often coincide. For example, in most states corresponding to the 'pinch-off' regime, the most beneficial action is often to increase both gate voltages.

Figure 4.4: Classification tools. (a) Example of blocks considered by the pre-classifier as corresponding to the 'single-electron transport' regime overlaid on the corresponding current map. The colour bar represents the number (M), out of nine, of sub-blocks which were not rejected by the pre-classification stage. **(b)** Blocks in **(a)**, displaying features corresponding to the 'single-electron transport' regime, overlaid on the corresponding current map. Inset: A block displaying bias triangles and the corresponding output value of the CNN binary classifier.

To train the DRL agent, we designed a reward function to ensure that the agent would learn to efficiently locate bias triangles. To this end, during training, the agent is rewarded for the detection of bias triangles and penalised for the number of blocks explored or measured in a single algorithm run, N. The reward $r = +10$ is assigned to the blocks exhibiting bias triangles. Other blocks are assigned $r = -1$. During training, the maximum number of blocks that could be measured in a given run, $N_{\max}$, is set to 300. If after $N_{\max}$ block measurements the agent had not found bias triangles, the algorithm is terminated and the agent is punished with $r = -10$. In other words, $N_{\max}$ determines how far from the starting block the agent can reach in gate voltage space, as it can only explore contiguous blocks.

We trained the duelling DQN (DRL decision agent) using the prioritised experience replay method[142] from a memory buffer. This method ensures that successful policy decisions are replayed more frequently in the DRL agent's learning process. The agent does not benefit from an ordered sequence of episodes during learning, yet it can learn from rare but highly successful policy decisions and it is less likely to settle in local minima of the decision policy. We trained the agent over 10000 episodes (algorithm runs) using the Adam optimiser[143], each time initialised in a random block for 4 different current maps which were previously recorded. The training takes less than an hour on a single CPU (Intel(R) Core(TM) i5-8500 CPU @ 3.00GHz).

4.4 Experimental results

We demonstrate the real-time ('online') performance of our algorithm in a double quantum dot device. The algorithm performance is evaluated according to the number of blocks explored in an algorithm run, N, which is equal to the number of blocks explored to successfully identify bias triangles unless $N = N_{\max}$, and according to the laboratory time spent in this task. For training and testing the algorithm's performance, we use different devices, both similar to the device shown in **Figure** 4.1**a**. We ran the DRL algorithm in two different regions of gate voltage space, I and II, which are centred in the coordinates from our super coarse tuning algorithm[62]. We ran the algorithm 10 times in each region. The DRL agent was initiated in a different block for every run, sampled uniformly at random. From these repeated runs, we can estimate the median $\bar{N}$ of the distribution of values of N obtained for a given region. We can also estimate (L, U), where L and U are the lower and upper deciles of the distribution. To identify bias triangles, the DRL agent required $\bar{N} = 40$ $(9, 104)$ for region I and $\bar{N} = 32$ $(10, 94)$ for region II. In both regions considered, our algorithm efficiently located bias triangles in a mean time of 30 minutes and, on one occasion, in less than 1 minute. This is an order of magnitude improvement in measurement efficiency compared to the laboratory time required to acquire a current map with the grid scan method, i.e. measuring the current while sweeping V_{B2} and stepping V_{B1}, which is approximately 5.5 hours with pixel resolution ($1\,\text{mV} \times 1\,\text{mV}$ resolution). This time corresponds to the measurement of the whole gate voltage window to be explored with no automatic bias triangle identification or any other computational overheads. This is the most time-consuming measurement strategy, but the most common approach until this work. The agent cannot move outside of the gate voltage window, nor can it measure the same block twice. In the

worst-case scenario, the algorithm will thus measure as many blocks as the traditional grid scan method with a small computational overhead. A grid scan of each 32 mV × 32 mV block with 1 mV × 1 mV pixel resolution takes 50 s. For the online runs of our algorithm, each block was measured and assessed for a median time of 23 s in region I and 26 s in region II. This demonstrates the low computational overhead required for the pre-classification, CNN classification, and forward pass of the DRL agent, as well as the efficacy of our random pixel measurements. Using a single CPU of a standard desktop computer, the algorithm is not limited by computation time. It can thus be run with the computing resources available in most laboratories.

Example trajectories of the agent within the gate voltage space give an insight into the transport properties that the agent has implicitly learnt from its environment. When initiated in a transport regime corresponding to pinch off (low current), the agent reduces the magnitude of the negative voltage applied to the gate electrodes, as human experts would do (**Figure** 4.5**a**). Conversely, when initiated in a transport regime corresponding to higher currents, the agent increases the magnitude of the negative voltage applied to the gate electrodes (**Figure** 4.5**b**). The policy thus leads to block measurements in the areas of gate voltage space where bias triangles are usually located. These areas can be anywhere within the defined parameter space.

We have performed an ablation study. Ablation studies are used to identify the relative contribution of different algorithm components. In this case, we aim to determine the benefit of using a DRL agent. We thus produced an algorithm in which the DRL decision agent was replaced with a random decision agent. We compared its performance with the DRL algorithm. The random agent selects an action, sampled uniformly and randomly. The quantum dot environment's (QDE's) action space is six-dimensional

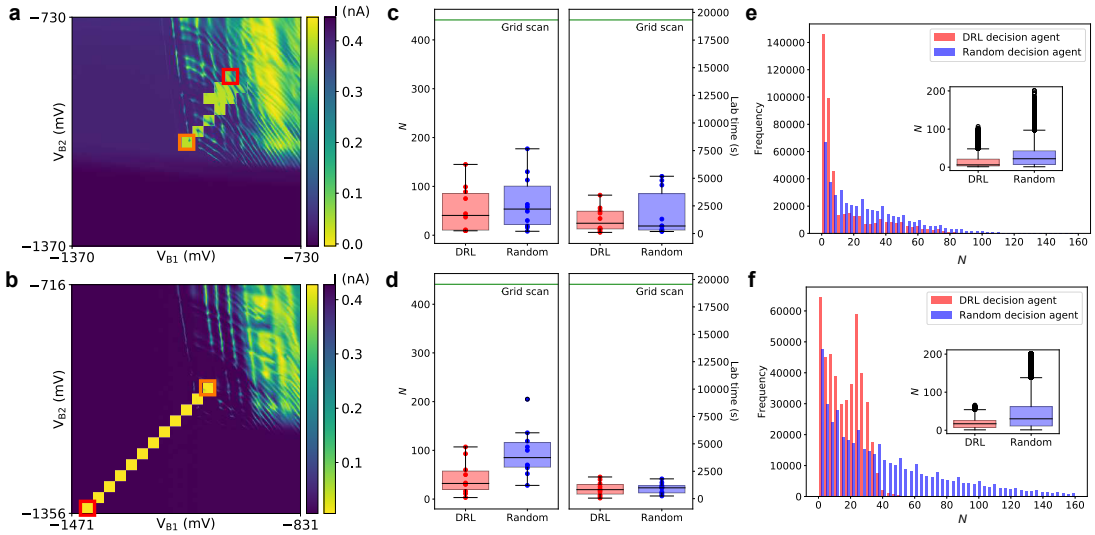


Figure 4.5: Performance benchmark. (a, b) Example trajectories of the DRL agent in gate voltage space. **a** (**b**) Corresponds to region I (II). The trajectories are indicated by inverting the colour scale of the current map for the blocks measured by the algorithm. The full current map measured by the grid scan method is displayed for illustrative purposes and it is not seen by the DRL agent. The red and orange squares indicate the start and end of the trajectory, respectively. **(c, d)** Real-time performance corresponds to the grid scan method (green line), the algorithm with a random decision agent (blue) and the algorithm with a DRL decision agent (red). The box plots indicate the laboratory time (right) and the corresponding number of blocks explored, N, (left) for regions of the gate voltage space I and II in **c** and **d**, respectively. The laboratory time and N are not proportional, as each block requires a different measurement and computational time depending on whether a block measurement is performed and processed by the CNN. The results of all 10 runs for both agents in each regime are plotted as points. The central line of the box plot corresponds to $\bar{N}$, while the upper and lower boundaries of the box display the upper ($Q3$) and lower ($Q1$) quartiles. The minimum and maximum whisker bars display ($Q1 - 1.5 \times IQR$) and ($Q3 + 1.5 \times IQR$) respectively, where IQR is the interquartile range. **(e, f)** Histograms of values of N for the random and DRL decision agents over 10 algorithm runs for each region, I (**e**) and II (**f**). This performance test was performed offline for better statistical convergence. The insets show the box plots, indicating the quartiles and $\bar{N}$ values for the DRL and random agents. In the inset, only the outlier points are plotted.

except in instances where the agent is in a state (block) along the edges (five-dimensional action space) and in the corners (four-dimensional action space) of the gate voltage window considered. This measurement strategy is similar to a random walk within the gate voltage space, but unlike a pure random walk strategy, it will not measure the same block twice. The random decision agent's measurement run will be terminated when the CNN classifies a block measurement as containing bias triangles. The random agent was initialised in the same random positions as the DRL agent so that a fair comparison could be made between their performances. We performed 10 runs of each algorithm in each of the two different regions of parameter space considered in this work, I and II (**Figure** 4.5 **c** and **d**). The DRL agent outperforms the random decision agent in the value of $\bar{N}$, and thus in the laboratory time required to successfully identify bias triangles. Note that the relation between $\bar{N}$ and the laboratory time is not linear, as high-resolution block measurements are only performed for each block classified as corresponding to the single-electron transport regime by the pre-classification stage.

In region II, the random agent requires $\bar{N}$ equal to 85 $(50, 143)$ which is approximately 2.6 times larger than the $\bar{N}$ corresponding to the DRL agent. The good performance of the random decision agent can be explained by its use of the pre-classifier, which makes the random search efficient. The random decision agent is an order of magnitude quicker than the grid scan method.

To test the statistical significance of the DRL agent's advantage, we have tested the performance of both algorithms in a much larger number of runs. To perform this statistical convergence test would have been too costly in laboratory time, so we used previously recorded current maps, which were measured by the grid scan method. We will call this performance test 'offline', as opposed to 'online' in the case of real-time

measurements. By initiating both agents 1024 times in each of the blocks in I and II, we obtained a histogram of the N blocks measured to successfully terminate the algorithm (see **Figure** 4.5**e** and **f** for I and II, respectively). We observe a higher number of runs for which the DRL algorithm performed fewer block measurements for successful termination. In region II, the DRL agent requires $\bar{N}$ of 17 $(2, 31)$, while the $\bar{N}$ for the random agent is 30 $(3, 101)$. Our results suggest that the DRL advantage is statistically significant. The two-tailed Wilcoxon signed rank test[144] allows us to make a statistical comparison of the two distributions corresponding to the DRL and the random agent. We have applied this test to the offline performance for regions I and II. The two-tailed Wilcoxon signed rank test yields a p-value < 0.001 for both regions. This means that the null hypothesis, stating there is no difference in the median performance between the two agents, can be rejected. In addition, the median of the differences $(\bar{N}_{\mathrm{DRL}} - \bar{N}_{\mathrm{Random}})$, estimated using the one-tailed Wilcoxon signed rank test, is less than zero. We can therefore confirm that the DRL agent offers a statistically significant advantage over the random agent.

To further illustrate the advantages of our algorithm, we compared its performance with a Nelder-Mead numerical optimisation method applied to achieve automatic tuning of quantum dots[133,57]. To ensure a fair comparison with our reinforcement learning method, our implementation of the Nelder-Mead optimisation was terminated when the CNN classified a block as exhibiting bias triangles in the same way as our DRL algorithm, i.e. when the output value of the CNN classifier was greater than 0.5. In the original implementation, stricter numerical stopping conditions must be met, thereby increasing the number of measurements performed before termination.

The Nelder-Mead, random decision, and DRL decision algorithms were compared offline.

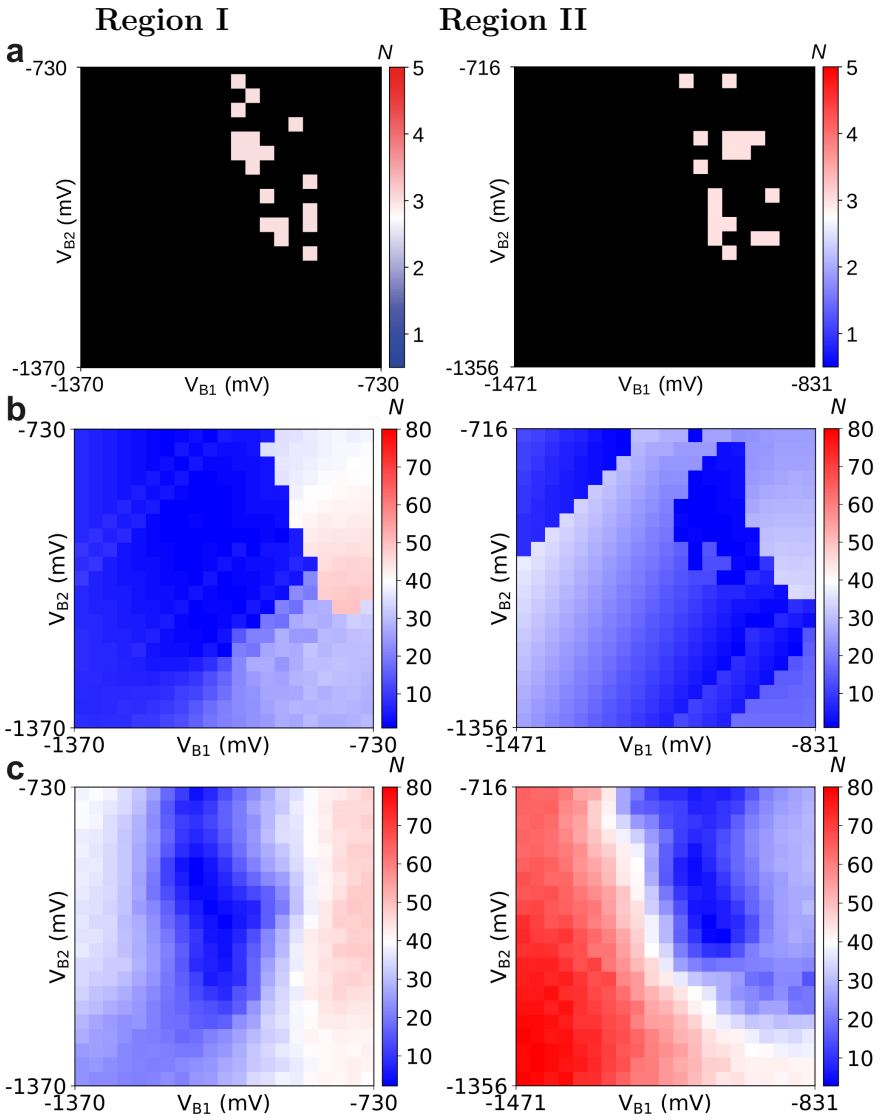


Figure 4.6: Offline performance distribution. The performance of different algorithms is evaluated by initiating an algorithm run in each block and estimating N for regions I and II. Black areas indicate that the algorithm failed when initiated at those blocks. **(a)** Performance distribution (heat-map) for the Nelder-Mead method, **(b)** the DRL decision agent, and **(c)** the algorithm with a random decision agent.

We have initiated the algorithms in each block within each gate voltage region and estimated $\bar{N}$, creating a performance distribution or heat map (**Figure** 4.6). We observe that large areas of gate voltage space which do not exhibit transport features correspond to large flat areas in the optimisation landscape and thus severely limit the Nelder-Mead method. Often the simplex was initiated in these areas and in those cases, the Nelder-Mead algorithm just repeatedly measured the area around the initial simplex. On other occasions, the algorithm moved away from the initial simplex but then became trapped in other areas of the parameter space in which transport features are not present. The method only succeeded in locating bias triangles when it was initiated in the double dot regime. The DRL decision agent's performance is non-uniform as the 'pinch-off' regime is less effectively characterised by the agent than the 'open' and 'single-electron transport' regimes. The performance of the random decision agent is also non-uniform, as it completes the tuning procedure more efficiently when initiated close to the target transport features.

The Nelder-Mead algorithm was also tested online under the same conditions as the DRL and random decision agents. None of 20 runs succeeded before reaching the predefined maximum number of measurements, $N_{\max}$, and thus the results are not presented alongside the online results of grid scan, random decision, and DRL decision algorithms in **Figure** 4.5. Other numerical optimisation methods, better suited to the task, could offer significantly better performances.

4.5 Discussion

We have demonstrated efficient measurement of a quantum dot device using reinforcement learning. We are able to locate bias triangles fully automatically, from a set of

gate voltages defined by a super coarse tuning algorithm[62], in a mean time of less than 30 minutes and in as little as one minute. Our approach gives a 10 times speed up in the median time required to locate bias triangles compared with grid scan methods. The approach is also less dependent on the transport regime in which the algorithm is initiated, compared to an algorithm based on a random agent and to a Nelder-Mead numerical optimisation method. We have also demonstrated the statistical advantage of a DRL decision agent over a random decision agent. Our DRL approach is also robust against featureless areas in the parameter space which limit other approaches. Our algorithm uses statistics calculated via pixel sampling to explore the transport landscape. This statistical state representation allows us to efficiently measure the transport regime (or the state of the environment in DRL terms) and avoid over-fitting during agent training. Other options for state representation that go beyond a statistical summary of current values could be considered. The measurement time remains, however, the dominant contribution in the time required to identify transport features. Fast readout techniques such as radio-frequency reflectometry can be used to reduce measurement times[145,99,53,146,147,148]. However, these techniques are better suited to the measurement of small gate voltage windows. A full tuning procedure could consist of a super coarse tuning algorithm, followed by this algorithm to locate bias triangles, and a fine-tuning algorithm such as the one described in Ref.[134]. Different pairs of bias triangles could be explored. Once a pair of bias triangles is chosen, now within a reduced gate voltage window, the fast readout can be optimised.

Our method is inherently flexible and modular such that it could be generalised to automate a variety of efficient measurement tasks. For example, the reward function could be modified so that the agent could learn to locate and score multiple bias triangles

within the current map. Furthermore, by retraining the CNN classifier and the DRL agent, the method would be able to locate different types of transport features. Our algorithm could also incorporate other gate electrodes by increasing the action space and retraining. This approach would significantly speed up super coarse, coarse and fine-tuning algorithms. We also expect DRL approaches to scale better than random searches as the dimensionality of the problem increases.

We envisage possible extensions of our approach using probabilistic reinforcement learning methods including Bayesian deep reinforcement learning[149,150] and model-based reinforcement learning[151], where the goal is to estimate the uncertainty when making a decision and incorporate domain knowledge into the reinforcement learning model. The resulting reinforcement learning model may be more efficient, especially when a limited amount of data is available.

An additional benefit of reinforcement learning is the capacity of the network's policy to be continuously updated. Thereby, the agent's policy can be updated in real time as the algorithm becomes familiar with a new device. This not only improves the general policy but also means that, over time, the pre-trained agent could learn the particularities of a specific device. To tune large quantum device arrays, due to the increasing dimensionality of the parameter space, DRL could offer a large advantage over conventional heuristic methods. Our algorithm can be implemented in arrays by considering double quantum dots independently[152] and compensating for the cross-talk. Our quantum dot environment and algorithmic framework offer a valuable resource to develop and test other algorithms and decision agents for quantum device measurement and tuning. Additionally, our duelling deep Q-network methods can be translated to further applications in experimental research.

Chapter 5

Bayesian methods for the automatic discovery of "sweet-spot" qubit operating regimes

"*There are things known and there are things unknown, and in between are the doors of perception*".

Aldous Huxley

Optimisation of spin-qubit quality factors.

Problem statement: Automatically identify coherence and control sweet spots in the operating regime of spin-qubits.

Solution: An algorithm which enables the automatic maximisation of the Rabi oscillation frequency or the number of coherent Rabi oscillations, with respect to system parameters including the magnetic field vector, and the voltages applied to device gates. Bayesian inference and Fourier analysis are used to automatically characterise these properties in a given setting and Bayesian optimisation is used to navigate the system parameter landscape to converge upon an optimum.

5.1 Introduction

While significant advancements have been made in the development of silicon-based quantum processors implementations[31,32,80,153,154], a pivotal challenge that remains is the optimisation of qubit properties for robust and high fidelity operation. Quantum dots have a large parameter space defining the tune-up and coherent control of a qubit, which corresponds to highly configurable qubit properties. However, optimising these parameters manually is labour-intensive, prone to human error, and becomes increasingly infeasible as the number of qubits grows. Automating this process is therefore imperative for the realisation of large-scale, error-resistant quantum computers.

The anisotropic nature of spin properties, specifically the g-factors, in valence band holes yields a complex landscape for the control and coherence of spin qubits[21]. This anisotropy gives rise to an electrostatic dependence of the hole spin coupling to the principle decoherence mechanisms, namely nuclear spin fluctuations[21] and charge noise[22]. This complicates scalability by introducing qubit properties which are non-uniform

across the quantum dot parameter space. However, it also enables the potential of a strategic advantage through the identification of operational "sweet spots"[155,33,22,26,156]. Within these sweet spot regimes, qubit control is optimal, and decoherence is minimised. This uniquely offers a mechanism to overcome the control-coherence trade-off. Theoretical analysis suggests that the exact operating points of these sweet spots are material and device-dependent, and due to fabrication imperfections, this makes the location difficult to model and predict for new devices[33,157]. The electrostatic potentials affect on the anisotropic g-factor dependency has origins in the itinerant and localised orbital currents through the nanostructure[158], in addition to the light-hole (LH) and heavy-hole (HH) state admixture, driven by a non-uniform strain profile.[159].

In this chapter, we introduce an algorithm that leverages Bayesian methods, specifically Bayesian Inference and Bayesian Optimisation, for the automatic optimisation of qubit properties[160] and identification of operational sweet spots in semiconductor quantum dots. The algorithm has an iterative two-step process to enhance this quality metric over the device's parameter space: first, it extracts a quality metric from measurements of the qubit, and second, it employs Bayesian Optimisation to propose a new set of candidate parameters. We demonstrate the effectiveness of the algorithm across different qubit parameterisations in two distinct semiconductor-based devices. By converging on optimal gate and field parameters for qubit operation, the algorithm represents a significant step toward the complete automation and optimisation of qubit control, thereby facilitating the development of more scalable and fault-tolerant quantum computing architectures.

5.2 Devices and measurement setup

We performed experiments in two different device architectures with two different types of spin qubits. The first was a Loss DiVincenzo hole spin qubit in a silicon fin field-effect transistor (finFET) similar to that in Camezind et. al.[161]. The second was a Singlet-Triplet hole spin qubit in a Ge/SiGe heterostructure planar device similar to that in Jirovec et. al.[27,162].

5.2.1 Rabi chevron measurements

The first experiment was conducted using a finFET device (**Figure** 5.1). Here, the device was operated in accumulation mode by applying negative voltages to the gate electrodes to form a hole double quantum dot. The sample was measured at $\approx 40\,\mathrm{mK}$ in a dilution refrigerator. A BasPI LNHR 927 DAC was used to apply the DC voltages to the gates and to apply a voltage bias across the device, from the source to the drain. High-frequency signals from a Tektronix AWG5208 were used to modulate the I and Q pulse envelopes of a microwave signal generated by a Keysight E8267D vector signal generator. The microwave envelopes and fast signals from the AWG were combined with the DC signals in an attenuated bias-tee and sent to the device.

We tune the DC voltages across the gate electrodes, including the three barrier gates (LB, CB, RB), and the two plunger gates (LP, RP), to a regime where we observe Pauli spin blockade (PSB). Particularly, we focus on a charge transition between states $(m+1, n+1)$ and $(m+0, n+2)$, such that the effective charge state occupancy is $(1, 1)$ and $(0, 2)$ respectively. We utilise the Pauli exclusion principle to initialise and read out the spin state of a Loss DiVincenzo qubit. To perform Rabi chevron measurements, the protocol consists of initialisation, manipulation, and readout (**Figure** 5.1). During

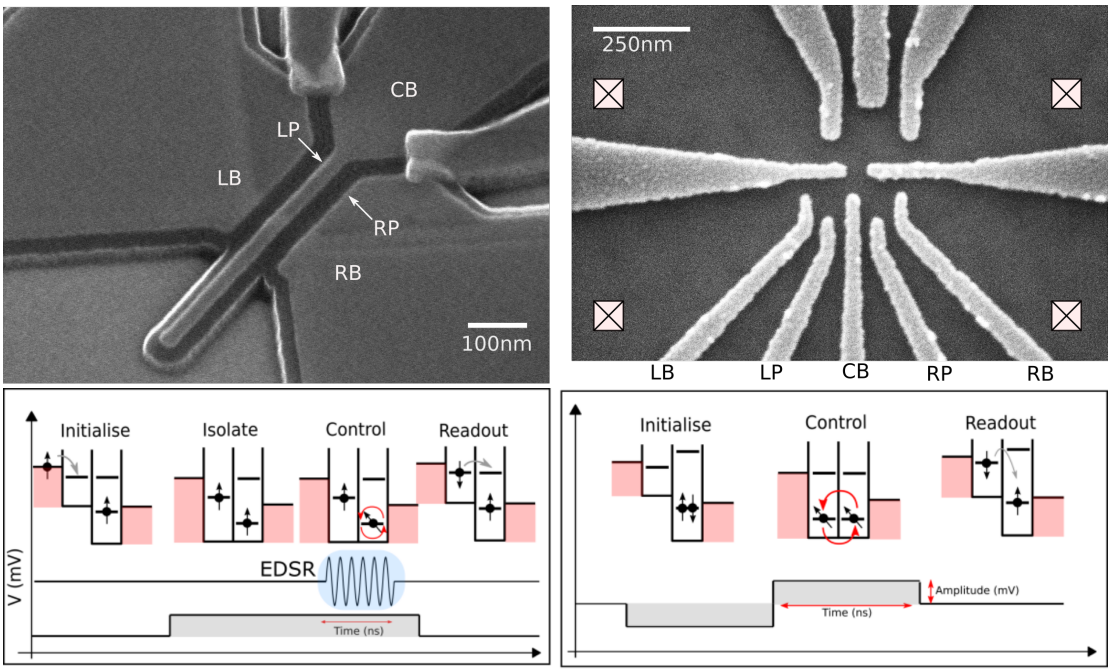


Figure 5.1: SEM images of devices, and pulse schemes for spin manipulation. Left: SEM of the Si finFET device with the left barrier gate (LB), central barrier gate (CB), right barrier gate (RB), left plunger gate (LP), and right plunger gate (RP) labelled. A double quantum dot with hole carriers is formed by tuning the voltages applied to these gates. Here, the dots will form in the fin-shaped channel which runs from the right to left barrier gates, perpendicular to the central barrier gate. The dots will form underneath the left and right plunger gates respectively, with the central barrier gate acting to control the exchange coupling between the dots. An in-plane external magnetic field is applied perpendicular to the fin, and DC measurements are performed by measuring the current from the source to drain through the fin channel. Hole spin movement and pulse scheme for performing EDSR, and PSB-based readout of a Loss-DiVencenzo qubit. The grey area represents a DC pulse applied to the plunger gates, to manipulate the charge state of the quantum dots, and the blue signal represents the high-frequency EDSR pulse used to manipulate the spin state of the qubit. The pulse scheme works by; initialising the quantum dots in the (1,1) ground spin state, performing EDSR using a variable frequency and duration of pulse, and performing readout using Pauli spin blockade. **Right**: SEM of the planar SiGe device, with the device gates labelled. A double quantum dot is formed under the left and right plunger gates respectively, and rf-reflectometry measurements are made using a tank circuit bonded to the ohmic contact of the charge sensor (see **Figure** 3.1). Hole spin movement and pulse scheme for performing EDSR based spin manipulation, and PSB-based readout of a Singlet-Triplet qubit. The grey areas represent the DC pulse applied to the device plunger gates, to manipulate the charge state of the double quantum dot. First, the system is reset in the (2,0) charge state, initialised to the (1,1) singlet state, left to evolve for some variable time, and then read out using a Pauli spin blockade mechanism by pulsing back to the (2,0) charge state regime.

initialisation, we pulse the plunger gates to shift the detuning and a hole tunnelling event changes the charge state from $(0, 2)$ to $(1, 1)$. The $(0, 2)$ state is a singlet due to the Pauli principle, thus the system is initialised by diabatically jumping into a spin singlet state with a $(1, 1)$ charge occupancy.

