

Department of Physics
Imperial College London

Optimizing Quantum Dynamics in the Age of Machine Learning

Frédéric Sauvage

Thesis submitted in partial fulfilment of the
requirements for the degree of
Doctor of Philosophy in Physics, August 2021

Supervised by Prof. Florian Mintert

The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a Creative Commons Attribution-Non Commercial 4.0 International Licence (CC BY-NC). Under this licence, you may copy and redistribute the material in any medium or format. You may also create and distribute modified versions of the work. This is on the condition that: you credit the author and do not use it, or any derivative works, for a commercial purpose. When reusing or sharing this work, ensure you make the licence terms clear to others by naming the licence and linking to the licence text. Where a work has been adapted, you should indicate that the work has been changed and describe those changes. Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.

Abstract

Optimizing the dynamics of quantum systems enables the design of high precision experiments and the development of quantum technologies. To date, such optimizations have been predominantly performed based on theoretical models and numerical simulations; given the size and intricacy of the systems that can now be controlled, this approach is reaching its limits. Optimizing dynamics based solely on experimental data provides a mean to exceed these limitations. However, the probabilistic nature of quantum measurements, combined with the relatively low repetition rates and high noise levels in many current experiments, render these optimizations uniquely difficult. Inspired by recent developments in the field of machine learning, we develop new methods for the efficient optimization of quantum dynamics in experimental situations.

First, we establish Bayesian optimization as a methodology well suited for these optimization problems. After thoroughly assessing its benefit on a paradigm problem of quantum optimal control, we refine the original framework to take into consideration the statistical features of quantum measurements. This allows to maximize the utility of each measurement data and results in enhanced convergence of the framework.

Going further, we investigate an aspect of optimization often ignored, namely the choice of the figure of merit. This figure, which for any optimization problem is identified as the quantity to be maximized, is not unique, thus leaving room for its refinement. After establishing criteria for adequate figures of merit, we show that improved figures, compared to several canonical ones, can be designed.

Finally, to fully reap the benefits of experimental optimizations, it is often desirable that not only one but many related optimizations are performed concurrently. To this intent, we propose two novel frameworks which are found to yield substantial improvements compared to existing methodologies.

Declaration

I hereby declare that the work presented in this thesis is my own work. Any information used from the work of others has been acknowledged and referenced appropriately.

Frédéric Sauvage,

October 4, 2021

Contents

Abstract	i
Introduction	1
I Background	6
1 Optimization tasks	7
1.1 Quantum optimal control	7
1.2 Variational quantum algorithms	12
1.3 Scope of the thesis	15
2 Methods	18
2.1 Optimization frameworks	19
2.1.1 Gradient methods	19
2.1.2 Gradient-free methods	23
2.1.3 Bayesian Optimization	25
2.2 Estimation of the figure of merit	31
2.2.1 Measuring quantum systems	32
2.2.2 Estimation by means of importance sampling	35
2.2.3 Application to the state fidelity	39
II Towards more efficient optimization	43
3 Bayesian optimization for efficient experimental optimization	44
3.1 Application and benchmarking in the context of ultracold–atoms systems	44
3.1.1 The superfluid to Mott insulator control problem	45

3.1.2	Results and discussions	49
3.2	Optimization with poor statistics	54
3.2.1	Bayesian optimization with binomial noise	55
3.2.2	The superfluid to Mott insulator transition revisited	61
3.2.3	Further results and outlook	65
3.3	Concluding remarks	68
4	Engineering the figure of merit	71
4.1	Desirable properties	73
4.2	Engineering and assessment of alternative figures of merit	76
4.2.1	Preparation of separable states	76
4.2.2	Preparation of stabilizer states	81
4.2.3	Alternative to the process fidelity	87
4.3	Concluding remarks	95
5	From one to many optimizations	97
5.1	Variational quantum algorithms with information sharing	98
5.2	Optimal control of continuous families of gates	105
5.2.1	Methodology	106
5.2.2	Results for problems of quantum gate engineering	111
5.2.3	Conclusion and outlook	118
5.3	Concluding remarks	120
Conclusion		121
Bibliography		122
Appendix A Bayesian Optimization		146
Appendix B Estimation		150
Appendix C Random unitaries and states		153
Appendix D Stabilizer states		159

List of Tables

5.1	Control of continuous families of 2– and 3–qubit gates.	116
-----	---	-----

List of Figures

1.1	Optimization tasks of interest.	17
2.1	Optimization based on experimental data.	18
2.2	Predictions with a Gaussian process.	28
2.3	Bayesian optimization (BO) of a toy problem.	30
2.4	Estimation with importance sampling.	36
3.1	The superfluid to Mott insulator (SF–MI) control problem.	46
3.2	Linear and optimized drivings for the SF–MI control problem	49
3.3	Benchmarking BO for the SF–MI control problem.	52
3.4	Comparison between binomial noise and its Gaussian approximation.	56
3.5	BO of a toy problem with single-shot measurements.	60
3.6	BO with poor statistics for the SF–MI control problem.	63
3.7	BO with poor statistics for state preparation problems with quantum circuits. .	66
4.1	Sketch of a good and a bad figure of merit.	72
4.2	Estimation accuracy for the state fidelity and local fidelity.	78
4.3	Comparison of two figures of merit for the SF–MI control problem.	80
4.4	Comparison of two figures of merit for tasks of GHZ state preparation.	85
5.1	Bayesian optimization with information sharing (BOIS).	101
5.2	Comparison of BO with different information sharing strategies.	103
5.3	Quantum optimal control of continuous families of targets.	109
5.4	Quantum optimal control of a family of single-qubit rotations.	113
5.5	Optimal controls and times for single-qubit rotations.	115

Introduction

Following decades of theoretical and experimental progress, quantum technology has matured to a widely recognized field, holding hopes for pervasive applications in a not-too-distant future. Quantum computers, albeit small and noisy, are now routinely accessed through the cloud [1], and quantum supremacy has recently been claimed [2, 3]. Already the ability to manipulate physical systems at elementary levels, whilst taking into consideration quantum mechanical effects, has permitted the realization of some of the most accurate measurement devices [4, 5]. These, in turn, can be employed to test further fundamental physic proposals and refine our understanding of the laws of nature [6, 7]. Still, more general applications of quantum technology are yet to be demonstrated.

The most anticipated applications – including the simulation of molecules and materials beyond the reach of classical computation [8], solving algebraic and graph problems [9, 10], and many more besides [11–13] – often presume the availability of flawless quantum devices. However, the current state of affairs, known as the noisy intermediate-scale quantum (NISQ) era [14], is rather characterized by relatively small quantum systems, subject to imperfect control and prone to errors. Crucially, due to the lack of systematic correction such errors accumulate, rendering any attempt of quantum computation or simulation quickly impractical. Ultimately, it is thought that these limitations will be overcome by means of quantum-error correction [15, 16] enabling the detection and correction of faulty operations faster than they occur. The advent of such fault-tolerant quantum devices is actively pursued, but is only expected in several years or decades, and requires steady improvement along the way. In the meantime, a new breed of quantum algorithms is being developed, aiming at exploiting the limited quantum resources that are already available, and taking into consideration their shortcomings to perform practical tasks. A significant part of this effort to develop future generations of quantum platforms and also to make the most of current technology can be seen through the lens of optimization.

A recurring theme in the improvement of quantum computers, and more generally within quantum experiments, is the necessity to devise adequate controls. Designing control protocols which allow to realize a specific quantum dynamics, as accurately and as fast as possible, forms the field of quantum optimal control (QOC). QOC [17, 18] proposes to formulate this design problem as an optimization: a figure of merit quantifying the success of a given control sequence is established, and controls are adjusted to maximize this figure. In particular, when applied to

systems where imperfections and experimental limitations need to be taken into account, QOC provides the tools to discover non trivial controls resulting in improved performances. Furthermore, due to its generic optimization approach, QOC has emerged as a versatile methodology which has been employed in virtually all implementations of quantum physics.

The optimization of (typically time-dependent) control fields has been predominantly performed based on theoretical models and numerical simulations. This approach has the advantage of making the most of cheap and abundant classical computing resources. However, it assumes a precise characterization of the system of interest and the ability to numerically simulate it. Given the increasingly complex quantum systems that we are gaining control over, it often becomes infeasible to find a sufficiently accurate model, let alone to simulate its dynamical behaviour. The design of control fields based solely on experimental observations, hence without resort to modeling, allows to circumvent such limitations. Still, performing optimizations based on measurements remains challenging. Due to the intrinsic probabilistic nature of quantum measurements, each step of the optimization requires many repetitions of the same experiment in order to accurately estimate the relevant figure of merit. This experimental effort, combined with the scarcity of near-term quantum resources, and the additional sources of noise found in any experimental setup, make such optimization uniquely demanding. Devising efficient optimization frameworks, robust to these various stochastic effects, which can identify close-to-optimal controls in realistic experimental conditions, would allow us to unlock the full range of possibilities offered by this measurement-based QOC approach.

In addition to these problems of accurate control, optimization is also at the core of the nascent field of variational quantum algorithms (VQAs) [19]. With error-corrected devices remaining years or decades away, it has become urgent to develop quantum algorithms which are tolerant to near-term technological limitations. For that purpose, VQAs propose to mitigate control errors and noise, found in current quantum devices, by means of optimization. Broadly speaking, a VQA consists of the preparation of a quantum state given a quantum circuit which parameters can be freely varied. These circuit parameters can then be iteratively refined to minimize a figure of merit characterizing the problem to be solved. Crucially, this figure is measured experimentally, thus permitting to directly probe quantum states which would have been intractable in classical computations. Additionally, resorting to optimization has the benefit of reducing the length of the quantum circuits involved, and compensating for some of the noise presents in NISQ devices. Due to these appealing features, combining advantages of both quantum and classical routines, many problems of interest have now been formulated in such variational forms [20–23].

For any VQA to be of practical use, it is crucial that the optimization of its circuit parameters can be performed in realistic experimental effort. Akin to problems of QOC based on experimental data, several challenges impeding the success of these optimizations have been identified

[22, 23]. To date, these have limited the largest experimental realizations of VQAs [24, 25] to address problems which could have been solved numerically. Overcoming these difficulties is required to further explore the potential offered by VQAs, with the hope that they could find practical applications behind what can be achieved classically.

The main goal of this thesis is to develop novel methods enabling the efficient optimization of quantum systems dynamics – depending on either analog control fields (in problems of QOC), or quantum circuits parameters (in problems of VQAs) – based on experimental data. In particular, we aim at leveraging advances from the field of machine learning (ML) [26] to facilitate these optimizations. ML provides a fresh perspective on optimization¹, along with practical tools. Among the flagship ML techniques, deep neural networks [27] have already made their way to the realm of quantum physics, and have been employed for problems of classification of quantum phases [28], of compact representation of quantum states [29, 30], and of control of quantum systems [31, 32]. However, typical ML algorithms often assume the availability of vast amount of data and processing power. Less studied, but more relevant to our purpose, is the field of probabilistic machine learning [33, 34]. Similar to ML, probabilistic ML provides the methodology to model and to predict based on observed data, but comes with the additional benefit of capturing the uncertainty entailed in such predictions. Such a feature is critical when data is scarce and noisy, as of interest here, and provides the basis for principled decision making as needed in an optimization context.

Outline of the thesis

Part I provides the necessary background for the thesis. In Chapter 1, we introduce the optimization tasks of interest, namely problems originating from the fields of QOC and VQAs. A summary of the problems which will be considered is provided in Figure 1.1. This is followed in Chapter 2 by a presentation of the methods permitting the optimization of the dynamics of a quantum system. First, we review widely used optimization frameworks, and also introduce Bayesian optimization from the probabilistic ML toolbox. Second, we detail elements of the statistical methodology necessary for the characterisation of quantum systems properties based on experimental data. The results of this thesis are then thematically grouped over three chapters forming Part II.

In Chapter 3, we study the choice of an adequate routine to perform optimization. In particular, we argue for the adoption of Bayesian optimization as particularly well-suited for problems of optimization based on experimental outcomes. First, its benefit is demonstrated for a paradigm control problem of the crossing of a quantum phase transition (in faster-than-adiabatic times). Next, we improve on the existing Bayesian optimization framework. By carefully integrating

¹”Learning” in machine learning invariably entails ”optimizing”.

statistical details of quantum measurements, we enable the framework to converge even in cases where only a few measurement repetitions (up to a single repetition) are performed at each optimization step. Overall, this allows optimization with reduced quantum resources, a particularly appealing feature in the current NISQ era where quantum data is scarce and expensive.

In Chapter 4, we study the importance of the choice of the figure of merit. This figure, which is identified as a measure of success guiding any problem of optimization, is not unique and thus can be refined. First, quantitative criteria for appropriate figures of merit are established. Then, these are used to study and devise more experiment-friendly figures, and we show in distinct examples that alternatives to figures commonly used could facilitate their corresponding optimization. As orthogonal to the choice of the optimization routine, these figures can provide an advantage to be combined with the improvements showcased in Chapter 3.

In Chapter 5, we take a step forward and consider the situation where many related optimizations need to be performed all together. As we will see, this situation is common in quantum technological applications, and calls for principled optimization methodologies. Rather than treating each of the many optimization tasks independently, we propose two frameworks aiming at exploiting their commonalities. A first framework is developed for families of VQAs such that measurement data can be shared in between each of the individual optimizations. A second framework is proposed to extend the QOC methodology from the optimization of a single target dynamics to the optimization of continuous families of such targets. Its training is performed assuming exact numerical simulations, and we discuss how it can also be optimized directly based on experimental data. Overall, the two schemes presented are shown to solve families of problems with significantly reduced effort compared to existing methods.

Overall, we address the problem of efficiently optimizing the dynamics of quantum systems from several angles, and devise novel methodologies which are found to surpass current approaches. The proposed frameworks are developed with general applications in mind, rather than being tied to too specific experimental details. The results presented are either obtained on public-access quantum chips or numerical simulations, in which case care is taken to incorporate imperfections and limitations mimicking experimental constraints. While demonstrated on relatively small system sizes and problems, we discuss and challenge the scalability of the methods presented along the way. Overall, it is our belief that the methods developed in this thesis can find application in a wide array of situations relevant to the current effort to scale quantum technology to practical problems.

The majority of the results presented in this thesis originate from:

- R. Mukherjee, FS, H. Xie, R. Löw, F. Mintert
Preparation of ordered states in ultra-cold gases using Bayesian optimization
New Journal of Physics 22, 075001 (2020)
- FS, F. Mintert
Optimal Quantum Control with Poor Statistics
PRX Quantum 1, 020322 (2020)
- C. Self, K. Khosla, A. Smith, FS, P. Haynes, J. Knolle, F. Mintert, M. Kim
Variational Quantum Algorithm with Information Sharing
npj Quantum Information 7, 116 (2021)
- S. Greenway, FS, K. Khosla, F. Mintert
Efficient Assessment of Process Fidelity
Physical Review Research 3, 033031 (2021)
- FS, F. Mintert
Optimal control of families of quantum gates
In preparation

Other works undertaken along the course of my doctoral research can be found in:

- A. Kuroś, R. Mukherjee, W. Golletz, FS, K. Giergiel, F. Mintert, K. Sacha
Phase diagram and optimal control for n -tupling discrete time crystal
New Journal of Physics 22, 095001 (2020)
- A. Pitchford, A. A. Rakhubovsky, R. Mukherjee, D. W. Moore, FS, D. Burgarth, R. Filip, F. Mintert
Optimal non-classical correlations of light with a levitated nano-sphere
arXiv preprint, 2006.15050 (2020)
- FS, S. Sim, A. A. Kunitsa, W. A. Simon, M. Mauri, A. Perdomo-Ortiz
FLIP: A flexible initializer for arbitrarily-sized parametrized quantum circuits
arXiv preprint, 2103.08572 (2021)

Part I

Background

Chapter 1

Optimization tasks

In this chapter, we introduce the optimization tasks of interest. These originate from the fields of quantum optimal control, reviewed in Section 1.1, and variational quantum algorithms, reviewed in Section 1.2. We conclude by defining more precisely the objective of this thesis in Section 1.3, and provide in Figure 1.1 a summary of concrete problems which will be investigated to assess the performances of the methods proposed in Part II.

1.1 Quantum optimal control

Technological advances in the control toolbox have enabled the realization of quantum physics experiments at unprecedented level of sophistication [35–39]. With this ability to precisely control physical systems at the elementary level, has emerged the question of how to design such controls optimally. Exploiting the possibility to vary the amplitude of a laser pulse in time was first envisioned in [40] to enhance the product of chemical reactions, and was experimentally demonstrated in [41, 42]. While originally intended for the tailoring of laser fields, the idea of optimizing any experimental degrees of freedom (which can be time-dependent or static) to achieve a specific objective has now been applied in many more situations, forming the field of quantum optimal control (QOC) [17, 18]. Due to its versatility, QOC has been recognized as a pillar supporting the development of quantum technology [43] to the benefit of, for instance, quantum computing, simulation and metrology experiments. In the following, we provide an overview of its current applications, before formalizing the general task of QOC.

State-of-the-art

Application to quantum computers. Digital quantum computers aim at storing and manipulating information encoded over register of qubits. A typical realization of a quantum computer relies on the implementation of few elementary operations, known as gates, acting on single qubit or pairs of qubits. When repeated and combined, such gates enable any operation

to be carried out on the whole register of qubits up to arbitrary precision, that is, permit universal quantum computing [44]. Any realization of a quantum computer strives for high speed and high accuracy gates which can be refined by means of QOC. QOC has been successfully applied to such quantum gate engineering problems over a wide array of platforms, such as, superconducting qubits [45, 46], trapped ions [47–49] or neutral atoms [50–52]. In particular, when control limitations, noise and other experimental imperfections are taken into account QOC enables the discovery of control protocols far from trivial but resulting in improved performances [53–56]. Most recent applications of QOC have considered operations acting on larger subsystem than the 2-level qubit model, including qudits or logical qubits [57–61], which are hoped to prove more robust to errors. Additionally, while QOC has primarily focused on the enhancement of gates operated individually, it has become of great interest to also improve on their realizations when acted in parallel [49, 62–64]. It is such ability to perform accurately and simultaneously many operations which will enable the scaling of quantum computing platforms.

Application to quantum simulators. As an appealing alternative to digital quantum computation, analogue quantum simulation [65] propose to utilize quantum resources toward the simulation of quantum physics, that is, without resorting to exponentially demanding numerical simulations. Broadly speaking, quantum simulation aims at imitating the behavior of a quantum system of interest by means of an auxillary quantum system that is easier to control and to measure. Among the many platforms suited for this purpose, ultra-cold atoms [66–68] and ions [69, 70], are often recognized as the most promising ones. These offer high levels of controllability, including the ability to realize tuneable potentials and interactions, complemented with the capability to resolve particles at the individual level. Taken together, these characteristics enable the synthesis and study of a vast range of physics. Still, realizations of such experiments remain demanding and prone to errors. In particular, the many steps involved in their preparation are often relatively slow, and can be improved by means of optimization. For instance, the trapping, cooling, displacement and loading of atoms (or ions) into a desired potential, have all been refined by means of QOC [71–76]. Overall, QOC offers the tools to improve on the preparation and manipulation of quantum states of increasing complexity, and ultimately to increase the quality and reach of current quantum simulators.

Application to quantum sensing. Finally, quantum sensing and metrology [77, 78] can also benefit from the design of better controls by means of QOC. Exploiting specificities of quantum mechanics such as entanglement, squeezing, and freedom in the measurements which can be performed can, in theory, enable precision measurement beyond what is possible classically [79, 80]. Still, maintaining such advantage in the presence of experimental noise requires great care. General quantum sensing protocols involve the creation of an initial quantum state sensitive to a physical quantity to be measured, its evolution, and a measurement of its final state. All these aspects can be refined by QOC to enhance the sensitivity of the system. In particular, the accurate preparation of initial entangled or squeezed states, despite experimental imperfections,

have been optimized in many different experiments [81–86]. Additionally, the driving of the system during its evolution, can also be tailored to limit detrimental environmental effects and to magnify the susceptibility of the system to the field to be measured [87–93].

Overall, QOC has found applications in many quantum experimental realizations. Despite the wide disparities in the practical details of the problems tackled, any QOC scheme can always be conceptualized as (i) the evolution of a quantum system under control fields which can be varied, (ii) the choice of a figure of merit quantifying the success of a given control, and (iii) the maximization of such figure. In the following subsection we review further this general formalism.

Formulation of the problem

A typical problem of optimal control is understood as the identification of a time-dependent control $f(t)$ in order to achieve a desired dynamical behavior of the system. In the realm of quantum physics, the evolution of a closed system is prescribed in terms of an ordinary differential equation, namely the Schrödinger equation (taking $\hbar = 1$)

$$i\frac{d}{dt}|\psi(t)\rangle = H(f(t))|\psi(t)\rangle, \quad (1.1)$$

which relates changes in the quantum state $|\psi(t)\rangle$, belonging to some Hilbert space \mathcal{H} , to the Hamiltonian operator H encoding the physical details of the controlled system.

In tasks of quantum state preparation, one aims at transforming a fixed initial state $|\psi(t=0)\rangle = |\psi_0\rangle$ into a target state $|\psi_{tgt}\rangle$ at later time $t = T$. For such purpose, the merit $F(f)$ of the control function f is measured in terms of the similarity between the final state $|\psi(T)\rangle$ under control, and the targeted one. This is often defined as the overlap $F(f) = |\langle\psi_{tgt}|\psi(T)\rangle|^2$ between the two states, which adopts a maximal value of $F = 1$ for identical prepared and target states. The corresponding task of QOC is the identification of the control f which maximizes the figure of merit F .

While it was previously assumed a fixed initial state $|\psi_0\rangle$, a more general problem of QOC aims at controlling the evolution of a system given an arbitrary initial state. Due to the linearity of Equation (1.1), the entire dynamics of a closed quantum system is uniquely described in terms of the unitary operator $U(t)$ that propagates any initial state $|\psi_0\rangle$ to the state $|\psi(t)\rangle = U(t)|\psi_0\rangle$ satisfying Equation (1.1). It can be shown that the evolution of such propagator follows the differential equation

$$i\frac{d}{dt}U(t) = H(f(t))U(t), \quad (1.2)$$

given initial conditional $U(t=0) = I$, with I being the identity operator acting on \mathcal{H} . Denoting U^{tgt} the target propagator to be realized at fixed time $t = T$, such QOC problem is defined

as the maximization of a similarity measure $F(f) = S(U^{tgt}, U(T))$ between the target and the realized propagators. Such objective is central to the engineering of faithful quantum computing gates (where the target U^{tgt} is the gate to be realized) and is commonly referred as a task of quantum gate engineering.