After initialisation, we begin the manipulation phase, isolating the qubit by applying fast voltage signals across the plunger gates to enter the Coulomb blockade. Driving with a microwave (MW) signal sent to the plunger gate enables spin manipulation; matching the signal frequency to the magnetic field-induced spin splitting causes electric dipole spin resonance (EDSR). By controlling the duration and phase of the MW signal, we induce Rabi oscillations, flipping the qubit's state on the Bloch sphere between the ground and excited spin states.

Following manipulation, in the readout stage we apply fast signals to the plunger gates, removing the Coulomb blockade, and utilising PSB for selective spin to charge state conversion. When repeated, this sequence generates an average current across the device corresponding to the qubit's spin state. In our system, the differential current with and without the MW signal indicates the mean relative spin state population and isolates the signal from the background current noise. To implement this, we incorporated a lock-in amplifier, chopping the MW signal at a specified frequency of 89.17 Hz, and a used pulse cycle repetition rate on the order of ≈ 1 MHz to obtain a high signal to noise ratio. The transport signal was then demodulated and digitised for data acquisition.

We vary the burst time ($0 - 100$ ns) and microwave driving frequency (range of 100 MHz) applied to the device plunger gate to measure a Rabi chevron (**Figure** 5.2). When the driving frequency is on resonance, the qubit state precesses between the ground and excited spin states, around an axis in the x-y plane controlled by the phase of the

MW signal, at the characteristic Rabi frequency. Off-resonant driving, however, causes a precession around an axis which is out of the x-y plane. The further the driving frequency is from resonance, the faster the oscillations, and the closer the precessional axis is to the z-axis (hence the signal amplitude decreases). The objective function of our optimisation procedure is the precessional frequency of the qubit state when driven at the resonance frequency. We must therefore measure oscillations over a range of driving frequencies and automatically extract the resonance and Rabi frequencies.

5.2.2 Single Rabi trace measurements

The second experiment was performed on a planar SiGe heterostructure device similar to the one described by Jirovec et. al.[27] (**Figure** 5.1). We control and readout the state of the singlet-triplet qubit in the sample at $\approx 50\,\mathrm{mK}$ in a dilution refrigerator. Gates G_{1-5}, including the three barrier gates (LB, CB, RB) and the two plunger gates (LP, RP), are used to define a double quantum dot, and gates G_{6-8} are used to define a charge sensing dot. The splitter gates G_9 and G_{10} are set to prevent current leakage between the device and sensor but maintain a capacitive coupling between them. We set the bias voltage across the device to $0\,\mathrm{mV}$. The ohmic of the charge sensing dot was connected to an impedance-matching network, with a $92\,\mathrm{pF}$ decoupling capacitor and $2.7\,\mathrm{H}$ inductor, to perform rf-reflectometry measurements. We used a Zurich Instruments UHFLI to generate an rf-tone to drive the circuit at a power of $-90\mathrm{dBm}$ and demodulate the reflected signal using an rf lock-in amplifier.

We manipulate and measure the oscillations of the qubit's state using high-frequency pulses, generated on an AWG (**Figure** 5.1). We perform initialisation and readout using the Pauli spin blockade principle, as described in section 5.2.1. Here, for the

readout, we integrate the reflected rf-signal from a coupled charge sensor and average multiple measurements. We thus obtain a signal which is proportional to the singlet state probability. We employed principal component analysis of the I and Q components of the demodulated signal, to select the first principal component and maximise the information in our signal[64]. Furthermore, to maintain the sensitivity of our readout, we use a charge sensor compensation algorithm[65]. This algorithm is initialised prior to the optimisation and works by calculating the effect of each device gate on the charge sensor Coulomb peak relative to the charge sensor plunger. This information is used to construct a set of linear virtual gates which perturb the electrostatic potential in the device, but use the charge sensor plunger gate to maintain a constant potential in the charge sensor, thereby maintaining a maximal sensitivity as the device gates are moved.

We can vary the amplitude of the control pulse (along the detuning axis), and qubit precessional time ($0 - 30\,\mathrm{ns}$), to measure the qubit state oscillations. We observe rotations driven by the g-factor difference induced by the spin-orbit interaction, and the driven by exchange interaction[27]. At large values of detuning, the g-factor difference driving mechanism becomes dominant, and the oscillation frequency is approximately independent of the detuning. Hence, we can measure oscillations in time at a single value of the detuning, and robustly calculate the quality factor relating to the x-rotations. These one-dimensional measurements are significantly faster to measure and process than their two-dimensional counterparts.

Our algorithm comprises two steps. The first step involves extracting the desired quality metric from qubit measurements. In the second step, Bayesian optimisation is employed to optimise the quality metric over a set of device parameters, including gate voltages, magnetic field strength, and driving power.

5.3 Quality metric analysis

We demonstrate two approaches for evaluating the qubit quality metrics. The first uses an analytical model for the measured qubit dynamics (section 5.3.1). The parameters are estimated using Bayesian inference and the quality metric is directly obtained from the relevant parameters. The second approach, based on the Fourier transform, is motivated for the case when no analytical model for the measured dynamics is available (section 5.3.2). Together, these two approaches provide a toolkit for evaluating the quality of qubits, catering to both well-understood systems where model-based predictions are viable, and to more complex scenarios that necessitate a non-parametric analysis of the dynamics. By implementing both techniques, we can ensure a more complete and accurate assessment of qubit performance.

5.3.1 Bayesian inference

The first component of our algorithm involves utilising Bayesian Inference to calculate a quality metric from the qubit measurements (refer to methods 3.2 for a full description). We formulate the problem as a probabilistic model where the goal is to infer the posterior distribution of the qubit properties given the measurement data. Bayesian inference (section 3.2) estimates a probabilistic distribution over the parameters of a model given data as evidence. From this distribution, we can extract the most probable values for each parameter. We use automatic differentiation variational inference (ADVI)[102], an efficient Bayesian inference algorithm which maps inference to an optimisation problem by approximating a high-dimensional Bayesian posterior with a simpler variational distribution.

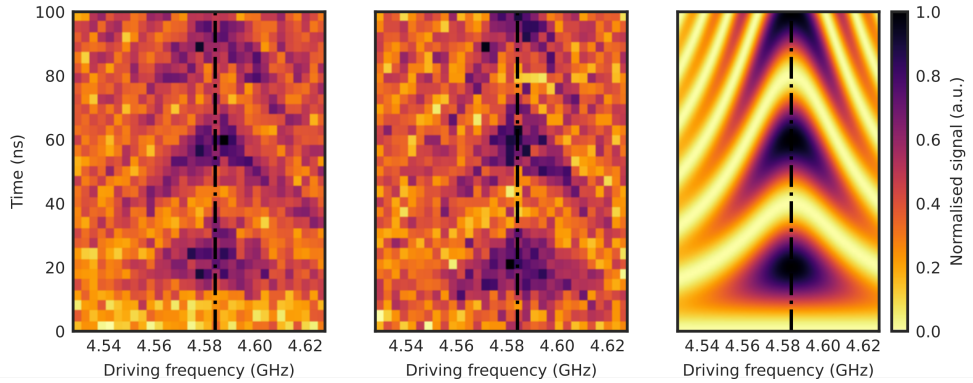


Figure 5.2: Bayesian inference of Rabi Chevron data. Left: Experimentally measured data of a Rabi Chevron on the Si finFET device. The driving frequency and manipulation time are swept and the normalised readout signal is recorded. This signal is the average demodulated signal difference recorded by measuring the current through the device, normalised between 0 and 1 for each rabi chevron measurement. **Centre** (**Right**): A simulation using the mean values of the posterior distributions which were inferred in real-time during the experiment using Bayesian inference, including (excluding) the inferred Gaussian noise level. The black line in each plot gives the position of the mean of the posterior distribution for the resonant frequency.

Parametric model for Rabi oscillations

We derive a parametric model for the Rabi chevron using rotating wave approximation and fit it to data obtained by the measurements described in section 5.2.1. The parameters for this model are the Rabi frequency, Larmor (or resonant) frequency, the phase of oscillations, the scale of measurement noise, and signal scaling factors.

We have a qubit with Hamiltonian $\hat{H}_q = -\frac{1}{2}\omega_q\sigma_z$ and electric dipole moment $\hat{d} = d\sigma_x$. The qubit interacts with an electric field of coherent light $E(t) = Ecos(\omega_d t)$ by the interaction Hamiltonian $\hat{H}_{int} = -\hat{E}\cdot\hat{d}$. The total Hamiltonian is therefore given by

$$\hat{H}_T = -\frac{\omega_q\sigma_z}{2} - E(t)\cdot\hat{d} \tag{5.1}$$

$$= -\frac{\omega_q\sigma_z}{2} - Acos(\omega_d t)\sigma_x \tag{5.2}$$

where $A = Ed$ quantifies how strong the interaction is.

If the interaction is turned off then we know the solution to the Hamiltonian is $|\psi\rangle = C_0 |0\rangle + C_1 |1\rangle$ and the time evolution would be captured by $|\psi(t)\rangle = C_0 e^{+i\omega_q t/2} |0\rangle + C_1 e^{-i\omega_q t/2} |1\rangle$[163]. We will therefore try and find solutions of the form

$$|\psi(t)\rangle = C_0(t) e^{+i\omega_q t/2} |0\rangle + C_1(t) e^{-i\omega_q t/2} |1\rangle . \tag{5.3}$$

Plugging this ansatz into the Schrödinger equation,

$$\frac{\partial |\psi(t)\rangle}{\partial t} = \left(\frac{-\omega_q \sigma_z}{2} - A cos(\omega_d t) \sigma_x \right) |\psi(t)\rangle \tag{5.4}$$

yields, two ordinary differential equations for $C_0(t)$ and $C_1(t)$,

$$C_0^{'} = iA cos(\omega_d t) e^{-i\omega_q t} C_1, \tag{5.5}$$

$$C_1^{'} = iA cos(\omega_d t) e^{+i\omega_q t} C_0. \tag{5.6}$$

These ODE's can be solved analytically by using the rotating wave approximation (RWA). Here, we expand $cos(\omega_d t) e^{\pm \omega_q t} = \frac{1}{2} \left(e^{i(\omega_d \pm \omega_q)t} + e^{-i(\omega_d \mp \omega_q)t} \right)$. We can now argue that fast timescale oscillations will average out in our regime of interest and, therefore, can ignore the fast rotation terms $e^{\pm i(\omega_d + \omega_q)t}$. Thus, we have

$$C_0^{'} = iA e^{-i(\omega_q - \omega_d)t} C_1, \tag{5.7}$$

$$C_1^{'} = iA e^{+i(\omega_q - \omega_d)t} C_0. \tag{5.8}$$

For a qubit initially in the ground state ($C_0(0) = 1, C_1(1) = 0$), the solutions are

$$C_0(t) = A\frac{exp(-i\Delta t/2)}{\Omega_R}\big(\Omega_R cos(\Omega_R t/2) + i\Delta sin(\Omega_R t/2)\big), \tag{5.9}$$

$$C_1(t) = iA\frac{exp(+i\Delta t/2)}{\Omega_R} sin(\Omega_R t/2), \tag{5.10}$$

where $\Delta = \omega_q - \omega_d$ and $\Omega_R = \sqrt{A^2 + \Delta^2}$. We can therefore calculate the dynamics of the system by evaluating the evolution of the population of the excited state $|1\rangle$

$$P_1(t) = |C_1(t)|^2 = \frac{A^2}{\Omega_R^2} sin^2(\Omega_R t/2) = \frac{f_{\text{Rabi}}^2}{\Omega^2} sin^2(\Omega t\pi), \tag{5.11}$$

where we have defined the rate of oscillation on resonance as the Rabi frequency, and converted to angular frequency $A = 2\pi f_{\text{Rabi}}$ (and $\Omega_R = 2\pi\Omega$).

For a spin qubit driven at by a microwave burst at f_{MW}, with a Larmor frequency of f_{Larmor} and Rabi frequency of f_{Rabi}, a measurement over time t_{burst} returns a noiseless current of

$$\begin{aligned}\mu = \alpha + &\frac{\beta f_{\text{Rabi}}^2 \exp\left(\frac{-t_{\text{burst}}}{T_2^*}\right)}{\Delta^2 + f_{\text{Rabi}}^2} \\ &\times \sin^2\left(\pi t_{\text{burst}}\sqrt{\Delta^2 + f_{\text{Rabi}}^2}\right),\end{aligned} \tag{5.12}$$

where $\Delta = |f_{\text{Larmor}} - f_{MW}|$, T_2^* is the Rabi decay rate, and α and β are scaling factors to transform the model from the qubit probability state space to the measured current signal. We fit the model to measured data using Bayesian inference and obtain a quality metric from the value of the Rabi frequency.

Bayesian inference noise classification

To eliminate regions of parameter space where no spin physics is observed, we evaluate the credibility of the inferred values. The inferred parameters are fed back into the model to produce a simulation of the measured data (**Figure** 5.2). Subsequently, the pixel-wise mean squared error (MSE) between the measured data and simulation is calculated. If the MSE exceeds a pre-determined threshold, the original measurement is considered noise and a quality metric of zero is returned. The threshold was determined empirically by analysing measurements, from the same and similar devices, and selecting a value that separated the signal from noise. The algorithm is outlined in pseudocode (Algorithm 1).

Algorithm 1: Objective function: **Bayesian inference and noise classification.**

Data: Experimental measurement.
Input: Prior distributions $p(\theta)$, analytical model $g(\theta)$, noise threshold B.
Output: Posterior probability distributions.
begin
  Set: control parameters x;
  Measure: y;
  Calculate: log-Norm between y and $g(\theta)$;
  Calculate: posterior probability distributions $p(\theta|d)$ using ADVI;
  Calculate: means of posterior distribution θ_{mean};
  Calculate: mean squared error mse between $g(\theta_{\text{mean}})$ and y;
  if $mse < B$;
  then
    Return: 0;
  else
    Return: $p(\theta|d)$;

5.3.2 Fourier analysis

In systems where there is no simple analytical model to describe the measurements, we use the discrete-time Fourier transform to obtain the frequency spectrum of the one-dimensional measurements defined in section 5.2.2. The Fourier transform has robust noise properties and is significantly faster than our Bayesian inference implementation. The quality metric is set to be the qubit quality factor, $Q = 2f_pT_2^*$ or the precessional frequency f_p. We select the dominant frequency component of the measured data, which corresponds to f_p, the precessional frequency. We extract the decay constant T_2^* from the full-width at half maximum of the peak representing the Rabi oscillations. This value tells us the dephasing noise in the rotating frame of the driven evolution[164].

Fourier analysis of damped oscillations

To analytically characterize the Rabi frequency and decay rate of the qubit's damped oscillations, we transform these parameters into the Fourier domain. This conversion allows us to approximate these values numerically by analyzing the Fourier spectrum derived from the experimental data.

The damped oscillation can be written as the real component of the sum of complex exponentials,

$$P_S(t) = e^{-t/T_2^*}\cos(\omega_0 t) \tag{5.13}$$

$$= \frac{1}{2}e^{-t/T_2^*}(e^{i\omega_0 t} + e^{-i\omega_0 t}). \tag{5.14}$$

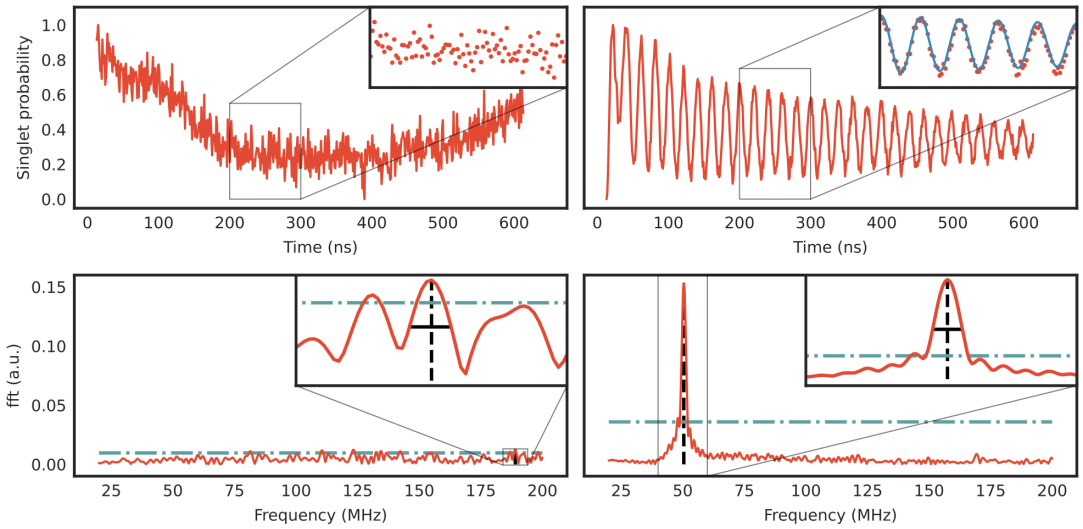


Figure 5.3: Fourier analysis of qubit state oscillations. We plot examples of measurements made by the algorithm and the corresponding Fourier analysis. **Upper Left (Right)**: The measured rf signal as a function of time, normalised between zero and one for each trace. Here we can see noise features (coherent oscillations). Inset, we show a scatter plot of a subset of the data (in addition to a simulated sinusoid with the frequency and decoherence extracted automatically by the algorithm). **Lower Left (Right)**: The Fourier transform of the measured signal (orange line), with an inset showing a subsection of the plot around the maximum frequency response. The algorithm automatically extracts the maximum frequency response (black dashed line) and the full width at half-maximum (black solid line) of the Lorentzian or tallest peak found. We plot the noise classification threshold (blue dotted-dashed line) for the half-maximum, which is given by four standard deviations above the mean of the frequency domain signal.

The Fourier transform is therefore

$$f(\omega) = \frac{1}{2}\int_0^\infty (e^{-at-i\omega t+i\omega_0 t} + e^{-at-i\omega t-i\omega_0 t})dt \tag{5.15}$$

$$\begin{aligned} &= \frac{1}{2}\left(\frac{1}{1/T_2^* + i(\omega_0+\omega)} + \frac{1}{1/T_2^* + i(\omega_0-\omega)}\right) \\ &- \frac{1}{2}\frac{1}{1/T_2^* + i(\omega_0+\omega)}\exp\left((-1/T_2^* + i(\omega+\omega_0))\tau\right) \\ &- \frac{1}{2}\frac{1}{1/T_2^* + i(\omega_0-\omega)}\exp\left((-1/T_2^* + i(\omega-\omega_0))\tau\right). \end{aligned} \tag{5.16}$$

This function has a sharp peak around $\omega = \omega_0$. Near this frequency we can, to a good approximation, discard the fast-rotating terms to give,

$$f(\omega) \approx \frac{1}{2}\left(\frac{1/T_2^* - i(\omega_0-\omega)}{(1/T_2^*)^2 + (\omega_0-\omega)^2}\right)\left(1 - e^{-\tau/T_2^*}e^{i(\omega_0-\omega)\tau}\right). \tag{5.17}$$

Therefore, the modulus of the Fourier transform is

$$|f(\omega)|^2 = \frac{1}{4}\frac{1}{(1/T_2^*)^2 + (\omega_0-\omega)^2}\left(1 - 2\cos\left((\omega-\omega_0)\tau\right)e^{-\tau/T_2^*} + e^{-2\tau/T_2^*}\right). \tag{5.18}$$

In the limit that $\tau \gg T_2^*$ (i.e. the maximum time of the measurement window is much larger than the dephasing time) the modulus of the Fourier transform is

$$|f(\omega)|^2 = \frac{1}{4}\frac{1}{(1/T_2^*)^2 + (\omega_0-\omega)^2}, \tag{5.19}$$

a Lorentzian where the maximum is at $\omega = \omega_0$ and the full-width of half maximum is T_2^*. Thus, we can extract the rate of decay using the full-width at half the maximum in the Fourier spectrum and calculate T_2^* (**Figure** 5.3).

Fourier analysis noise classification

To distinguish between spin physics and noise, we define a threshold of four standard deviations above the mean of the frequency domain representation. Frequency domain peaks generated by a noise distribution, assuming it is Gaussian, have a 99.994% statistical probability of falling within this threshold, so a signal above this is likely to represent the oscillation frequency of the qubit (**Figure** 5.3). If the half maximum of the tallest peak in the frequency spectrum does not reach this threshold, the measurement is classified as noise and the quality metric is set to zero. The algorithm is outlined in pseudocode (Algorithm 2).

Algorithm 2: Objective function: **Fourier analysis and noise classification.**

Data: Experimental measurement
Input: Frequency range, noise threshold B
Output: Dominant frequency, decay constant
begin
  Set: parameters x;
  Measure: y;
  Calculate: frequency domain representation $\Omega(\omega)$ of y using DTFT;
  Calculate the maximum frequency response, $\omega_{max} = argmax_{\omega}(\Omega)$;
  Calculate: the decay constant, T_2^*, using the full width of the peak at $\Omega(\omega_{max})/2$;
  if $\Omega(\omega_{max}) < (mean(\Omega(\omega)) + Bstd(\Omega(\omega)))$;
  then
    Return: 0;
  else
    Return: $Q = \omega_{max} * T_2^*$;

5.4 Bayesian optimisation of quality metrics

Once a reliable quality metric is established, the next step is to optimise this metric over the device's parameter space. We utilise Bayesian Optimisation, a probabilistic

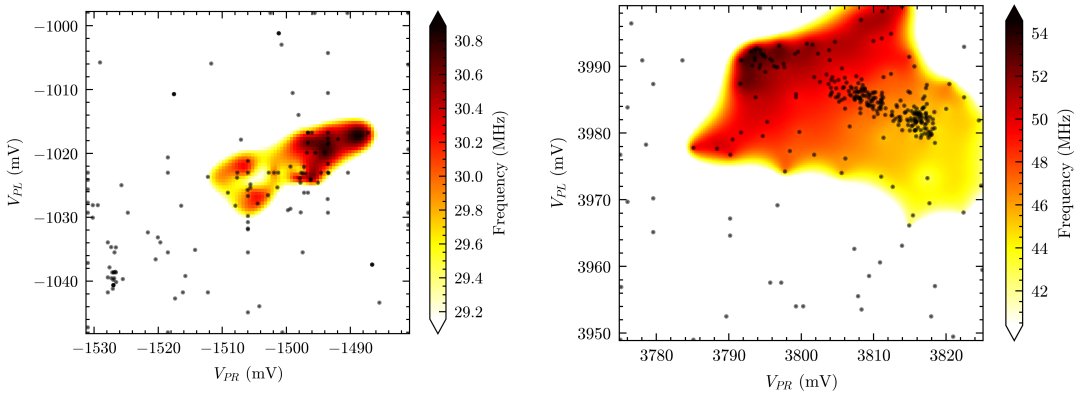


Figure 5.4: Gaussian Process Regression model of the variation of the Rabi frequency with respect to plunger gate voltages. Left: Points measured by the optimiser for the optimisation with respect to the plunger gate voltages (V_{PL} and V_{PR}) in the Si finFET device. We plot the Gaussian process regression model mean for the Rabi frequency as a function of the plunger gate voltages. **Right**: Points measured by the optimiser for the optimisation with respect to the plunger gate voltages (V_{PL} and V_{PR}) in the planar SiGe device. We plot the Gaussian process regression model mean for the Rabi frequency as a function of the plunger gate voltages.

model-based optimisation technique, to efficiently explore the high-dimensional parameter space. Our choice of Bayesian Optimisation is driven by its ability to reduce the number of required evaluations for convergence, as compared to alternative numerical optimisation methods such as Nelder-Mead. This reduction in the number of evaluations is crucial, as a significant portion of the algorithm's run-time is spent on measurements, and we aim to allocate more time towards computation to enhance overall algorithm efficiency. Furthermore, Bayesian Optimisation has the advantage of internally modelling measurement noise, making it more robust in experimental scenarios where measurement noise is inevitable.

An essential element of Bayesian Optimisation is the acquisition function, which guides the search for the next point to sample in the parameter space. In this study, we consider the Expected Improvement (EI) acquisition function, which offers an efficient balance between exploration and exploitation, with analytical tractability that lends itself to

high-dimensional optimisation problems[165]. Its ability to incorporate uncertainty and ensure asymptotic convergence to the global optimum makes it a robust and versatile choice in Bayesian optimisation.

To approximate the quality metric function, we employ Gaussian Processes (GPs) as a surrogate model[166] (**Figure** 5.4). GPs allow for a flexible, non-parametric model that can adapt to the complexity of the underlying quality metric. We use Gaussian process regression with a Matérn 5/2 kernel, which is distinguished by its smoothness, providing a more flexible fit than the squared exponential kernel while retaining computational efficiency[167]. Its capacity to capture long-range dependencies and non-stationary features makes it well-suited for modelling complex physical phenomena.

5.5 Experimental results

We present the results of the optimisation algorithm trialled in both devices, a Loss DiVincenzo qubit in Si finFET device, and a Singlet-Triplet qubit in a planar SiGe device.