More broadly, a QOC problem can always be formulated as the maximization of a figure of merit $F(f)$ depending on one (or several) time-dependent control function f . This figure F can be the similarity between states or propagators at final time T , as was just discussed, but could also quantify any other properties of the dynamical system, and could even depend on its trajectory over time $t \in]0, T[$. Additionally, when treating open quantum systems (that is when the system of interest interacts with an external environment), the representation of the system state as a pure state $|\psi(t)\rangle$ needs to be generalized to a mixed state $\rho(t)$ which evolution $\rho(t) = \Lambda(t)(\rho(t=0))$ is understood in terms of a quantum process $\Lambda(t)$ rather than the unitary evolution $U(t)$ which was described. While the dynamics of such systems are notoriously more involved than for closed systems, the exact same QOC principles, of establishing and maximizing a figure of merit F , equally apply.

With the notations introduced, a general problem of QOC is formulated as the maximization problem:

$$f^* = \arg \max_f F(f). \quad (1.3)$$

Equation (1.3) could be supplemented with constraints on the control, to reflect, for instance, bandwidth limitations or maximal values that it can adopt. Most often, however, we will enforce these constraints directly when parametrizing the control function $f(t)$, as we now discuss.

From functional to discrete optimization

The optimization problem over functions in Equation (1.3) is often found to be impractical, and is rather reformulated as an optimization over a discrete set of continuous parameters. This is achieved by parameterizing the control $f(t)$ in terms of a finite set of control parameters $\boldsymbol{\theta}$. For instance, it is common to decompose the control function as a weighted sum

$$f(t) = \sum_{i=1}^{P<\infty} \boldsymbol{\theta}_i \phi_i(t) \quad (1.4)$$

of a finite set $\{\phi_i(t)\}_{1 \leq i \leq P}$ of P basis functions. Typical choices of basis include the Fourier harmonics defined as

$$\begin{cases} \phi_i(t) = \cos(\omega_l t), & \text{if } i \text{ is odd} \\ \phi_i(t) = \sin(\omega_l t), & \text{if } i \text{ is even} \end{cases} \quad (1.5)$$

with respect to a finite set of frequencies ω_l . These frequencies can be chosen with respect to well-defined energy scales of the controlled system when possible, but could even be chosen randomly [94]. Alternatively, in the time-domain, piece-wise constant functions

$$\begin{cases} \phi_i(t) = 1, & \text{if } \frac{i-1}{T-1} \leq t < \frac{i}{T-1} \\ \phi_i(t) = 0, & \text{otherwise} \end{cases} \quad (1.6)$$

have also been frequently employed, as they permit faster numerical simulations of the controlled system. The parameterization in Equation (1.4) can further be composed with other functions [46, 94] to enforce, for instance, bounded values or boundary conditions for the control amplitudes, or to incorporate informed guesses on the optimal control functions.

More generally, any parameterization ensuring a well defined correspondence between control parameters and time-dependent control functions can be adopted. In particular, we will often resort to parameterization by means of cubic splines fitted to values of the controls taken at constant intervals, or even to neural networks, both as a flexible way to parameterize smooth control functions with a finite number of parameters. These will be described further in due course.

Given a choice of parameterization, the dependencies on the control function f can be replaced by dependencies on the control parameters $\boldsymbol{\theta}$. For instance, $F(f) \rightarrow F(\boldsymbol{\theta})$, resulting in a formulation

$$\boldsymbol{\theta}^* = \arg \max_{\boldsymbol{\theta}} F(\boldsymbol{\theta}) \quad (1.7)$$

of the general QOC problem to be solved. Despite the apparent simplicity of Equation (1.7), this will be the task that will occupy us for the rest of this thesis.

Taxonomy

The general QOC problem stated in Equation (1.3), or in its discrete form in Equation (1.7), has been approached with a wide range of methodologies [31, 46, 94–99]. It is not our intent to cover the peculiarities of each of these frameworks, still it is insightful to broadly categorize them.

Analytical and numerical optimizations. It is only in restricted cases that analytical solutions of the functional QOC problem expressed in Equation (1.3) can be obtained. Such solutions are most often derived by means of optimal control theory, typically based on the Pontryagin maximum principle [100], but remain limited to low dimensional problems (see for instance [97] for a recent review of optimal control theory applied to QOC). More general problems of QOC, however, evade such analysis, and rather require numerical optimizations over a finite number of parameters such as appearing in Equation (1.7).

Model–based and experiment–based QOC.¹ QOC is most often performed based on assumptions of a theoretical model describing the system under control. In such cases, the optimization of the control field can be performed based on numerical simulations of the model. Such approach bears the advantage of decoupling experimental realizations and numerical simulations, and of employing cheap and abundant classical computation resources. However, it necessitates a precise characterization of the controlled system in the first place, and is limited to problem sizes for which numerical simulations are possible. An alternative path toward QOC is to perform optimization directly based on experimental outcomes [98, 99], thus avoiding discrepancies between an idealized model and its physical realization. Furthermore, this approach does not restrict the optimizations to system sizes which can be numerically simulated. This, however, comes at the expense of more challenging optimization. In particular, the figure F needs to be estimated based on measurement outcomes (as opposed to the exact evaluations obtained in numerical simulations), which requires many rounds of preparations and measurements and limits the number of optimization iteration which can realistically be performed.

1.2 Variational quantum algorithms

In addition to these tasks of quantum control, optimization is also central to the field of variational quantum algorithms (VQAs). With fault-tolerant computers still years or decades away, VQAs aim at exploiting current quantum resources towards practical applications. In particular, it is hoped that some of the limitations imposed by current technology could be overcome by classical optimization. For that purpose, VQAs propose to embed the preparation of a quantum state into a classical optimization loop in charge of refining its preparation. Such hybrid quantum–classical frameworks were first motivated for problems of quantum chemistry, leading to the development of the variational quantum eigensolver (VQE) [20], and of combinatorial optimization, with the quantum approximate optimization algorithm (QAOA) [21]. VQAs have now been extended to a wide array of problems (see [22, 23] and references therein). Fast advances in the field have been facilitated by the development of powerful numerical libraries [101–103], allowing to quickly prototype new algorithms, and by an increased accessibility to cloud quantum computing resources [1]. Still, it remains the case that optimizing such algorithms at scale remains arduous, such that only proofs of concept, over relatively small system sizes (most often involving in between 1 to 10 qubits), have been experimentally demonstrated.

In the following, we present a general formulation of a VQA, and detail some of the main challenges faced when attempting to apply these algorithms to large problem sizes. We conclude

¹This dichotomy between model–based and experiment–based optimizations has found many names in the literature such as open–loop and closed–loop, or in–situ and ex–situ, or off–line and on–line optimizations.

this section by highlighting connections and differences with QOC. These are important as these two fields have grown separately, driven by different motivations, but still share many similarities.

Guiding principles

A general VQA problem consists of:

1. The mapping of a problem to be solved, to a synthetic qubit operator (the problem operator) H_{prob} which ground-state provides a solution to the problem.
2. The preparation of a quantum state $|\psi(\boldsymbol{\theta})\rangle$ by means of a parametrized quantum circuit $U(\boldsymbol{\theta})$ applied to a fixed initial state $|\psi_0\rangle$. This circuit is typically comprised of single- and two-qubit gates, which rotation angles form the set of parameters $\boldsymbol{\theta}$ that can be freely varied. Examples of parametrized quantum circuits can be found in Figure 1.1(c,d).
3. The estimation of the expectation value $\langle H_{prob} \rangle_{\boldsymbol{\theta}} = \text{Tr}[H_{prob}|\psi(\boldsymbol{\theta})\rangle\langle\psi(\boldsymbol{\theta})|]$ (or its gradient) of the problem operator with respect to the prepared state. This is achieved by means of repeated rounds of preparation of the same state $|\psi(\boldsymbol{\theta})\rangle$ and measurements of the constituents of H_{prob} .
4. The variation of the parameters $\boldsymbol{\theta}$ of the circuit in order to decrease the expectation value $\langle H_{prob} \rangle_{\boldsymbol{\theta}}$.

The steps (2-4), which form the outer-loop of classical optimization of the circuit parameters (often referred as the training of a VQA), are repeated until convergence of the algorithm. An optimization is deemed successful if the parameters $\boldsymbol{\theta}^*$ corresponding to the preparation of the ground state of H_{prob} (that is the state minimizing the expectation value $\langle H_{prob} \rangle_{\boldsymbol{\theta}}$) has been found. To be of practical interest, it is necessary that the mapping in step (1) can be performed in polynomial time (with respect to the size of the underlying problem to be solved), and that the number of constituents of H_{prob} , which needs to be estimated independently in step (3), scales at most polynomially with the problem size [104]. Ultimately, the success of any VQA lies in its ability to be trained in realistic conditions, for instance, in a limited number of rounds of state preparations, or in timescales during which the quantum device, used to perform the state preparations, does not have to be re-calibrated.

Since its recent inception, the field of VQAs has quickly developed and is now often perceived as a promising path toward a practical quantum advantage, that is, the ability to exploit quantum resources to solve problems of practical relevance (that could not have been solved by means of classical computations). On one hand, the freedom in the choice of the parameter values of the quantum circuits involved in any VQA permits the reduction of their depths compared with the lengthier (and fixed) circuits employed in non-variational quantum algorithms. Such

depth, which is the minimum number of gates that need to be sequentially performed, directly relates to the time required for a circuit to be executed. Given the limited coherence times² of NISQ devices, a shallow depth (thus entailing short execution times) is a prerequisite to obtain meaningful results. Additionally, errors known to plague NISQ devices can be, at least partially, mitigated during optimization [105]. For instance, systematic errors of over- or under-rotations in the gates implementation, could be compensated with correspondingly altered rotation parameters. For these reasons, VQAs are thought to be well-suited to overcome the shortcomings of NISQ devices, and already many quantum algorithms have been recast in such a variational form [106–108].

Despite these appealing properties, the training of VQAs in systems of more than a few qubits remains challenging. In addition to the fundamental limits imposed by the accuracy of current quantum devices, most of the challenges identified as limiting the scaling of VQAs can be understood from an optimization perspective. In particular, the Barren plateaus [109] phenomenon (this will be discussed further in Chapter 4) has been diagnosed as a generic issue preventing the optimization of a large variety of parameterized circuits when randomly initialized. More generally, the experimental effort in accurately estimating the expectation values of the operator H_{prob} in step (2) often becomes excessive, and, given the relatively low repetition rates of NISQ devices [110], quickly limits the number of optimization steps which can be performed [111]. Such issues, combined with the presence of local minima in the optimization landscape [112, 113], render the optimization of quantum circuits particularly challenging.

While intended for different purposes, VQAs are conceptually similar to QOC when performed based on experimental data. In both cases, a classical optimizer works in tandem with an experimental realization of a quantum system, to minimize a figure³ which is estimated experimentally. Along with their growth, the exact demarcation in between the two fields has become thinner. Still, a few conceptual differences deserve to be highlighted. First, in VQAs optimizations are typically performed at the gate level (that is, over rotation angles of a discrete set of gates) rather than at the pulse level (that is, over analog physical control fields) as is the case in QOC. Nonetheless, gates are always realized by means of physical pulses, such that one could reformulate any VQAs as an optimization over pulses rather than gates [114]. In particular, this allows to bypass the need for a decomposition of the state preparation in terms of gates to be applied sequentially, and could result in shorter preparation times. A second distinctive feature of VQAs, is that due to the specific dynamics induced by the structure of the quantum circuits, exact gradients of the figure of merit can often be estimated experimentally [115]. As such a significant body of research has focused on gradient-based methods for the optimization of VQAs. We discuss further these gradient-based methods and their limitations

²These are the timescales over which a quantum state remains sufficiently decoupled from its external environment.

³For consistency, we call this figure the figure of merit in both cases, albeit it is often referred as cost or objective function in the field of VQAs.

in Section 2.1.1. In any case, it is highly desirable to investigate more efficient optimization frameworks which could ease the training of VQAs.

1.3 Scope of the thesis

As alluded previously, the main goal of this research is to develop novel methodologies enabling the efficient optimization of quantum dynamics based on experimental outcomes. These can benefit both problems of experimental QOC and also VQAs. As such, the techniques presented in Part II will be evaluated on a variety of problems, sketched in Figure 1.1, drawing from these two fields.

While more challenging, experimental optimization allows, in the field of QOC, to circumvent the need for an accurate characterization of the system of interest, and to avoid prohibitively lengthy numerical simulations. This is of critical importance given the size and complexity of the quantum systems that are now controlled. For instance, modern quantum simulation experiments can manipulate hundreds to thousands of individual particles, making them already intractable numerically (at least, without resorting to approximations). Furthermore, in the context of quantum computers, it has become crucial to ensure accurate and simultaneous gate operations over large registers of qubits. However, addressing a targeted array of qubits often results in undesired effects (that are generically referred as cross-talk [116]) affecting other parts of the register. Given that such effects remain often hard to characterize precisely, experimental optimization can provide a viable approach towards the improvement of such intricate platforms. These optimizations are also central to the field of VQAs which bears the hope for a practical quantum advantage in a foreseeable future. More generally, the optimization approach considered in this thesis may be of particular relevance in the current effort of translating quantum physics to quantum technologies in a more industrial capability. This will involve gradually moving from extremely well controlled laboratories to mass production facilities, which calls for efficient and automated calibration and optimizations routines.

It remains the case that quantum dynamics optimizations have been predominantly performed based on models, which evidences the unique difficulties of experimental optimization. Given the several challenges which were identified in this chapter, we will strive to develop methods which should be

1. Efficient. Good solutions need to be obtained within realistic experimental effort. This is particularly relevant in complex experimental setups with low repetition rates.
2. Robust to noise. It is important to recognize that a figure of merit can only be estimated with finite accuracy, and that accurate estimates require the repetition of many

measurements. As such, it is key that the frameworks presented should be robust to noise.

3. Scalable. Ideally, the methods presented should benefit optimization problems defined over large system sizes. Still, most of the results that will be presented in this thesis are obtained based on relatively small problem sizes (either numerically simulated, or executed on NISQ devices). As such, it is important to assess that, at least in principle, the methods advocated here could be applied to larger problems.

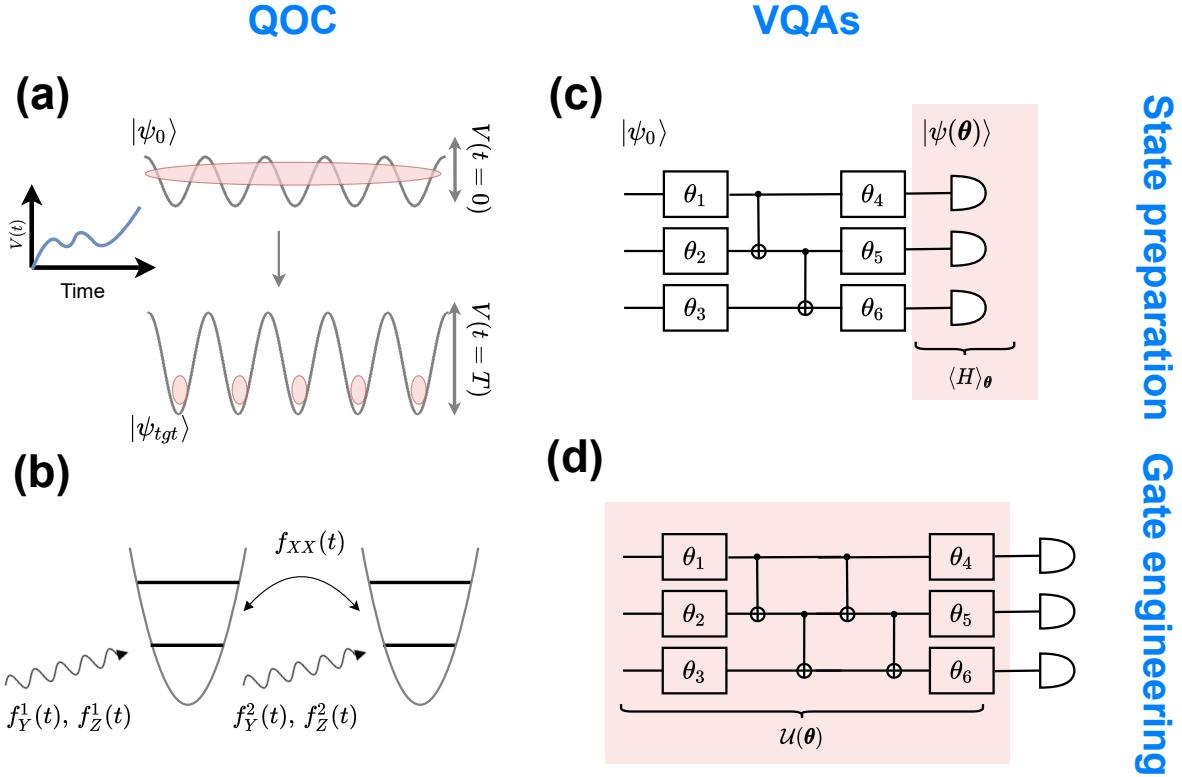


Figure 1.1: The optimization tasks considered in this thesis originate both from the field of quantum optimal control (QOC) and variational quantum algorithms (VQAs). (a-b) In QOC one aims at identifying the time-dependent controls which allow the realization of a desired dynamics at a time $t = T$. For instance in (a), the potential depth $V(t)$ is optimized such as to transform an initial state $|\psi_0\rangle$ into a desired target state $|\psi_{tgt}\rangle$. Alternatively in (b), the local $f_\sigma^i(t)$ and collective $f_{XX}(t)$ controls are jointly optimized such that the induced dynamics of the system resembles as much as possible a target gate U_{tgt} . In VQAs (c-d) the objects of optimization are the parameters (denoted θ_i) of a quantum circuit, and correspond to rotation angles of single or two-qubit gates. In a typical VQA (c) these angles are optimized such that the output state $|\psi(\theta)\rangle$, resulting on the action of the circuit onto an initial state $|\psi_0\rangle$, minimizes the expectation value $\langle H \rangle_\theta = \langle \psi(\theta) | H | \psi(\theta) \rangle$ of an operator H of interest. (d) Alternatively, these angles can be optimized such that the dynamics realized by the circuit $U(\theta)$ mimics a target unitary U_{tgt} . These are examples of optimization of quantum dynamics which are investigated in this thesis.

Chapter 2

Methods

Apart from a few exemplary problems for which analytical solutions can be derived, it is common to approach the optimization problems detailed in the previous chapter equipped with numerical routines, consisting of iterative updates of the parameters to be optimized. As depicted in Figure 2.1, a general framework of optimization can be broken down in two parts: (i) an optimization routine in charge of suggesting new parameters to be tried, (ii) the simulation of the quantum system which permits the evaluation of the figure of merit F to be maximized. When performed experimentally, this last part involves many repetitions of the same experiment in order to gather necessary measurements to the estimation of F . These two aspects of optimization and estimation are central to the success of any optimization task based on experimental data, and are detailed in this chapter.

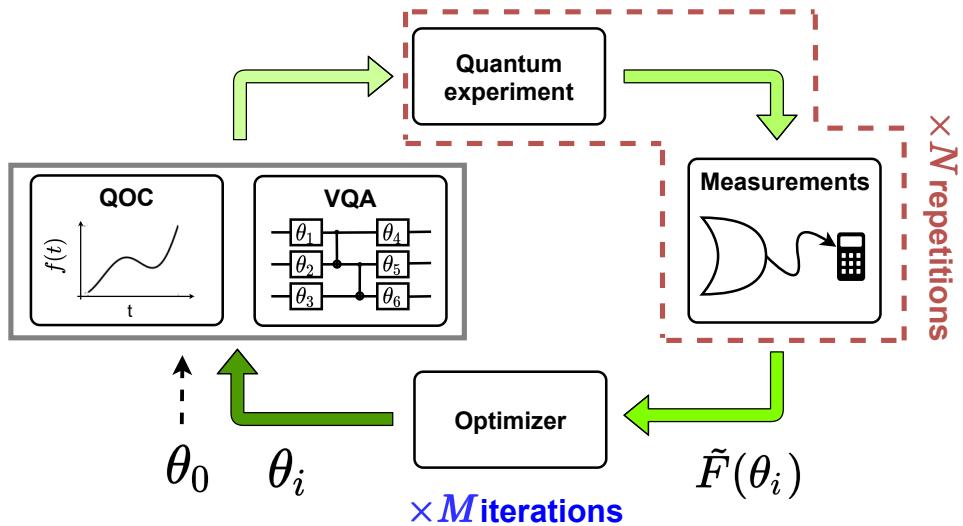


Figure 2.1: Sketch of a general optimization routine based on experimental data. At each iteration, the optimizer suggests a set of parameters θ_i to be tried. These parameters correspond to rotation angles of a quantum circuit in problems of VQAs, or permit to reconstruct a time-dependent control field $f(t)$ in problems of QOC. An estimate $\tilde{F}(\theta_i)$ of the figure of merit $F(\theta_i)$, for this choice of parameters, is obtained based on N rounds of experiments and measurements, and fed back to the optimizer. Such optimization step is repeated M times until convergence of the algorithm, or until an experimental budget has been exhausted.

Among the different optimization routines that are presented in Section 2.1, a particular emphasis is put on the description of Bayesian optimization (Section 2.1.3) which will enable the development of novel optimization frameworks studied in Chapter 3 and Chapter 5. Similarly, when presenting the statistical framework of estimation in Section 2.2, care is taken when deriving the expressions relating accuracy of the estimates and details of the figures to be estimated, as these will be central to the comparison and engineering of alternative figures of merit in Chapter 4.

2.1 Optimization frameworks

Along with the adoption of QOC as a practical tool to improve the control of quantum experiments and the development of VQAs, the importance of employing the right optimization routine has been quickly recognized, spurring the development of a rich variety of strategies [31, 46, 94–96, 115, 117–119]. Here, we review the most prominent methods grouped in a thematic way, highlighting their range of applicability and limit. Gradient-based methods (Section 2.1.1) provide a straightforward path toward optimization but are limited to well-behaved optimization landscapes. Gradient-free methods (Section 2.1.2) offer an alternative in cases where accurate estimation of the gradients is not possible, and are at the core of most optimizations that have been performed based on experimental outcomes. Also considered as gradient-free, but treated separately, we introduce Bayesian optimization (Section 2.1.3) which presents a distinct approach toward optimization, and which will be central to many of the results presented later.

In this section, we consider the general task of maximizing a figure of merit $F : \boldsymbol{\theta} \in \chi \subset \mathcal{R}^P \rightarrow \mathcal{R}$, with $\boldsymbol{\theta}$ the P -dimensional input parameters. In practice, these parameters would correspond to the tuneable control fields in QOC problems, or to parameterized-rotation angles in VQAs. Subscripts are used to index different sets of parameters, for instance as explored during the course of an optimization, while superscripts index the different elements of the same input parameter vector. Hence, $\boldsymbol{\theta}_i^{(j)}$ corresponds to the j -th element of the i -th set of input parameters. In experimental situations, only noisy evaluations of the figure of merit are accessible, and are denoted $y_i \approx F(\boldsymbol{\theta}_i)$, for input parameters $\boldsymbol{\theta}_i$.