5.5.1 A Loss DiVincenzo qubit in a Si finFET device

Here, we conducted independent trials of the algorithm in the Si finFET sample (section 5.2.1). In this instance, we optimise the Rabi frequency with respect to the plunger gate voltages over a $50\,\mathrm{mV} \times 50\,\mathrm{mV}$ range. For each trial of the algorithm, it successfully improved the optimal known value of the Rabi frequency, over 35 chevron measurements (**Figure** 5.5). Each measurement consisted of a 2D grid sweep (as described in section 5.2.1) of the microwave burst time and frequency (31×31 pixels), taking on the order of 850 s. The optimisation metric was evaluated using our proposed Bayesian

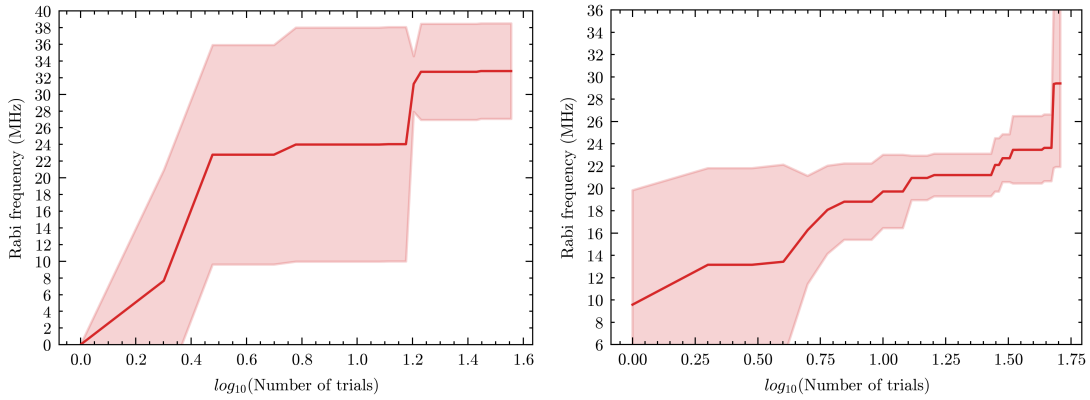


Figure 5.5: Optimisation of the Rabi frequency of a Loss-DiVincenzo qubit in a Silicon finFET device. Left: Convergence of the mean and standard deviation of the optimal Rabi frequency, during Bayesian optimisation, with respect to the number of measurements. Here, the Rabi frequency is optimised with respect to the left and right plunger gate voltages (V_{PL} and V_{PR}) over a $50\,\mathrm{mV} \times 50\,\mathrm{mV}$ range. **Right**: The convergence of the mean and standard deviation of the optimal Rabi frequency with respect to the number of measurements, optimising over the left ($10\,\mathrm{mV}$ range) and right ($10\,\mathrm{mV}$ range) plunger gate voltages, the central barrier ($10\,\mathrm{mV}$ range) gate voltage, the magnetic field strength ($120 - 140\,\mathrm{mT}$), and the power of the EDSR driving pulse ($(6 - 14\,\mathrm{dBm})$ at the signal generator output).

inference algorithm (section 5.3.1). The noise threshold method successfully classified the region of parameter space for the observation of Rabi chevron patterns, and the Bayesian optimiser clustered queries around this region. We also present the mean of the Gaussian process regression model, to illustrate the internal model developed by the black box Bayesian optimiser (**Figure** 5.4).

We then performed independent trials of the algorithm (**Figure** 5.5) while optimising the Rabi Frequency with respect to all the variable parameters. This included; the plunger gates; the central barrier gate; the magnetic field strength; and the power of the EDSR driving pulse. The algorithm sequentially improved the optimal value of the frequency on each independent trial, of 50 chevron measurements.

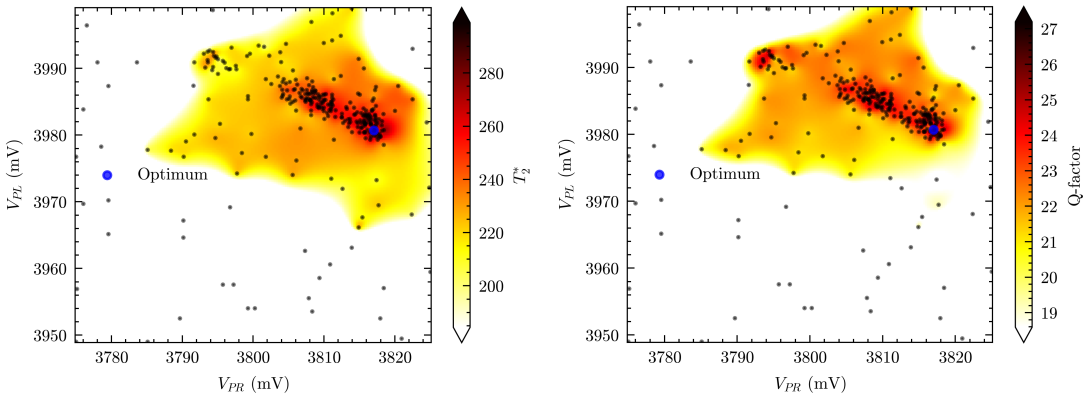


Figure 5.6: Coherence and quality factor sweet-spots. Points measured by the optimiser for the optimisation of the Q-factor with respect to the plunger gate voltages (V_{PL} and V_{PR}) in the planar SiGe device. The coherence sweet-spot coincides with the Q-factor sweet-spot for the device in this configuration. **Left**: We plot the Gaussian process regression model mean for the coherence time T_2^* as a function of the plunger gate voltages. **Left**: We plot the Gaussian process regression model mean for the quality factor Q as a function of the plunger gate voltages.

5.5.2 A Singlet-Triplet qubit in a planar SiGe device

We conducted independent trials of the algorithm of the Singlet-Triplet qubit in the planar SiGe device (section 5.2.2). As there is no simple phenomenological model to simulate the detuning dependence of the oscillations, we used the algorithm with Fourier analysis (section 5.3.2) to infer the Rabi frequency and decay rate, and thus calculate the quality metric. During 10 independent runs, we optimised the quality factor with respect to the plunger gate voltages (**Figure** 5.7). The algorithm automatically navigated to regions of the parameter space where x-rotations were observable and optimised this metric over the course of 300 trace measurements for each trial. The mean of the Gaussian process regression model, trained during the optimisation process, is plotted to demonstrate the information encoded in the model. A strong dependence of the Rabi frequency upon both plunger gate voltages is elucidated by the Gaussian process model (**Figure** 5.4). We also demonstrate the discovered, and learnt models, of the "sweet-

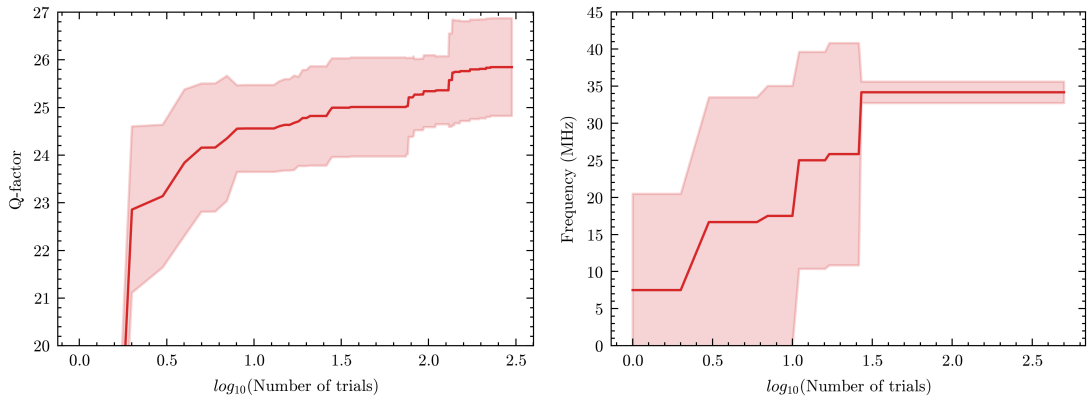


Figure 5.7: Optimisation of quality metrics of a Singlet-Triplet qubit in a SiGe heterostructure device. Left: Convergence the mean and standard deviation of the optimal Q-factor, during Bayesian optimisation, with respect to the number of measurements. Here, the Q-factor is optimised with respect to the left and right plunger gate voltages (V_{PL} and V_{PR}) over a 50 mV range for each gate. **Right**: The convergence of the mean and standard deviation of the optimal Rabi frequency with respect to the number of measurements, optimising over the plunger and barrier gate voltages (V_{PL}, V_{PR}, V_{BL}, V_{BC}, V_{BR}), using a 50 mV range for each gate.

spot" for both the coherence time T_2^*, and the Q-factor (**Figure** 5.6). A maximum quality factor of Q = 27.5 ($T_2^* = 247.61$ ns) was attained with a magnetic field strength of 3 mT. In all 10 independent runs, a quality factor of more than 20 was found in less than a minute. Each trace measurement took on the order of 2 s, with the data processing and quality metric calculation taking on the order of 1 ms.

We then proceeded to conduct 10 independent runs of the algorithm, with 300 trace measurements for each. This time optimising the precessional frequency with respect to the five gate voltages which define the electrostatic confinement potential of the double dot (**Figure** 5.7). These gate voltages include the left and right plunger gates, the left and right barrier gates, and the central barrier gate. Using an external magnetic field of 3 mT, the algorithm automatically found a region of the parameter space where a precessional frequency of 46.7 MHz could be induced.

5.6 Discussion

We have proposed an algorithm which can robustly and automatically characterise and optimise qubit quality metrics. We demonstrate two methods for the inference of qubit properties, which serve as the objective function for the optimisation. We believe that these methods, a model based and a model free instance, comprehensively cover the evaluation of qubit quality metrics in semiconductor quantum dot devices. Our algorithm could be trivially adapted in order to optimise quality metrics in a variety of quantum systems. The ability to automatically discover the optimal operating points is a vital step towards the construction of scalable quantum information processors with large quantum volumes. The automated procedure will function for an arbitrary number of parameters, however Bayesian optimisation becomes computationally inefficient above about 20 parameters. To scale this method to larger qubit systems, we recommend using a separate instance of the optimiser for each qubit or each unit cell of qubits.

The optimisation algorithm, coupled with Bayesian inference for the evaluation of the objective function, maximised the operational Rabi frequency in a Loss DiVincezo qubit. This qubit was tuned in a Si finFET device and manipulated all electrically using the Spin-Orbit interaction to drive EDSR. The algorithm successfully navigated to a region of gate voltage space where the Rabi oscillations were observable and constructed an internal model (the gaussian process regression model) of this region and the variation in Rabi frequency across it. This internal model could also be used for scientific investigation and discovery.

Using Fourier analysis, for the evaluation of the objective function, our algorithm automatically optimised the quality factor of a Singlet-Triplet qubit. The qubit was defined

in a Ge/SiGe planar device, and we manipulated the qubit state by leveraging the exchange interaction and g-tensor modulation mediated by spin-orbit coupling. The spin manipulation and decoherence properties were found to be tuneable, resulting in a variable precessional frequency and decay rate with respect to the electrostatic confinement potential. This variability allowed the algorithm to efficiently identify optimal operational points or "sweet spots" in the gate voltage parameter space, where the qubit quality factor, Q, was optimised. This optimal operational point, where $Q = 27.5$, was discovered at a field magnitude of only 3 mT.

In state-of-the-art devices, the coherence has been found to be limited by $1/f$ charge noise[22,153,154], however, the hyperfine noise contribution can also be substantial. A problem arises in that the "sweet spots" for these two decoherence processes do not necessarily align[154]. In this case, any black-box mechanism for automated optimisation may only discover the minima of the limiting decoherence mechanism, or a minima in the combined landscape which is not at the absolute sweet spot of either mechanism. This motivates further investigation, and the development of a theory driven automated strategy for the discovery and characterisation of the minima for both charge-noise induced and hyperfine induced decoherence.

Chapter 6

Robust optimal quantum control for high-fidelity gates with spin qubits

"*Anyone whose goal is 'something higher' must expect someday to suffer vertigo. What is vertigo? Fear of falling? No, Vertigo is something other than fear of falling. It is the voice of the emptiness below us which tempts and lures us, it is the desire to fall, against which, terrified, we defend ourselves*".

Milan Kundera

Pulse construction for noise robust and high-fidelity gate operations

Problem statement: The investigation and discovery of pulse envelope functions which perform high-fidelity and noise-robust unitary gate operations on spin-qubit systems.

Solution: The design of optimal control schemes for spin-qubit systems using a differentiable time-dependent Schrödinger solver, noise model, and pulse distortion model. This chapter demonstrates how differentiable models can be leveraged with stochastic gradient descent to effectively mitigate both coherent and incoherent errors in single- and two-qubit quantum gates, with a particular focus on spin qubits.

6.1 Introduction

Improving the fidelity of a quantum gate requires the systematic mitigation of both coherent and incoherent error sources. A prime example of coherent errors in driven single-qubit gates is spectator errors such as crosstalk. These can arise due to phenomena such as off-resonant driving, which occur due to microwave leakage or the utilisation of shared driving gates[80]. Other coherent error sources include the emergence of higher harmonics from non-linear driving[168], unexpected phase or frequency shifts[74], and the effects of counter-rotating drives like the Bloch Siegert shifts[169]. It is essential to factor in these non-linear driving effects when developing a theoretical optimal control framework. In contrast to coherent errors, incoherent errors in quantum systems arise from stochastic processes that lead to the loss of quantum information over time. A significant class of incoherent errors is associated with coloured noise sources, which have a non-uniform spectral distribution and are ubiquitous in realistic quantum systems.

Our work addresses the challenge of optimal quantum control, with specific applications in hole spin qubits, by presenting an automated, theory-driven approach for designing control pulses for high-fidelity gates. We lay an emphasis upon reducing coherent errors by introducing realistic constraints relating to the hardware and control stack and reducing incoherent errors by the design of robust optimal control pulses which are resistant to various sources of decoherence, such as coupling to a two-state fluctuator, or quasi-static errors in the local charge or magnetic environments. Crucial to both of these targets is the ability to efficiently compute gradients through a numerically sampled ensemble of trajectories, and corresponding fidelities. We focus on optimal control via stochastic gradient descent, using efficient automatic differentiation and optimisers enabled by modern machine learning tools.

We use JAX[170], a python based software package which enables automatic differentiation through all operations. Automatic differentiation is a computational technique that efficiently computes derivatives of functions by systematically applying the chain rule at the elemental operation level, enabling the calculation of gradients much faster compared to numerical or symbolic differentiation[171]. JAX also supports just-in-time compilation (JIT), which refers to the process of compiling Python functions into optimised machine code at runtime, thereby significantly improving the performance. Furthermore, JAX is vectorisable, meaning it can operate on and apply functions to entire arrays of data at once rather than performing operations on each element sequentially, which leverages the underlying hardware capabilities for parallel processing and results in substantial speedups. Lastly, JAX is GPU accelerated, indicating that it can utilise Graphics Processing Unit (GPU) hardware to perform computations much faster than a typical Central Processing Unit (CPU). The computational acceleration enabled by

JAX allows us to evaluate the time-evolution of the quantum system rapidly, and vectorise the fidelity calculation over a distribution of coloured noise samples and initial qubit states, as required for the optimisation of robust gates. The advantage of this method, with respect to previous gradient-based optimal control techniques, such as the original GRAPE proposal, is that we can add any modifications to our simulation, including vectorised ensemble averages, and still trivially and efficiently acquire gradients. These modifications can include any aspect of a realistic quantum device measurement setup, such as filtering effects, non-linear transmission functions, and experimental noise sources. The improved computational speed also enables us to do without approximations such as the rotating wave, or the white-noise spectra assumed by the Lindblad picture time evolution.

We demonstrate the utility of our approach with applications in a simulated hole-spin two-qubit processor[50], by designing robust and high fidelity single-, simultaneous single-, and two-qubit gates. The algorithm is general and can be applied to any controllable quantum system. However, the optimised quantum control schemes, are characteristic of the hardware and the control electronics, determined by the formulation of the Hamiltonian, the Hamiltonian's parameterisation, and the parameters defining the pulse distortion, such as the bandwidth, and the noise landscape. We use our method to design pulses which are optimally robust to pink-noise processes, analogous to coupling to a two-state fluctuator. Subsequently, we investigate the mitigation of cross-resonance and other coherent errors in performing simultaneous single-qubit gates. Lastly, we develop two-qubit CNOT gates which are robust with respect to a distribution of Hamiltonian system parameters, crucial for mitigating quasi-static errors caused by charge noise and hyperfine coupling. Together, these examples represent an advancement in the quest to

optimise spin qubit operations, addressing both the nuanced aspects of coherent errors and the challenges posed by environmental interactions.

6.2 Robust optimal quantum control

The fidelity of a quantum gate defines the closeness of the input-output state mapping to that of the ideal unitary operation. Given a quantum system in a well-defined initial state, and a Hamiltonian describing the dynamics of the evolution of the system under the time-dependent Schrödinger equation,

$$\hat{H}(t) = \hat{H}_0 + \hat{H}_C\left(\Omega(\theta, t)\right) + \hat{H}_{\text{noise}}(t) \tag{6.1}$$

with a drift term $\hat{H}_0$, a noise term $\hat{H}_{\text{noise}}$, and a control term $\hat{H}_C$ defined by a time dependent control signal $\Omega(\theta, t)$. The problem of optimal quantum control is to find a function $\Omega(\theta, t)$, parameterised by θ, which evolves the system to a desired target state or implements a target transformation. To implement a target transformation U_T, the ideal map can be calculated by minimising the loss function L, calculated from the Frobenius inner product, with respect to the control pulse $\Omega(\theta, t)$ [172],

$$L\big(\Omega(\theta, t)\big) = 1 - F_{avg} \tag{6.2}$$

$$F_{avg} = \frac{1}{M}\sum_{j}^{M}\left|\text{Tr}\left((U_T\Psi_j)^\dagger(U_C\Psi_j)\right)\right|^2 \tag{6.3}$$

where the control unitary transform $(U_C\Psi_j)$ is calculated by solving the time dependent Schrödinger equation. The fidelity is averaged over M initial states, j, which form a Haar measure over the state space. This averaging process utilises the concept of the Haar measure (Appendix C), which is a unique, translation-invariant measure on the

unitary group of quantum states[173]. For a single qubit, the gate fidelity is averaged over the six pure axial initial states, $\Psi(t=0)=\rho_j$, which form a Haar measure on the Bloch sphere. For the two-qubit subspace required for evaluating simultaneous single-, and two-qubit gates the Haar measure is well approximated by numerically sampling the two-qubit Clifford set[174] or using a Pauli twirling group[175].

In the instance of designing robust optimal control schemes, the fidelity F_{avg} averaged over the Haar measure, will also be averaged over sampled noise trajectories, stochastic instances of $\hat{H}_{noise}(t)$ or samples from a distribution of the parameters defining $\hat{H}_0(t)$, and this ensemble average will be maximised. There is, therefore, a computational balance, particularly with the two-qubit scenario, with respect to sampling initial states to improve the approximation of the Haar measure and sampling the stochastic noise process to improve the approximation of the dynamics of an ensemble subjected to environmental interactions.

6.3 Algorithm

The algorithm consists of a pulse distortion module which enables the mitigation of coherent errors; a numerically sampled stochastic noise module enabling robustness to incoherent errors; and an ODE module to solve the time evolution. These modules are constructed in JAX[170,178], so that all operations will be automatically differentiable, enabling optimisation of the pulse parameters via stochastic gradient descent. We also use just-in-time (JIT) compilation and vectorisation to compute the trajectories of the initial state set, and sampled stochastic processes in parallel. These accelerations can be employed on a CPU or a GPU and give a computational advantage over Qutip[176,177](**Figure** 6.1). Here, we show the scaling of the wall clock time with the

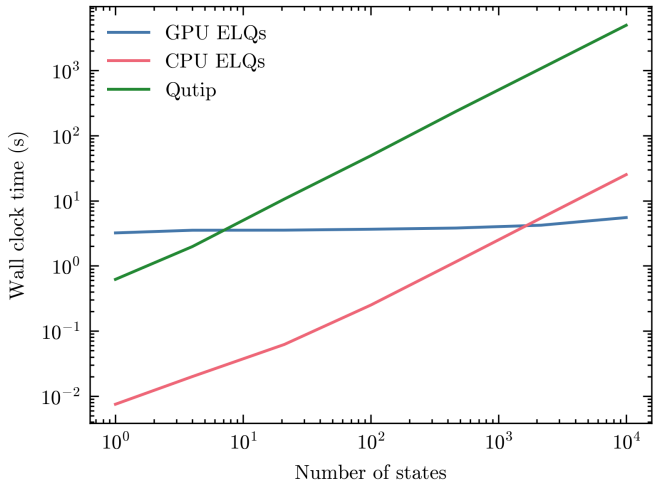


Figure 6.1: Wall-clock time comparison, for the scaling with the number of initial state vectors evaluated for the evolution of a two-spin quantum, parameterised by the values reported for hole spins in a Si finFET[50], system under a time-dependent Hamiltonian in the Schrödinger picture. Included is the scaling of our method (ELQ) compiled on a GPU, and CPU, compared to Qutip[176,177] running on a CPU.

number of initial state vectors for a simulation of the time evolution in the Schödinger picture.

We can now describe the general framework of our method, as depicted in **Figure** 6.2. The algorithm is initiated by defining a time-dependent pulse function $\Omega(\theta, t)$, which is parameterised by some variables θ, the system Hamiltonian $\hat{H}(t, x)$, the target Unitary U_T, and the pulse distortion function $f(y)$. The pulse function defines the mechanism with which we can influence the evolution of the quantum system (section 6.3.1), so the objective of the optimisation procedure will be to optimise the average gate fidelity F_{avg} with respect to the variables θ. Here, gate fidelity is averaged over numerically sampled stochastic noise processes and a set of initial states which approximate a Haar measure over the quantum system.

Within the main block, the pulse function is transformed to account for filter effects,

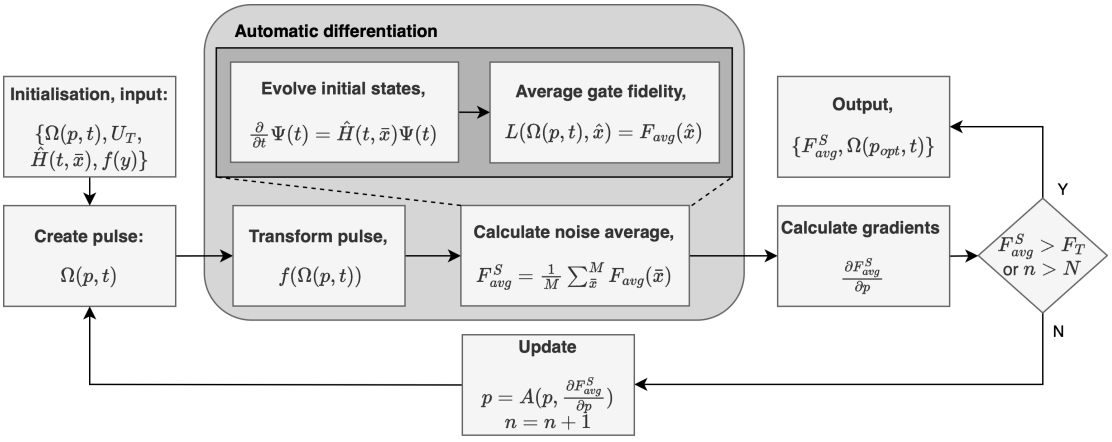


Figure 6.2: Algorithm flow diagram. The core function of the algorithm is to calculate the state fidelity, of the system evolved by the control pulse, averaged over a set of initial states and numerically sampled noise terms. The control pulse is subject to realistic experimental constraints via the pulse transformation function, including a non-linear transmission function and finite bandwidth filter function. The gradient of this average fidelity with respect to the parameters defining the control pulse, is then calculated using automatic differentiation in JAX. A gradient-based optimisation function then converges upon an optimal set of pulse parameters.

non-linear transmission functions, and other hardware constraints (section 6.3.1). The distorted pulse is then used to construct the control Hamiltonian $\hat{H}_C(t)$ and simulate the time-evolution of the system, for each of the initial states. The evolution is calculated by using an ODE solver for the Schrödinger picture with the drift Hamiltonian $\hat{H}_0(t)$ and stochastic process $\hat{H}_{noise}(t)$, to calculate the average gate fidelity and define the loss (equation 6.2). We map the problem to a standard ODE form[179] and solve using a Dormand-Prince's 8th order Runge-Kutta method[180], adapted to support complex coefficients. We can now calculate the gradient of the loss function with respect to the pulse parameters, $\partial L/\partial\theta$. A gradient-based optimisation algorithm (here we use the ADAM optimiser[181]), performs an update on pulse parameters. This process is continued until the target loss (L_T) is reached, or the number of iterations equals the defined maximum (N).

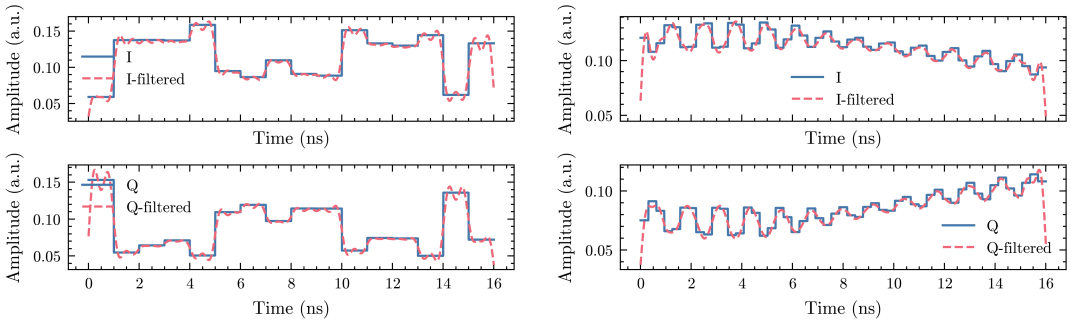


Figure 6.3: Pulse distortion. Examples of optimised pulse envelope, and the effect of the bandwidth, rise time, non-linear transmission, and amplitude constraints. In these examples, we use the transition bandwidth $b = 100\,\mathrm{MHz}$, the cut-off frequency $f_c = 2\,\mathrm{GHz}$, and the rise time $t_{rise} = 0.5\,\mathrm{ns}$ These pulses are optimised to perform an X-gate on the qubit in dot 1 using (**left**) 16, and (**right**) 51 parameters to define the PWC envelope functions for both the I and Q components.

6.3.1 Pulse distortions

In order for control pulses to be practically applicable, they must conform to a wide range of physical constraints. These constraints include signal rise times, finite bandwidths, and bounded amplitudes. Here, these effects are modelled and incorporated into the optimisation loop.

Bandwidth and rise time constraint

The AWG and filters in an experiment distort the pulse signal that is sent to the device. The signal is bounded to some finite value at the start and end of the pulse envelope, and a low-pass filter with finite bandwidth is applied. Central to this approach is the application of a low-pass filter through the convolution of the pulse envelope with the sinc function,

$$s[z] = \mathrm{sinc}\left(2f_c\left(z - \frac{Z-1}{2}\right)\right) \tag{6.4}$$

where f_c is the cut-off frequency, and N is the number of points in the filter. We calculate the filter length Z using its relation to the transition bandwidth $b \approx \frac{4}{Z}$. The convolutional kernel must be truncated to a finite length to be computationally useful. However, due to the asymptotic properties of the $1/x$ term this truncation leads to discontinuities that can introduce undesirable frequency artefacts. To counteract these effects, the truncated sinc function is multiplied by a Blackman window[83]. The Blackman window with a number of time steps Z is defined by:

$$w(z) = 0.42 - 0.5\cos\left(\frac{2\pi z}{Z-1}\right) + 0.08\cos\left(\frac{4\pi z}{Z-1}\right) \tag{6.5}$$

where $0 \leq z \leq Z-1$[182]. This results in an improved frequency response: by minimising oscillations in the pass band and improving attenuation in the stop band. To effectively implement constraints related to the device's rise time, we further introduce zero-padding to the windowed kernel function. The multiple of the number of zeroes added and the reciprocal of the computational sample rate corresponds to the device's rise time. This combined approach of convolution, windowing, and zero-padding offers a robust way to model and therefore mitigate pulse distortions.

Non-linear transmission function and amplitude constraint

We enforce a smooth and hard constraint on the pulse amplitude, and introduce a non-linearity to the transmission function, reproducing hardware and experimental constraints. Specifically, this accounts for the observed non-linear transmission function of drive amplitudes in spin qubits[74], which becomes particularly impactful when investigating the simultaneous driving of qubits. We achieve this by defining a phenomenological transformation[83], which is applied to the filtered pulses. The transformation consists

of the scaled sigmoid function

$$g(t,g) = 1 - \frac{1}{1+e^{gt}} \tag{6.6}$$

$$f\left(\Omega(t), l, u\right) = \left(2g\Big(\frac{\Omega(t) - (u+l)/2}{(u-l)/2}, 0.5\Big) - 1\right)\frac{u-l}{2} + \frac{u+l}{2} \tag{6.7}$$

where u, and l are the upper and lower bounds on the amplitude of the pulse envelope $\Omega(t)$.

Pulse parameterisation

The choice of pulse parameterisation remains a core component of optimal quantum control algorithms. Here, we will assess three parameterisations, namely PWC (Piece-wise Constant), the truncated Fourier series, and DRAG.

The PWC parameterisation is the most general and in principle can achieve the same results as any functional pulse. Defining the PWC parameterisation, we use P parameters, which give a piece-wise constant pulse with periods regulated by a sampling frequency given as $fs = P/T$, where T is the total gate time. This approach is more flexible than any given functional form and can therefore effectively discover novel control schemes. However, it has no bias towards smoothness which can introduce high-frequency Fourier terms and requires many parameters to approximate standard functional pulses.

The truncated Fourier series parameterisation can also approximate any functional pulse, given sufficient parameters. Additionally, it enforces a smoothness condition. However, the optimisation landscape becomes highly non-convex, as each parameter changes the whole time evolution.

The DRAG pulse is a popular choice[183] as it only requires two parameters, which are

tuned to reduce the spectral components of a pulse around the frequency which drives a primary leakage pathway. However, the Gaussian and Gaussian derivative functions do not reach zero, so have to be truncated, leading to significant discontinuities with high-frequency spectral components, long gate times which reduce the fidelity due to decoherence, or the use of a window function.

6.4 Applications, optimal control of hole spin qubits

Here, we choose to demonstrate our algorithm in a semiconductor device, where the spin of a hole is exploited as the quantum degree of freedom. Single qubit gates, for a hole spin confined in a semiconductor quantum dot, can be operated electrically using microwave control pulses characterised by their frequency, phase, and pulse envelopes of their in-phase and out-of-phase components. We therefore perform optimal control of this system by optimising these terms with respect to the average gate fidelity.

6.4.1 Hole spin qubit Hamiltonian

Time-independent Hamiltonian model

To create a simulation in a realistic regime, we use the parameters, measured experimentally in a fin field effect transistor[50]. We show the energy level diagram with respect to the detuning parameter and the magnetic field strength along the x-axis (**Figure** 6.4). We also show the conditional rotational Rabi fringes for the two-qubits, this provides the mechanism for engineering two-qubit CNOT gates. We use perturbation theory to derive an effective Hamiltonian in the spin basis, removing the $(2, 0)$ singlet state to second order. The time-independent spin Hamiltonian is derived in section 2.1.5. In

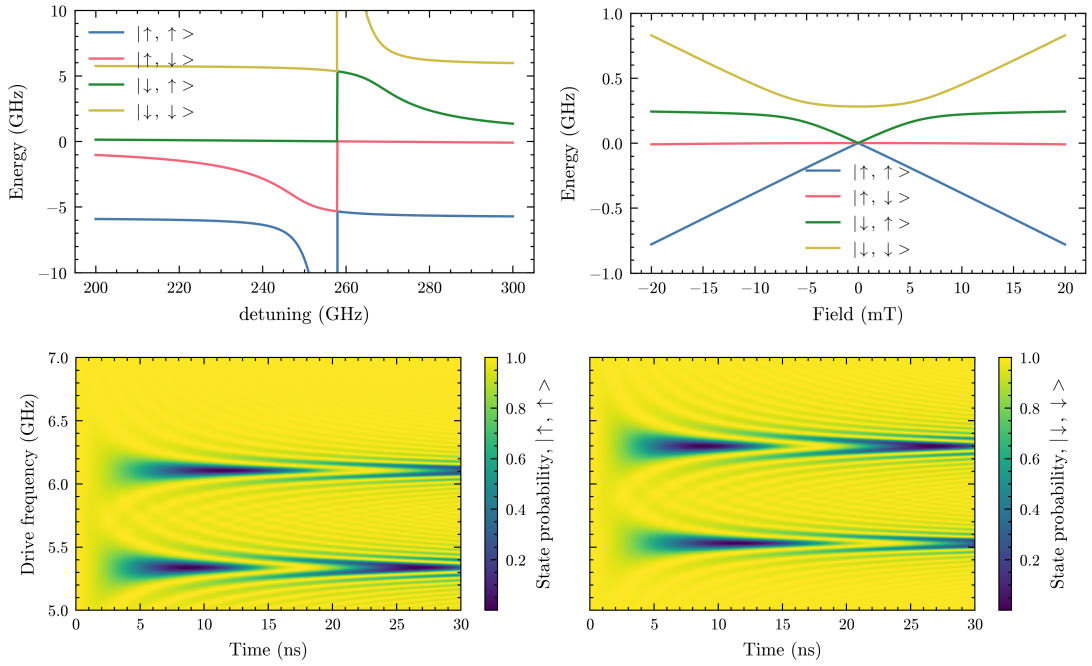


Figure 6.4: Energy diagram and controlled rotations. Upper: The Energies of the spin-states with respect to the detuning and the magnetic field along the x-axis. The Hamiltonian parameters used for these examples are; detuning $\epsilon = 482\,\mathrm{GHz}$, tunnel coupling $t_c = 5.61\,\mathrm{GHz}$, spin-orbit angle $\theta_{SO} = 0.41\pi$, $g_1 = 2.68$, $g_2 = 3.04$, $B = 146\,\mathrm{mT}$. **Lower**: Conditional rotational, Rabi Chevron for a driven two-qubit system initialised in the state $|\uparrow,\uparrow\rangle, |\uparrow,\downarrow\rangle, |\downarrow,\uparrow\rangle, |\downarrow,\downarrow\rangle$ respectively.

the basis for which the local spin Zeeman Hamiltonians are diagonal, it reads[50]:

$$\hat{H}_0 = \frac{\mu_B B}{2}\left(g_1\sigma_z^{(1)} + g_2\sigma_z^{(2)}\right) + \frac{\mathrm{J}_0}{4}\boldsymbol{\sigma}^{(1)}R_{\mathbf{n}}(-2\zeta)\boldsymbol{\sigma}^{(2)}, \tag{6.8}$$

where μ_B is the Bohr magneton, B is the norm of the magnetic field, g_1 and g_2 are the principal g-factors associated with spins 1 and 2, and $\sigma_z^{(1)}$ and $\sigma_z^{(2)}$ are the corresponding diagonal Pauli spin matrices, J_0 is the parameter of the spin exchange interaction and $R_{\mathbf{n}}(-2\zeta)$ is the matrix for the rotation in three dimensions by the angle -2ζ around vector $\mathbf{n}$. The latter arises from general spin-orbit coupling, either intrinsic or artificially implemented[50].

Control Hamiltonian model

Electric Dipole Spin Resonance (EDSR) offers a powerful technique for manipulating spin states in semiconductor quantum dots by leveraging the spin-orbit interaction[184], effectively creating a coupling between the spin states and the electric field. This coupling allows us to modulate the spin states using precisely tailored electric pulses. An IQ-modulated driving signal adds a time-dependent term to the Hamiltonian as follows:

$$\Omega(\theta,t) = I(\theta,t)\Big(v_1\sigma_x\otimes\mathcal{I} + v_2\mathcal{I}\otimes\sigma_x\Big) + Q(\theta,t)\Big(v_1\sigma_y\otimes\mathcal{I} + v_2\mathcal{I}\otimes\sigma_y\Big) \tag{6.9}$$

where v_1, and v_2 scale the signal to account for the lever arms, and the attenuation of qubits 1 and 2 respectively. The oscillating terms

$$I(\theta,t) = A(\theta,t)\sin(\omega t+\phi) \tag{6.10}$$

$$Q(\theta,t) = B(\theta,t)\cos(\omega t+\phi) \tag{6.11}$$

are defined by the envelope functions $A(t)$, and $B(t)$, the phase ϕ, and the frequency ω of the EDSR driving pulse that will be optimised in the following examples.

Noise Hamiltonian model

We develop two noise models which both operate by taking numerical samples to create an ensemble of quantum state trajectories, and we evaluate the gate fidelity by averaging over these ensembles, and over the Haar set of initial states:

i. The first model simulates the dephasing interaction which occurs due to the coupling of the quantum system with a two-state fluctuator. We generate an ensemble of noise trajectories using the interaction term:

$$\hat{H}^i_{\text{noise}}(t) = \eta^i_1(t)\Big(I \otimes \sigma_z\Big) + \eta^i_2(t)\Big(\sigma_z \otimes I\Big), \quad i = 1, 2, \ldots, N \tag{6.12}$$

where we take the stochastic noise processes $\eta^i_j(t)$ for dots $j = 1, 2$ respectively, to have a $1/f$ spectral density function. This pink noise spectrum captures the interaction of the system with a single two-state fluctuator[185,154]. In order to calculate samples from the spectral density function, we first construct a noise spectrum. To do this, we discretise the time domain and randomly generate samples from a Gaussian distribution. We then apply the discrete Fourier transform, creating a white noise frequency spectrum sample. This spectrum is filtered by the spectral density filter function, scaling the amplitude of each sample in the frequency domain by $f^{-1/2}$. The inverse Fourier transform is applied to the filtered frequency spectrum to generate a sample in the time domain of the $1/f$ stochastic noise process.

ii. We then model quasi-static noise, meaning noise that occurs on a time scale which is

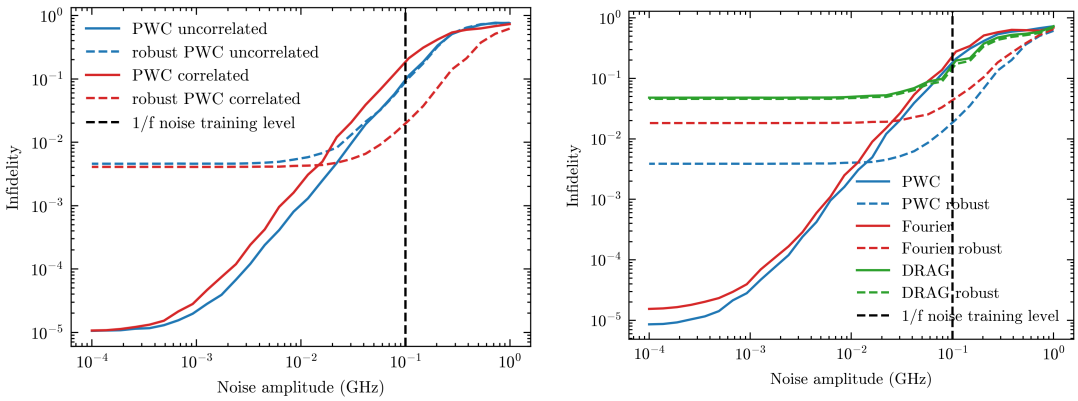


Figure 6.5: Robust 16 ns **single-qubit gates. Left**: Design of robust gates using PWC pulse envelopes (51 parameters for both I and Q components) in the presence of correlated and uncorrelated pink noise trajectories respectively. **Right**: The X-gate infidelity of the optimal pulse envelopes (51 parameters for both I and Q components) with respect to the correlated pink noise amplitude, for standard optimal pulses (solid lines), and robust optimal pulses (dotted lines).

slower than the gate time, such that it is modelled to be stationary throughout the simulated trajectory. We do this by generating a set of samples of a Hamiltonian parameter from a uniform distribution $U(a, b)$:

$$x_i \sim U(a, b), \quad i = 1, 2, \ldots, N \tag{6.13}$$

$$\hat{H}_0^i = h(x_i), \quad i = 1, 2, \ldots, N \tag{6.14}$$

to form an ensemble with time-independent Hamiltonian terms $\hat{H}_0^i$.

6.4.2 Applications

Noise robust single qubit X-gates

In our experiments, the primary focus is to use our algorithm to calculate the optimal pulse shape, frequency, and phase for a set of parameterisations. The quality of the generated quantum gates is evaluated based on the average gate fidelity over a Haar

measure and the distribution of stochastic noise processes. The target unitary, $U_T = \sigma_X \otimes I$, consists of an X-gate on qubit 1, with an identity performed on qubit 2. The optimised pulse sequence is obtained by iteratively updating the pulse parameters to maximise fidelity, considering various constraints such as bandwidth limitations, noise models, and other realistic experimental conditions.

To evaluate the effectiveness of our approach, we design noise robust optimal control pulses for single qubit X-gates. Here, we optimise pulses to minimise the infidelity averaged over sampled pink noise trajectories, for each of the initial states defining the Haar measure. We use PWC pulses with 51 parameters for the I and Q envelope functions respectively, as we find that this is the minimum number where adding parameters does not improve fidelities. We then assess the robustness of the resulting pulse to noise by plotting the mean infidelity in the presence of varying amplitudes of pink noise. Numerically sampled pink noise trajectories simulate the interaction between the system and a 2-state fluctuator with a correlated σ_z interaction term so that $\hat{H}_{\text{noise}}(t) = \eta_1(t)\sigma_z^{(1)} + \eta_2(t)\sigma_z^{(2)}$. The noise can be correlated between the two-qubits $\eta_1(t) \equiv \eta_2(t)$, completely uncorrelated $\eta_1(t) \neq \eta_2(2)$, or some mixture of the two. Here, we investigate these two limits and the affect they have on the achievable infidelity (**figure** 6.5). The correlated noise induces an inferior infidelity, compared to the uncorrelated noise terms. The algorithm is able to design a robust pulse envelope which significantly lowers the infidelity, at high-noise amplitudes, in the case of correlated noise. However, in the presence of purely uncorrelated noise, the robust pulse does not offer any improvement with respect to the pulse trained in the absence of noise. To investigate the further, we use our algorithm to design robust pulses in the presence of correlated noise, comparing piece-wise constant (PWC) and truncated Fourier series

with 51 parameters for I and Q respectively, and Gaussian DRAG pulses. Notably, the PWC pulse demonstrates superior resistance to noise-induced infidelity compared to the other methods. The results indicate that robust PWC and Fourier pulses offers a significant improvement in performance across a wide range of noise amplitudes. This underscores the potential of robust optimal control schemes in maintaining high-fidelity quantum gate operations in the presence of realistic noise environments.

Cross-talk and simultaneous X-gates

In this section, we analyse the effects of cross-talk and extraneous couplings on single-qubit gates and assess the performance of simultaneous single-qubit gates using our framework. The optimisation process aimed to minimise the infidelity of an X-gate on individual qubits in a two-qubit system, as well as the performance of simultaneous X-gates on both qubits. We optimise PWC pulses defined by 51 parameters for the I and Q components respectively, lasting for a duration of 16 ns.

We present (**Figure** 6.6) the infidelity of an X-gate on either qubit 1, $U_T = \sigma_X \otimes I$, or qubit 2, $U_T = I \otimes \sigma_X$, when the other qubit is transformed by an identity operation. The gate infidelity was evaluated by averaging over a two-qubit Haar measure and plotted against the iteration number during the training process. The infidelity of pulses optimised by averaging over only a single-qubit (1Q) Haar measure is reported by the red lines. Subsequently, the pulses were optimised by averaging the gate fidelity over the full two-qubit (2Q) Haar measure (blue lines). The results demonstrate a significant difference in the optimisation trajectories, with the two-qubit Haar measure optimisation yielding a more rapid convergence and lower infidelity, suggesting that it is able to learn to perform an identity on the second qubit by mitigating cross-resonance

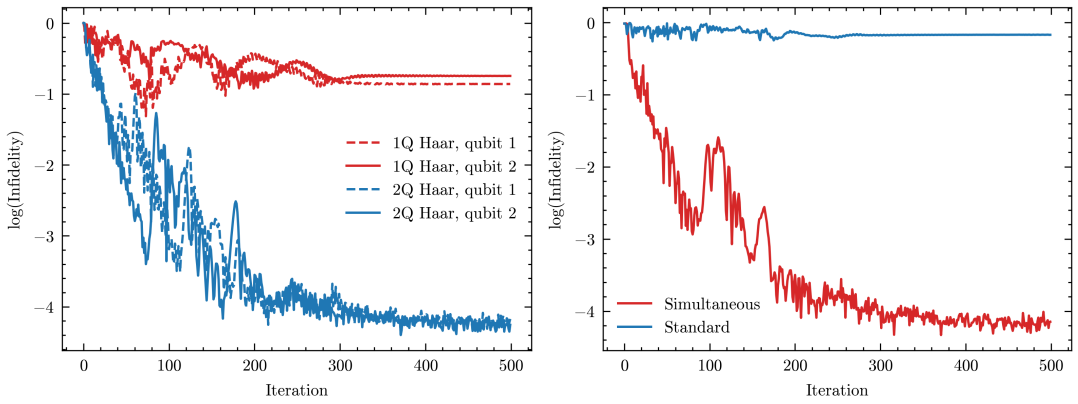


Figure 6.6: Optimisation convergence of single-, and simultaneous single-qubit gates. Left: X-gate infidelity averaged across a Haar measure of the two-qubit system (2Q Haar) for a target unitary with an X-gate on qubit 1 (dashed line) or qubit 2 (solid line) and an Identity operation on the other qubit. We plot the gate infidelity averaged in this way, against the iteration number during optimisation of a $16\,\mathrm{ns}$ PWC pulse parameterised by 51 parameters for the I and Q pulse envelopes. We first evaluate a pulse optimised by averaging the gate infidelity over just the one-qubit Haar measure (1Q Haar, red lines), and then by averaging the gate fidelity over the two-qubit Haar measure (1Q Haar, blue lines). **Right**: We show the gate infidelity of simultaneous X-gates averaged over the two-qubit Haar measure, against the iteration number during training. We compare the relative performance for pulses optimised while being evaluated to perform simultaneous X-gates (blue) as compared to pulses optimised to perform X-gates separately on each target qubit (red).

and extraneous two-qubit coupling interactions between the qubits even when one is nominally idle. This strongly motivates the use of a two-qubit Haar measure even when optimising and evaluating the performance of sequential single-qubit gates to ensure that an identity is performed on the second qubit, to mitigate cross-talk.

We then calculate the infidelity of simultaneous X-gates on both qubits ($U_T = \sigma_X \otimes \sigma_X$), averaged over the two-qubit Haar measure (**Figure** 6.6). Here, we contrast the simultaneous gate infidelity for pulses optimised to perform simultaneous X-gates (blue) with those optimised to perform X-gates separately on each qubit (red). The data reveal that pulses tailored for simultaneous operation exhibit a distinct advantage, converging to significantly lower infidelity than those optimised for individual qubit operations.

The analysis of single- and simultaneous single-qubit gates underscores the complexity inherent in optimising quantum control pulses for multi-qubit systems. It highlights the importance of considering the entire system's state space during optimisation to effectively mitigate cross-resonance and counter-act any two-qubit interactions.

Two-qubit CNOT gates

In quantum control systems, it is crucial to develop pulses that achieve high-fidelity gates and are also resilient to fluctuations in system parameters. To this end, we have engineered control pulses that maintain their fidelity across a distribution of system parameters rather than achieve an optimal performance at a single parameter value. We test this method on a two-qubit CNOT gate, under variations of the detuning parameter, and the magnetic field which are experimentally caused by flux noise and charge noise respectively. The CNOT gates are implemented by utilising the dependence of the resonance frequency of each qubit on the spin state of the other qubit[50,128](**Figure** 6.4).

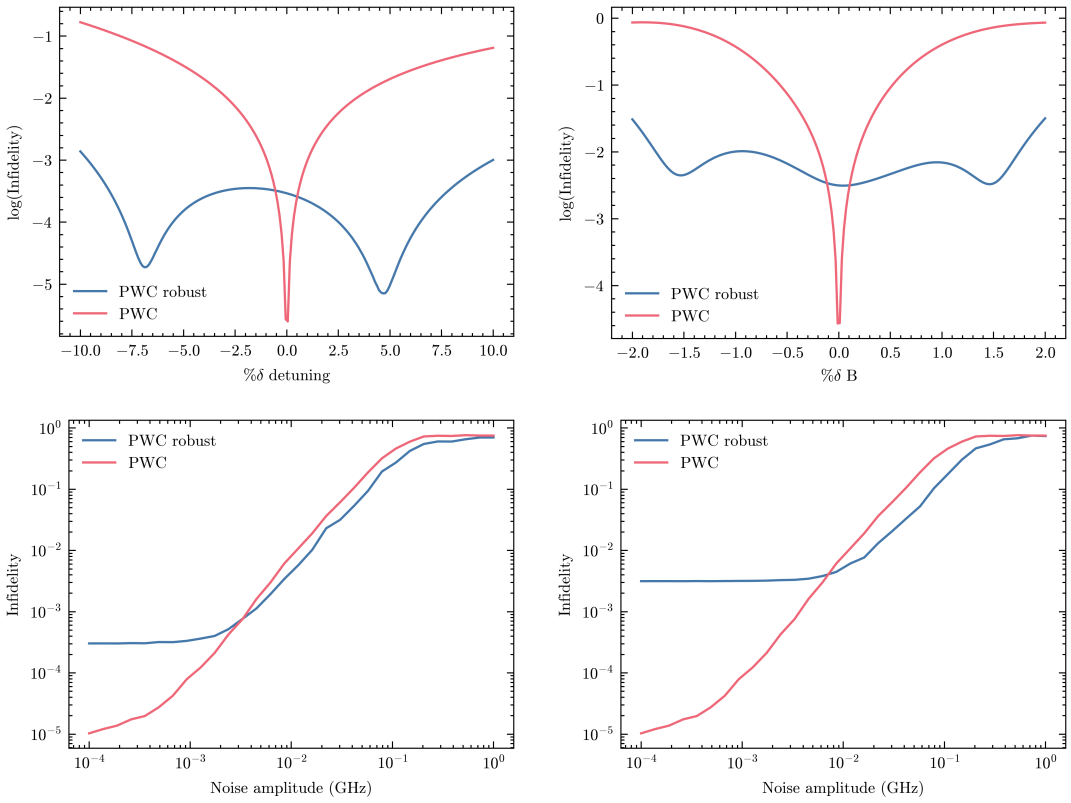


Figure 6.7: Infidelity of robust two-qubit CNOT gates. Left: We present the infidelity of a two-qubit CNOT gate as a function of the percentage error of detuning for robust (blue) and standard (red) optimised PWC pulses (**upper**). The robust pulse is trained by optimising the average gate fidelity over the distribution of the detuning, whereas the standard method only optimises the infidelity for a specific parameter. We then demonstrate that this robust pulse also provides protection with respect to high-frequency $1/f$ noise (**lower**). **Right**: The infidelity as a function of the percentage variation in magnetic field strength for standard and robust pulses. Here, the robust pulse is trained to minimise the average infidelity for samples of a distribution of magnetic field values (**upper**). We then demonstrate that this robust pulse also provides protection with respect to high-frequency $1/f$ noise (**lower**).

Control pulses were trained to be robust against a uniform distribution of Hamiltonian parameters, in order to mitigate quasi-static errors. The robustness criterion was to achieve the lowest possible infidelity averaged across the entire range of these parameters. This approach was compared to pulses optimised at a single sample point of the system parameter. **Figure** 6.7 showcases the infidelity of the resulting pulse over a range of system parameters. The robust pulses (blue lines) did not achieve the same low infidelity at the specific optimised parameter (0% error) as the standard pulses (red lines). However, they demonstrated significantly improved stability over the entire range of parameter values. The inset in each plot emphasises the stark contrast at the optimised parameter value where the standard pulse outperforms the robust pulse. Conversely, the robust pulse consistently maintains lower infidelity, demonstrating its effectiveness in an experimental scenario where exact parameter values cannot be guaranteed. Interestingly, we find that pulses trained to be robust to quasi-static errors are also demonstrably robust with respect to stochastic $1/f$ noise (**Figure** 6.7).

The results highlight the trade-off between achieving the lowest possible infidelity at a specific sample of the Hamiltonian and obtaining a more reliable and stable performance across a distribution of quasi-static errors. The robust training method is validated as it ensures that the quantum control pulses are not overly tailored to an idealised set of parameters but are effective under realistic, variable conditions. This approach is particularly advantageous for quantum computing applications where system parameters cannot be precisely controlled or may drift over time. The trade-off can be adjusted to meet specific experimental requirements by varying the width or type of the distribution in the optimisation process.

6.5 Discussion

In this work, we have presented a comprehensive and automated framework for designing optimal control schemes for quantum systems, with a particular focus on mitigating coherent and incoherent errors in hole-spin qubits. Our approach harnesses the power of modern machine learning tools, specifically automatic differentiation and stochastic gradient-based optimisers, to efficiently navigate the complex landscape of quantum control. Through the development of open-source code, we provide a valuable resource for the quantum computing community to design, optimise, and implement high-fidelity quantum gates. The methods described herein go beyond traditional approaches by encompassing a more realistic set of experimental constraints including bandwidth limitations, rise times, nonlinear transmission functions, and numerically sampled incoherent noise and error processes.

When numerically optimised pulses are subjected to distortions in post-processing or experiments, coherent errors are introduced, significantly limiting gate fidelities. To mitigate pulse distortions in experiments, it's crucial to integrate these constraints directly into the optimisation process. However, this integration and the computation of gradients through these transformations is a primary challenge in quantum optimal control. Conventional solutions have leaned on the inclusion of conditional expressions within the optimisation solver, and optimising numerically using gradients approximated via finite differences. Both these approaches are, however, computationally inefficient motivating our approach of using JAX-based automatic differentiation and optimisation via gradient descent. The coefficients used in applying these constraints, and therefore the resulting optimised pulses, are specific to each combination of device and control stack.

The robustness of a quantum gate is quantified by evaluating the averaged gate fidelity ($F_{robust} = \langle F_{avg} \rangle$) over a numerically sampled ensemble in the Schrödinger picture. Therefore, robust quantum gates can be designed by optimising the gate fidelity over such an ensemble average. We have presented two approaches for generating numerical samples to improve the robustness of quantum gates, which account for dynamically fluctuating noise interactions, and quasi-static errors. In the first method, we simulate the effect of a coupled two-state fluctuator on the quantum system, by numerically sampling $1/f$ noise trajectories and computing the time-evolution[185]. Secondly, we mitigate the effect of quasi-static errors by numerically sampling Hamiltonian parameters over some distribution, evolving the quantum system for each sample, and optimising the average gate fidelity over this ensemble.