2.1.1 Gradient methods

Gradient-based optimization holds a special appeal due to its simplicity. In its simplest flavour, it involves repeated steps of parameters updates in the direction of the steepest ascent of F . This corresponds to the iterative rule $\boldsymbol{\theta}_{i+1} = \boldsymbol{\theta}_i + \eta \nabla_{\boldsymbol{\theta}_i} F$, where the scalar $\eta > 0$ is the step

size, and where

$$\nabla_{\boldsymbol{\theta}} F = \left[\frac{\partial F}{\partial \boldsymbol{\theta}^{(1)}}, \dots, \frac{\partial F}{\partial \boldsymbol{\theta}^{(P)}} \right] \quad (2.1)$$

denotes the vector of gradients of F with respect to each of the parameters entries $\boldsymbol{\theta}^{(j)}$. For sufficiently small step sizes these updates ensure monotonic converge of F to a local maximum. More refined methods, typically resorting to higher order derivatives [119, 120] or adaptive step sizes [121] extend the simple iterative rule presented, and can result in improved convergence.

In addition to its conceptual simplicity, gradient optimization has proven to scale well to high-dimensional problems such as the optimizations of neural networks routinely performed over millions if not billions [122] of parameters. However, as is the case for any local optimization strategy, gradient ascent only guarantees convergence to a local minimum. Furthermore obtaining exact gradients (or unbiased estimates when these need to be estimated experimentally) is not always possible. In such situation, these gradients need to be approximated by means of finite differences. In the following, we review the different flavours of gradient schemes which have commonly been employed when optimizing quantum controls and circuits. We first take a detour to discuss gradients obtained based on numerical simulations of quantum dynamics. This will help us to contrast this numerical situation to the case when gradients need to be assessed experimentally, which is presented afterwards.

Exact gradients in numerical simulations

Obtaining the gradients of numerically simulated quantum systems is at the core of many QOC frameworks. The now widely adopted, gradient-ascent pulse engineering (GRAPE) algorithm [95] took a first step in this direction by providing the methodology permitting the efficient computation of approximate gradients in cases when the control fields are discretized in times (that is, with the piece-wise constant parameterization of Equation (1.6)). Such methodology was subsequently extended to the evaluation of exact gradients [118], and higher order derivatives [120]. The more general situation, where the optimized control fields can be taken to be continuous functions (still parameterized by a finite number of parameters), has now also been solved in [46]. As based on numerical simulations, the application of all the methods just mentioned is limited to system sizes which can be numerically evolved. Still, recent work [123] has shown that it was also possible to derive and to exploit gradients obtained for larger system sizes when treated by means of approximate evolution.

In parallel with the development of these specialized frameworks, the field of machine learning (ML) – where gradients are prominently featured – has matured and now offers flexible tools which are also well suited for the study and optimization of quantum dynamics. In particular auto-differentiation (AD) frameworks provide efficient methods (benefiting from careful implementations and leveraging GPUs capacities) to obtain gradients in numerical computa-

tion. In a nutshell, AD [124, 125] develops the machinery to systematically extend an arbitrary numerical computation to permit the evaluation of its gradients, at the cost of only a modest computational overhead. Such capacity is now integrated into any modern ML libraries [126, 127], and is central to the training of almost all ML architectures where it is known as back-propagation. Still, AD is not limited to the training of neural networks but enables differentiation over virtually any computational programs. For instance, it has been applied recently to numerical routines encountered in the realm of quantum physics, such as the contraction of tensor networks [128, 129], the eigendecomposition of Hermitian operators [130], the simulation of parameterized quantum circuits [101, 102], and the discretized evolution of quantum systems [131–134]. Latest progress in the field [135] have improved the implementation of AD for numerical solvers of ordinary differential equation (such as the Schrödinger equation in Equation (1.1)), thus broadening its applicability to problems of QOC. Resorting to AD tools, rather than more specialized QOC frameworks, has the additional advantage of facilitating its integration into larger optimization frameworks; this will be explored in Section 5.2.

Overall, it is the case that any quantum dynamics which can be numerically simulated could also be optimized by means of gradient methods. In contrast, assessing exact gradients based on experimental outcomes is only possible in certain situations that are now discussed.

Exact gradients in experimental situations

Outside the realm of numerically simulated systems, exact gradients are known to be accessible only for specific problems [115, 136–139]. For instance, this is the case when (i) a quantum state is evolved by means of a quantum circuit (or follows an equivalent dynamics) with parameters corresponding to single-qubit rotation angles, and when (ii) the figure of merit F is taken to be the expectation value of an operator with respect to the output state of the circuit. In such situation, the gradients of F can be reformulated in the form [115, 136]

$$\frac{\partial F}{\partial \boldsymbol{\theta}^{(j)}} = \frac{F(\boldsymbol{\theta} + s\mathbf{e}_j) - F(\boldsymbol{\theta} - s\mathbf{e}_j)}{2\sin(s)}, \quad (2.2)$$

that is, in terms of the figure F evaluated for the shifted parameters $\boldsymbol{\theta} \pm s\mathbf{e}_j$, with scalar shift $s > 0$, and where \mathbf{e}_j denotes the unit vector with all entries equal to zero except $\mathbf{e}_j^{(j)} = 1$.

Equation (2.2), which is often referred as the parameter-shift rule, ensures that if F can be estimated for arbitrary parameters, so can be its gradients. It should be noted that, as estimated on a finite number of measurements, evaluations of F are not exact but rather stochastic. Still assuming that estimates of F are unbiased, that is, are equal to the true value of F in average, it follows that gradients estimates based on Equation (2.2) are also unbiased. It is in such sense that should be understood the "exactness" of the parameter-shift rule: stochastic but unbiased estimates of the true gradients.

This possibility to assess gradients experimentally enables the optimization of parameterized quantum circuits by means of stochastic gradient ascent, akin to the way that neural networks are trained¹. This brings hope that similar success, as attained in training large neural networks, could also be achieved when optimizing parameterized quantum circuits. Still, it remains to ascertain if such optimization approach can truly permit the scaling of VQAs. First, when Equation (2.2) is repeated over all the P entries of the gradient vectors, it entails evaluations of F for a number $2m$ of different parameter values, which quickly renders experimental estimations of the full gradient vector demanding. Second, it should be stressed that Equation (2.2) relies on the implicit assumption that the implemented quantum circuits behave as expected. However, the presence of cross-talk, control errors, or any additional uncontrolled terms, resulting from an experimental realization of the circuit, signifies that Equation (2.2) should be treated as an approximation. To which extent this approximation is accurate enough remains to be studied further. Taken together, these two reasons may explain the disparities between experimental realizations of VQAs, which have mostly been performed without the parameter-shift rule, and the many numerical proofs of concept of VQAs, which often assume exact gradients. As such, it is desirable to consider alternatives to the parameter-shift rule for the training of VQAs, and, also for any other situation where Equation (2.2) (or similar rules [137–139]) does not apply.

Finite-differences gradients

In the general case where the exact gradients cannot be decomposed in terms of measurable quantities, they still can be approximated by means of the finite differences

$$\frac{\partial F}{\partial \theta^{(j)}} \approx \frac{F(\theta + \varepsilon \mathbf{e}_j) - F(\theta - \varepsilon \mathbf{e}_j)}{2\varepsilon}, \quad (2.3)$$

given \mathbf{e}_j the unit vector previously defined, and $\varepsilon > 0$ a scalar perturbation. In the limit of vanishing perturbation $\varepsilon \rightarrow 0$, Equation (2.3) recovers the definition of exact gradients. As such, it would be tempting to take ε to be as small as can be resolved by the experimental control apparatus. While such intuition would be true for exact evaluations of F , in the case of finite errors in the values of F , originating for instance from finite sampling, the overall variance of Equation (2.3) scales as ε^{-2} . Hence, in practice it is desirable to identify (often empirically) values of ε which are small enough to avoid too large approximation errors, but large enough to ensure reasonable statistical errors.

As before, approximating the full gradient vector requires evaluating $2m$ times the figure F . This unfavourable scaling in the number of parameters m can be avoided by perturbing the

¹In the context of neural networks, numerical gradients are exact, but are evaluated on partial batch of data randomly sampled yielding stochastic deviations from the true gradients (which would have been assessed on the whole dataset).

parameters simultaneously rather than individually. This insight forms the basis of the Simultaneous Perturbation Stochastic Ascent (SPSA) algorithm [140] which proposes to approximate the whole gradient vector in terms of a single finite difference

$$\nabla_{\boldsymbol{\theta}} F \approx \frac{F(\boldsymbol{\theta} + \varepsilon \boldsymbol{\Delta}) - F(\boldsymbol{\theta} - \varepsilon \boldsymbol{\Delta})}{2\boldsymbol{\Delta}}, \quad (2.4)$$

along a global direction $\boldsymbol{\Delta} \in \mathcal{R}^P$ which elements are randomly drawn at each iteration of the algorithm. In such case, the evaluation of Equation (2.4) only requires evaluations of F for the two set of parameters $\boldsymbol{\theta} \pm \varepsilon \boldsymbol{\Delta}$ per iteration. Furthermore, SPSA prescribes for both the magnitude of the perturbation ε and the step size η , taken in the approximate gradient direction, to be decreased during the course of the optimization. Given an appropriate decreasing schedule (see [141] for guidelines), SPSA presents similar convergence guarantees than a finite-difference scheme performed over each individual parameter, but offers a practical advantage in terms of the total number of function evaluations required. Furthermore, as opposed to Equation (2.2) SPSA does not rely on any assumption about the underlying dynamics of the system. As such, it has been used in many experimental optimizations [142–144], and we will regard it as a competitive benchmark when assessing the merit of the methods that we develop in Chapter 3.

2.1.2 Gradient-free methods

While gradient methods only consider local updates of the parameters, gradient-free methods permit a broader range of optimization strategies. Among such methods, Nelder-Mead and genetic algorithms have traditionally been applied both to problems of QOC and VQAs. Additionally, techniques inherited from the field of ML have recently been explored.

Nelder-Mead

Nelder-Mead algorithm provides a simple, off the shelf optimization routine for black box optimization. In a nutshell, Nelder-Mead works by maintaining a simplex of $P + 1$ affinely independent set of parameters $\{\boldsymbol{\theta}_i\}_{1 \leq i \leq P+1}$ which are sequentially updated according to a deterministic rule. At each iteration, the worst set of parameters (based on evaluations of the figure of merit), is replaced by its partial reflection to the centroid of the remaining P points. Such strategy is usually fast in converging to local minima of the problem, and, due to its simplicity it has frequently been used in problems of QOC [45, 71, 76, 89, 94, 145, 146] and in experimental implementation of VQAs [20]. For these reasons, we will also consider it as a benchmark in Chapter 3.

Genetic algorithms

Differential evolution and other types of genetic algorithms [147], belong to a class of optimization heuristics inspired by natural evolution. They have been used since the infancy of QOC [40, 42] and are still drawing interest both for problems of QOC and VQAs [117, 148]. Most of these algorithms rely on maintaining a collection $\{\theta_i\}$ (named in this context population) of candidate solutions, each corresponding to a distinct set of parameters, which are evolved according to a set of stochastic rules. Such evolution typically involves (i) a probabilistic selection of the most appropriate candidates, at each round, based on evaluations of their figures of merit, (ii) crossover between different candidates solutions, and (iii) the introduction of additional element of randomness. The main appeal of genetic algorithms, compared to any local optimization, lies in its ability to explore more globally the parameter space. In particular, in the first iterations of the algorithm, candidates solutions remain widespread over the whole parameter space, and only converge after many iterations. As such, these strategies are particularly well suited for problems of optimization with optimization landscapes presenting many local extrema [117]. Given that these algorithms provide a distinct approach compared to gradient-based methods, and that implementations of genetic algorithms are plenty, we will also include them for benchmarking purposes in Chapter 3.

Reinforcement learning

Following recent breakthroughs in problems of artificial intelligence [149, 150], reinforcement learning (RL) [151] has emerged as a powerful paradigm with broad applicability to many domains of science. Roughly speaking, RL envisions the task of learning as repeated interactions between a learner and an unknown environment, along which the learner is able to iteratively refine its strategy (that is how it interacts with the environment) in order to achieve a specific objective. Rather than a monolithic framework, RL has developed in a wide array of techniques, many of which are now based on neural networks. In the realm of quantum physics, these have been employed for the design of quantum experiments [152], problems of quantum error correction [32, 153–155] and tasks of QOC, including quantum state preparation [31, 54] and quantum gate engineering [55, 155].

To date, only a few results have demonstrated the applicability of RL to problems of QOC or VQAs based on experimental data (with the exception of [152, 155]). Such applicability may be limited by the large amount of training data that RL frameworks typically require. Still, it should be acknowledged that due to the flexibility of the RL principles, these can be applied to problems exceeding the scope of this thesis. In particular, RL is naturally suited for tasks of control where partial observations of the system under control can be obtained during its evolution [32], or for problems involving both continuous and combinatorial optimizations,

such as finding the optimal parameters and circuit structures in problems of VQAs [156]. Furthermore, bridging the gap between numerical simulations and real world experiments is an active field of research [157], and may soon bring the advances easing the applicability of RL in experimental situations. Given that any implementation of RL requires a significant amount of care and expertise, we will not actively consider RL frameworks as benchmarks for the methods presented in Chapter 3, but will provide qualitative comparison whenever is possible.

2.1.3 Bayesian Optimization

We now turn our attention to an optimization technique originating from the field of probabilistic machine learning [34], namely Bayesian optimization (BO). BO [158–160] has been recognized as a prime choice for black-box optimization, that is, when the evaluation of the function to be optimized is expensive and potentially noisy. As such, it has been applied to demanding problems such as the optimization of the structure of neural networks [161], problems of robotics [162], or the design of new molecules [163], and may be perfectly suited for problems of experimental optimizations. Still, at the beginning of this thesis, applications of BO to the realm of quantum physics were scarce [99].

In sharp contrast to the methodologies previously mentioned, BO advocates the construction of a surrogate model of the function F to be maximized, in order to accelerate its optimization. Given limited and noisy evaluations of F , such model can only be approximate, and understanding how likely it is to deviate from F , that is understanding the uncertainty entailed by the modeling approach, is key to its appropriate use. Probabilistic modeling provides the methodology to capture such uncertainty. Broadly speaking, the model considered in BO relies on an initial distribution $p(f)$ over conceivable model functions f , that is refined (by means of Bayesian inference) to be consistent with the set of observations \mathbf{y} which have been obtained. The resulting distribution $p(f|\mathbf{y})$ allows to formulate the choice of the next parameters to be evaluated as a Bayesian decision problem. These two aspects of modeling and deciding under uncertainty form the basis of an iterative optimization procedure, and are now described.

Modeling functions with Gaussian processes

Prior distribution. In the spirit of Bayesian inference, a first step toward probabilistic modeling is the definition of a prior distribution. Here, this needs to be done over model functions² f , rather than over a discrete set of random variables. Furthermore as the exact nature of the function F to be optimized is unknown, this choice can only be made based on

²Here, f always refers to the functions modeling F , and not to the control fields to be optimized in problems of QOC (discussed in the previous chapter).

general considerations. For problems of control of physical systems it seems reasonable to expect well-behaved functions without wild unnatural oscillations or discontinuities. In practice, such general assumptions can be formulated in terms of a Gaussian process distribution [164] denoted $p(f)$, which technicalities are deferred to Appendices A.1 and A.2. In essence, $p(f)$ allows to favor smooth and regular functions f to describe the unknown F . As this distribution does not yet incorporate any observed data it is therefore referred as the prior distribution.

Likelihood. The next step consists in updating this prior distribution $p(f)$ to take into consideration the evaluations of F that have been acquired. For that purpose, observations and model need to be connected. At an arbitrary step D of the optimization, D of such evaluations would have been obtained, and we denote the vector of observations $\mathbf{y} = [y_1, \dots, y_D]$, which corresponds to the noisy values of F for input parameters $\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_D$. In the general situation where the exact details of the noise in the observed data is unknown, it is convenient to approximate it by means of a phenomenological Gaussian noise. More exactly, it is typically posited that $y_j = f(\boldsymbol{\theta}_j) + \varepsilon_j$, where the noise terms ε_j are assumed to be independently drawn from the same Gaussian distribution, with zero mean and variance σ_N^2 , denoted $p(\varepsilon_j) = \mathcal{N}(0, \sigma_N^2)$.

For a noisy evaluation y_i obtained for input $\boldsymbol{\theta}_i$, these assumptions yield a likelihood

$$\begin{aligned} p(y_i|f(\boldsymbol{\theta}_i)) &= \mathcal{N}(f(\boldsymbol{\theta}_i), \sigma_N^2) \\ &= \frac{1}{\sigma_N \sqrt{2\pi}} \exp\left[-\frac{(y_i - f(\boldsymbol{\theta}_i))^2}{2\sigma_N^2}\right] \end{aligned} \quad (2.5)$$

to observe y_i given a model function f (with the explicit density function of the Gaussian distribution provided in the second line). Given independence of the noise terms, the likelihood of the full body of observations \mathbf{y} to be observed given a model f , is the multivariate Gaussian distribution

$$\begin{aligned} p(\mathbf{y}|f) &= \mathcal{N}(\mathbf{f}, \sigma_N^2 I) \\ &= \frac{1}{\sigma_N \sqrt{2\pi}} \frac{\exp\left[-(\mathbf{y} - \mathbf{f})^T (\sigma_N^2 I)^{-1} (\mathbf{y} - \mathbf{f})\right]}{\sqrt{(2\pi)^D |\sigma_N^2 I|}} \end{aligned} \quad (2.6)$$

with mean vector $\mathbf{f} = [f(\boldsymbol{\theta}_1), \dots, f(\boldsymbol{\theta}_D)]$, which is the vector of values taken by f for the parameters $\{\boldsymbol{\theta}_i\}$, and with covariance matrix $\sigma_N^2 I$, where I denotes the $D \times D$ identity. Again we explicitly provide the expression of the probability density function (of this now multivariate Gaussian distribution) in the second line, where $|A|$ denotes the determinant of the matrix A . For noiseless data, $\sigma_n = 0$, this reduces to the delta function $p(\mathbf{y}|f) = \delta_{\mathbf{y}=\mathbf{f}}$. In the more general case where $\sigma_n > 0$, the noise variance σ_N^2 is often not known beforehand but can be fitted to the observed data (Appendix A.3).

Posterior distribution and predictions. Having connected data \mathbf{y} to model f (Equa-

tion (2.6)), we can now proceed to express the posterior distribution $p(f|\mathbf{y})$ by means of the Bayes rule

$$p(f|\mathbf{y}) = \frac{p(f)p(\mathbf{y}|f)}{p(\mathbf{y})}, \quad (2.7)$$

which features, in addition to the terms already defined, the marginal likelihood $p(\mathbf{y}) = \int d\mathbf{f} p(\mathbf{f})p(\mathbf{y}|\mathbf{f})$ that normalizes the overall distribution. In general, this term is notoriously difficult to evaluate as resulting from a D -fold integral over the values that can take each of the variables $f(\boldsymbol{\theta}_i)$. However, given the choice of the prior Gaussian process $p(f)$ and of the Gaussian likelihood $p(f|\mathbf{y})$, Equation (2.7) admits an analytical expression. In particular, the posterior distribution of the value $f(\boldsymbol{\theta}^*)$ taken by the model at arbitrary parameter $\boldsymbol{\theta}^*$ also follows a Gaussian distribution:

$$p(f(\boldsymbol{\theta}^*)|\mathbf{y}) = \mathcal{N}(\mu_f(\boldsymbol{\theta}^*), \sigma_f^2(\boldsymbol{\theta}^*)), \quad (2.8)$$

with mean $\mu_f(\boldsymbol{\theta}^*)$ and variance $\sigma_f^2(\boldsymbol{\theta}^*)$ (both depending implicitly on \mathbf{y}). The detailed expressions of these two functions are provided in Equation (A.5) (of the Appendix A.4), where it is highlighted that their evaluations entail a computational effort scaling as D^3 , that is, growing cubically with the number D of observations available.

The distribution in Equation (2.8) is central to BO as it allows to make prediction for arbitrary inputs $\boldsymbol{\theta}^*$ (even in the absence of corresponding observations), and is thus often referred as the predictive distribution. Notably, rather than a single deterministic value $f(\boldsymbol{\theta}^*)$ (as would be predicted by a non-probabilistic ML model), it provides a full probabilistic account of the values that the unknown function could take.

Illustration. This predictive distribution is illustrated in Figure 2.2 for a dataset of 10 observations (red dots) and a one-dimension parameter θ . Given the body of 10 observations, Equation (2.8) (and its explicit expression in Equation (A.5)) is evaluated and depicted in terms of its mean $\mu_f(\theta)$ (plain dark blue) and the 95% confidence interval centered around the mean which, given the Gaussian nature of Equation (2.8), has bounds $\mu_f(\theta) \pm 1.96 \times \sigma_f(\theta)$. In addition, the full probability distributions $p(f(\theta)|\mathbf{y})$ for specific values of the parameters $\theta = 0.2$ and 2.8 are also reported (in shaded gray and in units chosen to ease their visualizations). In particular, for $\theta = 0.2(2.8)$ the predictive distribution exhibits a small(large) confidence interval corresponding to the presence(absence) of observations in the vicinity of the parameter value. In other words, it can be seen that the model uncertainty naturally grows when further away from the training data points. This ability to capture modeling uncertainty is central to the choice of the next parameters to be probed, that we now discuss.

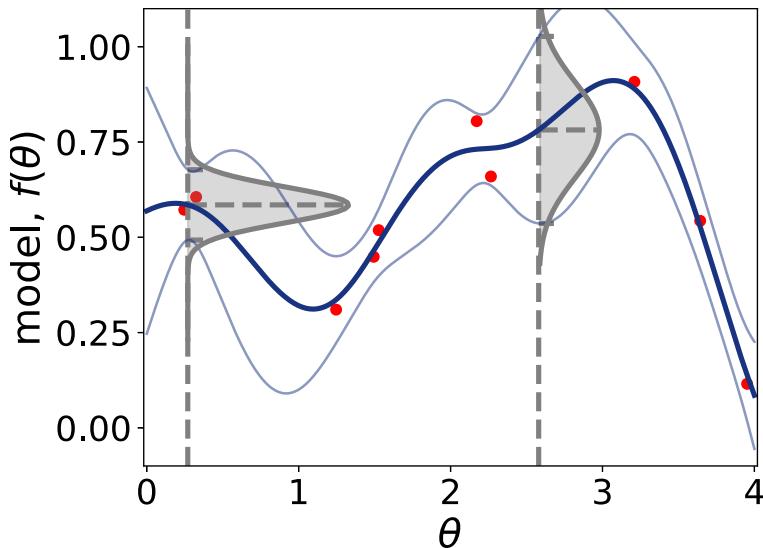


Figure 2.2: From [165]: Gaussian process model fitted to a vector \mathbf{y} of 10 observations (red dots). The predictive distribution $p(f(\theta)|\mathbf{y})$ from Equation (2.8) is represented by its mean (thick blue line), and a 95% confidence interval (light blue around the mean). In addition, the full probability distributions are plotted for values of $\theta = 0.2$ and 2.6 (in shaded gray).