Focusing on a two-qubit hole spin system, we have successfully designed optimal control for single-, simultaneous single-, and two-qubit gates. The optimal control pulses significantly reduce coherent errors and incorporate robustness to incoherent errors, a critical impediment in the path to reliable quantum computing. In the first example we illustrate the necessity of optimising specifically for simultaneous single-qubit gates, and of computing average gate fidelities over a two-qubit Haar measure even for standard single-qubit gates in order to mitigate unwanted two-qubit coupling interactions and cross-resonant driving. We then demonstrate the design of optimal control pulses which are robust with respect to the interaction of the system with a two-state fluctuator, by numerically sampling $1/f$ noise trajectories. Lastly, we find that two-qubit CNOT gates optimised to be robust with respect to quasi-static errors also demonstrate robustness with respect to a $1/f$ stochastic noise process. The discovery of robust control pulses which provide resistance to both high- and low-frequency noise processes is notable, and

deserved of further investigation.

Future work will aim to expand upon the versatility of our framework to encompass a wider array of qubit systems and control scenarios. Additionally, we will explore the integration of our methodology with error correction protocols to further enhance the robustness of quantum operations. The ultimate goal remains to facilitate the realisation of scalable and fault-tolerant quantum computing, and our research marks a significant stride towards that horizon.

Chapter 7

A differentiable digital twin: Using Neural CDEs to simulate quantum systems

"*Is there no way out of the mind?*"

Sylvia Plath

Quantum system emulation and optimal control

Problem statement: Noisy quantum system emulation and pulse optimisation in the absence of a known Hamiltonian model.

Solution: Learning a model for the dynamics of a noisy quantum system using Neural Controlled Differential Equations. These models are universal function approximators with a structural bias to systems with dynamics captured by a set of unknown controlled differential equations. This differentiable model can be used for optimal quantum control as in chapter 6. Building on recent advances in machine learning and optimal quantum control, this framework bridges the gap between physical experiments and incomplete theoretical models.

7.1 Introduction

Coherent and incoherent noise sources both contribute to errors in quantum circuits but in fundamentally different ways. The challenge in reducing errors spans the dual space of noise sources in quantum circuits: coherent (unitary) and incoherent (statistical). Coherent noise arises from factors like imperfect calibration of gates and non-linear transfer functions, while incoherent noise sources are more stochastic, influenced by random fluctuations in qubit control parameters and coupling to external degrees of freedom. The behaviour of these noise types is complex; for instance, errors due to coherent noise grow quadratically with the number of gates, unlike those caused by incoherent noise, which grow linearly[186].

Optimal control schemes are defined to optimally limit coherent errors. GRAPE (Gradient Ascent Pulse Engineering) has been a cornerstone in the quantum control community for optimising quantum systems through gradient descent methods[187]. It gained favour

for its analytical approach to gradient calculations, providing computational benefits over prior techniques reliant on finite differences. However, modifying GRAPE to suit realistic system constraints has remained a significant challenge. While differentiable analytic models like GRAPE have made significant strides in optimising quantum control schemes, they often fall short of capturing the full complexity of real-world quantum systems. One critical limitation is the insufficient understanding and modelling of various noise sources that influence quantum behaviour, due to experimental conditions and environmental interactions. This deficiency leads to a gap between theoretical predictions and experimental results, limiting the efficacy of control methods derived from such models. This underscores the need for more comprehensive approaches that can incorporate realistic constraints and unknown transformations in quantum systems.

The field of simulating and controlling quantum systems has seen the influx of various machine learning techniques to tackle its intrinsic challenges. Lin et al.[188] exploited Long Short-Term Memory Recurrent Neural Networks (LSTM-NNs) combined with the bootstrap method to simulate the long-time dynamics of open quantum systems. Their work highlights the high accuracy and computational efficiency of LSTM-based approaches, particularly for the time evolution of quantum states. Similarly, Flurin et al.[189] implemented a recurrent neural network to reconstruct quantum dynamics based on physical observations, without specifying the a priori rules of quantum mechanics. The study demonstrates the robustness of neural networks in parameter estimation, tomography, and quantum control, emphasising their practicality in real-time implementations. Choi et al.[190], contributed a package they termed QNODE, which uses a latent Neural Ordinary Differential Equation (NODE) to simulate closed and open quantum systems. QNODE was able to extrapolate beyond its training region while

abiding by quantum-mechanical laws like Heisenberg's uncertainty principle, demonstrating the promise of latent NODEs for understanding quantum phenomena. Neural ODEs are highly expressive, capable of approximating complex functional relationships, and inherently differentiable, making them amenable to gradient-based optimisation algorithms. However, the evolution of ODEs is completely determined by their initial conditions so we cannot incorporate data that changes as the system evolves.

In contrast to these methods, our approach aims to utilise Neural Controlled Differential Equations (Neural CDEs)[178], which offer a promising path towards a learnable and differentiable digital twin. This method harnesses the mature field of controlled differential equations[178] to enable the simulation of dynamic systems which evolve under a set of ordinary differential equations subject to a continuously varying control field. The integration of Neural CDEs promises several advantages. Primarily, the evolution of the system is no longer determined completely by the initial conditions. Instead, control inputs modulate the behaviour of the system, enabling the model to capture the time evolution of quantum systems under some dynamic control. These models are computationally efficient, with a speedup that is supported by modern advances in machine learning computation, such as just-in-time compilation, vectorisation, and GPU acceleration[178]. They use neural networks, universal function approximators capable of encapsulating the non-linear dynamics inherent in quantum mechanical systems, but also contain an inductive bias towards dynamics captured by a set of CDEs. Neural CDEs also exhibit robustness against noise and errors owing to their continuous nature, an attribute vital for capturing the complexities of quantum mechanics where coherent and incoherent noise factors have a significant impact[178]. Lastly, a pivotal advantage is the inherent differentiability of these models, which facilitates the seamless applica-

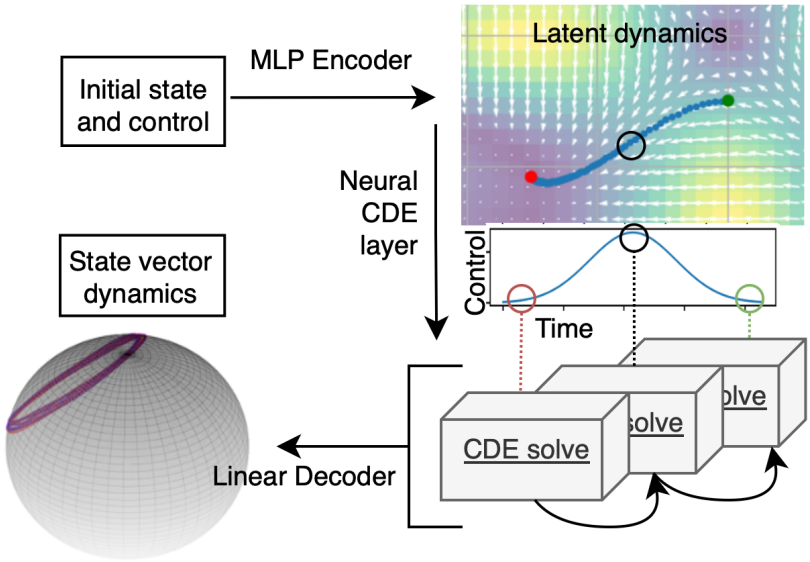


Figure 7.1: Neural CDE algorithm schematic for emulating the evolution of quantum systems with variable control pulses. The encoder module, consisting of a learnt function, parameterised by a neural network (Multi Layer Perceptron (MLP)), which maps from the state vector space to an initial state in the latent state space. The initial value of the control term is also embedded in this latent state space by the encoder. The Neural CDE layer transforms the state, in the latent state space, using learnt dynamics consisting of a mapping function (parameterised by a neural network) embedded within a differential equation solver. This mapping function uses the current state of the system in latent space, and the gradient of the variable control term, to calculate some time evolution in the latent space. The decoder module is a learnt linear function which maps the final state in the latent space, back to the original vector state space. The decoder module then outputs the state evolution as a time series of the expectation values of the Pauli operators.

tion of gradient-based optimisation techniques for designing optimal control schemes via methods such as stochastic gradient descent.

7.2 Neural Controlled Differential Equations

Neural CDEs extend the traditional neural network architecture by incorporating continuous-time dynamics. In this paradigm, neural networks are used to approximate the evolution of a system governed by a set of CDEs. The general form of a controlled differential

equation for a n dimensional state z of a system is

$$\frac{dz}{dt} = A(t)z + B(t)X, \tag{7.1}$$

where A is an $n \times n$ matrix, B is an $n \times v$ matrix for some function $X : \mathbb{R} \to \mathbb{R}^v$ which is the control. Specifically, Neural CDEs are composed of an encoder-decoder structure where an input state is mapped to a latent space by a learnt encoder function, the state is transformed in the latent space by a network approximating the dynamics of the CDE, and a learnt decoder function maps the final state from the latent space back to the original space. The parameters defining the neural networks of the Neural CDEs are optimised during training, fine-tuning the model to predict the evolution of the state of a system under varying control schemes (**Figure** 7.1).

Neural CDEs operating in an auto-encoder structure, construct a map $x \mapsto y$ by learning functions ζ_θ, f_θ, and a linear map l_θ, such that:

$$z(0) = \zeta_\theta(t_0, x_0), \qquad \text{Encode} \tag{7.2}$$

$$\frac{dz}{dt}(t) = f_\theta\left(z(t), \frac{dX}{dt}(t)\right), \qquad \text{transform CDE} \tag{7.3}$$

$$y(t) \approx l_\theta(z(t)), \qquad \text{Decode} \tag{7.4}$$

where z is a hidden state, and we let $x = ((t_0, x_0), \ldots, (t_n, x_n))$ with $t_i \in \mathbb{R}$ and $x_i \in \mathbb{R}^v$. We can interpolate to make $X : [t_0, t_n] \mapsto \mathbb{R}^{1+v}$ such that $X(t_i) = (t_i, x_i)$. Importantly, X includes time and therefore has $v + 1$ dimensions.

7.3 Network architecture

The functions ζ_θ (encoder), f_θ (CDE function), and linear map l_θ (decoder) can be parameterised by neural networks. We can use the memory-efficient adjoint method[178] and automatic differentiation frameworks to train these networks to approximate the target transformations. The size of these networks is tabulated below:

Layer (type:depth-idx)	**Non-linearity**	**Param #**
NeuralCDE	–	–
Encode:	–	–
Linear: 1-2	relu	24
CDE function: 1-1	–	–
Linear: 2-1	relu	1,152
Linear: 2-2	tanh	2,064
Decode:	–	–
Linear: 1-3	relu	27
Total params:	–	3,267

7.4 Simulating quantum systems

7.4.1 Singlet-Triplet qubits

For a quantum system with a small number of states, the time evolution of the dynamics of the system can be simulated classically. Given the Hamiltonian of the system $H(t)$, the time evolution of the state $|\psi(t)\rangle$ is governed by the time-dependent Schrödinger equation. For a Singlet-Triplet qubit the effective Hamiltonian of the qubit subsystem, in the $\{|S\rangle, |T_0\rangle\}$ basis, is:

$$H = \begin{pmatrix} -J(\epsilon) & \frac{\Delta g \mu_B B}{2} \\ \frac{\Delta g \mu_B B}{2} & 0 \end{pmatrix} \tag{7.5}$$

where $J(\epsilon)$ is the detuning dependent exchange energy, Δ_g is the difference in the g-factors of the two spins, and B is the external magnetic field. All of the elements in the Hamiltonian are electrically tunable. Pulsing on the detuning, ϵ, influences $J(\epsilon)$ as well as the ratio between $J(\epsilon)$ and $\Delta g \mu_B B$, which in turn determines the rotation axis tilted by an angle $\theta = \arctan(\frac{\Delta g \mu_B B}{J(\epsilon)})$ from the Z-axis. Hence, for large detuning $\theta \approx \frac{\pi}{2}$ corresponding to X-rotations, while for small detuning $\theta \approx 0$ enabling Z-rotations.

We determine the angle and frequency of the rotations[27], using the tunnel coupling t_c, as follows

$$J(\epsilon) = |\frac{\epsilon}{2} - \sqrt{\frac{\epsilon^2}{4} + 2t_c^2}| \tag{7.6}$$

$$\theta = \arctan(\frac{\Delta g \mu_B B}{J(\epsilon)}) \tag{7.7}$$

$$\omega = \frac{1}{\hbar}\sqrt{J(\epsilon)^2 + (\Delta g \mu_B B)^2}. \tag{7.8}$$

7.4.2 Model training

We train the Neural CDE on trajectories using variable control fields for the Singlet-Triplet qubit in the Schrödinger picture (**Figure** 7.2). We parameterise the Hamiltonian using data from a hole spin singlet-triplet qubit in a SiGe device[27]. We generate training data with our Hamiltonian model by applying random control fields (values of the detuning) and simulating the evolution of the state of the singlet-triplet qubit in the computational basis. We add Gaussian noise, with a variance of 0.01, to the expectation values of $\{\sigma_X, \sigma_Y, \sigma_Z\}$ and record this data to train the Neural CDE. Here, the loss function is the MSE of the time series of the $\{\sigma_X, \sigma_Y, \sigma_Z\}$ projections of the qubit state. The Neural CDE therefore learns to output the evolution of these expectation values

over time under a variable detuning control field. To improve generalisability we train the network on time-dependent controls generated via: random Fourier terms; random constant pulses; random pink noise trajectories; and random white noise trajectories. In total, we used 5000 trajectories of 30 ns and trained the model for 2000 epochs on a Nvidia GP102 GPU.

7.4.3 Results

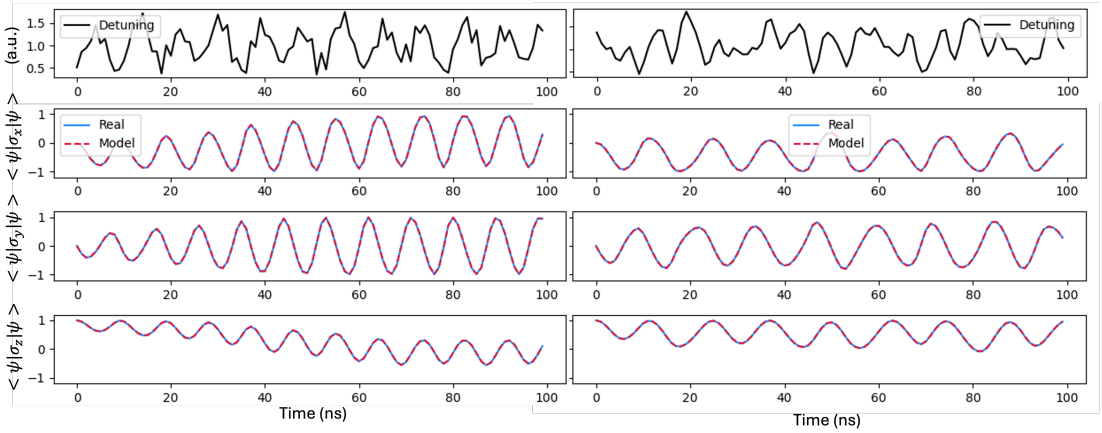


Figure 7.2: Dynamics comparison: The Pauli operator expectations, for the time dynamics as calculated by the Neural CDE (dashed red), and the time-dependent Schrödinger equation (solid blue) under a detuning pulse with random Gaussian and Fourier features (black line).

We now examine how accurately the Neural CDE can reproduce the dynamics of the physics model, and compare the computational efficiency of this method to the time-dependent Schrödinger evolution. To evaluate the learnt dynamics of the Neural CDE, we evolve a Singlet-Triplet qubit system using the time-dependent Schrödinger equation and compare the resulting trajectories (**Figure** 7.2). In all cases, we observe a close correspondence for the first 60 ns. This is a significant result, given that all the training data consisted of 30 ns trajectories. This performance beyond the scope of the training

indicates that the Neural CDE has learnt a general model for the dynamics of the quantum system, even with the measurement noise added to the training data.

We now generate a two-dimensional plot to show the effect of the control amplitude on the time evolution predicted by the Neural CDE. We vary the time and amplitude of a square detuning control pulse and report the projection of the state-vector onto the z-axis, again comparing to the results given by the evolution of the dynamics predicted by solving the time-dependent Schrödinger equation (**Figure** 7.3). Here we use the JAX code, and the implicit just-in-time compilation and vectorisation, developed for the previous chapter to accelerate the evaluation of the time-dependent Schrödinger equation. We also sweep the time-step duration of the ODE solver and report the mean-squared-error in the final state prediction averaged over this two-dimensional plot for both the Neural CDE and the Schrödinger equation. We then plot the wall-clock time required to achieve a certain level of accuracy. The Neural CDE offers a computational speedup with respect to the evolution calculated using the time-dependent Schrödinger equation.

7.5 Optimal quantum control

The fidelity of a quantum gate defines the closeness of the input-output state mapping to that of the ideal unitary operation. The problem of optimal quantum control is to find a function $u(t)$ which evolves the system to a desired target state or implements a target transformation. For a transition to a target state $|\psi_{target}\rangle$, the solution can be

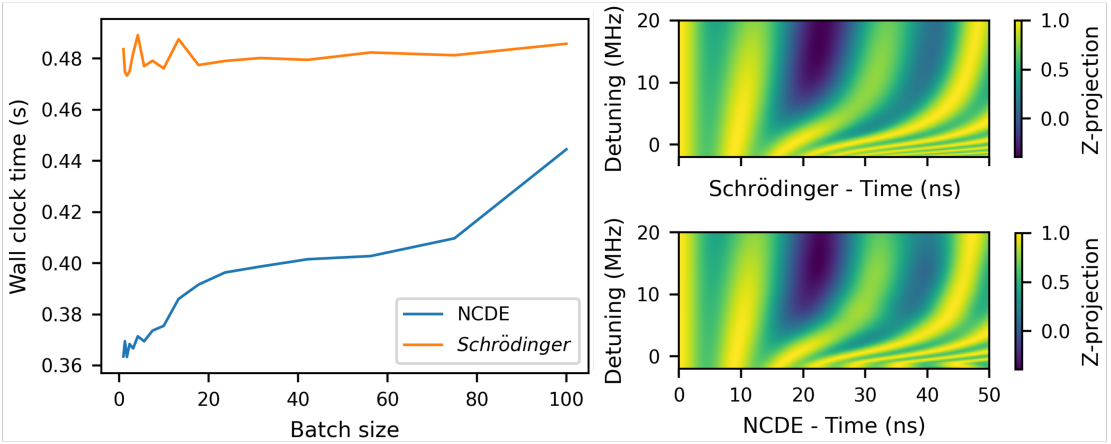


Figure 7.3: Comparison to time-dependent Schödinger equation. Left: The wall-clock time as a function of the batch size evaluated for the evolution of the Schödinger equation (orange), and calculated using the NCDE (blue). **Right:** A comparison of the time evolution of a singlet-triplet qubit, as a function of the time and detuning amplitude of the pulse envelopes for the Schödinger model (top) and the Neural CDE model (bottom).

found by minimising the loss function,

$$L\Big(u(t)\Big) = 1 - \frac{1}{M}\sum_{i}^{M} |\left\langle \psi_{\text{target}} \right| U\Big(u(t)\Big) \left| \psi_{\text{i}} \right\rangle|^2 \tag{7.9}$$

$$= 1 - \frac{1}{M}\sum_{i}^{M} |\left\langle \psi_{\text{target}} \right| |l_\theta\Big(z(\psi_i, t)\Big)\rangle|^2 \tag{7.10}$$

where $z(\psi_i, t)$ is found by solving the CDE $\frac{dz}{dt}(t) = f_\theta\left(z(t), \frac{du}{dt}(t)\right)$ with $z(0) = \zeta_\theta(t_0, x_0)$. The fidelity of the process mapping is calculated by averaging the state fidelity over a set of M initial states, $\Psi(t=0) = |\psi_{\text{initial}}\rangle$, which define a Haar measure for the system. In most cases of interest, an analytical solution to this control problem does not exist, hence optimal control solutions must be found numerically. We, therefore, propose a method to numerically determine optimal control schemes for implementing high-fidelity quantum gates by minimising L with respect to $u(t)$. Here, the values of the weights and biases θ in the networks $\{l_\theta, f_\theta, \zeta_\theta\}$ can be frozen and then the gradients of the

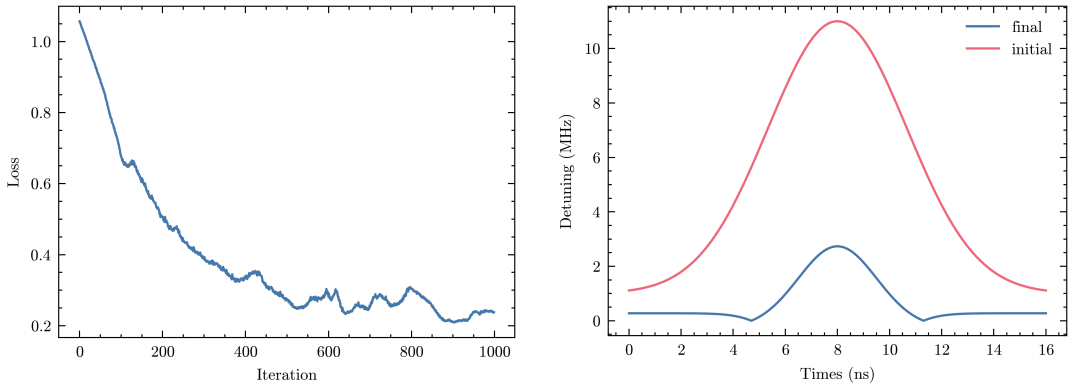


Figure 7.4: Quantum optimal control with Neural Controlled Differential Equations. We use automatic differentiation through the learnt NCDE model and the ADAM optimiser to converge upon an optimal control pulse for the $X_{\pi/2}$ target unitary operation with respect to the ground initial state. **Left:** We plot the convergence of the loss function with sequential iterations of the optimisation loop. **Right:** We plot the initial guess for the pulse envelope and the envelope optimised using our method (final).

parameterised function $u(t)$ calculated, to learn optimal control schemes via stochastic gradient descent. The gradient of a loss function L with respect to the control input $u(t)$ can be expressed as:

$$\frac{\partial L}{\partial u} = \frac{\partial L}{\partial y} \cdot \frac{\partial y}{\partial z} \cdot \frac{\partial z}{\partial u}, \tag{7.11}$$

where $\frac{\partial L}{\partial y}$ represents the gradient of the loss with respect to the output y, $\frac{\partial y}{\partial z}$ is the gradient of the output with respect to the internal state z, obtained during the decoding process, $\frac{\partial z}{\partial u}$ denotes the gradient of the internal state with respect to the control input u, captured through the transformation CDE. This approach enables the optimisation of control inputs by iteratively adjusting u to minimise the loss function L, thereby learning the optimal control pulses for the system. Here, we perform an

$$X_{\pi/2} = \begin{bmatrix} \frac{1}{\sqrt{2}} & -i\frac{1}{\sqrt{2}} \\ -i\frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{bmatrix} \tag{7.12}$$

gate, with respect to the ground initial state, by optimising the parameters (a, b, c) of a Gaussian pulse envelope

$$f_{control}(t) = a + be^{-(t-t_{max}/2)^2/(2c^2)}. \tag{7.13}$$

We calculate the loss function by evaluating the mean squared error between the prediction of the final state of by the Neural CDE and the expectation values for the Pauli operators of the target state (**Figure** 7.4). We use the ADAM optimiser[181] to converge upon the optimal control pulse for implementing the target gate in this system.

7.6 Discussion

In this chapter, we have delved into the concept of a differentiable digital twin for quantum systems, an approach that leverages Neural Controlled Differential Equations for optimal quantum control. This methodology not only lends an advantage to computational efficiency and precision but also introduces a novel paradigm for simulating and controlling quantum systems. By framing quantum systems within the context of a differentiable digital twin, we highlight the utility of iterative refining and adaptation of control strategies in response to dynamic changes and uncertainties inherent in quantum systems. This approach highlights a critical advancement over traditional methods, offering a more flexible and accurate model for predicting quantum behaviour.

Despite these promising results, we acknowledge certain limitations in our study, principally the need for testing and evaluation on real experiments, and the comparison of optimal control solutions to other methodologies and qubit systems. These challenges offer fertile ground for future research, which could explore training Neural CDEs on a

real device and testing the developed control strategies on the same device.

Chapter 8

Conclusion

"*Our blue shadows are enormous. We move in a gigantic, joyous world. A tall pillar on the corner is tightly swathed in wet canvases, across which a paintbrush has scattered coloured whirlwinds*".

Vladimir Nabokov

This work set out to develop methodologies to automate the control and optimise the performance of spin qubits in semiconductor quantum dot devices. I have developed three key results during the course of this research, contributing algorithms which in turn: efficiently locate a potential qubit operational regime within the device gate voltage space; automatically optimise the qubit properties; and design optimal control pulses for implementing high-fidelity qubit gate operations. In this section, I will discuss the general applicability of such methods towards the automation of scientific experimentation, specifically in the domain of semiconductor spin qubits. Following this, I will delineate the potential limitations of the methods presented in this thesis, and outline some problems that have been beyond its scope. Lastly, I propose areas for future investigation.

8.1 Reflections on automated scientific control and discovery

The methods utilised in the main chapters of this thesis have significant technical implications. Firstly, deep reinforcement learning emerges as an effective approach for estimating gradients through black-box physical phenomena. The estimated gradients can be leveraged to learn strategies to enable the automation of the discovery of desired physical regimes. This is exemplified by its successful application in automatically identifying bias triangles. Secondly, Bayesian inference and experimental design prove resilient against experimental noise. They play a crucial role in minimising experimental overhead and offer a reliable analysis of uncertainty in parameter estimates. Thirdly, Bayesian optimisation stands out as an effective black-box optimisation method, particularly adept at handling time-intensive and noisy experimental data. Contemporary machine learning programming tools, such as JAX, greatly enhance the utility and im-

pact of physics simulations. Their vectorisation, just-in-time compilation, and GPU compatibility significantly reduce computation time, thus expanding the horizon of possible investigations. Additionally, their automatic differentiation features facilitate the theoretical optimisation of physical phenomena, such as pulse envelopes for quantum gates. Efficient computation of gradients also enables parameter inference via stochastic gradient descent, effective in cases such as Hamiltonian parameter estimation. Lastly, Neural CDEs offer a powerful paradigm for training differentiable models of physical systems. These networks can be pre-trained using simulations and then fine-tuned on real devices. The advantage of this method is twofold: it allows the networks to be updated as device parameters shift over time, and it maintains the capabilities of automatic differentiation. Consequently, this enables the optimisation of control schemes through stochastic gradient descent.