Choice of the next parameter

We now detail how the predictive probability distribution (Equation (2.8)) is employed when deciding which set of parameters $\boldsymbol{\theta}_{D+1}$ needs to be evaluated next. One could choose it where the model adopts a maximal mean value μ_f . However, as based on a finite number of evaluations, it should be acknowledged that this model can only be approximate and is likely to miss some important features of the underlying function F , especially at the beginning of the optimization when only a few evaluations have been obtained. Thus, it is also of interest to evaluate F where the uncertainty in the model is high, that is where the standard deviation σ_f^2 is large.

This trade-off³ between exploiting the model to accelerate the original optimization problem, and exploring new regions of the parameter space to refine the model, can be formulated in terms of the utility $\alpha(\boldsymbol{\theta})$ in evaluating F given the set of parameters $\boldsymbol{\theta}$. In the context of BO such utility function is called an acquisition function. Given the choice of an acquisition function, the next set of parameters to be evaluated is chosen where the acquisition function adopts its maximal value:

$$\boldsymbol{\theta}_{D+1} = \arg \max_{\boldsymbol{\theta}} \alpha(\boldsymbol{\theta}), \quad (2.9)$$

that is, where the parameters are deemed to be the most useful.

Among the several acquisition functions encountered in the literature [159], we limit the discussion to the upper confidence bound (UCB) which was found to be the most adequate for

³It is often referred as exploitation-exploration trade-off.

the problems that are considered in Part II. This acquisition function is defined as:

$$\alpha_{UCB}(\boldsymbol{\theta}) = \mu_f(\boldsymbol{\theta}) + k\sigma_f(\boldsymbol{\theta}), \quad (2.10)$$

where the scalar $k \geq 0$ balances the bias toward exploration of the parameter space (for large values of k) or exploitation (for small value of k). In practice we choose the value of k to be linearly decreasing (from an initial value k_0 to a final value of 0) during the course of the optimization. Such decreasing schedule encourages exploration at the beginning of the optimization, but as the optimization progresses it limits the choice of the new parameters to be evaluated to the most promising ones (according to the model). An adequate value of k_0 (typically in the range [1, 10]) can be chosen by means of a few trial-and-errors.

It should be highlighted that, as based exclusively on the surrogate model, the auxiliary optimization problem in Equation (2.9) is much simpler than the original problem of maximizing F . In particular, it does not require any additional evaluations of F . Finally, once the set of parameter $\boldsymbol{\theta}^{D+1}$ maximizing Equation (2.9) is identified, the corresponding evaluation of F is performed and its result is added to the body of observations. Given this new dataset the model (that is, the predictive distribution in Equation (2.8)) is refined.

Overall, repetitions of these steps of modeling and deciding the next parameters to be evaluated, that have been described, form the basis of BO. This is now illustrated on a toy problem.

Iterations of Bayesian optimization

A whole run of BO involves repeated cycles of updates of the surrogate model, to include the latest observations obtained, and choices of the next set of parameters to be evaluated. This is illustrated in Figure 2.3 for the optimization of a toy function

$$F(\theta) = \sin^2\left(\frac{1}{2}\sin\left(3\theta + \frac{9}{10}\right) + \frac{3\theta}{2} + \frac{9}{20}\right), \quad (2.11)$$

depending on a one-dimension parameter θ . In this simple example, Gaussian noise, with zero mean and 5% standard deviation, is added to the evaluations of F to imitate experimental noise.

Similar to Figure 2.2, the predictive distribution of Equation (2.8) is depicted, in the top panels of Figure 2.3, in terms of its mean value (thick blue) and a 95% confidence interval (shades of blue). This is to be compared to the true value of F (depicted in dashed red). The model is built for the first time, in Figure 2.3(a), after a stage of initialization where F is evaluated for 10 values of θ randomly sampled. Already, the model captures some of the most salient features of the unknown function, but misses the central peak. Nonetheless, this lack of certainty is reflected in wide confidence intervals, especially when further away from any observation.

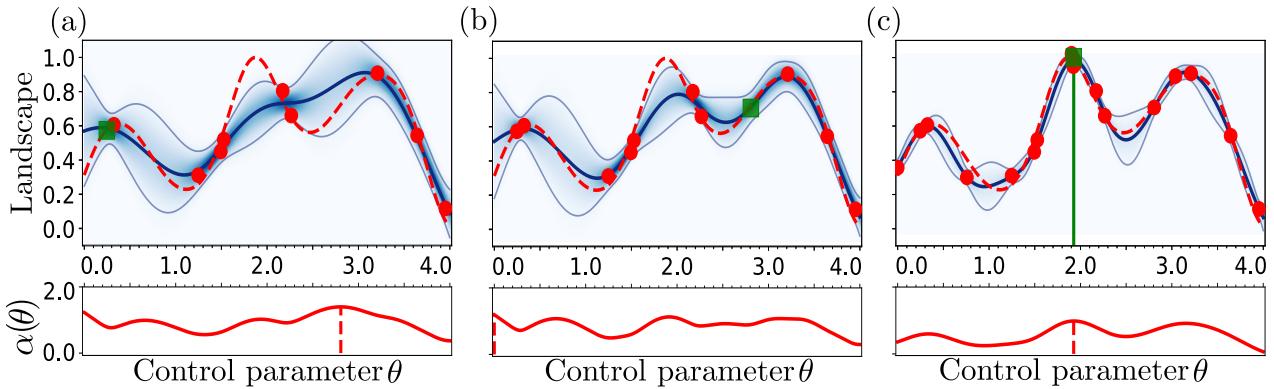


Figure 2.3: From [165]: Bayesian optimization (BO) of the toy function in Equation (2.11) after (a) 10, (b) 11 and (c) 20 noisy evaluations. The top panel depicts the underlying function F to be maximized (dashed red), and the available measurements (red circles and green square for the last measurement obtained). Similar to Figure 2.2, the model is represented by its mean (solid blue line) and a 95% confidence interval (solid gray). The lower panel depicts the acquisition function α (Equation (2.10) with $k = 4$ for (a) and (b) and $k = 0$ for (c)), the maximum of which determines the parameter to be evaluated next.

The acquisition function in Equation (2.10) is plotted in each of the bottom panels. At the initial stage of Figure 2.3(a) this function is found to be maximized for $\theta \approx 2.8$. Consequently, F is measured for this new parameter, and its noisy evaluation (depicted in green square in Figure 2.3(b)) is added to the body of observations. This yields an updated model which is in closer agreement with the unknown function F , and a reduced confidence interval in the vicinity of to the new data point. After 10 iterations of BO, Figure 2.3(c), the algorithm has identified the optimal parameter (vertical green line) with high precision. As can be seen in this last step, given the 10 additional evaluations which have been obtained over the course of the optimization, the model replicates accurately F for promising regions of the parameter space (that is, around the two largest local maxima), but BO has avoided the effort of too many evaluations in other regions which were identified as of less interest.

In summary, BO offers a distinct alternative to more traditional optimization approaches. Its ability to select observations in a principled way over the entire parameter space may be particularly well suited to situations where data is scarce and convexity of the optimization landscape is not guaranteed. This will be assessed in Section 3.1 where a thorough comparison to the other optimization routines presented in this section is performed. Furthermore BO, has the additional advantage of integrating probabilistic elements of data acquisition in its modeling approach. This was presented here under the common assumption of Gaussian noise (Equation (2.6)), but can be refined further to take into consideration more precise knowledge of the measurement statistics, in order to increase the overall efficiency of BO. This will be demonstrated in Section 3.2. Lastly, it should be highlighted that the model, as expressed in Equation (2.8), allows to make predictions based on the whole body of observations which are available. As such it always benefits from the inclusion of more data. This feature will

be central to the framework presented in Section 5.1, where it is shown that measurements obtained over the course of distinct optimizations can be shared, to enhance the individual optimizers performances.

This concludes our review of the algorithms which can be applied to optimize the dynamics of quantum systems. We now proceed to discuss the second aspect which needs to be considered when performing optimization based on experimental data, namely the estimation of the figure of merit F .

2.2 Estimation of the figure of merit

Any optimization relies on the ability to evaluate a figure of merit at each iteration of the routine (Figure 2.1). For numerically simulated systems this figure is readily available. However, when dealing with experimental outcomes one can only estimate its value based on a finite set of measurements. Quantum mechanics postulates that the process of measuring a quantum system is intrinsically probabilistic, and that measurements only reveal partial information about the underlying state measured. As such, it becomes necessary to repeat many measurements to accurately characterize a quantum state.

If not treated carefully, the experimental effort entailed by these measurement repetitions can quickly become prohibitive. This issue is magnified in the case of optimization, where the estimation of the figure of merit needs to be performed, not only once, but at each iteration. A naive approach towards estimating properties of a quantum state consists of performing tomography, that is, fully reconstructing the unknown state, and subsequently evaluating any desired property of the state based on such reconstruction. In general, the measurement effort necessitated by such approach scales exponentially with the size of the system, making the tomography of states with $n > 10$ qubits already very demanding. Still, it is often the case that one does not require the full knowledge of the unknown state, but rather only aims at assessing some specific properties. In particular, in this thesis we will always be concerned with figures of merit taken to be expectation values of operators with respect to the unknown state. In such situations, more efficient estimation schemes have been devised [166–168].

After reviewing elements of the quantum measurement formalism in Section 2.2.1, we detail in Section 2.2.2 a sampling scheme [166, 167] enabling the estimation of figures of merit. This will permit us to establish general expressions relating details of the figure to be estimated and accuracy in its estimation. Finally, these expressions are put into practice in Section 2.2.3 for the case of the state fidelity, a figure of interest for many tasks of quantum state preparation. The methodology developed in this chapter will serve us as a basis for the comparison and design of experiment-friendly figures of merit in Chapter 4.

2.2.1 Measuring quantum systems

Quantum measurements

A quantum measurement is defined as a set of measurement operators $\{M_m\}$ satisfying the completeness relation $\sum_m M_m^\dagger M_m = \mathcal{I}$. Quantum mechanics postulates that measuring a (potentially mixed) state ρ in this setting results in an outcome m with probability $p_m(\rho) = \text{Tr}[M_m^\dagger M_m \rho]$, which is known as the *Born's rule* [44]. Projective measurements, which are the most common type of measurements encountered in experimental realizations of quantum physics, restrict the measurement operators M_m to be orthogonal projectors, that is, operators which satisfy the relations $M_m M_{m'} = \delta_{m,m'} M_m$ and $M_m^\dagger = M_m$, in which case $p_m(\rho) = \text{Tr}[M_m \rho]$. To such operators can be associated an observable M admitting a decomposition

$$M = \sum_m \lambda_m M_m, \quad (2.12)$$

with real values λ_m .

For a given state ρ , measuring the observable M is understood as assessing its expectation value, defined as $\text{Tr}[M^\dagger \rho]$ and that we denote⁴

$$\langle M \rangle_\rho \equiv \text{Tr}[M^\dagger \rho]. \quad (2.13)$$

This expectation can be recast in terms of measurement probabilities:

$$\langle M \rangle_\rho = \sum_m \lambda_m \text{Tr}[M_m \rho] = \sum_m \lambda_m p_m(\rho), \quad (2.14)$$

showing its operational meaning as the average value of the random variable which takes a value of λ_m for measurement outcome m .

In addition to its first statistical moment (the average, or expectation value), the second moment (the variance) of this measurement random variable is defined as

$$\begin{aligned} \Delta_\rho^2[M] &= \sum_m p_m(\rho) [\lambda_m - \langle M \rangle_\rho]^2 \\ &= \langle M^2 \rangle_\rho - \langle M \rangle_\rho^2. \end{aligned} \quad (2.15)$$

This variance quantifies fluctuations in the measurement values and is bounded in

$$0 \leq \Delta_\rho^2[M] \leq \frac{(\lambda_{\max} - \lambda_{\min})^2}{4}, \quad (2.16)$$

⁴We will often resort to this shorthand notation in the rest of this thesis.

where $\lambda_{\max}(\lambda_{\min})$ denotes the largest(lowest) eigenvalue of M . These fluctuations vanish when the underlying state measured is an eigenstate of M , and are maximal for any state in equal superposition of the lowest and highest eigenvectors of M .

Operators of the form Equation (2.12) are called *physical observables* as they can be related to physical observations, permitting the estimation of their expectation values. In particular, the figures of merit that are presented and employed in this thesis will always be expectation values of such operators, and we now detail how they are estimated in practice.

Estimation of expectation values of physical observables

Without exact knowledge of the measured state ρ one can only estimate, rather than exactly evaluate, expectation values. This is done by first estimating the unknown probabilities $p_m(\rho)$, appearing in Equation (2.14), in terms of the outcomes frequencies $f_m(\rho) = N_m/N$, where N_m is the number of measurement outcomes m recorded over a total number of N repetitions. Consequently, an estimate \tilde{M}_ρ of the (true) expectation value $\langle M \rangle_\rho$ is obtained as

$$\tilde{M}_\rho = \sum_m \lambda_m f_m(\rho). \quad (2.17)$$

Importantly, this estimate is unbiased, that is, its average equals the true expectation value:

$$E[\tilde{M}_\rho] = \langle M \rangle_\rho, \quad (2.18)$$

where $E[\cdot]$ denotes an average taken over the probably distribution of the measurement outcomes $p(N_1, \dots, N_S)$. Additionally, the variance of this estimate is given by

$$\Delta_\rho^2[\tilde{M}] = E[(\tilde{M}_\rho - \langle M \rangle_\rho)^2] = \frac{\Delta_\rho^2[M]}{N}, \quad (2.19)$$

which depends on the variance of a single measurement outcome defined in Equation (2.15), and on the number of measurement repetitions N . For bounded eigenvalues λ_m , this variance scales as N^{-1} , and as expected vanishes with increased N . That is, in the limit of an infinite number of measurements one can estimate expectation values of M up to arbitrary precision. However, in the realistic situations discussed in this thesis, this sampling number N is taken finite entailing non-vanishing estimation errors $\Delta_\rho^2[\tilde{M}] > 0$.

Estimation with restricted measurement settings

In principle, any Hermitian operator $M = M^\dagger$ can be recast in the form of Equation (2.12) and thus could be estimated as was just described. In practice, however, measurements are

physical processes which need to be engineered, and experimental restrictions put constraints on the sets of the measurements $\{M_m\}$ which can be performed accurately. In particular, as entangling operations are often unreliable compared to local operations, it is common to only consider local measurements. These are defined as the projectors $M_m = M_{m_1}^{(1)} \otimes \dots \otimes M_{m_n}^{(n)}$ resulting from a tensor product of projectors $M_{m_i}^{(i)}$ acting locally on each of the n individual parts of the system. To make the distinction clear between general observables $M = M^\dagger$ and the observables corresponding to measurements which can accurately be performed, we will refer to the later explicitly as *experimental observables* and denote them $\{W_i\}$ in the rest of the thesis.

Assuming that the observable of interest M can be decomposed as a weighted sum over D experimental observables W_i such as

$$M = \sum_{i=1}^D \alpha(i) W_i, \quad (2.20)$$

with real coefficients $\alpha(i) \neq 0$, one can proceed by first assessing the individual expectation values $\langle W_i \rangle_\rho$ in terms of the estimates $\tilde{W}_{i,\rho}$ (obtained as previously discussed), and by reconstructing the estimate

$$\tilde{M}_\rho = \sum \alpha(i) \tilde{W}_{i,\rho} \quad (2.21)$$

of the expectation value $\langle M \rangle_\rho$. This estimate is guaranteed to be unbiased by virtue of the linearity of the expectation, that is,

$$\begin{aligned} E[\tilde{M}_\rho] &= \sum_{i=1}^D \alpha(i) E[\tilde{W}_{i,\rho}] = \sum_{i=1}^D \alpha(i) \langle W_i \rangle_\rho \\ &= \langle M \rangle_\rho. \end{aligned} \quad (2.22)$$

Given access to a set of experimental observables $\{W_i\}$ spanning the space of linear operators, one can always decompose an arbitrary observable M according to Equation (2.20) and estimate it. Followed naively this approach entails the estimation of as much as $D = d^2$ individual terms. Already, for system sizes of $N = 10$ qubits one would need to estimate $4^{10} \approx 10^6$ of such terms and it would seem, at first glance, hopeless to consider larger systems. This argument counting, however, fails in recognizing the importance of the distributions of the weights $\alpha(i)$. Furthermore, this strategy does not prescribe how to allocate measurements over the different operators W_i . In particular, it would seem natural to allocate more measurements to the terms contributing the most to the decomposition Equation (2.20). Building up on these intuitions, we now detail an estimation scheme following the treatment presented in [166, 167].

2.2.2 Estimation by means of importance sampling

With the goal of building better estimates in mind, we start this section by formalizing the notion of efficient estimation, then provide details of an estimation scheme (called importance sampling) based on Monte Carlo sampling which can be shown to be efficient. This will permit us to derive analytical expressions relating the details of a general observable to be measured, to the variance of its estimate. These expressions will be applied in the next section to more specific observables.

In the following, it is assumed the existence of a fixed orthonormal basis of d^2 experimental observables $\{W_i\}$, such that $\text{Tr}[W_i^\dagger W_j] = \delta_{i,j}$. In this basis, an arbitrary observable M admits the decomposition

$$M = \sum_{i=1}^D \alpha(i) W_i, \text{ with } \alpha(i) = \text{Tr}[W_i^\dagger M] \neq 0. \quad (2.23)$$

Efficient estimation

The accuracy in an estimate \tilde{X} of the true quantity X is quantified by the variance $\Delta^2[\tilde{X} - X]$ in the estimation error $\tilde{X} - X$. That is, high(low) estimation accuracy corresponds to low(high) variance. For unbiased estimates (as systematically considered in this thesis), such that $E[\tilde{X}] = X$, this variance in the estimation error equals the variance of the estimate itself, $\Delta^2[\tilde{X} - X] = \Delta^2[\tilde{X}]$. An alternative version of accuracy is sometimes expressed as the probability $p(|\tilde{X} - X| > \varepsilon)$ for the absolute value of the estimation error to be larger than a quantity $\varepsilon > 0$ [166, 167, 169]. However, we find the use of the variance, as a single quantity to gauge estimation accuracy, to be more adequate as it leads to more compact and easy-to-interpret expressions. Nonetheless, one can relate these two quantities by means of the Chebyshev's inequality, which specifies that, for an unbiased estimate \tilde{X} with variance $\Delta^2[\tilde{X}]$,

$$p(|\tilde{X} - X| \geq \varepsilon) \leq \frac{\Delta^2[\tilde{X}]}{\varepsilon^2}. \quad (2.24)$$

Overall, an (unbiased) estimate is deemed better than another (unbiased) one if, for the same number N of measurements, it presents a lower variance.

Importance sampling

To allow for freedom in the allocation of the number N_i of measurements to be performed for each observable W_i , and to avoid issues in non integer values of N_i (for instance, when the

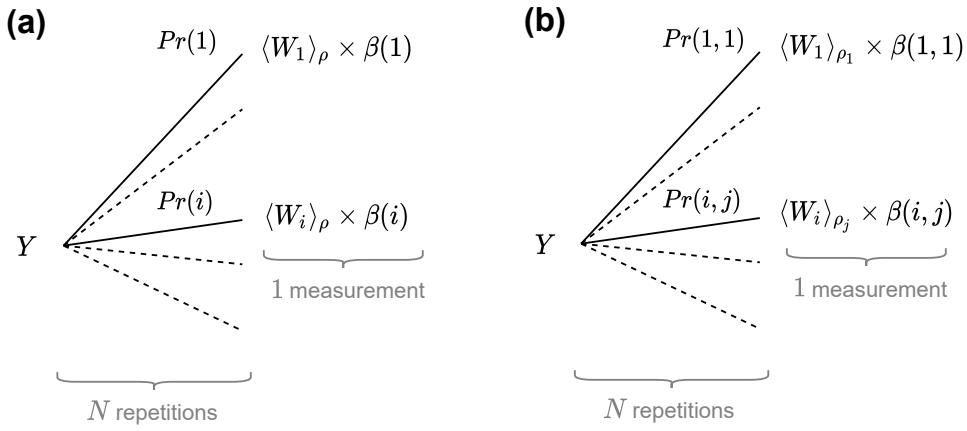


Figure 2.4: Estimation with importance sampling. (a) Estimation of the expectation value $\langle M \rangle_\rho$ of an operator $M = \sum \alpha(i) W_i$, with respect to a fixed state ρ . Each sample corresponds to (i) the random choice of the experimental observable W_i to be measured according to an importance sampling distribution $Pr(i)$, and (ii) a single measurement of W_i which outcome is rescaled by a factor $\beta(i)$. An estimate of the expectation value $\langle M \rangle_\rho$ is obtained by averaging the sampled values over N repetitions of (i-ii). Both the importance distribution and the scaling factors can be chosen to modify the statistics of this estimate. The choices resulting in unbiased estimates with minimal variance are provided in Equation (2.26) and Equation (2.31). (b) This scheme can be extended to situations where both the observable W_i and the state ρ_j to be measured can be varied. The importance distribution $Pr(i, j)$ and the scaling factors $\beta(i, j)$ can be adapted to this extended setup.

total number N of measurements is smaller than the number D of experimental observables to be measured), we follow the statistical treatment presented in [166, 167]. As sketched in Figure 2.4(a) each measurement involves (i) sampling an index $i \in \{1, \dots, D\}$ according to an importance distribution $Pr(i)$, (ii) performing a single measurement of the corresponding observable W_i , and rescaling its outcome value by $\beta(i)$. The random variable resulting from steps (i-ii) is denoted Y_ρ and an estimate of $\langle M \rangle_\rho$ can be obtained by averaging the values taken by $Y_\rho^{(s)}$ over N measurement repetitions such as

$$\tilde{M}_\rho = \frac{1}{N} \sum_{s=1}^N Y_\rho^{(s)}, \quad (2.25)$$

with $s = 1, \dots, N$ indexing the measurement repetitions.

Notably, the importance distribution $Pr(i) > 0$ (with $\sum Pr(i) = 1$) specifies how to allocate (in average) measurements to a particular observable W_i , while the terms $\beta(i)$ ensure consistency of the outcomes of the sampling scheme with the desired expectation value. In particular it is necessary to choose

$$\beta(i) = \frac{\alpha(i)}{Pr(i)} \quad (2.26)$$

to ensure that Y_ρ (and thus also its average \tilde{M}_ρ) is unbiased. Indeed, for the choice of $\beta(i)$ in

Equation (2.26), we obtain

$$E[Y_\rho] = \sum_{i=1}^D Pr(i)\beta(i)\langle W_i \rangle_\rho = \langle M \rangle_\rho. \quad (2.27)$$

The variance of Y_ρ (which effectively follows a mixture distribution) can be recast as

$$\begin{aligned} \Delta^2[Y_\rho] &= E[Y_\rho^2] - E[Y_\rho]^2 \\ &= \sum_{i=1}^D Pr(i)E[(\beta(i)W_i)^2] - \langle M \rangle_\rho^2 \\ &= \sum_{i=1}^D \frac{\alpha^2(i)}{Pr(i)} \langle W_i^2 \rangle_\rho - \langle M \rangle_\rho^2, \end{aligned} \quad (2.28)$$

where the second line results from the decomposition of the expectation $E[Y_\rho^2]$ over the different choices of measurement indices and from replacing $E[Y_\rho]$ according to Equation (2.27). The last line follows from the choice of $\beta(i)$ in Equation (2.26).

Finally, it should be noted that the estimate \tilde{M}_ρ in Equation (2.25), which is averaged over N outcomes of Y_ρ , has an expected value $E[\tilde{M}_\rho] = \langle M \rangle_\rho$ and a variance $\Delta^2[\tilde{M}_\rho] = \Delta^2[Y_\rho]/N$. That is, both Y_ρ and \tilde{M}_ρ are unbiased estimates of $\langle M \rangle_\rho$, with variances $\Delta^2[Y_\rho]$ and $\Delta^2[\tilde{M}_\rho]$ only differing by a factor N . As such when discussing estimate accuracy of $\langle M \rangle_\rho$ we may refer to one or the other.

Allocation of the measurements

Having chosen the scaling coefficients ensuring unbiased estimates in Equation (2.26), it is left to decide how measurements should be allocated over the observables W_i . The most common distribution advocated in the literature [166, 167] is given as

$$Pr(i)^* = \frac{\alpha^2(i)\langle W_i^2 \rangle_\rho}{\sum_i \alpha^2(i)\langle W_i^2 \rangle_\rho}. \quad (2.29)$$

The variance of the estimate corresponding to this choice of distribution is obtained by inserting Equation (2.29) into Equation (2.28) resulting in the expression

$$\Delta^2[Y_\rho] = D \sum_i \alpha^2(i)\langle W_i^2 \rangle_\rho - \langle M \rangle_\rho^2, \quad (2.30)$$

which relates the variance of the estimate to details of the observable M of interest, which appears through the coefficients $\alpha(i)$ and the number D of non vanishing coefficients in Equation (2.23).

To explicitly evaluate Equations (2.29) and (2.30) it remains to assess the values of the statistics $\langle W_i^2 \rangle_\rho = \text{Tr}[(W_i^2)^\dagger \rho]$. As we will see, in many situations of interest – which involve a choice of operator W_i such that W_i^2 is proportional to the identity, this is discussed in Section 2.2.3 – these values do not depend on the unknown state ρ , and can be directly evaluated. Still, in the more general case where the expectations $\langle W_i^2 \rangle_\rho$ depend on ρ , one can still approximate them. For instance, these could be approximated by the maximum or, the average, of $\langle W_i^2 \rangle_\rho$ over random states, which can be evaluated (this is discussed further in due course).

In addition to the popular choice of importance probability distribution given in Equation (2.29), we derive in Appendix B.1 the probability distribution which explicitly minimizes the variance in Equation (2.28), and obtain an optimal choice of

$$Pr(i)^* = \frac{|\alpha(i)\sqrt{\langle W_i^2 \rangle_\rho}|}{\sum_i |\alpha(i)\sqrt{\langle W_i^2 \rangle_\rho}|}, \quad (2.31)$$

for the probabilities, which yields an variance

$$\Delta^2[Y_\rho] = \left(\sum_i |\alpha(i)\sqrt{\langle W_i^2 \rangle_\rho}| \right)^2 - \langle M \rangle_\rho^2. \quad (2.32)$$

This variance is guaranteed to be lower or equal compared to Equation (2.30). That is, one should in principle favour the distribution in Equation (2.31) compared to the choice of Equation (2.29). We note that, in some of the examples that we are going to explore both these probability distributions coincide with the simple uniform distribution $Pr(i) = 1/D$, and thus would result in the same variance. In such situation, we will use the formula in Equation (2.30) which is easier to manipulate. In any case, given Equation (2.28) (and Equations (2.30) and (2.32) for specific choices of importance distribution) we are now in position to evaluate the variance of estimates corresponding to arbitrary observables.

Extension in case of varied observables and varied input states

The methodology developed so far can be extended to situations involving several experimental observables W_i , and, also several states ρ_j to be measured (illustrated in Figure 2.4(b)). For instance, a more general general figure F is given in terms of the decomposition

$$F = \sum_{i,j} \alpha(i,j) \langle W_i \rangle_{\rho_j} \quad (2.33)$$

over expectation values $\langle W_i \rangle_{\rho_j}$ of varied observables and states, indexed by i and j respectively. For such figure F , the importance sampling strategy and formulas reported can simply be extended to take into consideration the changes $i \rightarrow (i, j)$, $\langle \cdot \rangle_\rho \rightarrow \langle \cdot \rangle_{\rho_j}$ and $D \rightarrow D_{ij}$, with D_{ij} now denoting the number of non-null coefficients $\alpha(i, j)$. For instance, the importance

distribution in Equation (2.29) becomes

$$Pr(i, j)^* = \frac{\alpha^2(i, j) \langle W_i^2 \rangle_{\rho_j}}{\sum_{i,j} \alpha^2(i, j) \langle W_i^2 \rangle_{\rho_j}}, \quad (2.34)$$

and the resulting variance Equation (2.30) transforms to

$$\Delta^2[Y] = D_i D_j \sum_{i,j} \alpha^2(i, j) \langle W_i^2 \rangle_{\rho_j} - \langle F \rangle^2. \quad (2.35)$$

This extended scenario is investigated further in Section 4.2.3.

In summary, we have detailed a scheme for the estimation of general physical observables. The overall methodology, relying on formalizing the choice of the measurement to be taken as a sampling procedure, is depicted in Figure 2.4 and the explicit expressions of the scaling coefficients and of the sampling probabilities were provided in Equation (2.26) and Equation (2.29) (or Equation (2.31)) respectively. The variances of the resulting estimates were derived in Equation (2.30) (or Equation (2.32)). These expressions relate details of the decomposition of an observable in a specific basis, to its estimation accuracy (variance of its estimate).

Before proceeding further, it should be highlighted that we have assumed the ability to seamlessly vary the measurement observable W_i (and also the initial state ρ_j in the extended setup which was just discussed) at each step of the sampling procedure. In practice, these changes may incur a significant experimental overhead, such that it would become sensible to repeat several times the same measurement. The sampling scheme presented could be extended to also take into consideration a number $m > 1$ of measurements to be performed for each of the measurement settings sampled (for instance, see Appendix A [170]), but this will not be investigated in this thesis. Furthermore, we note that the general approach towards estimation which is adopted here (and sometimes referred as direct estimation) only considers measuring experimental observables which explicitly appear in the decomposition of the observable of interest (Equation (2.23)). Recent work [168] has also advocated the use of measurements performed in random basis, but is out of the scope of this thesis.

2.2.3 Application to the state fidelity

To exemplify the formulas obtained in the previous section, we now examine the task of estimating the state fidelity in the context of n -qubit systems. We start by detailing the basis of experimental observables which is chosen – this basis will be used many times in the rest of the thesis when investigating qubit systems – before defining further the state fidelity. Then, we apply the methodology developed, and derive the variance of the estimates of the state fidelity for several families of target states.

Basis of Pauli observables

The basis of experimental observables considered is taken to be formed of tensor products of single-qubit Pauli operators (X , Y and Z) and identity (I). Denoting the set $\mathcal{P} = \{X, Y, Z, I\}$, these Pauli observables are defined as the d^2 distinct operators of the form

$$\{W_i = \frac{1}{\sqrt{d}} P_{j_1} \otimes \dots \otimes P_{j_n} | P_j \in \mathcal{P}\}, \quad (2.36)$$

with scaling factor $1/\sqrt{d}$ ensuring their unit norms ($\text{Tr}[W_i^\dagger W_i] = 1$). Notably, such observables only involve local measurements (over each of the n qubits) which are deemed more reliable than measurements requiring entangling operations. Indeed, for state-of-the-art quantum computing platforms, two-qubit gates are typically one or two orders of magnitudes less accurate than single-qubit rotations. Given that we will often resort to such basis $\{W_i\}$ of observables, it is worth detailing further its properties.

First, any of the Pauli observables $W_i = W_i^\dagger$ is Hermitian, and square to the (scaled) identity $W_i^2 = I_n/d$ (with I_n the identity acting on n -qubit states). Second, the set $\{W_i\}$ forms an orthonormal basis for the space of linear operators. That is, given any operators W_i and W_j , $\text{Tr}[W_i^\dagger W_j] = \delta_{i,j}$. Third, any operator W_i admits an eigendecomposition

$$W_i = \sum_{k=1}^{d=2^n} \lambda_k^i |\lambda_k^i\rangle\langle\lambda_k^i| \quad (2.37)$$

with eigenvalues λ_k^i only adopting values of $\pm 1/\sqrt{d}$, and with eigenvectors $|\lambda_k^i\rangle$ which are tensor products of the eigenvectors of the individual Pauli matrices (or identities) composing to W_i . It follows that their expectation values are bounded in

$$-1/\sqrt{d} \leq \langle W_i \rangle_\rho = \text{Tr}[W_i^\dagger \rho] \leq 1/\sqrt{d}, \quad (2.38)$$

for any state ρ .

State fidelity for problems of quantum state preparation

As discussed in Chapter 1, any task of quantum state preparation requires the definition of a similarity measure between a realized state ρ and a target state $\sigma_{tgt} = |\psi_{tgt}\rangle\langle\psi_{tgt}|^5$. This can be assessed by means of the state fidelity \mathcal{F} defined as

$$\mathcal{F}(\rho) = \text{Tr}[\sigma_{tgt}\rho], \quad (2.39)$$

⁵In tasks of quantum state preparation this target is most often taken to be a pure state, as considered here, but could equivalently be taken as a mixed state.

with values bounded in $[0, 1]$ and with maximal value of $\mathcal{F}(\rho) = 1$ attained only when the prepared state ρ matches exactly its target σ_{tgt} (up to a global phase). Such features, and other appealing properties detailed in [44, 171], make the state fidelity a popular choice when comparing quantum states.

The fidelity in Equation (2.39) can be understood as the expectation value $\langle \sigma_{tgt} \rangle_\rho$ of the projector $\sigma_{tgt} = |\psi_{tgt}\rangle\langle\psi_{tgt}|$ with respect to the unknown state ρ , such that we can apply the methodology previously developed. For that purpose, we first decompose σ_{tgt} in the basis $\{W_i\}$, which yields the d^2 coefficients

$$\alpha_{tgt}(i) = \text{Tr}[W_i^\dagger \sigma_{tgt}] \in [-\frac{1}{\sqrt{d}}, \frac{1}{\sqrt{d}}]. \quad (2.40)$$

For pure target states, these coefficients satisfy the equality

$$\sum_{i=1}^{d^2} \alpha_{tgt}^2(i) = 1, \quad (2.41)$$

that follows from the purity rule [44] $\text{Tr}[\sigma_{tgt}^2] = 1$ (in case of a mixed target state this equality becomes $\text{Tr}[\sigma_{tgt}^2] \leq 1$).

Next, Equation (2.41) can be used to simplify the estimation variance given in Equation (2.30). This yields the expression

$$\Delta^2[\tilde{\mathcal{F}}] = \frac{1}{N} \left(\frac{D}{d} - \mathcal{F} \right), \quad (2.42)$$

relating the variance of an estimate $\tilde{\mathcal{F}}$ of the fidelity \mathcal{F} (obtained over N measurements), to the number D of non vanishing coefficients $\alpha_{tgt}(i) \neq 0$ in Equation (2.40). In particular, this variance does not depend on the state ρ .

Given the constraints on the coefficients $\alpha(i)$ provided in Equations (2.40) and (2.41), it follows that $D \geq d$ (and by definition of the coefficients $D \leq d^2$). Hence, we obtain bounds on the variance of the fidelity estimate (with respect to an arbitrary target state) given by:

$$\frac{1 - \mathcal{F}}{N} \leq \Delta^2[\tilde{\mathcal{F}}] \leq \frac{d - \mathcal{F}}{N}. \quad (2.43)$$

We now discuss specific scalings of this variance for several families of target states of interest.

As can be seen in Equation (2.42), for any target state which can be decomposed over $D = d = 2^n$ (to be compared to the maximal number of $D = d^2$) Pauli observables, the variance of the fidelity estimates does not scale with the dimension d of the target. That is, it would require the same experimental effort $N \approx \varepsilon^{-2}$ to reach an estimate accuracy of $\Delta^2[\tilde{\mathcal{F}}] = \varepsilon^2$, independently of the size of the system considered. Target states satisfying such decompositions include the computational basis states (that are tensor products of qubits in either the $|0\rangle$ or $|1\rangle$ state), or

more generally, any other tensor product of eigenvectors of the Pauli operators⁶.

More interestingly, any stabilizer states [15] (Appendix D) can also be decomposed over $D = d$ elements of the Pauli basis of observables (with coefficients $\alpha(i) = \pm 1/\sqrt{d}$), thus ensuring that estimates of the fidelity with fixed accuracy can be obtained in constant effort, independently of d . These states include the Greenberger-Horne-Zeilinger(GHZ) state and other graph states that are commonly used for quantum information and computing purposes. This may have come as a surprise, given that such states can exhibit non-trivial entanglement [172]. Still, given their particular structure it is now known [166, 167] that they can be probed, and thus could be optimized, in a scalable way.

Other families of states of interest include the W states defined as

$$|\psi_W\rangle = \frac{1}{\sqrt{n}}(|10\dots0\rangle + |01\dots0\rangle + \dots + |00\dots1\rangle) \quad (2.44)$$

which are known to display a different form of entanglement compared to the GHZ state previously mentioned [173], and which can be decomposed over (at most) $D = n^2d$ operators W_i [166]. In such situation, the variance in Equation (2.42) scales as n^2 .

However, for more general target states with support over $\sim d^2$ elements of the operator basis $\{W_i\}$, the variance of the fidelity estimates will grow as the dimension d of the underlying state space, that is exponentially with the system size n . This quickly precludes the accurate estimation of the corresponding fidelities (given a bounded number of measurement repetitions N).

Overall, these results show that the methodology established in this section enables to quickly assess the estimation accuracy of arbitrary observables. We will resort to it, and develop it further, in Chapter 4, in order to compare different figures of merit, in light of their estimation accuracy.

With this ability to estimate and to optimize, that have been developed in this chapter, we are now in position to improve on problems of optimization of quantum dynamics, and now proceed to present the results of this thesis.

⁶It should be noted that for such target states, it is more sensible to perform projective measurements directly onto these states, rather than decomposing them in terms of Pauli observables.

Part II

Towards more efficient optimization

Chapter 3

Bayesian optimization for efficient experimental optimization

Central to the success of any task of optimization is the choice of an adequate optimizer. In this chapter, we investigate the benefits of Bayesian optimization (BO) to perform optimization based on experimental data. BO was suggested in Section 2.1.3 as a promising optimization routine in case when data is noisy and scarce. This situation is common in quantum physics, and particularly prominent in complex experimental setups such as ultracold-atoms quantum simulators.

To demonstrate the potential benefits of BO, we apply it to a paradigm control problem found in the realm of many-body physics: the crossing of a quantum phase transition in faster-than-adiabatic times. This problem is investigated in Section 3.1 to thoroughly benchmark BO against more traditional optimizers. In Section 3.2, we extend the existing BO framework to the poor statistics regime, that is, when the figure of merit is estimated based on very few measurement repetitions. To assess the benefit of this novel approach, we first revisit the task of the quantum phase transition. Later, the framework is applied to two problems of quantum state preparation with parameterized quantum circuits, one performed on numerical simulations and the other one on public-access IBM quantum chips [1]. In all cases, we report an improved convergence of the method proposed compared to other routines commonly employed (these were detailed in Section 2.1). Most of the results and discussions presented in this chapter can be found in [165, 174].

3.1 Application and benchmarking in the context of ultracold-atoms systems

Due to their excellent controllability, ultracold-atoms experiments have been recognized as a platform of choice for the simulation of complex quantum many-body systems [67, 68, 175]. Typical experimental routines involve many steps of preparation, which are followed by further

dynamical transformations of the initial state. All these can be refined by means of QOC [71–74]. In particular, many of the control protocols that are involved often rely on adiabatic or empirical schemes [176–179] which, despite their intuitive appeals, are lengthy and prone to errors (due to detrimental and prolonged interactions with the environment). Devising faithful control protocols over shorter timescales can thus directly enhance the quality of the corresponding simulations, and can extend the reach of physics which could be probed accurately.

Still, optimizing such protocols remains challenging. Due to the many-body nature of the dynamics typically involved, performing accurate numerical simulations of the systems is often not possible, leaving only the option of measurement–based optimization advocated here. Additionally, ultracold atoms simulations typically exhibit low repetition rates, such that one cannot assume access to a large amount of experimental data. These unique challenges in optimizing the control of these experiments call for well-crafted optimization methods and are identified as a situation for which BO may be particularly suitable.

To illustrate the potential of BO in this context, we apply it to an optimal control task of transitioning from a superfluid to a Mott insulator phase (SF–MI transition) in faster-than-adiabatic times. As a well studied problem [176, 180, 181], it will serve us as a perfect example for establishing the merits of BO. Furthermore, preparing many-body states with spatial correlations resembling solid-state matter, such as the Mott insulating state and also many more states [177, 179], is of great interest to the condensed matter community. More generally, the task of QOC, that is addressed here, corresponds to the optimization of a ground-state to ground-state transition in short times, which has broad applicability. For instance, it could permit the acceleration of adiabatic quantum computing schedules [182] or the discovery of swift particles shuttling protocols [183, 184].

In Section 3.1.1 we describe the underlying physics of the system considered and introduce the optimal control problem to be tackled. Later, in Section 3.1.2, we discuss an example of optimal control protocol discovered with BO, and present results of a thorough comparison between BO and the alternative optimization routines which were detailed in Section 2.1.

3.1.1 The superfluid to Mott insulator control problem

The controlled system

The Bose-Hubbard model, that we consider here, is a widely studied model in many-body physics. Although conceptually simple, it cannot be mapped to a single particle problem and contains a transition between a superfluid (SF) and a Mott insulating (MI) quantum phases [180]. For a system of N bosonic atoms in a lattice with L sites, the Bose-Hubbard Hamiltonian

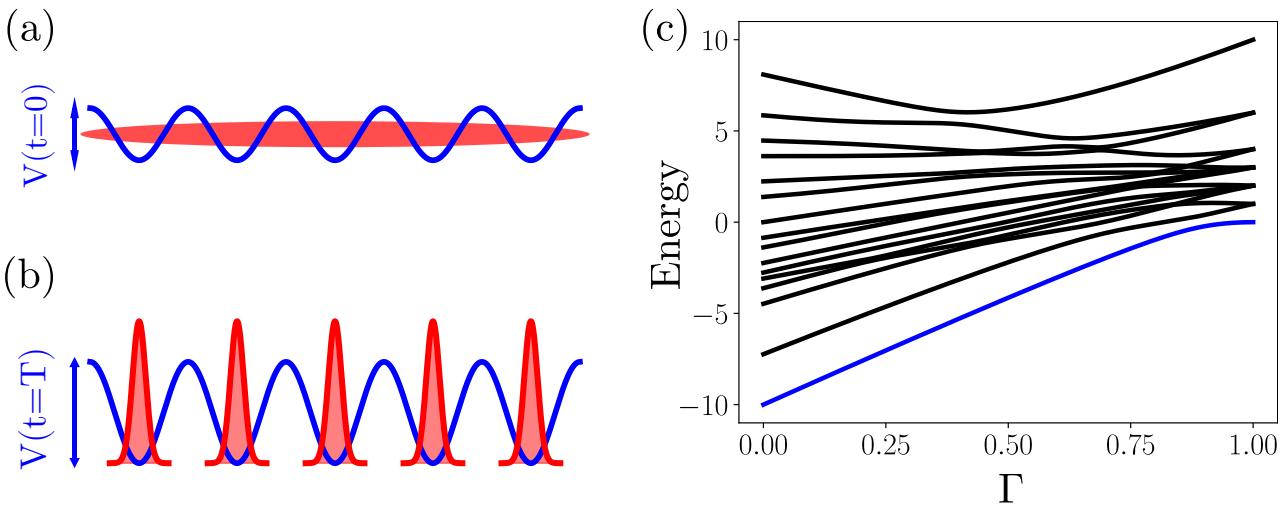


Figure 3.1: From [165]. Illustration of the superfluid to Mott insulator (SF–MI) control problem. The system considered here is a one–dimension bosonic chain (a–b) described by Equation (3.1). We aim at driving the system from an initial superfluid (a) to a Mott insulating phase (b) at time $t = T$ by dynamically changing the depth of the optical potential $V(t)$ (blue). (c) Energy spectrum of the Hamiltonian in Equation (3.2) for different values of the control parameter Γ , and with ground state energies highlighted in blue.

is given as

$$H(t) = -J(t) \sum_{i=1}^L (b_i b_{i+1}^\dagger + b_{i+1} b_i^\dagger) + \frac{U(t)}{2} \sum_{j=1}^L \hat{n}_j (\hat{n}_j - 1), \quad (3.1)$$

where $b_i (b_i^\dagger)$ denotes the annihilation (creation) operator of a boson at site i and $\hat{n}_i = b_i^\dagger b_i$ denotes the number operator. These bosonic operators follow the usual commutation relation, $[b_i^\dagger, b_j] = \delta_{ij}$. The first term of the r.h.s. in Equation (3.1) describes the tunneling of bosons between neighboring sites with energy $J(t)$, and favours the delocalization of the atoms over the lattice. The second term of the r.h.s captures the on–site interactions between two (or more) bosonic atoms with strength $U(t)$.

As sketched in Figure 3.1(a–b), such model can be realized in terms of bosonic particles in an optical lattice potential (in blue) [66, 68, 175]. Dynamical changes in the depth of the lattice potential $V(t)$ allow to vary both the energies $J(t)$ and $U(t)$ in time. The tunneling $J(t)$ changes depending on the energy barrier between neighboring lattice sites, while variations in the on-site interaction $U(t)$ result from changes in the confinement of the atoms at each site. The exact expressions relating the depth of the lattice potential to the interaction and tunneling energies can be found, for instance, in [181]. To make the optimization problem more tractable, we introduce a single dimensionless quantity to be controlled, defined as $\Gamma(t) = U(t)/(U(t) + J(t))$ such that $\Gamma(t) \in [0, 1]$ (for the positive values of U and J that are considered here). Given this

new quantity, the Hamiltonian in Equation (3.1) is recast as

$$H(t) = -(1 - \Gamma(t)) \sum_{i=1}^L (b_i b_{i+1}^\dagger + b_{i+1} b_i^\dagger) + \frac{\Gamma(t)}{2} \sum_{j=1}^L \hat{n}_j (\hat{n}_j - 1). \quad (3.2)$$

Depending on the energy ratio U/J , this Hamiltonian admits two distinct ground states, that we now describe. In the limit where $\Gamma(t) = 0$ ($U \ll J$), the interaction term vanishes and the ground state of each atom is delocalized over the entire lattice. This superfluid (SF) ground state is described as a superposition of different Fock states

$$|\psi_{SF}\rangle \propto \left(\sum_i^L b_i^\dagger \right)^N |\mathbf{0}\rangle, \quad (3.3)$$

where $|\mathbf{0}\rangle$ is the many-body vacuum state in the Fock basis. This SF state is characterized by a well-defined global phase but variance $\Delta^2[\hat{n}_i] > 0$ in the number operator.