8.2 Reflections on spin qubits

Spin qubits in quantum dots, particularly hole spin qubits in silicon, represent a significant area of advancement in quantum computing, offering both unique advantages and specific challenges. Following are some reflections, exploring the advantages and disadvantages of a silicon spin-based quantum computing technology stack.

Spin qubits are highly compatible with foundry fabrication processes, making them a practical choice for scalable quantum computing. They have a small device footprint, which is advantageous for dense integration. They can also be interfaced with cryo-CMOS control electronics to enable multiplexed driving[191] and readout[192] of multiple qubits, a crucial feature for complex quantum computations. Another remarkable advantage is the potential for "hot" spin qubits, which can operate at higher temperatures

(around 4K), significantly reducing the resources and infrastructure required for the construction and operation of a quantum computing system. This higher temperature operation increases the cooling power at the device stage, allowing for more robust driving and readout capabilities. Furthermore, spin qubits are known for their fast operation and relatively high coherence times[17]. Hole spin qubits in silicon, in particular, stand out due to several distinct benefits. They exhibit a smaller hyperfine interaction, attributable to the low number of nuclear charges in enriched silicon and the p-orbital characteristics of valence shell hole vacancies. This aspect is crucial for maintaining coherence in quantum states. Hole spin qubits demonstrate strong electrical tunability of spin qubit properties[154]. This tunability is a valuable asset for precise control of qubit behaviour. Moreover, the high spin-orbit interaction in hole spin qubits enables all-electrical ultra-fast spin operations, a key requirement for efficient quantum computing.

Despite these advantages, spin qubits face several challenges. The tune-up process of devices can be complex and demanding. They are highly susceptible to charge noise due to the localisation of quantum states and the lack of inherent noise protection mechanisms. Furthermore, the small length scales involved pose a challenge in separately reading out individual quantum dots. There is also the critical challenge of achieving a high on-off ratio for two-qubit interactions, essential for effective quantum gate operations.

In summary, while spin qubits, and specifically hole spin qubits in silicon, present a promising pathway in the development of quantum computers, their full potential can only be realised by overcoming certain inherent challenges, particularly related to improvements in device fabrication, modelling, automation and control.

8.3 Summary of results and applications

The first key result, presented in "**Deep Reinforcement Learning for Efficient Measurement of Quantum Devices**" demonstrates an effective method for automatically discovering bias triangle charge transport features in a planar GaAs device. I employ a statistical current metric to inform the DRL agent, based on duelling deep Q-networks, and use a computer vision algorithm, a Convolutional Neural Network, to classify bias triangles. This method yields a 10x speedup in measurement time alone, compared to the grid-scan which is conventionally used in manual tune-ups and is fully automatic. I also demonstrate an advantage from using the DRL agent, as compared to a random decision agent, or a Nelder-Mead numerical optimiser. The DRL agent can update its policy in real-time and can be trained to achieve other objectives, such as incorporating more gates into the optimisation or discovering other physical features.

The framework presented in Chapter 5, "**Bayesian methods for the automatic discovery of 'sweet-spot' qubit operating regimes**", demonstrably classifies regimes within the device parameter space where spin qubit procession is observable, and optimises the qubit properties, such as the Rabi frequency, within this regime. Here, I have successfully employed the algorithm in hole-spin systems, optimising the Rabi frequency in Si finFET Loss DiVincenzo qubit, and discovering a "sweet-spot" for the Q-factor in a planar Si/SiGe Singlet-Triplet qubit. These optimal operational regimes are discovered automatically using a Bayesian framework which is robust with respect to noisy measurements and model uncertainty.

I have contributed an algorithm, and associated open source software suite, in "**Robust optimal quantum control for high-fidelity gates with spin qubits**" to develop

high-fidelity quantum gates. The algorithm enables the design of optimal quantum control schemes, with modules which effectively mitigate both coherent and incoherent gate errors. The code implementation uses vectorisation, and just-in-time compilation to reduce the wall-clock time of simulations of the time dynamics of quantum systems, in either the Schrödinger or the Lindblad picture. JAX also enables all functions to be automatically differentiable, and this feature is utilised to calculate gradients through the simulation and optimise control schemes using stochastic gradient descent. Incoherent errors are mitigated, by optimising the average fidelity with respect to numerically sampled noise trajectories, or distributions of system parameters. Coherent errors are also reduced, by including pulse distortions and similar effects within the optimisation loop. I include examples which show these advantages within a hole-spin qubit system, applied to single-, simultaneous-single, and two-qubit gates.

Chapter 7, **"A differentiable digital twin: Using Neural CDEs to simulate quantum systems"**, extends the work of Chapter 6 by introducing a novel method for training a differentiable model of a quantum system using Neural CDEs. This provides a mechanism to bridge the gap between incomplete physics models and real experiments, using a universal function approximator with an inductive bias to emulating evolving systems governed by controlled differential equations, such as driven quantum systems. These models can then be leveraged, by using stochastic gradient descent, to optimise control schemes to achieve high-fidelity quantum gates.

8.4 Limitations and future research

In the concluding reflections of this thesis, it is essential to acknowledge and examine the limitations of the methods presented across its chapters, along with potential strategies

for future enhancements.

The methodologies applied in the first two approaches (Chapters 4 and 5) are characterised by their black-box nature. The first method, the DRL agent, learns a model implicitly before deployment, while the second, Bayesian optimisation, adapts and learns during the deployment phase. An integration of these approaches could potentially offer a more comprehensive strategy. However, a key limitation here is the absence of physics-based models within these algorithms. Incorporating such models could lead to more efficient strategies, though it necessitates a deeper understanding of the device's physics. Furthermore, this integration calls for devices with fewer defects and more consistency with theoretical models, thereby underscoring the need for improved fabrication techniques.

In contrast, the third method employs a model-based approach (Chapter 6). The primary challenge with this method lies in the discrepancies between the theoretical model and real-world experimental conditions. While refining the models and employing cleaner devices can mitigate this issue to some extent, additional strategies are required to bridge the gap fully. In this context, two approaches are recommended. First, employing a numerical optimiser, such as Bayesian optimisation, coupled with the randomised benchmarking-based ORBIT optimisation metric[193], could fine-tune the control envelopes suggested by the algorithm on actual devices. Second, implementing a machine-learning paradigm to implicitly learn the discrepancies between simulation and experiment offers a promising avenue of research. This could involve using a deep reinforcement learning agent or a physics-inspired method. Particularly, Neural Controlled Differential Equations (Neural CDEs) are advantageous (Chapter 7) due to their inherent inductive bias towards modelling systems governed by ordinary differen-

tial equations, akin to closed or open quantum systems. However, this method requires extensive training of the neural networks and has yet to be demonstrated on a real experiment.

In conclusion, while the methods presented in this thesis mark significant advancements in the field of quantum computing, their limitations also pave the way for future research and development. Integrating physics-based models, refining fabrication processes, and employing advanced machine-learning techniques are crucial steps toward developing more accurate and efficient quantum computing technologies.

Bibliography

[1] Hugh Everett. "Relative state" formulation of quantum mechanics. *Reviews of Modern Physics*, 29(3):454–462, 1957.

[2] David Deutsch (University of Oxford). Quantum theory, the Church–Turing principle and the universal quantum computer. *The Royal Society A*, 400(1818), 1985.

[3] P.W. Shor. Algorithms for quantum computation: discrete logarithms and factoring. pages 124–134, 2002.

[4] Francesco A. Evangelista and Victor S. Batista. Editorial: Quantum computing for chemistry. *Journal of Chemical Theory and Computation*, 19(21):7435–7436, 11 2023.

[5] Amine Zeguendry, Zahi Jarir, and Mohamed Quafafou. Quantum machine learning: A review and case studies. *Entropy*, 25(2), 2023.

[6] 40 years of quantum computing. *Nature Reviews Physics*, 4(1):1–1, 2022.

[7] Daniel Loss, David P DiVincenzo, and P DiVincenzo. Quantum computation with quantum dots. *Phys.\ Rev.~A*, 57(1):120–126, 1997.

[8] John Clarke and Frank K. Wilhelm. Superconducting quantum bits. *Nature*, 453 (7198):1031–1042, 2008.

[9] Morten Kjaergaard, Mollie E. Schwartz, Jochen Braumüller, Philip Krantz, Joel I.-J. Wang, Simon Gustavsson, and William D. Oliver. Superconducting qubits: Current state of play. *Annual Review of Condensed Matter Physics*, 11(1):369–395, 2020.

[10] A. Steane. The ion trap quantum information processor. *Applied Physics B: Lasers and Optics*, 64(6):623–642, 1997.

[11] Iulia Georgescu. Trapped ion quantum computing turns 25. *Nature Reviews Physics*, 2(6):278–278, 2020.

[12] Karen Wintersperger, Florian Dommert, Thomas Ehmer, Andrey Hoursanov, Johannes Klepsch, Wolfgang Mauerer, Georg Reuber, Thomas Strohm, Ming Yin, and Sebastian Luber. Neutral atom quantum computing hardware: performance and end-user perspective. *EPJ Quantum Technology*, 10(1):32, 2023.

[13] A. P. Nizovtsev, S. Ya Kilin., F. Jelezko, T. Gaebal, I. Popa, A. Gruber, and J. Wrachtrup. A quantum computer based on NV centers in diamond: Optically detected nutations of single electron and nuclear spins. *Optics and Spectroscopy (English translation of Optika i Spektroskopiya)*, 99(2):233–244, 2005.

[14] L. M. K. Vandersypen, H. Bluhm, J. S. Clarke, A. S. Dzurak, R. Ishihara, A. Morello, D. J. Reilly, L. R. Schreiber, and M. Veldhorst. Interfacing spin qubits in quantum dots and donors—hot, dense, and coherent. *npj Quantum Information*, 3(1):34, 2017.

[15] P. F. Braun, X. Marie, L. Lombez, B. Urbaszek, T. Amand, P. Renucci, V. K. Kalevich, K. V. Kavokin, O. Krebs, P. Voisin, and Y. Masumoto. Direct observation of the electron spin relaxation induced by nuclei in quantum dots. *Physical Review Letters*, 94(11):1–4, 2005.

[16] Andrea Morello, Jarryd J. Pla, Patrice Bertet, and David N. Jamieson. Donor spins in silicon for quantum technologies. *Advanced Quantum Technologies*, 3(11): 2000005, 2020.

[17] Guido Burkard, Thaddeus D. Ladd, Andrew Pan, John M. Nichol, and Jason R. Petta. Semiconductor spin qubits. *Rev. Mod. Phys.*, 95:025003, Jun 2023.

[18] R. Hanson, L. P. Kouwenhoven, J. R. Petta, S. Tarucha, and L. M.K. Vandersypen. Spins in few-electron quantum dots. *Reviews of Modern Physics*, 79(4):1217–1265, 2007.

[19] Donald E. Savage, Mark Friesen, Max G. Lagally, Mark A. Eriksson, Robert Joynt, Paul Rugheimer, and Daniel W. van der Weide. Practical design and simulation of silicon-based quantum-dot qubits. *Physical Review B - Condensed Matter and Materials Physics*, 67(12):4, 2003.

[20] M. G. Borselli, K. Eng, R. S. Ross, T. M. Hazard, K. S. Holabird, B. Huang, A. A. Kiselev, P. W. Deelman, L. D. Warren, I. Milosavljevic, A. E. Schmitz, M. Sokolich, M. F. Gyure, and A. T. Hunter. Undoped accumulation-mode Si/SiGe quantum dots. *Nanotechnology*, 26(37):1–4, 2015.

[21] Yinan Fang et al. Recent advances in hole-spin qubits. *Mater. Quantum. Technol.*, 3:012003, 2023.

[22] N. Piot, B. Brun, V. Schmitt, S. Zihlmann, V. P. Michal, A. Apra, J. C. Abadillo-Uriel, X. Jehl, B. Bertrand, H. Niebojewski, L. Hutin, M. Vinet, M. Urdampilleta, T. Meunier, Y. M. Niquet, R. Maurand, and S. De Franceschi. A single hole spin with enhanced coherence in natural silicon. *Nature Nanotechnology*, 17(10):1072–1077, 2022.

[23] M. L. V. Tagliaferri, P. L. Bavdaz, W. Huang, A. S. Dzurak, D. Culcer, and M. Veldhorst. Impact of valley phase and splitting on readout of silicon spin qubits. *Phys. Rev. B*, 97:245412, Jun 2018.

[24] Zhanning Wang, Abhikbrata Sarkar, S. D. Liles, Andre Saraiva, A. S. Dzurak, A. R. Hamilton, and Dimitrie Culcer. Electrical operation of hole spin qubits in planar mos silicon quantum dots, 2023.

[25] M. Pioro-Ladrière, T. Obata, Y. Tokura, Y. S. Shin, T. Kubo, K. Yoshida, T. Taniyama, and S. Tarucha. Electrically driven single-electron spin resonance in a slanting zeeman field. *Nature Physics*, 4(10):776–779, 2008.

[26] Ognjen Malkoc, Peter Stano, and Daniel Loss. Charge-noise-induced dephasing in silicon hole-spin qubits. *Phys. Rev. Lett.*, 129:247701, Dec 2022.

[27] Daniel Jirovec, Andrea Hofmann, Andrea Ballabio, Philipp M. Mutter, Giulio Tavani, Marc Botifoll, Alessandro Crippa, Josip Kukucka, Oliver Sagi, Frederico Martins, Jaime Saez-Mollejo, Ivan Prieto, Maksim Borovkov, Jordi Arbiol, Daniel Chrastina, Giovanni Isella, and Georgios Katsaros. A singlet-triplet hole spin qubit in planar ge. *Nature Materials*, 20(8):1106–1112, 2021.

[28] Nico W. Hendrickx, William I. L. Lawrie, Maximilian Russ, Floor van Riggelen, Sander L. de Snoo, Raymond N. Schouten, Amir Sammak, Giordano Scappucci, and Menno Veldhorst. A four-qubit germanium quantum processor. *Nature*, 591 (7851):580–585, 2021.

[29] Marko Milivojević. Electrical control of the hole spin qubit in si and ge nanowire quantum dots. *Phys. Rev. B*, 104:235304, Dec 2021.

[30] Ke Wang, Gang Xu, Fei Gao, He Liu, Rong-Long Ma, Xin Zhang, Zhanning Wang, Gang Cao, Ting Wang, Jian-Jun Zhang, Dimitrie Culcer, Xuedong Hu, Hong-Wen Jiang, Hai-Ou Li, Guang-Can Guo, and Guo-Ping Guo. Ultrafast coherent control of a hole spin qubit in a germanium quantum dot. *Nature Communications*, 13 (1):206, 2022.

[31] Jørgen Holme Qvist and Jeroen Danon. Anisotropic g-tensors in hole quantum dots: Role of transverse confinement direction. *Phys. Rev. B*, 105:075303, Feb 2022.

[32] L. Bellentani, M. Bina, S. Bonen, A. Secchi, A. Bertoni, S. P. Voinigescu, A. Padovani, L. Larcher, and F. Troiani. Toward hole-spin qubits in Si p-mosfets within a planar cmos foundry technology. *Phys. Rev. Appl.*, 16:054034, Nov 2021.

[33] Stefano Bosco, Bence Hetényi, and Daniel Loss. Hole spin qubits in Si finfets with fully tunable spin-orbit coupling and sweet spots for charge noise. *PRX Quantum*, 2:010348, Mar 2021.

[34] Cécile X. Yu, Simon Zihlmann, JoséC. Abadillo-Uriel, Vincent P. Michal, Nils Rambal, Heimanu Niebojewski, Thomas Bedecarrats, Maud Vinet, Étienne Dumur, Michele Filippone, Benoit Bertrand, Silvano De Franceschi, Yann-Michel Niquet, and Romain Maurand. Strong coupling between a photon and a hole spin in silicon. *Nature Nanotechnology*, 18(7):741–746, 2023.

[35] Shigenobu Toné, Toshikazu Kubo, Takashi Ohyama, Takahiro Kusakabe, and Yohsuke Minatogawa. Quantitation of mRNA using in vitro RNA amplification and Northern hybridization. *Analytical Biochemistry*, 284(2):420–422, 2000.

[36] L. I. Glazman and M. Pustilnik. Coulomb blockade and Kondo effect in quantum dots. *New Directions in Mesoscopic Physics (Towards Nanoscience)*, pages 93–115, 2003.

[37] R. C.C. Leon, C. H. Yang, J. C.C. Hwang, J. Camirand Lemyre, T. Tanttu, W. Huang, K. W. Chan, K. Y. Tan, F. E. Hudson, K. M. Itoh, A. Morello, A. Laucht, M. Pioro-Ladrière, A. Saraiva, and A. S. Dzurak. Coherent spin control of s-, p-, d- and f-electrons in a silicon quantum dot. *Nature Communications*, 11 (1):1–7, 2020.

[38] Justin K. Perron, M. D. Stewart, and Neil M. Zimmerman. A new regime of Pauli-spin blockade. *Journal of Applied Physics*, 119(13), 2016.

[39] M. Veldhorst, J. C.C. Hwang, C. H. Yang, A. W. Leenstra, B. De Ronde, J. P. Dehollain, J. T. Muhonen, F. E. Hudson, K. M. Itoh, A. Morello, and A. S. Dzurak. An addressable quantum dot qubit with fault-tolerant control-fidelity. *Nature Nanotechnology*, 9(12):981–985, 2014.

[40] K. C. Nowack, F. H.L. Koppens, Yu V. Nazarov, and L. M.K. Vandersypen. Coherent control of a single electron spin with electric fields. *Science*, 318(5855): 1430–1433, 2007.

[41] Filip K. Malinowski, Frederico Martins, Peter D. Nissen, Edwin Barnes, Łukasz Cywiński, Mark S. Rudner, Saeed Fallahi, Geoffrey C. Gardner, Michael J. Manfra, Charles M. Marcus, and Ferdinand Kuemmeth. Notch filtering the nuclear environment of a spin qubit. *Nature Nanotechnology*, 12(1):16–20, 2017.

[42] Pascal Cerfontaine, Tim Botzem, Julian Ritzmann, Simon Sebastian Humpohl, Arne Ludwig, Dieter Schuh, Dominique Bougeard, Andreas D. Wieck, and Hendrik Bluhm. Closed-loop control of a GaAs-based singlet-triplet spin qubit with 99.5% gate fidelity and low leakage. pages 2–7, 2019.

[43] Xin Zhang, Elizaveta Morozova, Maximilian Rimbach-Russ, Daniel Jirovec, Tzu-Kan Hsiao, Pablo Cova Fariña, Chien-An Wang, Stefan D. Oosterhout, Amir Sammak, Giordano Scappucci, Menno Veldhorst, and Lieven M. K. Vandersypen. Universal control of four singlet-triplet qubits. 2023.

[44] Dimitrije Stepanenko, Mark Rudner, Bertrand I. Halperin, and Daniel Loss. Singlet-triplet splitting in double quantum dots due to spin-orbit and hyperfine interactions. *Physical Review B - Condensed Matter and Materials Physics*, 85 (7), 2012.

[45] E. Ferraro, M. Fanciulli, and M. De Michielis. Gate fidelity comparison in semiconducting spin qubit implementations affected by control noises. *Journal of Physics Communications*, 2(11), 2018.

[46] Maximilian Russ and Guido Burkard. Three-electron spin qubits. *Journal of Physics: Condensed Matter*, 29(39):393001, aug 2017.

[47] E. A. Laird, J. M. Taylor, D. P. Divincenzo, C. M. Marcus, M. P. Hanson, and A. C. Gossard. Coherent spin manipulation in an exchange-only qubit. *Physical Review B - Condensed Matter and Materials Physics*, 82(7):1–6, 2010.

[48] J. M. Taylor, V. Srinivasa, and J. Medford. Electrically protected resonant exchange qubits in triple quantum dots. *Physical Review Letters*, 111(5):1–5, 2013.

[49] Steven M. Girvin and Kun Yang. *Modern Condensed Matter Physics*. Cambridge University Press, 2019.

[50] Simon Geyer, Bence Hetényi, Stefano Bosco, Leon C. Camenzind, Rafael S. Eggli, Andreas Fuhrer, Daniel Loss, Richard J. Warburton, Dominik M. Zumbühl, and Andreas V. Kuhlmann. Two-qubit logic with anisotropic exchange in a fin field-effect transistor. 2022.

[51] Jan Klos, Fabian Hassler, Pascal Cerfontaine, Hendrik Bluhm, and Lars R. Schreiber. Calculation of tunnel couplings in open gate-defined disordered quantum dot systems. *Phys. Rev. B*, 98:155320, Oct 2018.

[52] T. A. Baart, P. T. Eendebak, C. Reichl, W. Wegscheider, and L. M. K. Vandersypen. Computer-automated tuning of semiconductor double quantum dots into the single-electron regime. *Applied Physics Letters*, 108(21):213104, 05 2016. ISSN 0003-6951.

[53] C. Volk, A. M. J. Zwerver, U. Mukhopadhyay, P. T. Eendebak, C. J. van Diepen, J. P. Dehollain, T. Hensgens, T. Fujita, C. Reichl, W. Wegscheider, and L. M. K. Vandersypen. Loading a quantum-dot based "qubyte"register. *npj Quantum Information*, 5(1):29, 2019.

[54] C. J. van Diepen, P. T. Eendebak, B. T. Buijtendorp, U. Mukhopadhyay, T. Fujita, C. Reichl, W. Wegscheider, and L. M. K. Vandersypen. Automated tuning of inter-dot tunnel coupling in double quantum dots. *Applied Physics Letters*, 113 (3):033101, 07 2018. ISSN 0003-6951.

[55] A. R. Mills, M. M. Feldman, C. Monical, P. J. Lewis, K. W. Larson, A. M. Mounce, and J. R. Petta. Computer-automated tuning procedures for semiconductor quantum dot arrays. *Applied Physics Letters*, 115(11):113501, 09 2019. ISSN 0003-6951.

[56] T.-K. Hsiao, C.J. van Diepen, U. Mukhopadhyay, C. Reichl, W. Wegscheider, and L.M.K. Vandersypen. Efficient orthogonal control of tunnel couplings in a quantum dot array. *Phys. Rev. Appl.*, 13:054018, May 2020.

[57] Sandesh S. Kalantre, Justyna P. Zwolak, Stephen Ragole, Xingyao Wu, Neil M. Zimmerman, M. D. Stewart, and Jacob M. Taylor. Machine learning techniques for state recognition and auto-tuning in quantum dots. *npj Quantum Information*, 5(1):6, 2019.

[58] Justyna P. Zwolak, Thomas McJunkin, Sandesh S. Kalantre, J.P. Dodson, E.R. MacQuarrie, D.E. Savage, M.G. Lagally, S.N. Coppersmith, Mark A. Eriksson, and Jacob M. Taylor. Autotuning of double-dot devices in situ with machine learning. *Phys. Rev. Appl.*, 13:034075, Mar 2020.

[59] Julian D. Teske, Simon Sebastian Humpohl, René Otten, Patrick Bethke, Pascal Cerfontaine, Jonas Dedden, Arne Ludwig, Andreas D. Wieck, and Hendrik Bluhm. A machine learning approach for automated fine-tuning of semiconductor spin qubits. *Applied Physics Letters*, 114(13):133102, 04 2019. ISSN 0003-6951.

[60] J. Darulová, S.J. Pauka, N. Wiebe, K.W. Chan, G.C Gardener, M.J. Manfra, M.C. Cassidy, and M. Troyer. Autonomous tuning and charge-state detection of gate-defined quantum dots. *Phys. Rev. Appl.*, 13:054005, May 2020.

[61] R. Durrer, B. Kratochwil, J.V. Koski, A.J. Landig, C. Reichl, W. Wegscheider, T. Ihn, and E. Greplova. Automated tuning of double quantum dots into specific charge states using neural networks. *Phys. Rev. Appl.*, 13:054019, May 2020.

[62] H. Moon, D. T. Lennon, J. Kirkpatrick, N. M. van Esbroeck, L. C. Camenzind, Liuqi Yu, F. Vigneau, D. M. Zumbühl, G. A. D. Briggs, M. A. Osborne, D. Sejdinovic, E. A. Laird, and N. Ares. Machine learning enables completely automatic tuning of a quantum device faster than human experts. *Nature Communications*, 11(1):4161, 2020.

[63] V. Nguyen, S. B. Orbell, D. T. Lennon, H. Moon, F. Vigneau, L. C. Camenzind, L. Yu, D. M. Zumbühl, G. A. D. Briggs, M. A. Osborne, D. Sejdinovic, and N. Ares. Deep reinforcement learning for efficient measurement of quantum devices. *npj Quantum Information*, 7(1):100, 2021.

[64] Barnaby van Straaten, Federico Fedele, Florian Vigneau, Joseph Hickie, Daniel Jirovec, Andrea Ballabio, Daniel Chrastina, Giovanni Isella, Georgios Katsaros, and Natalia Ares. All rf-based tuning algorithm for quantum devices using machine learning. 2022.

[65] Joseph Hickie, Barnaby van Straaten, Federico Fedele, Daniel Jirovec, Andrea Ballabio, Daniel Chrastina, Giovanni Isella, Georgios Katsaros, and Natalia Ares. Automated long-range compensation of an rf quantum dot sensor. 2023.

[66] Peter W. Shor. Scheme for reducing decoherence in quantum computer memory. *Phys. Rev. A*, 52:R2493–R2496, Oct 1995.