In the opposite limit where $\Gamma(t) = 1$ ($U \gg J$), the hopping between adjacent sites is suppressed and the ground state of the system consists of localized particles that minimize the interaction energy. For a commensurate filling of N/L atoms per lattice site, this Mott insulating (MI) ground state is a product of local Fock states over each lattice site:

$$|\psi_{MI}\rangle \propto \prod_i^L (b_i^\dagger)^{N/L} |\mathbf{0}\rangle. \quad (3.4)$$

In this case, the global phase coherence of the atomic wave function is lost and the variance in the number operator vanishes. To provide further insights on the energies at play in between these two extreme values of $\Gamma = 0$ or $\Gamma = 1$, we now study the eigenspectrum of Equation (3.2). In the rest of this section, we restrict our attention to the control of a chain with $L = 5$ lattice sites, unit filling ($L = N = 5$) and periodic boundary conditions.

Many-body energy spectrum

The energies E_n of the many-body eigenstates $|n\rangle$ of H are obtained by means of exact diagonalization for varied values of Γ , and are displayed in Figure 3.1(c). Given the size and translation invariance of the system considered, we obtain a total of 16 distinct eigenvalues, with further degeneracy at $\Gamma = 1$. The ground energies of the spectrum are highlighted in blue, ranging from the SF ground energy on the left ($\Gamma = 0$) to the MI ground energy on the right ($\Gamma = 1$) with specific values of $E_0(\Gamma = 0) = -2N$ and $E_0(\Gamma = 1) = 0$ respectively. As we will see, the control task corresponds to a ground to ground state transition (from left to right), which from an energy perspective, is complicated by the reduced energy gap around $\Gamma \approx 0.9$

and the presence of multiple avoided crossings in the spectrum. We now discuss further the exact details of the QOC task.

Optimal control problem

We aim at identifying the optimal control function $\Gamma^{opt}(t)$ which transforms an initial SF state into a targeted MI state, at time $t = T$. Furthermore, we impose boundary conditions of $\Gamma(t = 0) = 0$ and $\Gamma(t = T) = 1$, such that the initial and target state should be ground states of the Hamiltonian of the system at initial and final times. In addition to this general objective of state preparation, to fully specify the QOC problem we still need to define the figure of merit to be maximized, the duration T allowed for the transformation, and the finite set of control parameters $\boldsymbol{\theta}$ which will be optimized over.

First, the control $\Gamma(t)$ is parametrized in terms of a vector $\boldsymbol{\theta}$ of $P = 10$ parameters $\boldsymbol{\theta}^{(j)} \in [0, 1]$. Each parameter corresponds to the value adopted by the control at fixed time steps $\Delta t = T/(P + 1)$. That is, $\Gamma(t = j \times \Delta t) = \boldsymbol{\theta}^{(j)}$ with $(j = 1, \dots, 10)$. The values of the control at intermediate times are obtained by fitting a cubic spline to these P values and the two boundary conditions $\Gamma(0) = 0$ and $\Gamma(T) = 1$. Such parametrization ensures that any set of control parameters $\boldsymbol{\theta}$ defines a unique control function $\Gamma(t)$ that varies smoothly.

Second, we consider the fidelity (with respect to the target MI state) as the figure of merit to be maximized. It is defined as $\mathcal{F}_{MI}(\boldsymbol{\theta}) = |\langle \psi_{MI} | \psi_{\boldsymbol{\theta}}(T) \rangle|^2$, where $|\psi_{\boldsymbol{\theta}}(T)\rangle$ denotes the state realized at time $t = T$ given the control parameters $\boldsymbol{\theta}$. While the state fidelity is a natural choice for any quantum state preparation problem, experimentally assessing its value may not always be possible. Hence, a second choice of figure of merit will be explored, and discussed in due course.

Lastly, the duration T of the protocol is defined with respect to the quantum speed limit (QSL) which is given in [76, 185] for this problem as $T_{QSL} = \pi/\Delta$, with Δ being the minimum energy gap between the ground and first excited state (this gap is evaluated based on the numerical energy values that were found in Figure 3.1(c)). Such QSL [186] theoretically bounds the minimum time necessary for a perfect preparation of the MI state. However, there is no guarantee that perfect preparation can be achieved in such time, especially given the restricted family of control functions (parametrized by a finite number of 10 parameters) that are optimized over. Importantly, the durations T that will be considered are taken to be close to T_{QSL} , that is shorter than the adiabatic timescales $T_{adiab} \gg 1/\Delta$ of the system, which are also defined with respect to the energy gap Δ .

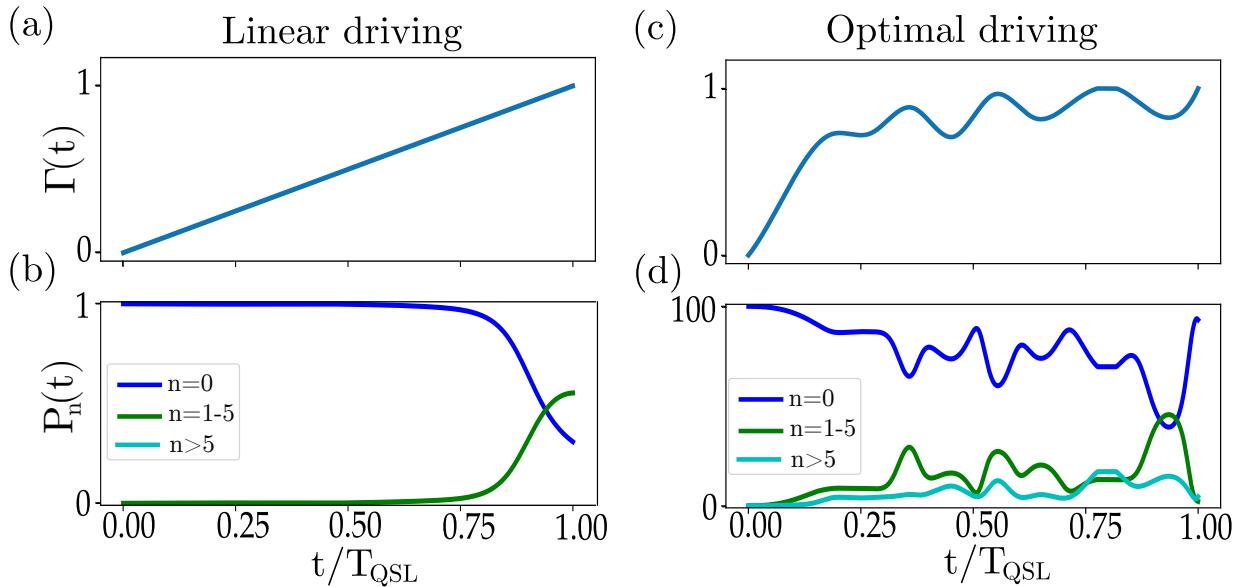


Figure 3.2: From [165]. Two control strategies for the SF–MI transition are presented: a linear driving of the system (a), and a driving optimized with BO (c). The evolution of the population of the instantaneous eigenstates (defined in the main text) is reported for the linear (b) and the optimized (d) controls. For clarity, these populations are grouped by ground state only (blue), first to fifth (green), and higher (turquoise) energy eigenstates. In each case, the final ground state populations $P_0(T_{QSL})$ is equal to the fidelity of the realized state with respect to the target MI state.

3.1.2 Results and discussions

Before proceeding to a more systematic study of BO, we start by presenting the results obtained for a single optimization problem, defined for $T = T_{QSL}$ and assuming that the optimizer has access to the exact values of the fidelity at each iteration. In this situation, BO consistently converges in a few hundred of iterations to values of the fidelity of $\mathcal{F}_{MI} \approx 90\%$. The optimized control is plotted in Figure 3.2(c), where it can be compared to a more naive strategy, depicted in Figure 3.2(a), consisting in a linear increase of the values of Γ from 0 to 1.

To understand the dynamics involved, it is insightful to examine the squared overlaps between the system state $|\psi(t)\rangle$ and the instantaneous eigenvectors $|n(t)\rangle$, which are obtained by repeated eigendecomposition of the Hamiltonian at varied time t . These overlaps $P_n(t) = |\langle\psi(t)|n(t)\rangle|^2$, that are called instantaneous populations, are reported in Figure 3.2(b),(d) for the linear control and the optimized control respectively. Given that the MI state is the ground state of the Hamiltonian of the system at final time $T = T_{QSL}$, the ground state instantaneous population $P_0(T_{QSL})$ is effectively equal to the fidelity \mathcal{F}_{MI} of the realized state with respect to the MI one, and can be assessed directly from the plots.

It can be seen, in Figure 3.2(b), that in case of linear driving, most of the system remains in the ground state until around $t = 0.75 \times T_{QSL}$. However, at later time, for values of Γ

inducing a smaller energy gap in between the ground and first eigenstate (as can be observed in Figure 3.1(b)), excitations to the excited instantaneous eigenstates (green curve) quickly build up. Notably, the population in these excited states is found to monotonically increase, resulting in a final fidelity of only $\mathcal{F}_{MI} \approx 30\%$. In contrast, as seen for the optimized driving in Figure 3.2(d), part of the initial state gets quickly promoted to excited eigenstates, but, revert to the ground state over the course of the dynamics. This ability to recover a large overlap, found to be $\mathcal{F}_{MI} \approx 90\%$, with the MI state is attributed to the taylored oscillations in the optimized control which are seen in Figure 3.2(c). Similar patterns were recently reported [187] – for a comparable control problem but with larger system size – suggesting a potential generalization of the drivings found by BO.

Finally, we also assess the minimal time that it would have taken for a linear driving of the system to yield a similar value of the fidelity to the one obtained with the optimized controls. This minimal time is estimated by repeating simulations of the system with linearly increased control over varied durations T . We find that it would have taken around 7 times longer for the linear protocol to reach a value $\mathcal{F}_{MI} \approx 90\%$ of the fidelity. Hence, we saw that optimized controls can yield similar fidelity than simpler control protocols, but, in shorter timescales. In experimental conditions this would permit to limit errors occurring due to lengthy interactions with the environment. These results highlight the appeal of QOC, provided that these optimized controls can be efficiently discovered in the first place.

Benchmarking methodology

To appreciate further the appeal of BO, it is appropriate to compare it to other optimization routines commonly employed. Hence, in addition to BO we will also study the differential evolution (DE), from the family of evolutionary techniques, the Nelder–Mead (NM), and the Simultaneous Perturbation Stochastic Ascent (SPSA) algorithms, which were all reviewed in Section 2.1. Both NM and DE are implemented using Scipy [188], SPSA is implemented following [140], and BO relies on the implementation which can be found in [189]. Additionally, a simple random search (RANDOM) is also included in the list of the optimizers considered. It relies on randomly sampling a new set of parameters at each iteration and will serve us as a benchmark. When design choices (often known as hyperparameters) in the implementation of the algorithms investigated are known to affect their performances, we do our best effort to select them adequately.

The performance of each of the optimizers mentioned is judged by its ability to reach low infidelity, defined as $\mathcal{I} = (1 - \mathcal{F}_{MI})$, in as few iterations as possible. Furthermore, since any optimization relies on the choice of random initial parameters, and may be affected by further stochastic effects, we repeat each optimization 30 times and systematically report the median

values (displayed by means of filled symbols) and interquartile ranges (displayed by means of shaded regions) when presenting the results in Figure 3.3¹.

To be instructive, the comparison between BO and the other optimizers will be performed over different problem configurations that are now described. First, the duration T is known to have a strong influence on the complexity of many QOC tasks. For instance, it was reported that close to the QSL, quantum state preparation problems often become harder [31, 117], with hardness characterized by the presence of many local extrema in the optimization landscape. As we aim at finding optimal controls in relatively short timescales, we will investigate three different situations corresponding to a duration of $T = 0.5 \times T_{QSL}$, T_{QSL} , and $1.2 \times T_{QSL}$.

Additionally, the convergence of any optimizer may be tied to the choice of the figure of merit employed. Hence, in addition to the fidelity that was discussed, we will also resort to a second figure of merit. With $\Delta^2[\hat{n}_i]$ the variance in the occupation number \hat{n}_i at site i , this second figure is defined as the average $\langle \Delta^2[\hat{n}_i] \rangle_{i=1,\dots,L}$ over each lattice site. In principle, this figure can be assessed in ultracold atoms experiments permitting site resolution [190]. Furthermore, as this variance is minimized(maximized) for the MI(SF) state, it could serve as a figure of merit for each of the corresponding problem of state preparation. As such, this average variance is deemed of particular interest, and we will also perform optimization with this new figure. Finally to make it more experimentally realistic, we will limit its evaluation to be based on a finite number $N = 1000$ of repeated measurements, thus introducing sampling noise in its values, in contrast to the exact values of the fidelity.

Benchmarking results

The results obtained for the 6 configurations that has just been described (corresponding to 3 durations and 2 figures of merit) are displayed in Figure 3.3. As can be seen, over most of the different configurations investigated, BO (orange curve) exhibits faster convergence towards low infidelity compared to any other alternative. We now discuss more in-depth these results starting with the case where exact values of the fidelity are available to each optimizer (first row).

For the short evolution time of $T = 0.5 \times T_{QSL}$ (first column), it can be seen that most of the optimizers converge to what seems to be the global optimum in circa 100 iterations. In particular, the two local optimizers (NM and SPSA) exhibit the fastest convergence towards the control solution. However, the maximum fidelity reached is only of the order of 50%. These findings of easy optimization (that is with a landscape seemingly devoid of local extrema) for durations shorter than the QSL are in accordance with the results presented in [31].

¹We follow this methodology for all the optimization results that are reported in this thesis.

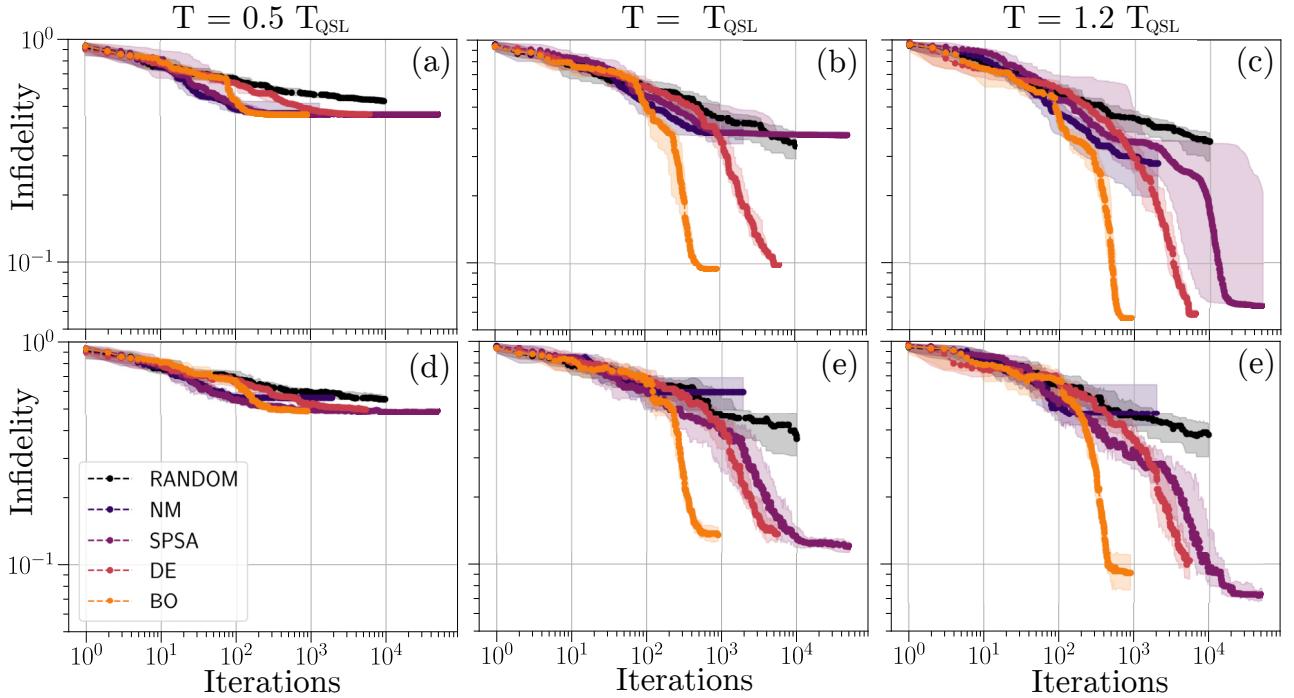


Figure 3.3: From [165]. Benchmarking of BO (orange) against 4 other optimization strategies (colours in legend and described further in main text) for the SF–MI control problem. The convergence of each optimizer is reported in terms of the infidelity $\mathcal{I} = 1 - \mathcal{F}_{MI}$ as a function of the number of iteration steps (both in log scales). The comparison is performed for different control durations of $T = 0.5 \times T_{QSL}$ (first column), $T = T_{QSL}$ (middle column) and $T = 1.2 \times T_{QSL}$ (last column), and also for different figures of merit with either the exact evaluation of the fidelity \mathcal{F}_{MI} (first row), or a more experiment–realistic figure (second row) estimated with a finite number of measurements, and defined in the main text. In all cases, the median (filled symbols) and interquartile range (shaded regions) are evaluated over 30 repetitions of the same task.

For longer durations $T \geq T_{QSL}$, however, only DE and BO consistently find solutions with infidelities below 10%, with BO converging in roughly an order of magnitude less iterations. The presence of many local minima is evidenced by the performance of the SPSA optimization routine. For $T = T_{QSL}$ the 30 optimization runs of SPSA get trapped in local minima with a low fidelity $\mathcal{I} \approx 50\%$. The situation is improved for $T = 1.2 \times T_{QSL}$, but closer inspection of the interquartile intervals (depicted as shaded regions) shows that more than one fourth of the SPSA runs remain trapped until up to 4×10^4 iterations. This feature of the optimization landscape was also noticed in [31] and explains the substantial performance gap between the global optimizers (such as DE and BO) and the local optimizers (such as NM and SPSA) which are studied. Furthermore, it can be seen that for the increased duration $T = 1.2 \times T_{QSL}$ the performances of the optimizers are globally ameliorated. This suggests that finding the right balance between short control durations, thus limiting accumulated errors, and longer times, which facilitate the identification of the optimal controls, may play an important role when identifying appropriate experimental controls.

The results for the optimizations performed based on the noisy evaluations of the second figure of merit are reported in Figure 3.3(second row). In this more experimentally realistic situation, NM is found to underperform the simple random strategy, highlighting the extreme sensitivity of NM to noise. Surprisingly, for the case $T \geq T_{QSL}$ performances of SPSA are improved. This can be partially explained by the ability of SPSA, which is originally designed for situations of stochastic gradients [140], to benefit from noise in the evaluation of the figure of merit in order to escape local minima. Akin to the noiseless scenario, for $T \geq T_{QSL}$ BO is found to converge substantially faster than any of the other optimization strategies studied. Overall the final fidelities attained have dropped by roughly 5%, compared to the noiseless situation, indicating the additional difficulty in performing optimization with noisy observations.

Discussion

These results aim at encouraging the use of BO as a practical tool when optimizing the dynamics of complex quantum systems. By resorting to an extensive comparison, permitted by the use of numerical simulations of limited system sizes, we were able to characterize the potential of BO, and found it to outperform the other alternative investigated in the most relevant situations, that is when $T \geq T_{QSL}$. Ultimately, we argue for the adoption of BO in experimental setups, rather than based on the numerical simulations employed here. Still, it is expected that the results reported, especially when incorporating sampling noise, should be representative of the difficulties that any optimization will encounter in realistic situations.

During the course of this thesis, BO has gained in popularity and has been applied to several experimental realizations of QOC and VQAs [191–193], confirming further its appeal. We note

however, that in all these examples only a few 100s of iterations were performed, compared to the maximum of 1000 that was studied here. Such limited numbers of iterations can be explained by the experimental effort required per iteration, which typically involves many repeated measurements. A large number of sampling repetitions ensures low statistical variance when estimating the figure of merit, which in general facilitates the optimization, but, comes at the expense of an increased experimental effort at each iteration. In the following section we explore further the importance of this sampling number, and extend BO to also operate in regimes where this number is taken small.

3.2 Optimization with poor statistics

Most of the optimization results presented in problems of QOC and VQAs feature the convergence of a figure of merit as a function of a number of optimization iterations. When optimizations are performed based on experimental outcomes, it is important to recognize that the experimental effort associated with any of these iterations depends on the desired accuracy in the estimates of the figure of merit. Good estimates, that are estimates with high statistical accuracy, are obtained by repeating the same experiment many times. This is typically the regime in which operate experimental optimizations [45, 56, 191, 192], where this number of repetitions (the sampling number) is often found to be in the range of a few hundreds up to more than tens of thousands of repetitions per optimization step. This raises several questions that we aim at addressing in this section: can one still be able to perform optimization with poor statistical estimates, that is, with estimates obtained with very few measurement repetitions? And if so, is it beneficial when convergence is judged based on the overall experimental effort required, rather than the bare number of iterations?

In the following, we study the importance of the choice of the sampling number N when performing optimization. In Section 3.2.1, we motivate and present a framework extending BO to operate in the poor statistics regime, up to the case of single-shot measurements (that is, when only a single $N = 1$ measurement repetition is performed). This is achieved by incorporating precise details of measurement statistics, as found in typical quantum physics experiments. This new framework is then judged based on its ability to find good solutions with minimal experimental effort. For that purpose, in Section 3.2.2, we revisit the SF–MI transition problem and show enhanced convergence of the proposed framework. This is confirmed, in Section 3.2.3, on additional problems of state preparation with parametrized quantum circuits. In all cases, the proposed framework is compared to a traditional implementation of BO and to the SPSA algorithm, and is found to require one or two orders of magnitude less experimental repetitions than its competitors.

3.2.1 Bayesian optimization with binomial noise

Recall from Section 2.1.3 the main steps involved in building a (surrogate) model in BO. First, prior assumptions $p(f)$ on the functions f – modeling the unknown figure of merit F to be maximized – are formulated in terms of a Gaussian process distribution. Next, this prior distribution is conditioned, by means of Bayes rules, on the vector $\mathbf{y} = [y_1, \dots, y_D]$ of observations available. These observations correspond to noisy evaluations of F for input parameters $[\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_D]$. Relating these noisy observations to the underlying model is achieved by specifying the likelihood $p(\mathbf{y}|f) = p(\mathbf{y}|\mathbf{f})$ of observing \mathbf{y} given a function f which adopts values of $\mathbf{f} = [f(\boldsymbol{\theta}_1), \dots, f(\boldsymbol{\theta}_D)]$.

As was discussed in Section 2.1.3, the likelihood effectively contains assumptions about the noise induced by the physical measurement process. For instance, the likelihood

$$P(\mathbf{y}|f) = \mathcal{N}(\mathbf{f}, \sigma_N^2 I), \quad (3.5)$$

which is employed in traditional BO, and which appeared in Equation (2.6), carries the assumption of independent and identically distributed (i.i.d.) additive Gaussian noise. Such likelihood allows for an analytical treatment of the Gaussian process. that is, it yields an analytical expression of the posterior distribution (given in Equation (2.8)), but, may not always be justified. In particular, we will see that the assumptions entailed by Equation (3.5) are too strong when the values of F are estimated based on a small number of measurement repetitions. While a general figure of merit F is function of several experimental observables – or equivalently a function of probabilities of outcomes of several measurements – we first elaborate on the case of a single observable.

Poor statistics of a single observable

Consider a quantum measurement with binary outcomes as involved, for instance, when estimating the expectation value of a tensor product of Pauli operators. Given the two possible outcomes, the expected value of an operator M associated with such measurement can be evaluated based solely on the probability of observing one of the outcomes (the probability of observing the other outcome directly follows), and we call the outcome corresponding to the largest eigenvalue of M a success. Hence, estimating the expectation value of M only requires to estimate the probability of observing a success, which is denoted $p(\boldsymbol{\theta})$ and which would depend on some control parameters $\boldsymbol{\theta}$. This probability can be estimated in terms of the frequency $\tilde{n} = n/N$ of recording n success outcomes over a total of N repetitions. This number n follows

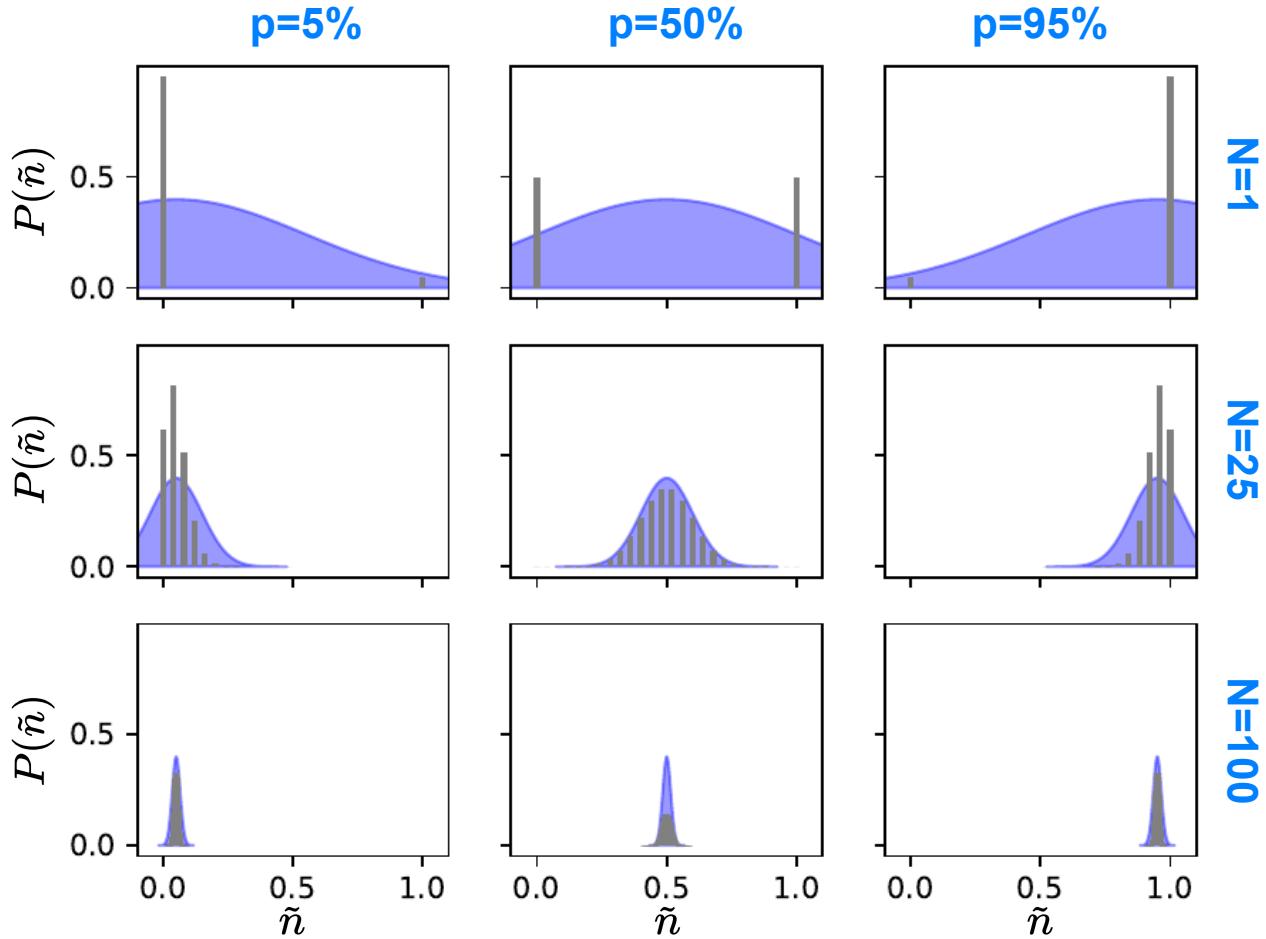


Figure 3.4: Comparison between the statistics of an estimate \tilde{n} obtained according to the true binomial distribution (grey) of the measurement outcomes (Equation (3.6)) and its approximation (blue) in terms of a Gaussian distribution (Equation (3.7)). The probabilities $p(\tilde{n})$ are reported for different number of repetitions with values $N = 1$ (first row), 25 (second row), 100 (last row) and for different values of the underlying probability of success with values $p = 0.05$ (first column), 0.5 (second column) and 0.95 (last column). Qualitatively, it is only in the case of enough repetitions $N \gtrsim 100$ that the Gaussian approximation becomes justified. This motivates us to extend BO to take into consideration the true nature of the binomial measurement outcomes, rather than resorting to the approximate Gaussian distribution, commonly employed.

the binomial distribution

$$P(n|p(\boldsymbol{\theta})) = \binom{N}{n} (p(\boldsymbol{\theta}))^n (1 - p(\boldsymbol{\theta}))^{N-n}, \quad (3.6)$$

with average value $\langle n \rangle = Np(\boldsymbol{\theta})$ and variance $\Delta^2[n] = Np(\boldsymbol{\theta})(1 - p(\boldsymbol{\theta}))$.

Given a sufficiently large number N of measurement repetitions, and provided that the probability of success is sufficiently different from $p(\boldsymbol{\theta}) = 0$ and $p(\boldsymbol{\theta}) = 1$, it is known that the distribution of the frequencies \tilde{n} is well approximated by the Gaussian distribution

$$\mathcal{N}\left(p(\boldsymbol{\theta}), \frac{p(\boldsymbol{\theta})(1 - p(\boldsymbol{\theta}))}{N}\right). \quad (3.7)$$

A rule of thumb for Equation (3.7) to be a good approximation is that for

$$Np(\boldsymbol{\theta}) \geq 5, \text{ and for } N(1 - p(\boldsymbol{\theta})) \geq 5 \quad (3.8)$$

to be simultaneously satisfied. In other cases, in particular for small number of repetitions N , this Gaussian approximation is not justified anymore. That is, the distribution of the discrete and bounded values that \tilde{n} can adopt will deviate strongly from the Gaussian distribution in Equation (3.7).

This can be qualitatively observed in Figure 3.4 where the exact distributions (in grey) of the values of \tilde{n} obtained according to the true binomial nature of the measurements (Equation (3.6)) and its approximation (in blue) in terms of a Gaussian distribution (Equation (3.7)) are compared for different sampling number N and different underlying success probabilities p (dropping the dependency to $\boldsymbol{\theta}$). For the minimum number of repetitions $N = 1$ (first row), it can be seen that the Gaussian distribution systematically fails in approximating the true distribution of the estimates \tilde{n} . Similarly, inconsistencies between the two distributions are observed for an increased number $N = 25$ of repetitions (second row) and values of the probabilities $p = 5\%$ and 95% . It is only in the case of larger number of repetitions $N = 100$ (last row) than the two distributions qualitatively agree.

Additionally, it is also the case that the variance $\Delta^2[\tilde{n}] = p(\boldsymbol{\theta})(1 - p(\boldsymbol{\theta}))/N$ of the estimate \tilde{n} depends on the underlying probability $p(\boldsymbol{\theta}) \in [0, 1]$, and varies in the range $[0, 1/(4N)]$. This is in contrast with the assumption of identical noise contained in the likelihood in Equation (3.5). Again, these inconsistencies between the assumed and true estimates statistics are exacerbated when the number N of measurement repetitions is taken small. As we are interested in exploring the use of BO in the regime of small N , it seems unproductive to rely on the Gaussian noise assumption. This motivates us to extend BO to incorporate details of the binomial distribution as appearing in Equation (3.6).

Modeling with binomial noise

In order to achieve an accurate modeling of the statistical noise found in the estimates of a figure of merit, several refinements of the original framework of BO need to be undertaken. A full account of the technical implementation details required can be found in [174](Sec. II), and here we only highlight the main conceptual aspects.

1. As we aim at modelling the outcome probabilities $p(\boldsymbol{\theta})$, it is crucial that the values of the model remain bounded in the interval $[0, 1]$. While functions f following a Gaussian process distribution do not naturally satisfy such condition, their values can be squashed to $[0, 1]$ by composing f with any monotonic function satisfying $\pi : \mathcal{R} \mapsto [0, 1]$. In practice, we resort to the cumulative distribution function of a standard normal distribution $\pi(x) = \int_{-\infty}^x dy \exp(-y^2/2)/\sqrt{2\pi}$, resulting in modeling probabilities $g(\boldsymbol{\theta}) = \pi(f(\boldsymbol{\theta})) \in [0, 1]$ of the true probabilities $p(\boldsymbol{\theta})$.
2. Measurement data is included by means of the proper binomial likelihood. Equation (3.6) only specifies the probability $P(n|\pi(f(\boldsymbol{\theta}))$ to observe n successes for fixed parameters $\boldsymbol{\theta}$, and needs to be extended to take into consideration a whole dataset of measurements outcomes $\mathbf{n} = [n_1, \dots, n_D]$ obtained for parameters $\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_D$. Given the independence of the measurement outcomes obtained for varied parameters $\boldsymbol{\theta}$, it follows that the likelihood of observing \mathbf{n} given a model f is the product

$$P(\mathbf{n}|f) = \prod_{j=1}^D P(n_j|\pi(f(\boldsymbol{\theta}_j))), \quad (3.9)$$

which will be used.

3. The predictive distribution resulting from the changes listed in Items 1 and 2 does not admit an analytical solution anymore, but is rather obtained by means of the Laplace approximation (detailed in [164]), which permits its efficient evaluation.
4. Items 1 to 3 only consider the modeling of a single measurement probability p . A general figure of merit $F = \mathcal{Q}(p_1, \dots, p_k)$ is function of several measurement probabilities p_k which are individually modeled in terms of g_j (following the guidelines in Items 1 to 3). The overall model g of F is taken to be $g = \mathcal{Q}(g_1, \dots, g_k)$.

Other aspects of BO, such as the choice of the next parameters to be evaluated (Section 2.1.3), remain unchanged. Despite our claim of accurate modeling of measurement noise, it should be stressed that the practical implementation of the framework requires the numerical routine listed in Item 3. This consists in an approximation permitting efficient use of the model, rather than fundamental assumptions on data acquisition. As such, it is expected that despite this approximation, the overall scheme would perform better than the traditional BO framework, especially for small number of repetitions. This will be verified in the following examples. The

new flavour of BO is referred as BO with binomial modelling (and sometimes simply binomial BO) as opposed to a traditional BO with Gaussian modeling of the noise. To assess the correct implementation of Items 1 to 4, we start by applying the framework to a one-dimension optimization problem, for which visual inspection of the model is possible.

Illustrative example

Here, we revisit the problem of maximization of the toy function defined in Equation (2.11). This function, which is depicted in Figure 3.5 (dashed red line in the top panels), takes value $F(\theta) \in [0, 1]$ for $\theta \in [0, 4]$, and thus provides a synthetic example of a measurement probability to be maximized. This probability $F(\theta)$ could correspond, for instance, to the measurement of a single Pauli observable with respect to a state prepared with control parameter θ . In order to validate the ability of the framework to optimize in the poor statistics regime, we investigate the extreme case where estimates of F are obtained based on a single ($N = 1$) measurement repetition. These measurements (red dots in the top panel) are randomly sampled with probability $F(\theta)$ to take a value of 1 (and probability $1 - F(\theta)$ to take a value of 0). Despite this very limited amount of measurement information, we aim at quickly identifying the parameter θ^* which maximizes F .

Different stages of BO are plotted in Figure 3.5 after that 30 (a), 31 (b) and 100 (c) single-shot measurements have been taken. The surrogate model of BO with binomial modelling, is depicted in the top panels of each of the subfigures, in terms of its mean value (blue), a 95% confidence interval (gray contours) and the actual probability density (in shades of blue). The acquisition function corresponding to each of the three models is displayed (red curve) in the lower plots.

We first consider the initialization stage, depicted in Figure 3.5(a), when the surrogate model is built for the first time based on 30 single-shot measurements which have been taken for random values of the parameter θ . As can be seen, the model already replicates (in average) some of the main features of the unknown function F , but given the very limited data, fails to distinguish in between local and global maxima. Close inspection of the model confirms the proper implementation of binomial BO. First, as expected, its values are well behaved and bounded in the interval $[0, 1]$. As would also be the case for the Gaussian modeling of noise, it can be observed that the width of the confidence interval increases when further away from any observations. More subtly, this width also widens for probabilities identified to be close to 50% (for instance, around $\theta = 3.5$), that is, when the variance in the measurement data is expected to be maximal and thus when binomial measurement data is the less informative. Such feature would not have been captured by the Gaussian likelihood, and indicates an adequate implementation of the framework.

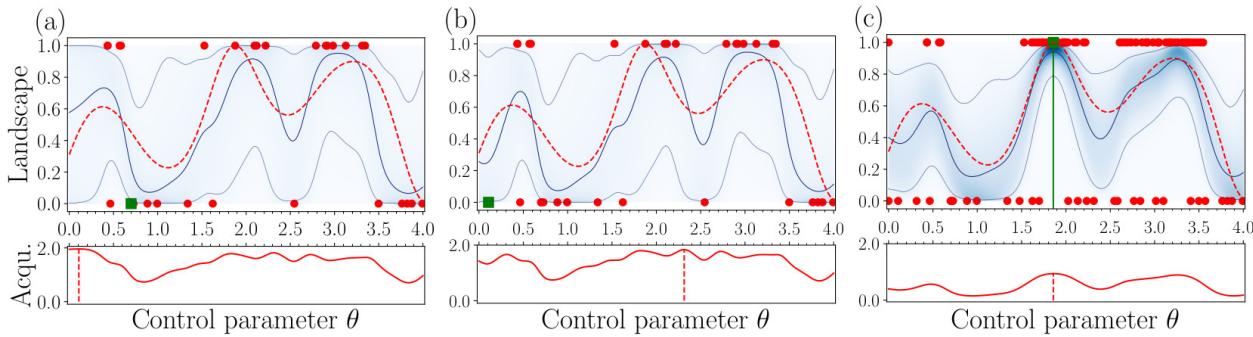


Figure 3.5: From [174]. We revisit the optimization of the toy function $F(\theta) \in [0, 1]$ defined in Equation (2.11) and depicted in the top panels (dashed red lines). This function serves the purpose of a synthetic measurement probability to be maximized. Each evaluation of F consists in a single measurement repetition ($N = 1$). These measurements (red dots) are randomly drawn such that they take a value of 1 (or 0) with probability $F(\theta)$ (or $1 - F(\theta)$). Three stages of BO with binomial modeling after (a) 30, (b) 31, and (c) 100 single-shot measurements have been obtained. In all three cases, the surrogate model is represented in the top panel in terms of its mean value (solid blue line), a 95% confidence interval (solid gray line) and its actual probability density (shades of blue). Additionally, the lower panel depicts the acquisition function (plain red line), which maximum (dashed vertical line) determines the next parameter to be evaluated. Details of this acquisition function are the same as used in Figure 2.3.

As is usual in BO, the next parameter to be evaluated is chosen where the acquisition function (bottom panel) is maximized (red vertical line). At the initial step in Figure 3.5(a), this acquisition is maximized close to $\theta = 0.1$. The resulting measurement of 0, obtained for this parameter, is added to the dataset of observations, and can be seen in Figure 3.5(b) (green square). Given this additional measurement, the model now exhibits a closer fit to the true function F , and a reduced confidence interval (that is, an increased confidence in its prediction) close to the newly acquired data. Still, as this additional binary data only reveals very limited information, this increase in confidence remains slender.

During the course of the optimization (from left to right), the surrogate model becomes a better approximation of the actual landscape especially in the vicinity of the two most probable maxima but, as desired, the algorithm avoids the unnecessary effort that would be required to approximate the landscape well in other domains. After 70 iterations, the algorithm has identified the parameter $\theta_f \approx 1.86$ (indicated by the vertical green line in Figure 3.5(c)) as optimal. This choice of parameter yields a value of $F(\theta_f) = 1 - 4 \times 10^{-3}$, close to the true maxima of $F = 1$, and was found given a total number of 100 measurements that would have only been enough to assess the value of F with a resolution of 0.01 for a single parameter.

Already, this allows us to appreciate the advantage of resorting to a small number of measurements over many different parameter values, rather than concentrating too many measurements on a single parameter value. Exploring further this toy problem, we find similar performances for BO given a number of $N = 10(25)$ measurement repetitions after around $S = 25(15)$

evaluations. These would have required a total of $N_r = 250(375)$ individual measurements, substantially larger than the total number $N_r = 100$ that was required for the single-shot, $S = 1$, case. For this simplistic model, it would thus seem that reducing N is beneficial for algorithmic convergence when judged based on the overall experimental effort necessitated (that is, based on the total number N_r of measurements required). To confirm such assertion, and to study more systematically the improvement enabled by the new flavour of BO, we now revisit the SF–MI control task.

3.2.2 The superfluid to Mott insulator transition revisited

Control problem

As was the case in Section 3.1.1, the system under control is described by the Bose-Hubbard Hamiltonian, which was defined in Equation (3.2), with periodic boundary conditions and with a system size of $B = 5$ bosons and $L = 5$ sites. Recall that our aim is to find an optimal control of the system permitting the transformation of a superfluid (SF) initial state to a Mott insulating (MI) state at time $t = T$. The control function $\Gamma(t)$ to be optimized is here parameterized in terms of $P = 5$ parameters (with cubic spline interpolation as detailed before), and the final time allowed for the transition is taken to be $T = 1.5 \times T_{QSL}$ (with the quantum speed limit T_{QSL} defined in Section 3.1.1).

Rather than considering the state fidelity as the figure to be maximized, we resort to an alternative figure of merit, defined in terms of local observables that could be experimentally probed (provided single-site resolution of the lattice [190]). Denoting the probability to observe exactly one atom at site i by p_i , we will strive to maximize the average unit filling

$$\mathcal{U} = \frac{1}{L} \sum_{i=1}^L p_i, \quad (3.10)$$

over the L lattice sites. As does the fidelity, this figure yields a maximal value only for deterministic unit filling of each site, that is, when the target MI state has been realized. As such it is equally well qualified as a guide for the optimization problem at hand. In contrast to the state fidelity, however, states with an atom distribution similar to the MI state yield a close to optimal value, and we found this new figure of merit to be advantageous in terms of scalability and statistical accuracy, but, defer its study to Section 4.2.1. Hence, in the following we aim at maximizing Equation (3.10), and we now detail how this figure can be related to experimental measurement data.

Given the (discrete) translation invariance of the system Hamiltonian in Equation (3.2) and of the superfluid initial state in Equation (3.3), any local observable is also invariant under

lattice-site translation. In particular, the probability of unit filling is independent of the site location i , such that $p_i = \mathcal{U}$ in Equation (3.10). Assuming independence of the measurements at each lattice site, it follows that for a single ($N = 1$) measurement, the number $n_1 \in [0, L]$ of sites observed (over the whole lattice) with a single atom follows a binomial distribution with success probability \mathcal{U} and number L of trials. For N repeated measurements, the total number $n \in [0, NL]$ of sites observed with a single atom has thus a binomial likelihood of

$$P(n|\mathcal{U}) = \binom{NL}{n} (\mathcal{U})^n (1 - \mathcal{U})^{NL-n}, \quad (3.11)$$

given the unit-filling probability \mathcal{U} and corresponding to a number NL of trials. This is the likelihood used when implementing the binomial flavour of BO. Note that the assumption of independence in the measurement outcomes at each lattice site is an approximation. In particular, close to the MI state, measurements over each site will become increasingly correlated. This approximation allows us to avoid modeling each probability p_i individually (which can be done for the system size studied but would require extensive computational effort for larger systems), and, as we will see, already yields significantly increased optimization performances.

To assess the merit of binomial BO, we will compare it both to BO with Gaussian modeling (that is, a traditional implementation of BO) and to the SPSA algorithm which was found earlier to be robust, and even to improve, in the presence of noise. The efficiency of any optimizer is judged by its ability to converge towards good solutions, that is, to low filling errors defined as $\mathcal{E} = 1 - \mathcal{U}$, with minimal experimental effort. This effort is quantified as the total number N_r of experimental runs (consisting of individual cycles of preparation, evolution, and measurements) that would have been required. This number takes into account both the number N of measurement repetitions and the number S of optimization steps performed. As usual, any given optimization is repeated 30 times and the statistics, such as the median and the interquartile intervals, assessed over these repetitions are reported. Finally, we will also compare the results of optimizations performed with finite number of measurement repetitions to the results of optimizations without noise, that is, performed given exact values of the figure of merit. In the latter case, all the optimizations converge reliably to the same solution yielding a minimum filling error $\mathcal{E}_{min} \approx 1.6\%$ which will be reported.