[67] Rajeev Acharya, Igor Aleiner, Richard Allen, Trond I. Andersen, Markus Ansmann, Frank Arute, Kunal Arya, Abraham Asfaw, Juan Atalaya, Ryan Babbush, Dave Bacon, Joseph C. Bardin, Joao Basso, Andreas Bengtsson, Sergio Boixo, Gina Bortoli, Alexandre Bourassa, Jenna Bovaird, Leon Brill, Michael Broughton, Bob B. Buckley, David A. Buell, Tim Burger, Brian Burkett, Nicholas Bushnell, Yu Chen, Zijun Chen, Ben Chiaro, Josh Cogan, Roberto Collins, Paul Conner, William Courtney, Alexander L. Crook, Ben Curtin, Dripto M. Debroy, Alexander Del Toro Barba, Sean Demura, Andrew Dunsworth, Daniel Eppens, Catherine Erickson, Lara Faoro, Edward Farhi, Reza Fatemi, Leslie Flores Burgos, Ebrahim Forati, Austin G. Fowler, Brooks Foxen, William Giang, Craig Gidney, Dar Gilboa, Marissa Giustina, Alejandro Grajales Dau, Jonathan A. Gross, Steve Habegger, Michael C. Hamilton, Matthew P. Harrigan, Sean D. Harrington, Oscar Higgott, Jeremy Hilton, Markus Hoffmann, Sabrina Hong, Trent Huang, Ashley Huff, William J. Huggins, Lev B. Ioffe, Sergei V. Isakov, Justin Iveland, Evan Jeffrey, Zhang Jiang, Cody Jones, Pavol Juhas, Dvir Kafri, Kostyantyn Kechedzhi, Julian Kelly, Tanuj Khattar, Mostafa Khezri, Mária Kieferová, Seon Kim, Alexei Kitaev, Paul V. Klimov, Andrey R. Klots, Alexander N. Korotkov, Fedor Kostritsa, John Mark Kreikebaum, David Landhuis, Pavel Laptev, Kim-Ming Lau, Lily Laws, Joonho Lee, Kenny Lee, Brian J. Lester, Alexander Lill, Wayne Liu, Aditya Locharla, Erik Lucero, Fionn D. Malone, Jeffrey Marshall, Orion Martin, Jarrod R. McClean, Trevor McCourt, Matt McEwen, Anthony Megrant, Bernardo Meurer Costa, Xiao Mi, Kevin C. Miao, Masoud Mohseni, Shirin Montazeri, Alexis Morvan, Emily Mount, Wojciech Mruczkiewicz, Ofer Naaman, Matthew Neeley, Charles Neill, Ani Nersisyan, Hartmut Neven, Michael Newman, Jiun How Ng, Anthony Nguyen, Murray Nguyen, Murphy Yuezhen Niu,

Thomas E. O'Brien, Alex Opremcak, John Platt, Andre Petukhov, Rebecca Potter, Leonid P. Pryadko, Chris Quintana, Pedram Roushan, Nicholas C. Rubin, Negar Saei, Daniel Sank, Kannan Sankaragomathi, Kevin J. Satzinger, Henry F. Schurkus, Christopher Schuster, Michael J. Shearn, Aaron Shorter, Vladimir Shvarts, Jindra Skruzny, Vadim Smelyanskiy, W. Clarke Smith, George Sterling, Doug Strain, Marco Szalay, Alfredo Torres, Guifre Vidal, Benjamin Villalonga, Catherine Vollgraff Heidweiller, Theodore White, Cheng Xing, Z. Jamie Yao, Ping Yeh, Juhwan Yoo, Grayson Young, Adam Zalcman, Yaxing Zhang, Ningfeng Zhu, and Google Quantum AI. Suppressing quantum errors by scaling a surface code logical qubit. *Nature*, 614(7949):676–681, 2023.

[68] Jonathan J. Burnett, Andreas Bengtsson, Marco Scigliuzzo, David Niepce, Marina Kudra, Per Delsing, and Jonas Bylander. Decoherence benchmarking of superconducting qubits. *npj Quantum Information*, 5(1):54, 2019.

[69] C. H. Yang, K. W. Chan, R. Harper, W. Huang, T. Evans, J. C. C. Hwang, B. Hensen, A. Laucht, T. Tanttu, F. E. Hudson, S. T. Flammia, K. M. Itoh, A. Morello, S. D. Bartlett, and A. S. Dzurak. Silicon qubit fidelities approaching incoherent noise limits via pulse engineering. *Nature Electronics*, 2(4):151–158, 2019.

[70] Reed W. Andrews, Cody Jones, Matthew D. Reed, Aaron M. Jones, Sieu D. Ha, Michael P. Jura, Joseph Kerckhoff, Mark Levendorf, Seán Meenehan, Seth T. Merkel, Aaron Smith, Bo Sun, Aaron J. Weinstein, Matthew T. Rakher, Thaddeus D. Ladd, and Matthew G. Borselli. Quantifying error and leakage in an encoded si/sige triple-dot qubit. *Nature Nanotechnology*, 14(8):747–750, 2019.

[71] Irina Heinz and Guido Burkard. Crosstalk analysis for single-qubit and two-qubit gates in spin qubit arrays. *Phys. Rev. B*, 104:045420, Jul 2021.

[72] Irina Heinz and Guido Burkard. Crosstalk analysis for simultaneously driven two-qubit gates in spin qubit arrays. *Phys. Rev. B*, 105:085414, Feb 2022.

[73] Maximilian Rimbach-Russ, Stephan Philips, Xiao Xue, and Lieven Vandersypen. Simple framework for systematic high-fidelity gate operations. *Quantum Science and Technology*, 8, 09 2023.

[74] Brennan Undseth, Xiao Xue, Mohammad Mehmandoost, Maximilian Rimbach-Russ, Pieter T. Eendebak, Nodar Samkharadze, Amir Sammak, Viatcheslav V. Dobrovitski, Giordano Scappucci, and Lieven M.K. Vandersypen. Nonlinear response and crosstalk of electrically driven silicon spin qubits. *Phys. Rev. Appl.*, 19:044078, Apr 2023.

[75] Stefania Lazăr, Quentin Ficheux, Johannes Herrmann, Ants Remm, Nathan Lacroix, Christoph Hellings, Francois Swiadek, Dante Colao Zanuz, Graham J. Norris, Mohsen Bahrami Panah, Alexander Flasby, Michael Kerschbaum, Jean-Claude Besse, Christopher Eichler, and Andreas Wallraff. Calibration of drive nonlinearity for arbitrary-angle single-qubit gates using error amplification. *Phys. Rev. Appl.*, 20:024036, Aug 2023.

[76] Yuval Baum, Mirko Amico, Sean Howell, Michael Hush, Maggie Liuzzi, Pranav Mundada, Thomas Merkh, Andre R.R. Carvalho, and Michael J. Biercuk. Experimental deep reinforcement learning for error-robust gate-set design on a superconducting quantum computer. *PRX Quantum*, 2(4), November 2021. ISSN 2691-3399.

[77] Murphy Yuezhen Niu, Sergio Boixo, Vadim Smelyanskiy, and Hartmut Neven. Universal quantum control through deep reinforcement learning. *AIAA Scitech 2019 Forum*, pages 1–21, 2019. ISSN 2056-6387.

[78] Eric Hyyppä, Antti Vepsäläinen, Miha Papič, Chun Fai Chan, Sinan Inel, Alessandro Landra, Wei Liu, Jürgen Luus, Fabian Marxer, Caspar Ockeloen-Korppi, Sebastian Orbell, Brian Tarasinski, and Johannes Heinsoo. Reducing leakage of single-qubit gates for superconducting quantum processors using analytical control pulse envelopes, 2024.

[79] F. Setiawan, Peter Groszkowski, and Aashish A. Clerk. Fast and robust geometric two-qubit gates for superconducting qubits and beyond. *Phys. Rev. Appl.*, 19: 034071, Mar 2023.

[80] Stephan G. J. Philips, Mateusz T. Mądzik, Sergey V. Amitonov, Sander L. de Snoo, Maximilian Russ, Nima Kalhor, Christian Volk, William I. L. Lawrie, Delphine Brousse, Larysa Tryputen, Brian Paquelet Wuetz, Amir Sammak, Menno Veldhorst, Giordano Scappucci, and Lieven M. K. Vandersypen. Universal control of a six-qubit quantum processor in silicon. *Nature*, 609(7929):919–924, 2022.

[81] Navin Khaneja, Timo Reiss, Cindie Kehlet, Thomas Schulte-Herbrüggen, and Steffen J. Glaser. Optimal control of coupled spin dynamics: design of nmr pulse sequences by gradient ascent algorithms. *Journal of Magnetic Resonance*, 172(2): 296–305, 2005.

[82] Frederik F Floether, Pierre de Fouquieres, and Sophie G Schirmer. Robust quantum gates for open systems via optimal control: Markovian versus non-markovian dynamics. *New Journal of Physics*, 14(7):073023, jul 2012.

[83] Yao Song, Junning Li, Yong-Ju Hai, Qihao Guo, and Xiu-Hao Deng. Optimizing quantum control pulses with complex constraints and few variables through autodifferentiation. *Phys. Rev. A*, 105:012616, Jan 2022.

[84] Christiane P. Koch, Ugo Boscain, Tommaso Calarco, Gunther Dirr, Stefan Filipp, Steffen J. Glaser, Ronnie Kosloff, Simone Montangero, Thomas Schulte-Herbrüggen, Dominique Sugny, and Frank K. Wilhelm. Quantum optimal control in quantum technologies. strategic report on current status, visions and goals for research in europe. *EPJ Quantum Technology*, 9(1):19, 2022.

[85] David L. Craig, Natalia Ares, and Erik M. Gauger. Differentiable master equation solver for quantum device characterisation, 2024.

[86] Xiaotong Ni, Hui-Hai Zhao, Lei Wang, Feng Wu, and Jianxin Chen. Integrating quantum processor device and control optimization in a gradient-based framework. *npj Quantum Information*, 8(1):106, 2022.

[87] Mohamed Abdelhafez, Brian Baker, András Gyenis, Pranav Mundada, Andrew A. Houck, David Schuster, and Jens Koch. Universal gates for protected superconducting qubits using optimal control. *Phys. Rev. A*, 101:022321, Feb 2020.

[88] Nicolas Wittler, Federico Roy, Kevin Pack, Max Werninghaus, Anurag Saha Roy, Daniel J. Egger, Stefan Filipp, Frank K. Wilhelm, and Shai Machnes. Integrated tool set for control, calibration, and characterization of quantum devices applied to superconducting qubits. *Phys. Rev. Appl.*, 15:034080, Mar 2021.

[89] H. Jirari. Implementation of one-qubit gates using optimal control theory for a dissipative system. *The European Physical Journal B*, 94(10):217, 2021.

[90] D. Fernández-Fernández, Yue Ban, and G. Platero. Quantum control of hole spin qubits in double quantum dots. *Phys. Rev. Appl.*, 18:054090, Nov 2022.

[91] Julian David Teske, Pascal Cerfontaine, and Hendrik Bluhm. qopt: An experiment-oriented software package for qubit simulation and quantum optimal control. *Phys. Rev. Appl.*, 17:034036, Mar 2022.

[92] Pascal Cerfontaine, Tim Botzem, Julian Ritzmann, Simon Sebastian Humpohl, Arne Ludwig, Dieter Schuh, Dominique Bougeard, Andreas D. Wieck, and Hendrik Bluhm. Closed-loop control of a gaas-based singlet-triplet spin qubit with 99.5% gate fidelity and low leakage. *Nature Communications*, 11(1):4144, 2020.

[93] Dylan J. Gorman, Kevin C. Young, and K. Birgitta Whaley. Overcoming dephasing noise with robust optimal control. *Phys. Rev. A*, 86:012317, Jul 2012.

[94] Rong-Long Ma, Ao-Ran Li, Chu Wang, Zhen-Zhen Kong, Wei-Zhu Liao, Ming Ni, Sheng-Kai Zhu, Ning Chu, Cheng-Xian Zhang, Di Liu, Gang Cao, Gui-Lei Wang, Hai-Ou Li, and Guo-Ping Guo. Single spin qubit geometric gate in a silicon quantum dot, 2023.

[95] J. Fransson and M. Råsander. Pauli spin blockade in weakly coupled double quantum dots. *Phys. Rev. B*, 73:205333, May 2006.

[96] Akira Fujiwara, Hiroshi Inokawa, Kenji Yamazaki, Hideo Namatsu, Yasuo Takahashi, Neil Zimmerman, and Stuart Martin. Single electron tunneling transistor with tunable barriers using silicon nanowire metal-oxide-semiconductor field-effect transistor. *Applied Physics Letters*, 88:053121–053121, 02 2006.

[97] Florian Vigneau, Federico Fedele, Anasua Chatterjee, David Reilly, Ferdinand Kuemmeth, M. Fernando Gonzalez-Zalba, Edward Laird, and Natalia Ares. Probing quantum devices with radio-frequency reflectometry. *Applied Physics Reviews*, 10(2), February 2023. ISSN 1931-9401.

[98] M Esterli, R M Otxoa, and M F Gonzalez-Zalba. Small-signal equivalent circuit for double quantum dots at low-frequencies. *Appl. Phys. Lett.*, 114(25):253505, June 2019.

[99] F. J. Schupp, F. Vigneau, Y. Wen, A. Mavalankar, J. Griffiths, G. A. C. Jones, I. Farrer, D. A. Ritchie, C. G. Smith, L. C. Camenzind, L. Yu, D. M. Zumbühl, G. A. D. Briggs, N. Ares, and E. A. Laird. Sensitive radiofrequency readout of quantum dots using an ultra-low-noise SQUID amplifier. *Journal of Applied Physics*, 127(24):244503, 06 2020. ISSN 0021-8979.

[100] Natalia Ares. Machine learning as an enabler of qubit scalability. *Nature Reviews Materials*, 6(10):870–871, 2021.

[101] Mario Krenn, Jonas Landgraf, Thomas Foesel, and Florian Marquardt. Artificial intelligence and machine learning for quantum technologies. *Phys. Rev. A*, 107: 010101, Jan 2023.

[102] Alp Kucukelbir, Dustin Tran, Rajesh Ranganath, Andrew Gelman, and David M. Blei. Automatic differentiation variational inference. 2016.

[103] Solomon Kullback and Richard A Leibler. On information and sufficiency. *The annals of mathematical statistics*, 22(1):79–86, 1951.

[104] Diederik P Kingma and Max Welling. Auto-encoding variational bayes, 2022.

[105] Leon C. Camenzind, Simon Geyer, Andreas Fuhrer, Richard J. Warburton, Dominik M. Zumbühl, and Andreas V. Kuhlmann. A hole spin qubit in a fin field-effect transistor above 4 kelvin. *Nature Electronics*, 5(3):178–183, 2022.

[106] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518 (7540):529–533, 2015. ISSN 14764687.

[107] David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharshan Kumaran, Thore Graepel, Timothy Lillicrap, Karen Simonyan, and Demis Hassabis. A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science*, 362(6419):1140–1144, 2018. ISSN 10959203.

[108] David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016. ISSN 14764687.

[109] Moritz August and José Miguel Hernández-Lobato. Taking Gradients Through Experiments: LSTMs and Memory Proximal Policy Optimization for Black-Box Quantum Control. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 11203 LNCS: 591–613, 2018. ISSN 16113349.

[110] Thomas Fösel, Petru Tighineanu, Talitha Weiss, and Florian Marquardt. Reinforcement Learning with Neural Networks for Quantum Feedback. *Physical Review X*, 8(3):31084, 2018. ISSN 21603308.

[111] Marin Bukov, Alexandre G.R. Day, Dries Sels, Phillip Weinberg, Anatoli Polkovnikov, and Pankaj Mehta. Reinforcement Learning in Different Phases of Quantum Control. *Physical Review X*, 8(3):31086, 2018. ISSN 21603308.

[112] Han Xu, Junning Li, Liqiang Liu, Yu Wang, Haidong Yuan, and Xin Wang. Generalizable control for quantum parameter estimation through reinforcement learning. *npj Quantum Information*, 5(1):1–8, 2019. ISSN 20566387.

[113] Sahar Daraeizadeh, Shavindra P. Premaratne, and A. Y. Matsuura. Designing high-fidelity multi-qubit gates for semiconductor quantum dots through deep reinforcement learning. (1), 2020.

[114] Steven Herbert and Akash Sengupta. Using Reinforcement Learning to find Efficient Qubit Routing Policies for Deployment in Near-term Quantum Computers. 1:1–13, 2018.

[115] Pantita Palittapongarnpim, Peter Wittek, Ehsan Zahedinejad, Shakib Vedaie, and Barry C. Sanders. Learning in quantum control: High-dimensional global optimization for noisy quantum dynamics. *Neurocomputing*, 268:116–126, 2017. ISSN 18728286.

[116] Zheng An and D. L. Zhou. Deep reinforcement learning for quantum gate control. *Epl*, 126(6), 2019. ISSN 12864854.

[117] Riccardo Porotti, Dario Tamascelli, Marcello Restelli, and Enrico Prati. Coherent transport of quantum states by deep reinforcement learning. *Communications Physics*, 2(1), 2019. ISSN 23993650.

[118] Jonas Schuff, Lukas J. Fiderer, and Daniel Braun. Improving the dynamics of quantum sensors with reinforcement learning. *New Journal of Physics*, 22(3), 2020. ISSN 13672630.

[119] Tingwu Wang, Xuchan Bao, Ignasi Clavera, Jerrick Hoang, Yeming Wen, Eric Langlois, Shunshi Zhang, Guodong Zhang, Pieter Abbeel, and Jimmy Ba. Benchmarking Model-Based Reinforcement Learning. pages 1–25, 2019.

[120] Peng Wei, Na Li, and Zairong Xi. Open quantum system control based on reinforcement learning. *Chinese Control Conference, CCC*, 38:6911–6916, 2019. ISSN 21612927.

[121] Xun Gao and Lu Ming Duan. Efficient representation of quantum many-body states with deep neural networks. *Nature Communications*, 8(1):1–5, 2017. ISSN 20411723.

[122] Ariel Barr, Willem Gispen, and Austen Lamacraft. Quantum Ground States from Reinforcement Learning. 107(2001):1–19, 2020.

[123] Dong Ling Deng. Machine Learning Detection of Bell Nonlocality in Quantum Many-Body Systems. *Physical Review Letters*, 120(24):240402, 2018. ISSN 10797114.

[124] Giuseppe Carleo and Matthias Troyer. Solving the quantum many-body problem with artificial neural networks. 606(February):602–606, 2017.

[125] Vegard B. Sørdal and Joakim Bergli. Deep reinforcement learning for quantum Szilard engine optimization. *Physical Review A*, 100(4):1–9, 2019. ISSN 24699934.

[126] M. *et al.* Veldhorst. A two-qubit logic gate in silicon. *Nature*, 526:410–414, 2015. ISSN 14764687.

[127] W. *et al.* Huang. Fidelity benchmarks for two-qubit gates in silicon. *Nature*, 569, 2019. ISSN 14764687.

[128] D. M. Zajac, A. J. Sigillito, M. Russ, F. Borjans, J. M. Taylor, G. Burkard, and J. R. Petta. Resonantly driven CNOT gate for electron spins. *Science*, 359(6374): 439–442, 2018. ISSN 10959203.

[129] S. *et al.* Tarucha. A quantum-dot spin qubit with coherence limited by charge noise and fidelity higher than 99.9%. *Physical Review Letters*, 13, 2018. ISSN 2.

[130] J. *et al.* Medford. Self-consistent measurement and state tomography of an exchange-only spin qubit. *Nat. Nanotechnol*, 8, 2013. ISSN 9.

[131] J. Darulová, S. J. Pauka, N. Wiebe, K. W. Chan, G. C. Gardener, M. J. Manfra, M. C. Cassidy, and M. Troyer. Autonomous tuning and charge state detection of gate defined quantum dots. pages 1–19, 2019.

[132] Justyna P. Zwolak, Sandesh S. Kalantre, Xingyao Wu, Stephen Ragole, and Jacob M. Taylor. QFlow lite Dataset: A Machine-Learning approach to the charge states in quantum dot experiments. *PLoS ONE*, 13(10):1–18, 2018. ISSN 19326203.

[133] Justyna P. Zwolak, Thomas McJunkin, Sandesh S. Kalantre, J.P. Dodson, E.R. MacQuarrie, D.E. Savage, M.G. Lagally, S.N. Coppersmith, Mark A. Eriksson, and Jacob M. Taylor. Autotuning of Double-Dot Devices In Situ with Machine Learning . *Physical Review Applied*, 13(3):1–9, 2020. ISSN 2331-7019.

[134] N. M. van Esbroeck, D. T. Lennon, H. Moon, V. Nguyen, F. Vigneau, L. C. Camenzind, L. Yu, D. M. Zumbühl, G. A. D. Briggs, D. Sejdinovic, and N. Ares. Quantum device fine-tuning using unsupervised embedding learning. pages 1–11, 2020.

[135] Renato Durrer, Benedikt Kratochwil, Jonne V. Koski, Andreas J. Landig, Christian Reichl, Werner Wegscheider, Thomas Ihn, and Eliska Greplova. Automated tuning of double quantum dots into specific charge states using neural networks. 1:1–9, 2019.

[136] D. T. Lennon, H. Moon, L. C. Camenzind, Liuqi Yu, D. M. Zumbühl, G. A.D. Briggs, M. A. Osborne, E. A. Laird, and N. Ares. Efficiently measuring a quantum device using machine learning. *npj Quantum Information*, 5(1):1–8, 2019. ISSN 20566387.

[137] Yann Lecun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521 (7553):436–444, 2015. ISSN 14764687.

[138] Leon C. Camenzind, Liuqi Yu, Peter Stano, Jeramy D. Zimmerman, Arthur C. Gossard, Daniel Loss, and Dominik M. Zumbühl. Hyperfine-phonon spin relaxation in a single-electron GaAs quantum dot. *Nature Communications*, 9(1):1–6, 2018. ISSN 20411723.

[139] Leon C. Camenzind, Liuqi Yu, Peter Stano, Jeramy D. Zimmerman, Arthur C. Gossard, Daniel Loss, and Dominik M. Zumbühl. Spectroscopy of Quantum Dot Orbitals with In-Plane Magnetic Fields. *Physical Review Letters*, 122(20):207701, 2019. ISSN 10797114.

[140] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. OpenAI Gym. pages 1–4, 2016.

[141] Ziyu Wang, Tom Schaul, Matteo Hessel, Hado Van Hasselt, Marc Lanctot, and Nando De Freitas. Dueling Network Architectures for Deep Reinforcement Learning. *33rd International Conference on Machine Learning, ICML 2016*, 48:1995–2003, 2016.

[142] Tom Schaul, John Quan, Ioannis Antonoglou, and David Silver. Prioritized experience replay. *4th International Conference on Learning Representations, ICLR 2016 - Conference Track Proceedings*, pages 1–21, 2016.

[143] Diederik P. Kingma and Jimmy Lei Ba. Adam: A method for stochastic optimization. *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings*, pages 1–15, 2015.

[144] Frank Wilcoxon. Individual Comparisons by Ranking Methods. *Biometrics Bulletin*, 1(6):80, 1945. ISSN 00994987.

[145] Alessandro Crippa, Romain Maurand, Dharmraj Kotekar-Patil, Andrea Corna, Heorhii Bohuslavskyi, Alexei O. Orlov, Patrick Fay, Romain Laviéville, Sylvain Barraud, Maud Vinet, Marc Sanquer, Silvano De Franceschi, and Xavier Jehl. Level Spectrum and Charge Relaxation in a Silicon Double Quantum Dot Probed by Dual-Gate Reflectometry. *Nano Letters*, 17(2):1001–1006, 2017. ISSN 15306992.

[146] N. Ares, F. J. Schupp, A. Mavalankar, G. Rogers, J. Griffiths, G. A.C. Jones, I. Farrer, D. A. Ritchie, C. G. Smith, A. Cottet, G. A.D. Briggs, and E. A. Laird. Sensitive Radio-Frequency Measurements of a Quantum Dot by Tuning to Perfect Impedance Matching. *Physical Review Applied*, 5(3):1–6, 2016. ISSN 23317019.

[147] Damaz De Jong, Jasper Van Veen, Luca Binci, Amrita Singh, Peter Krogstrup, Leo P. Kouwenhoven, Wolfgang Pfaff, and John D. Watson. Rapid Detection of Coherent Tunneling in an InAs Nanowire Quantum Dot through Dispersive Gate Sensing. *Physical Review Applied*, 11(4):1, 2019. ISSN 23317019.

[148] M. Jung, M. D. Schroer, K. D. Petersson, and J. R. Petta. Radio frequency charge sensing in InAs nanowire double quantum dots. *Applied Physics Letters*, 100(25): 1–5, 2012. ISSN 00036951.

[149] N. *et al.* Vlassis. Bayesian reinforcement learning. Reinforcement learning. pages 359–386, 2012.

[150] K. *et al.* Azizzadenesheli. Efficient exploration through bayesian deep q-networks. *2018 Workshop (ITA) IEEE.*, pages 1–9, 2018.

[151] A. *et al.* Ayoub. Model-based reinforcement learning with value-targeted regression. *PMLR.*, pages 463–474, 2020.

[152] G. A. *et al.* Oakes. Automatic virtual voltage extraction of a 2xN array of quantum dots with machine learning. 2020.

[153] Aaron J. Weinstein, Matthew D. Reed, Aaron M. Jones, Reed W. Andrews, David Barnes, Jacob Z. Blumoff, Larken E. Euliss, Kevin Eng, Bryan H. Fong, Sieu D. Ha, Daniel R. Hulbert, Clayton A. C. Jackson, Michael Jura, Tyler E. Keating, Joseph Kerckhoff, Andrey A. Kiselev, Justine Matten, Golam Sabbir, Aaron Smith, Jeffrey Wright, Matthew T. Rakher, Thaddeus D. Ladd, and Matthew G. Borselli. Universal logic with encoded spin qubits in silicon. *Nature*, 615(7954): 817–822, 2023.

[154] N. W. Hendrickx, L. Massai, M. Mergenthaler, F. Schupp, S. Paredes, S. W. Bedell, G. Salis, and A. Fuhrer. Sweet-spot operation of a germanium hole spin qubit with highly anisotropic noise sensitivity, 2023.

[155] Zhanning Wang, Elizabeth Marcellina, Alex. R. Hamilton, James H. Cullen, Sven Rogge, Joe Salfi, and Dimitrie Culcer. Optimal operation points for ultrafast, highly coherent ge hole spin-orbit qubits. *npj Quantum Information*, 7(1):54, 2021.

[156] Miguel J. Carballido, Simon Svab, Rafael S. Eggli, Taras Patlatiuk, Pierre Chevalier Kwon, Jonas Schuff, Rahel M. Kaiser, Leon C. Camenzind, Ang Li, Natalia Ares, Erik P. A. M Bakkers, Stefano Bosco, J. Carlos Egues, Daniel Loss, and Dominik M. Zumbühl. A qubit with simultaneously maximized speed and coherence, 2024.

[157] Zhanning Wang, Abhikbrata Sarkar, S. D. Liles, Andre Saraiva, A. S. Dzurak, A. R. Hamilton, and Dimitrie Culcer. Electrical operation of hole spin qubits in planar mos silicon quantum dots. *Phys. Rev. B*, 109:075427, Feb 2024.

[158] J. van Bree, A. Yu. Silov, M. L. van Maasakkers, C. E. Pryor, M. E. Flatté, and P. M. Koenraad. Anisotropy of electron and hole g tensors of quantum dots: An intuitive picture based on spin-correlated orbital currents. *Phys. Rev. B*, 93: 035311, Jan 2016.

[159] John Rooney, Zhentao Luo, Lucas E. A. Stehouwer, Giordano Scappucci, Menno Veldhorst, and Hong-Wen Jiang. Gate modulation of the hole singlet-triplet qubit frequency in germanium. 2023.

[160] Peter Stano and Daniel Loss. Review of performance metrics of spin qubits in gated semiconducting nanostructures. *Nature Reviews Physics*, 4(10):672–688, 2022.

[161] Leon C. Camenzind, Simon Geyer, Andreas Fuhrer, Richard J. Warburton, Dominik M. Zumbühl, and Andreas V. Kuhlmann. A hole spin qubit in a fin field-effect transistor above 4 kelvin. *Nature Electronics*, 5(3):178–183, 2022.

[162] Daniel Jirovec, Philipp M. Mutter, Andrea Hofmann, Alessandro Crippa, Marek Rychetsky, David L. Craig, Josip Kukucka, Frederico Martins, Andrea Ballabio, Natalia Ares, Daniel Chrastina, Giovanni Isella, Guido Burkard, and Georgios Katsaros. Dynamics of hole singlet-triplet qubits with large g-factor differences. *Phys. Rev. Lett.*, 128:126803, Mar 2022.

[163] Mahdi Naghiloo. Introduction to experimental quantum measurement with superconducting qubits, 2019.

[164] Fei Yan, Simon Gustavsson, Jonas Bylander, Xiaoyue Jin, Fumiki Yoshihara, David G. Cory, Yasunobu Nakamura, Terry P. Orlando, and William D. Oliver. Rotating-frame relaxation as a noise spectrum analyser of a superconducting qubit undergoing driven evolution. *Nature Communications*, 4(1):2337, 2013.

[165] Ziyu Wang and Nando de Freitas. Theoretical analysis of bayesian optimisation with unknown gaussian process hyper-parameters. 2014.

[166] Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian processes for machine learning.* Adaptive computation and machine learning. MIT Press, 2006. ISBN 026218253X.

[167] Marc G. Genton. Classes of kernels for machine learning: a statistics perspective. *J. Mach. Learn. Res.*, 2:299–312, mar 2002. ISSN 1532-4435.

[168] P. Scarlino, E. Kawakami, D. R. Ward, D. E. Savage, M. G. Lagally, Mark Friesen, S. N. Coppersmith, M. A. Eriksson, and L. M. K. Vandersypen. Second-harmonic coherent driving of a spin qubit in a si/sige quantum dot. *Phys. Rev. Lett.*, 115: 106802, Sep 2015.

[169] Judit Romhányi, Guido Burkard, and András Pályi. Subharmonic transitions and bloch-siegert shift in electrically driven spin resonance. *Phys. Rev. B*, 92:054422, Aug 2015.

[170] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018.

[171] Davan Harrison. A brief introduction to automatic differentiation for machine learning, 2021.

[172] Mark D. Bowdrey, Daniel K.L. Oi, Anthony J. Short, Konrad Banaszek, and Jonathan A. Jones. Fidelity of single qubit maps. *Physics Letters A*, 294(5): 258–260, 2002.

[173] Michael A. Nielsen. A simple formula for the average gate fidelity of a quantum dynamical operation [rapid communication]. *Physics Letters A*, 2002.

[174] A. D. Córcoles, Jay M. Gambetta, Jerry M. Chow, John A. Smolin, Matthew Ware, Joel Strand, B. L. T. Plourde, and M. Steffen. Process verification of two-qubit quantum gates by randomized benchmarking. *Phys. Rev. A*, 87:030301, Mar 2013.

[175] Michael R. Geller and Zhongyuan Zhou. Efficient error models for fault-tolerant architectures and the pauli twirling approximation. *Phys. Rev. A*, 88:012314, Jul 2013.

[176] J.R. Johansson, P.D. Nation, and Franco Nori. Qutip: An open-source python framework for the dynamics of open quantum systems. *Computer Physics Communications*, 183(8):1760–1772, 2012.

[177] J.R. Johansson, P.D. Nation, and Franco Nori. Qutip 2: A python framework for the dynamics of open quantum systems. *Computer Physics Communications*, 184 (4):1234–1240, 2013.

[178] Patrick Kidger, James Morrill, James Foster, and Terry Lyons. Neural Controlled Differential Equations for Irregular Time Series. *Advances in Neural Information Processing Systems*, 2020.

[179] Patrick Kidger. *On Neural Differential Equations*. PhD thesis, University of Oxford, 2021.

[180] P. J Prince and J. R. Dormand. High order embedded Runge–Kutta formulae. *J. Comp. Appl. Math*, 7(1):67–75, 1981.

[181] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. 2017.

[182] F.J. Harris. On the use of windows for harmonic analysis with the discrete fourier transform. *Proceedings of the IEEE*, 66(1):51–83, 1978.

[183] M. Werninghaus, D. J. Egger, F. Roy, S. Machnes, F. K. Wilhelm, and S. Filipp. Leakage reduction in fast superconducting qubit gates via optimal control. *npj Quantum Information*, 7(1):14, 2021.

[184] Alessandro Crippa, Romain Maurand, Léo Bourdet, Dharmraj Kotekar-Patil, Anthony Amisse, Xavier Jehl, Marc Sanquer, Romain Laviéville, Heorhii Bohuslavskyi, Louis Hutin, Sylvain Barraud, Maud Vinet, Yann-Michel Niquet, and Silvano De Franceschi. Electrical spin driving by g-matrix modulation in spin-orbit qubits. *Phys. Rev. Lett.*, 120:137702, Mar 2018.

[185] Baker Shalak, Christophe Delerue, and Yann-Michel Niquet. Modeling of spin decoherence in a si hole qubit perturbed by a single charge fluctuator. *Phys. Rev. B*, 107:125415, Mar 2023.

[186] Guanru Feng, Joel J. Wallman, Brandon Buonacorsi, Franklin H. Cho, Daniel K. Park, Tao Xin, Dawei Lu, Jonathan Baugh, and Raymond Laflamme. Estimating the coherence of noise in quantum control of a solid-state qubit. *Phys. Rev. Lett.*, 117:260501, Dec 2016.

[187] Navin Khaneja, Timo Reiss, Cindie Kehlet, Thomas Schulte-Herbrüggen, and Steffen J. Glaser. Optimal control of coupled spin dynamics: design of nmr pulse sequences by gradient ascent algorithms. *Journal of Magnetic Resonance*, 172(2): 296–305, 2005.

[188] Kunni Lin, Jiawei Peng, Feng Long Gu, and Zhenggang Lan. Simulation of open quantum dynamics with bootstrap-based long short-term memory recurrent neural network. *The Journal of Physical Chemistry Letters*, 12(41):10225–10234, 10 2021.

[189] E. Flurin, L. S. Martin, S. Hacohen-Gourgy, and I. Siddiqi. Using a recurrent neural network to reconstruct quantum dynamics of a superconducting qubit from physical observations. *Phys. Rev. X*, 10:011006, Jan 2020.

[190] Matthew Choi, Daniel Flam-Shepherd, Thi Ha Kyaw, and Alán Aspuru-Guzik. Learning quantum dynamics with latent neural ordinary differential equations. *Phys. Rev. A*, 105:042403, Apr 2022.

[191] Jeroen Petrus Gerardus Van Dijk, Bishnu Patra, Sushil Subramanian, Xiao Xue, Nodar Samkharadze, Andrea Corna, Charles Jeon, Farhana Sheikh, Esdras Juarez-Hernandez, Brando Perez Esparza, Huzaifa Rampurawala, Brent R. Carlton, Surej Ravikumar, Carlos Nieva, Sungwon Kim, Hyung-Jin Lee, Amir Sammak, Giordano Scappucci, Menno Veldhorst, Lieven M. K. Vandersypen, Edoardo Charbon, Stefano Pellerano, Masoud Babaie, and Fabio Sebastiano. A scalable cryo-cmos controller for the wideband frequency-multiplexed control of spin qubits and transmons. *IEEE Journal of Solid-State Circuits*, 55(11):2930–2946, 2020.

[192] Yatao Peng, Andrea Ruffino, Tsung-Yeh Yang, John Michniewicz, Miguel Fernando Gonzalez-Zalba, and Edoardo Charbon. A cryo-cmos wideband quadrature receiver with frequency synthesizer for scalable multiplexed readout of silicon spin qubits. *IEEE Journal of Solid-State Circuits*, 57(8):2374–2389, 2022.

[193] J. Kelly, R. Barends, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. G. Fowler, I.-C. Hoi, E. Jeffrey, A. Megrant, J. Mutus, C. Neill, P. J. J. O'Malley, C. Quintana, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, T. C. White, A. N. Cleland, and John M. Martinis. Optimal quantum control using randomized benchmarking. *Phys. Rev. Lett.*, 112:240504, Jun 2014.

[194] Clement H. Wong, M. A. Eriksson, S. N. Coppersmith, and Mark Friesen. High-fidelity singlet-triplet $s-T_-$ qubits in inhomogeneous magnetic fields. *Phys. Rev. B*, 92:045403, Jul 2015.

[195] J. R. Schrieffer and P. A. Wolff. Relation between the anderson and kondo hamiltonians. *Phys. Rev.*, 149:491–492, Sep 1966.

[196] Sergey Bravyi, David P. DiVincenzo, and Daniel Loss. Schrieffer–wolff transformation for quantum many-body systems. *Annals of Physics*, 326(10):2793–2826, October 2011. ISSN 0003-4916.

[197] Sigurdur I. Erlingsson, J. Carlos Egues, and Daniel Loss. Energy spectra for quantum wires and two-dimensional electron gases in magnetic fields with rashba and dresselhaus spin-orbit interactions. *Phys. Rev. B*, 82:155456, Oct 2010.

[198] M. M. Salomaa. Schrieffer-wolff transformation for the anderson hamiltonian in a superconductor. *Phys. Rev. B*, 37:9312–9317, Jun 1988.

[199] Rukhsan Ul Haq, Sachin Satish Bharadwaj, and Towseef Ali Wani. An explicit method for schrieffer-wolff transformation, 2019.

[200] Rukhsan Ul Haq and Keshav Singh. A systematic method for schrieffer-wolff transformation and its generalizations, 2020.

[201] I. Roth, R. Kueng, S. Kimmel, Y.-K. Liu, D. Gross, J. Eisert, and M. Kliesch. Recovering quantum gates from few average gate fidelities. *Phys. Rev. Lett.*, 121: 170502, Oct 2018.

[202] R. B. A. Adamson and A. M. Steinberg. Improving quantum state estimation with mutually unbiased bases. *Phys. Rev. Lett.*, 105:030406, Jul 2010.

He was the most inconsiderable creature in that swarming mass of mankind which for a brief space occupied the surface of the earth; and he was almighty because he had wrenched from the chaos the secret of its nothingness.
W. Somerset Maugham

Appendix A

Singlet-Triplet qubit Hamiltonian

Here I demonstrate the construction of a Hamiltonian which describes the physics of a Singlet-Triplet hole spin qubit in a semiconductor quantum dot device[50]. I consider a tunnel-coupled double quantum dot in the presence of an external magnetic field and include a strong spin-orbit interaction (SOI). The system can be modelled by the Hamiltonian

$$H_{\text{tot}} = H_{\text{orb}} + H_Z + H_{\text{SO}} \tag{A.1}$$

where H_Z is the Zeeman Hamiltonian and H_{SO} describes the SOI[50]. The orbital part H_{orb} reads

$$H_{\text{orb}} = \epsilon \left|S_{20}\right\rangle \left\langle S_{20}\right| + 2t_C\big(\left|S_{11}\right\rangle \left\langle S_{20}\right| + \left|S_{20}\right\rangle \left\langle S_{11}\right| \big) \tag{A.2}$$

where $|S_{20}\rangle$ is the singlet with double occupancy in the left dot which is detuned by an amount ϵ from the $(1,1)$ configuration, t_C is the tunnel coupling strength in the one-particle picture, and $|S_{11}\rangle$ is the singlet with one hole in each dot. Diagonalising yields the hybridised singlet states

$$\left|S_E\right\rangle = \cos\left(\frac{\Omega}{2}\right)\left|S_{20}\right\rangle + \sin\left(\frac{\Omega}{2}\right)\left|S_{11}\right\rangle \tag{A.3}$$

$$\left|S_G\right\rangle = \cos\left(\frac{\Omega}{2}\right)\left|S_{11}\right\rangle - \sin\left(\frac{\Omega}{2}\right)\left|S_{20}\right\rangle \tag{A.4}$$

where $\Omega = \arctan\left(\sqrt{\frac{2t_C}{\epsilon}}\right)$ is the orbital mixing angle and the corresponding energies are $E(S_{E/G}) = \epsilon/2 \pm \sqrt{\epsilon^2/4 + 2t_C^2}$[194]. In the following, we can consider a regime where the excited singlet is far detuned, as is common in qubit operation, and may be neglected. Defining $|S\rangle \equiv |S_G\rangle$, we may focus on the four-dimensional space spanned by the states $|S\rangle, |T+\rangle, |T_0\rangle$ and $|T-\rangle$, where $|T_{0,\pm}\rangle$ denote the triplet states in the $(1,1)$ configuration. For a magnetic field of magnitude B and angle θ as measured from the

plane defined by the heterostructure, the Zeeman Hamiltonian H_Z reads,

$$H_Z = \frac{\mu B}{2}\Bigg(\sin(\theta)\Big(\Sigma_\pm \pm g_\perp^+ \left|T_\pm\right\rangle \left\langle T_\pm\right| + g_\perp^-(\left|S\right\rangle \left\langle T_0\right| + \left|T_0\right\rangle \left\langle S\right|)cos\left(\frac{\Omega}{2}\right)\Big) \\ + \frac{cos(\theta)}{\sqrt{2}}\Big(\Sigma_\pm(g_\parallel^+ \left|T_0\right\rangle \left\langle T_\pm\right| \mp g_\parallel^- \left|S\right\rangle \left\langle T_\pm\right| cos\left(\frac{\Omega}{2}\right) + H.c.)\Big)\Bigg) \quad \text{(A.5)}$$

where $g_{j\pm} = g_{jL} \pm g_{jR}$.

Finally, we consider a SOI parameterised by an in-plane spin-orbit vector, $\mathbf{t}_{\mathrm{SO}} = (t_x, t_y, t_z)$, where all entries are real. Such in-plane spin-flip tunnelling terms stem from the cubic Rashba SOI, while this type of SOI does not induce out-of-plane terms t_z. This interaction arises due to the mixing of heavy and light hole states in hole spin systems. Consequently, the spin-orbit Hamiltonian H_{SO} reads,

$$H_{\mathrm{SO}} = \sin\left(\frac{\Omega}{2}\right)\Big(i\sqrt{2}t_z \left|T_0\right\rangle \left\langle S\right| + \Sigma_\pm(t_x \pm it_y)\left|T_\pm\right\rangle \left\langle S\right| + H.c.\Big) \quad \text{(A.6)}$$

Given the reduced basis, $\left|S_E\right\rangle, \left|S_G\right\rangle, \left|T_+\right\rangle, \left|T_-\right\rangle, \left|T_0\right\rangle$

$$H_Z = \frac{\mu B}{2}\begin{pmatrix} \epsilon & 2t_c & 0 & 0 & 0 \\ 2t_c & g_\perp^- \sin\theta\cos^2\frac{\Omega}{2} & -g_\parallel^- \frac{\cos\theta}{\sqrt{2}}\cos\frac{\Omega}{2} & +g_\parallel^- \frac{\cos\theta}{\sqrt{2}}\cos\frac{\Omega}{2} & g_\perp^- \sin\theta\cos\frac{\Omega}{2} \\ 0.0 & -g_\parallel^- \frac{\cos\theta}{\sqrt{2}}\cos\frac{\Omega}{2} & +g_\perp^+ \sin\theta & 0 & g_\parallel^+ \frac{\cos\theta}{\sqrt{2}} \\ 0.0 & +g_\parallel^- \frac{\cos\theta}{\sqrt{2}}\cos\frac{\Omega}{2} & 0 & -g_\perp^+ \sin\theta & g_\parallel^+ \frac{\cos\theta}{\sqrt{2}} \\ 0.0 & g_\perp^- \sin\theta\cos\frac{\Omega}{2} & g_\parallel^+ \frac{\cos\theta}{\sqrt{2}} & g_\parallel^+ \frac{\cos\theta}{\sqrt{2}} & -g_\perp^- \sin\theta \end{pmatrix}, \quad \text{(A.7)}$$

$$H_{SO} = \sin\left(\frac{\Omega}{2}\right)\begin{pmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & t_x + it_y & tx - it_y & i\sqrt{2}t_z \\ 0 & t_x - it_y & 0 & 0 & 0 \\ 0 & t_x + it_y & 0 & 0 & 0 \\ 0 & -i\sqrt{2}t_z & 0 & 0 & 0 \end{pmatrix}, \quad \text{(A.8)}$$

where $\Omega = \arctan\left(\sqrt{2}\frac{t_c}{\epsilon}\right)$, t_c is the tunnel coupling, $t_{SO} = (t_x, t_y, t_z)$ the spin-orbit vector, ϵ the detuning, B the magnetic field strength, and θ the field's vector angle from the plane.

This Hamiltonian is crucial for the simulation and understanding of the physics of the singlet-triplet hole spin experiments conducted in this work, particularly in Chapter 5.

Appendix B

The Schrieffer-Wolff transformation

The study of quantum systems often involves dealing with complex Hamiltonians that describe the dynamics of the system. In many cases, these Hamiltonians can be separated into a dominant part and a smaller perturbation. The Schrieffer-Wolff transformation is a powerful tool that allows us to derive such an effective Hamiltonian[195]. The transformation is especially useful in the study of strongly correlated electron systems, where the electron-electron interaction cannot be treated perturbatively[196]. This has found applications in various areas of condensed matter physics, including the study of the Kondo effect[195], quantum magnetism[197], and high-temperature superconductivity[198]. We use it here to derive, to first order, an effective Hamiltonian without the (2,0) charge basis state from a simulation of the dynamics of a system of coupled Loss DiVincenzo qubits, as reported in Chapter 6.

In essence, the Schrieffer-Wolff transformation is a canonical transformation that is used to derive an effective Hamiltonian that describes the low-energy physics of a system. The transformation involves the calculation of a generator function, an anti-Hermitian operator that brings the Hamiltonian into a block-diagonal form. However, the calculation of this generator function can be challenging, especially for complex systems. The following section delineates a systematic method for calculating the Schrieffer-Wolff generator function. This method provides a step-by-step procedure that can be applied to a wide range of systems. We start by separating the Hamiltonian into diagonal and off-diagonal components. We then use the commutation relation to derive a set of equations for the elements of the generator matrix.

Principles

First, let's consider a Hamiltonian H that includes both diagonal and off-diagonal terms:

$$H = H_0 + V \tag{B.1}$$

We can write H_0 in terms of a set of operators α_i that satisfy the commutation relations

$$[\alpha_i, \alpha_j^\dagger] = \delta_{ij} \tag{B.2}$$

The off-diagonal part of the Hamiltonian can be written in terms of a set of operators β_i that satisfy the same commutation relations as the α_i operators.

Now, we can perform a unitary transformation on the Hamiltonian H using the transformation operator

$$U = e^S \tag{B.3}$$

where S is an anti-Hermitian matrix[196]. The effective Hamiltonian obtained in this way is a perturbation series in powers of the off-diagonal terms, with the leading term being the diagonal part of the Hamiltonian. This means that the effective Hamiltonian can often be truncated after a few terms, and the resulting approximation can be used to study the system. This transformation yields a new Hamiltonian H', given by

$$H' = UHU^\dagger = H_0 + V' + \mathcal{O}(V^2) \tag{B.4}$$

where V' is the transformed off-diagonal part of the Hamiltonian, given by

$$V' = [S, H_0] + \frac{1}{2}[S, [S, V]] \tag{B.5}$$

The idea behind the Schrieffer-Wolff transformation is to choose S in such a way that V' is small and can be treated as a perturbation[195]. To achieve this, we can choose S to be of the form

$$S = \sum_{i,j} c_{ij} \alpha_i^\dagger \alpha_j \tag{B.6}$$

where the coefficients c_{ij} are chosen to satisfy the condition

$$[H_0, S] = [H_0, \sum_{i,j} c_{ij} \alpha_i^\dagger \alpha_j] = -V \tag{B.7}$$

Substituting this form of S into the expression for V' and using the commutation relations for the α_i and β_i operators, we obtain

$$V' = -[H_0, S] = \sum_{i,j} c_{ij} (\beta_i^\dagger \alpha_j - \alpha_i^\dagger \beta_j) \tag{B.8}$$

Note that V' is now proportional to the off-diagonal terms $\beta_i^\dagger \alpha_j - \alpha_i^\dagger \beta_j$. Since these terms are typically small compared to the diagonal terms, we can treat V' as a perturbation and obtain an effective Hamiltonian for the system that includes only diagonal terms plus a perturbation in V'. If our Hamiltonian H has diagonal terms H_0 and off-diagonal terms V, we can propose a generator matrix S for the Schrieffer-Wolff transformation[199,200]. This generator matrix S will have the same structure as $\eta = [H_0, V] = \sum_{i,j} (\beta_i^\dagger \alpha_j - \alpha_i^\dagger \beta_j)$, with coefficients such that $-V = [S, H_0] = \sum_{i,j} c_{i,j} [\eta_{i,j}, \alpha_i^\dagger \alpha_j]$.

Generator matrix

The generator matrix, S, is an off-diagonal anti-hermitian matrix, that can be written as

$$S = \begin{bmatrix} 0 & \gamma_{12} & \cdots & \gamma_{1n} \\ -\gamma_{12} & 0 & \cdots & \gamma_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ -\gamma_{1n} & -\gamma_{2n} & \cdots & 0 \end{bmatrix}. \tag{B.9}$$

We can calculate the coefficients γ_{ij} by using the commutator relation, $-V = [S, H_0]$. The left-hand side of the equation is $-V$, which is an off-diagonal matrix with elements $-\beta_{ij}$. The commutator $[S, H_0] = SH_0 - H_0S$ will be an off-diagonal matrix with elements given by $(\alpha_{jj}\gamma_{ij} - \alpha_{ii}\gamma_{ij})$ for $i \neq j$. Setting $-V = [S, H_0]$, we can equate the elements of the matrices on both sides of this equation to get, $\alpha_{jj}\gamma_{ij} - \alpha_{ii}\gamma_{ij} = -\beta_{ij}$. Solving for γ_{ij}, we find:

$$\gamma_{ij} = \frac{-\beta_{ij}}{\alpha_{jj} - \alpha_{ii}}. \tag{B.10}$$

This gives the coefficients γ_{ij} of the generator matrix S. Note that these coefficients satisfy $\gamma_{ij} = -\gamma_{ji}$, as required for S to be anti-Hermitian.

Appendix C

The Haar Measure

Quantum gate fidelity is a crucial metric for assessing the performance of quantum gates, which are the fundamental building blocks of quantum algorithms. The fidelity provides a measure of how closely the operation of a quantum gate matches an ideal unitary transformation. To evaluate the average performance of a quantum gate across the entire state space, it is necessary to average its fidelity over a set of initial states. This averaging process utilises the concept of the Haar measure, which is a unique, translation-invariant measure on the unitary group of quantum states[173]. We utilise the Haar measure to calculate the average gate fidelity in order to derive optimal quantum control pulse schemes for the operation of spin qubits in chapter 6.

The Haar measure provides a uniform distribution over the group of unitary matrices, ensuring that every possible quantum state is equally likely to be sampled. Formally, for a quantum system described by a Hilbert space of dimension d, the Haar measure over the unitary group $U(d)$ allows for the integration over the group such that:

$$\int_{U(d)} f(U)\, d\mu(U) = \int_{U(d)} f(VUV^\dagger)\, d\mu(U), \tag{C.1}$$

where f is any integrable function over $U(d)$, and V is any fixed unitary matrix. This equation illustrates the invariance of the Haar measure under unitary transformations.

The average gate fidelity of a quantum gate can be computed by averaging the fidelity over the Haar measure[201]. If $\mathcal{E}$ represents the quantum channel corresponding to the real operation of the gate, and U is the ideal unitary operation, the fidelity between $\mathcal{E}$ and U for a state ρ is given by:

$$F(\rho, U, \mathcal{E}) = \left(\mathrm{Tr}\sqrt{\sqrt{U\rho U^\dagger}\mathcal{E}(\rho)\sqrt{U\rho U^\dagger}} \right)^2. \tag{C.2}$$

The average gate fidelity is then obtained by integrating this fidelity over all pure states

with respect to the Haar measure:

$$\bar{F}(U,\mathcal{E}) = \int_{U(d)} F(\rho, U, \mathcal{E})\, d\mu(U). \tag{C.3}$$

There are various methods for estimating the Haar measure for two-qubit systems including the Clifford state set[174], mutually unbiased bases[202], and Pauli twirling[175].

The set of states that can be generated by the action of the Clifford group on a standard initial state, is referred to as the Clifford state set. This set exhibits symmetries that are highly representative of the entire state space, albeit not uniformly as in the Haar measure. Specifically, the Clifford state set forms a unitary 2-design, a set of unitaries that mimics the statistical properties of the Haar measure up to the 2-nd moment. The Clifford group for a 2-qubit system consists of 11520 elements and is generated by the following gates, the Hadamard gate (H), the Phase gate (S), and the Controlled-NOT gate (C_{NOT})[174].

MUBs are sets of orthonormal bases in which the probability distribution of measurement outcomes in one basis provides no information about the outcomes in another[202].

The Pauli twirling technique simplifies the process of approximating the Haar measure by replacing a general noise channel with a "twirled" channel that is easier to handle computationally[175]. The Pauli twirling group for 2 qubits, consists of 16 two-qubit operators constructed from tensor products of the single-qubit Pauli matrices (X, Y, Z) and the identity matrix (I).

We demonstrate the performance of each of these methods in approximating the fidelity calculated using the full Clifford group (**Figure** C.1). Here, we perform the time evolution of a two-qubit quantum system under an imperfect unitary U_C and calculate the average fidelity with respect to a target unitary U_T over a set of initial states which approximate the Haar measure. We then calculate the dependence of the infidelity on the number of sampled initial states from the Clifford state set and evaluated the error of these fidelity estimates (**Figure** C.1).

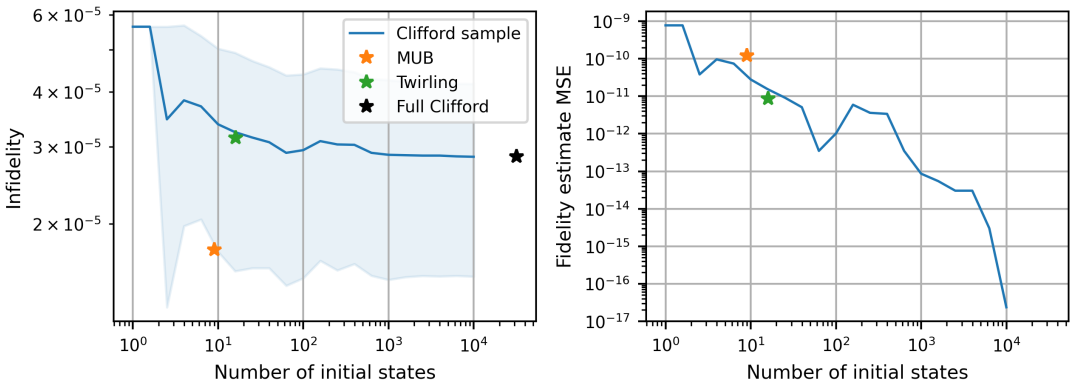


Figure C.1: Comparison of two-qubit gate infidelity and Haar measure approximation accuracy across different sampling methods. A simulated two qubit system is evolved under an imperfect unitary U_C and the average fidelity with respect to a target unitary U_T is evaluated over a set of initial states which approximate the Haar measure. **Left**: The average infidelity of a two-qubit quantum gate as a function of the number of initial states sampled from the mutually unbiased bases (MUB), the Pauli twirling group, and a numerical sampling of Clifford states. **Right**: The mean squared error (MSE) of the fidelity estimates relative to those obtained from the full Clifford state set, showcases the variance in estimation precision with an increasing number of initial states.