Results

Figure 3.6(a) shows the results obtained with BOs with binomial and Gaussian modeling, for a varied number of measurement repetitions N in between 1 to 100 (colours in legend). All the optimizations are limited to a number $S = 2000$ of iterations. As can be seen, binomial BO (filled symbols) exhibits significantly faster convergence than its Gaussian counterpart (empty ones). In particular, with $N = 25$ or fewer repetitions the Gaussian version fails to improve

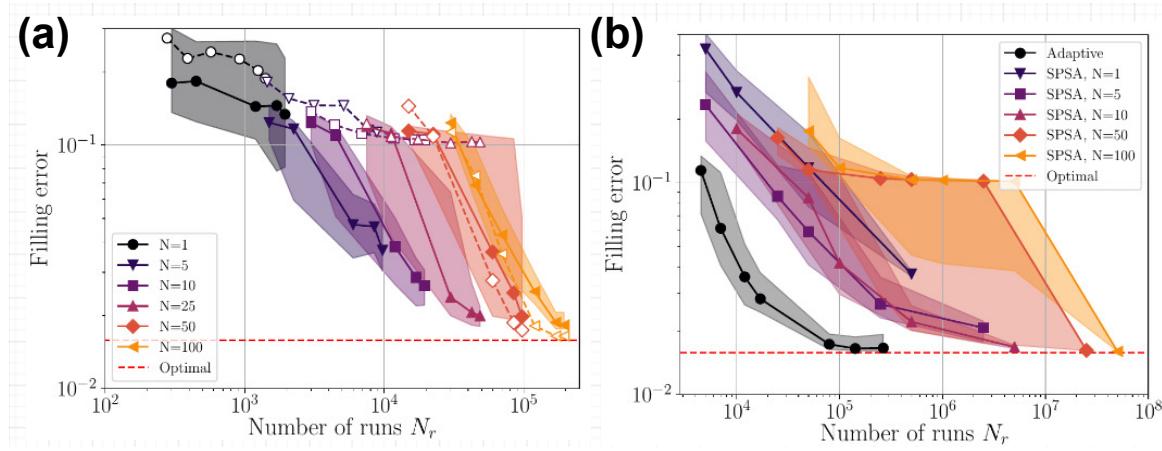


Figure 3.6: From [174]. Optimization with poor statistics (or equivalently, with small number $N \leq 100$ of measurement repetitions per iteration) for the SF-MI control problem. The convergence of each optimization strategy is judged in terms of the filling errors \mathcal{E} that it can achieve given the total number of runs N_r that would have been required. The median results (symbols) and interquartile intervals (shaded eras) are obtained over 30 repetitions of the same problem. The best control solution found for this task results in an error $\mathcal{E}_{min} \approx 1.6\%$ (horizontal dashed red line). (a) Results obtained with BO with binomial (filled symbols) or Gaussian (open ones) noise modeling are reported for different number N of measurement repetitions (colours in the legend). (b) Binomial BO with the adaptive strategy described in the main text (black) is compared to SPSA optimization for different numbers N of measurement repetitions (colours in legend).

beyond the mediocre value of $\mathcal{E} = 10\%$ of the filling error, while the binomial one converges steadily towards the minimum \mathcal{E}_{min} (vertical dashed line). It is only in cases with more frequent measurement repetitions ($N = 50$ and $N = 100$ in Figure 3.6(a)) that the Gaussian version manages to converge towards the low filling error \mathcal{E}_{min} , with convergence rates at par, or slightly better than binomial BO. Overall, this confirms our intuition that a detailed description of the measurement noise is critical when N is taken small, such that despite the approximations needed for the practical implementation of binomial BO, it is found to substantially outperform the original version of BO. Given this improved ability to optimize with poor statistics, we now discuss further the choice of an optimal number N of measurement repetitions.

It can be seen in Figure 3.6(a) that, given a fixed target error $\mathcal{E} > \mathcal{E}_{min}$, resorting to smaller N is in general favorable (except for the extreme case with $N = 1$). It thus seems to be always preferable to explore many parameters with poor statistics rather than trying to estimate too accurately the figure of merit for a few ones. Reducing the number N of repetitions of the same measurement, however, implies an additional computational overhead. As discussed in Section 2.1.3, the numerical complexity in building the models used in BO scales as S^3 , with the number S of observations. For a fixed total number $N_r \propto SN$ of measurements a reduction of N implies an increase in S . That is, whereas fundamentally optimizations with the fewest repetitions seem to perform best, these necessitate an increased computational effort which

may slow down the optimization.

One option to overcome this increased computational requirement would be to scale the computational resources in accordance to the number S of observations expected. Alternatively, one could devise an optimization schedule ensuring that this computational effort never becomes too excessive. When studying optimizations with large total number $N_r > 10^5$ of measurements, we find it practical to start the optimizations with few measurement repetitions (in this case $N \leq 10$), which as seen in Figure 3.6(a) permits the rapid identification of good solutions in parameter space. As the search circles in closer to the optimal solution, and as the observations accumulate, one can restrict the subsequent steps of optimization to a smaller domain. This is achieved by reducing the parameter space around the best parameters encountered, and by dropping the data outside this reduced domain, thus attenuating the numerical effort. As the search approaches a high-fidelity solution, subsequent optimizations can be performed with an increased number of repetitions N in which case we resort to Gaussian modeling, as the latter was found to display similar convergence than the former for large N , but, is numerically faster.

Results for this adaptive schedule are presented in Figure 3.6(b) (black curve). The first 4 points of the curve pertain to the first stage of optimization with a small number of repetitions (here $N = 10$). Binomial modeling is there used for up to $S = 1500$ steps. After this, the parameter space is reduced around the 5% best set of parameters explored, and the number of repetitions is increased. Results of this second round for 1000 extra iterations steps and several choices for the increased number of repetitions $N = 50, 100, 200$ are shown as the remaining three points of the curve. This strategy converges towards the supposed minimum \mathcal{E}_{min} after a total number of $N_r \approx 10^5$ runs. Overall, we find this ad-hoc strategy to be appropriate when exploring the advantage offered by binomial BO in the limit of large number N_r of total repetitions.

Finally, results obtained with SPSA optimization for different numbers of measurement repetitions are depicted in colour in Figure 3.6(b). As can be seen, the convergence with SPSA is substantially slower than with the adaptive schedule of BO. The optimizations with SPSA using $N = 10$ and $N = 100$ repetitions manage to converge close to \mathcal{E}_{min} , but only did so after one, or two, orders of magnitude more runs than with the adaptive method (or even without the adaptive strategy, as can be seen when comparing these results with Figure 3.6(a)).

Overall, we found significantly improved convergences with binomial BO. For a limited total number $N_r < 10^4$ of repetitions, it permits the identification of solutions already close to the true minimum, when a traditional implementation of BO and the SPSA algorithm only improve marginally compared to randomly generated control parameters. This regime is relevant to many complex experimental setups where, as opposed to the numerical simulations employed here, only a limited number of repetitions is possible. For instance, ultracold-atoms experi-

ments often exhibit preparation times of the order of 0.1 second to more than 10s of seconds², thus making optimizations necessitating more than 10^4 repetitions either demanding or unrealistic. Nonetheless, even for increased number of repetitions, binomial BO compares favorably with the other alternatives that were studied, and can be extended further by means of the adaptive strategy presented. Finally, we note that additional results showcasing the robustness of binomial BO with respect to a noisy preparation of the initial SF state are presented in [174].

3.2.3 Further results and outlook

Having demonstrated the applicability of binomial BO on a problem of QOC, we now aim at exploring further its potential when optimizing parameterized quantum circuits. For that purpose, two tasks of quantum state preparation by means of the circuits depicted in Figure 3.7(a,b) are studied. The first example is explored based on thorough numerical simulations, while the second one is performed on public-access IBM quantum chips. In both cases, the figure of merit to be maximized is taken to be the fidelity \mathcal{F} , between realized and target states, which is decomposed over a Pauli basis of observables (defined in Section 2.2.3). Compared to the previous example, this figure now depends on several experimental observables which each requires an individual model of its corresponding measurement probability (discussed in Item 4). This will allow us to verify the ability of the framework to cope with this more general situation.

We first consider the problem of preparing a three-qubit Greenberger–Horne–Zeilinger (GHZ) state $|\Psi_1\rangle = (|000\rangle + |111\rangle)/\sqrt{2}$ given the parameterized circuit depicted in Figure 3.7(a). The goal is to find suitable angles $\theta_j \in [0, 2\pi]$ such that the circuit maps the initial state $|000\rangle$ into the GHZ state $|\Psi\rangle$. In Figure 3.7(b) we compare optimizations relying on BO with binomial (filled symbols) and Gaussian modeling (empty symbols), for varied numbers of repetitions ranging from $N = 1$ to $N = 1000$ (these repetitions are performed per observable and per iteration). Results are reported in terms of the infidelity $\mathcal{I} = 1 - \mathcal{F}$ as a function of the number N_r of repetitions, that is, the total number of circuit executions. In the case of $N = 1$ measurement (depicted in black), the optimizations based on Gaussian modeling mostly fail, whereas binomial modeling already yields infidelities in the percent regime in a couple of thousands circuit executions. More generally, for up to $N = 100$ measurement repetitions we find binomial BO to systematically outperform its Gaussian counterpart, and it is only for larger $N = 1000$ repetitions, that the two flavours become at par.

With this enhanced capability of binomial modeling to converge even with poor statistics of the fidelity, one can appreciate the advantage of choosing small values N of measurement repetitions in order to accelerate the optimization convergence. As can be seen in Figure 3.7(c), in any

²See for instance Fig.3.8 [194] for details of a typical experimental sequence, or [99, 176, 191, 195, 196] for timescales in different setups.

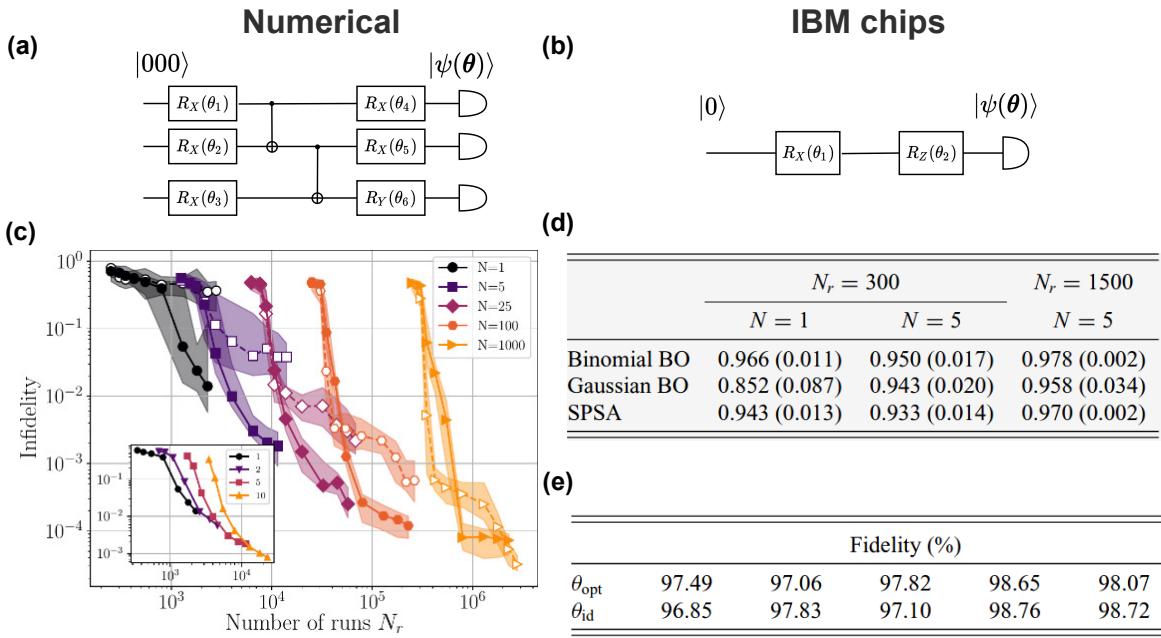


Figure 3.7: From [174]. Two tasks of quantum state preparation with parameterized quantum circuits (a,b). In each case, the aim is to identify the optimal circuit angles θ_i permitting the realization of a predefined target state. (c) The optimization results for the state preparation of a GHZ state with the circuit (a) are plotted in terms of the infidelities achieved as a function of the total number N_r of circuit runs which were required. Results are systematically obtained based on 30 repetitions of the same configuration and the median results (filled symbols) and interquartile intervals (shaded regions) are plotted. The Gaussian (empty symbols and dashed lines) and binomial (filled symbols and plain lines) flavours of BO are compared for different number of measurement repetitions N (colours in legend). (d,e) Results for the preparation of a single-qubit state performed on publicly accessible IBM quantum chips with the circuit (b). (d) Median fidelity results (and their standard deviations in parenthesis) given varied total numbers N_r of circuit executions and numbers N of measurement repetitions. The optimizations are performed with Gaussian and binomial BOs, and also SPSA. (e) Comparison of the fidelities measured for the optimized angles θ_{opt} found by BO (given $N_r = 1500$ and $N = 5$) and for the ideal (that is, in case of perfect realization of the circuit) solution of the problem θ_{id} , over 5 different optimization runs.

case where there is comparable data for a fixed budget of N_r circuit executions (or for a fixed infidelity \mathcal{I}), a choice of the smallest N available results in the lowest infidelities (or the smallest number of circuit executions). Such effect can be verified further in the inset of Figure 3.7(c) where a comparison of binomial BO for a varied number of repetitions N in between 1 and 10 is provided. These results strengthen our intuition, built on the previous two examples, that it is often preferable to resort to poorly resolved measurement data taken over many parameters, rather than highly accurate data over only a few points.

Further results for this task of GHZ state preparation can be found in [174]. In particular we study the convergence of the framework in the presence of additional noise, such as the inclusion of random rotations and also errors in the measurement readouts. It is found that

despite these additional sources of noise, the framework keeps converging, albeit at a slower pace, and remains more competitive than the other alternatives studied. Instead of presenting such results – based on theoretical modeling and simulation of the noise – into more details, we rather focus our attention on a problem of optimization performed with uncharacterized sources of noise.

As a final demonstration, we study a problem of state preparation, performed on public-access IBM quantum chips [1], aiming at realizing a single-qubit target state

$$|\Psi_2\rangle = \cos(\pi/8)|0\rangle + \exp[-i\pi/4]\sin(\pi/8)|1\rangle, \quad (3.12)$$

using the circuit depicted in Figure 3.7(c). Such problem admits a solution $\theta_{id} = (\pi/4, \pi/4)$ in case of flawless implementation of the circuit. However, due to control errors commonly featured in NISQ devices, the true optimal parameters θ_{opt} permitting the accurate realization of the state in Equation (3.12) may differ from such ideal solution. Furthermore, when performed directly on the chip, the identification of these optimal parameters is complicated by stochastic errors in the execution of the circuits and in the measurement readouts (in addition to the statistical noise arising from finite measurements).

Results of optimizations performed with a total number of circuit executions limited to 1500 runs, but with final results verified on a large enough number of 200 000 measurements (ensuring statistical errors of the order of 10^{-3}) are presented in Figure 3.7(d,e). The fidelities estimated with the parameters optimized with binomial BO, Gaussian BO and SPSA, are reported in Figure 3.7(d). These are obtained over varied configurations consisting of a total of either $N_r = 300$ (with either $N = 1$ or 5 measurement repetitions) or 1500 (with $N = 5$ repetitions) circuit executions. The optimizations performed for each of these 3 configuration are repeated in between 5 and 15 times, such that we can report statistics accurate enough to be distinguished. Medians and their standard errors are reported in Figure 3.7(d). Consistently with the findings based on numerical simulations, we find binomial BO to systematically yield the highest fidelities. Furthermore, for the total of $N_r = 300$ circuit executions it is also verified that optimizations with a single $N = 1$ repetition per measurement perform better than with 5 measurement repetitions.

Finally in Figure 3.7, we report the results for the 5 individual runs of optimizations performed with binomial BO given $N_r = 1500$ circuit executions (corresponding to the results reported in the top right entry in Figure 3.7(d)). The fidelities found for the parameters θ_{opt} optimized by BO are compared to the fidelities found for the idealized solution θ_{id} . These two quantities are systematically estimated at the end of any of the BO runs (which were performed over a time window of one week). Over these 5 runs the fidelities for the optimal and idealized parameters are found comparable with deviations in the sub-percent range. Such deviations are smaller than the intrinsic fluctuations of the device, which can be characterized in terms of the standard

deviation of 0.92% in the fidelities which are estimated based on the fixed parameters θ_{id} , but at different times. The overall average obtained for both sets of parameters is almost identical with values of 98.2%, consistent with the readout errors reported on the chips which, at the time of the study, were ranging between 1% and 5%.

Overall, these results of optimization performed on the IBM chips are found consistent with our earlier conclusions based on numerical simulations, and establish further binomial BO as a very frugal alternative to other optimization algorithms. Despite the simplicity of this last example, it is already remarkable that close-to-optimal parameters are identified in a total number of circuits executions ($N_r = 1500$) that is typically used in a single step of optimization of parametrized quantum circuits [56, 144, 192, 197].

3.3 Concluding remarks

The ability to find close-to-optimal solutions solely based on limited experimental data can advance technological developments and precision experiments on a wide range of quantum physics platforms. It offers a very resource-efficient pathway towards the optimal use of existing quantum hardware comprised, for instance, of tens to hundreds of superconducting qubits, or hundreds to thousands of atoms. Given the availability of cheap resources for classical computation, it is essential to leverage them as much as possible to support the limited capabilities of near-term quantum hardware. For that purpose, we have explored the use of (and have improved on) BO.

For the SF–MI transition task presented in Section 3.1, the performance of BO was found superior to the other optimization routines which were studied. This was attributed to the ability of BO to select the control parameters to be probed in a principled way balancing local improvement and global exploration of potentially rewarding domains of the parameter space. In situations where the presence of local minima (as observed for the task of the SF–MI transition close to the quantum speed limit) impedes the success of optimizations, BO reached better solutions than the local optimizers studied, while remaining more efficient than the global optimizer that was considered. Given that this situation is commonly found in many problems of QOC and VQAs [31, 112, 113, 117], this encourages the adoption of BO as an appealing alternative to traditional optimization strategies. In Section 3.2, we have studied the effect of varying the number of measurement repetitions performed at each step of the optimization. While such number has traditionally been kept large to permit accurate estimation of the figure of merit, here we have rather considered the opposite regime characterized by poor statistical estimates resulting from a small number of repetitions. By incorporating precise details of the measurement statistics, as obtained in typical quantum experiments, we endowed BO with the ability to operate even in the presence of very noisy data. Given this new ability, it was found

that choosing a small number of measurement repetitions was in general favorable for enhanced convergence, in stark contrast with current practices.

Having showcased the benefits of BO, it is also in order to stress out its limitations. In the several examples studied, the number of parameters was taken to be around $P = 10$ parameters. Such number could be enough to parameterize a couple of experimental control fields, but may fall short in situations where many experimental degrees of freedom could be optimized over, or in large scale implementations of parameterized quantum circuits. In practice, directly optimizing over high-dimension parameter spaces, with $P \gtrsim 30$, is challenging. Still, several methods extending BO to this regime have been proposed and could be incorporated. These often rely on the optimization over a reduced parameter subspace, which can either be learnt [198] or randomly selected [199]. Going even further, this could form the basis of an iterative strategy. For instance, in a problem of QOC one would optimize over perturbations of the control field, parameterized by a small number $P' \ll P$ of parameters, which could be iteratively resampled [146]. The success of such iterative procedure could be made experimentally feasible when used in conjunction with the low data requirement of BO.

It should be highlighted that the conclusion that was drawn regarding the use of a small number of measurement repetitions was based on the assumption that changes in the control parameters do not incur a large experimental overhead. When this overhead becomes significant, compared to the time required to repeat the same experiment, it should be taken into consideration. In this case, convergence of the algorithms should be studied in terms of the experimental time required, rather than the number of experimental repetitions as was done here. Given specific experimental timescales, the results presented in Section 3.2 could directly be recast in such form.

Going further, we remark that the probabilistic modeling approach adopted in BO offers room for even more refinement. Akin to the methodology developed for binomial BO one could also include other sources of noise, in addition to the statistical noise resulting from finite sampling. For instance, values of readout errors, which are often well characterized, could be directly incorporated into the models by means of altering the likelihoods involved. More generally, the field of probabilistic machine learning is advancing at fast pace and has a lot to offer to the characterization and optimization of quantum dynamics. Faster implementations of the models used have been developed for exact [200] or approximative [201] inferences with Gaussian processes. These can accelerate the algorithms studied here, especially when the number of iterations becomes substantial. Alternative probabilistic models such as ensembles of neural networks [202] could also provide a substitute for Gaussian processes in case of high-dimension parameter space and large number of iterations.

Finally, it should be acknowledged that the system sizes optimized over in this chapter were taken relatively small, such that a thorough study of BO was possible. Ultimately, we argue

for the adoption of the methods developed in larger experimental situations. As such, it is of interest to identify potential roadblocks in the scaling of the techniques presented and, more generally, in the scaling of any experiment-based optimization, to large system sizes. For instance, the binomial flavour of BO, which takes into consideration the statistics of projective quantum measurements, requires to model each of the measurement probabilities involved individually. This can quickly become overwhelming for figures of merit, such as the fidelity, which often need to be decomposed in terms of a large number of experimental observables. More generally, difficulties of accurate estimation of the fidelity, and concentration of its values to a narrow range, are known to complicate optimization tasks defined over large Hilbert spaces. Such detrimental effects are discussed in the next chapter, where we argue that these could be overcome, or at least mitigated, by means of engineered figures of merit.

Chapter 4

Engineering the figure of merit

Any optimization, be it a task of QOC or a VQA, relies on iterative evaluations of a figure of merit which guides an optimizer toward the solution of a problem. Importantly, for a given problem the choice of this figure is not unique, thus leaving freedom in its design. Choosing the right figure could result in accelerated convergence or even could enable an optimization which otherwise would have failed. For instance, recent works [203, 204] have resorted to ad-hoc figures of merit intended to overcome (or at least to limit) issues related to Barren plateaus [109] in the training of VQAs. These Barren plateaus, which arise due to the concentration of the figure of merit to almost identical values over most of the parameter space, are not limited to problems of VQAs but are rather relevant to any problem of optimization of quantum dynamics. In addition to such issues, it is also important to consider the experimental effort which is required to estimate a figure of merit up to desired accuracy during the course of an optimization. As we will see, both of these aspects need to be carefully taken into consideration when performing optimization over large Hilbert spaces.

In this chapter, we discuss the engineering of more experiment-friendly figures of merit. We start in Section 4.1 by listing the desirable properties of these figures, and by establishing quantitative criteria allowing their study. These criteria are taken general enough to be relevant to any optimization performed based on experimental outcomes. Their use is illustrated in Section 4.2.1, where an alternative to the state fidelity is investigated for problems of preparation of separable quantum states. This alternative was introduced in [204] but was only discussed in the context of Barren plateaus in qubit systems. The discussion is extended here to also encompass a study of its estimation properties and its application to boson systems. This is followed in Section 4.2.2 by the engineering of improved figures of merit for the preparation of stabilizer states. Finally, in Section 4.2.3 we study the 0-fidelity proposed in [170] as an alternative to the process fidelity, and which can be used for task of quantum gate engineering. In all cases, the alternative figures which are investigated in this chapter are shown to outperform their alter egos in light of the criteria developed.

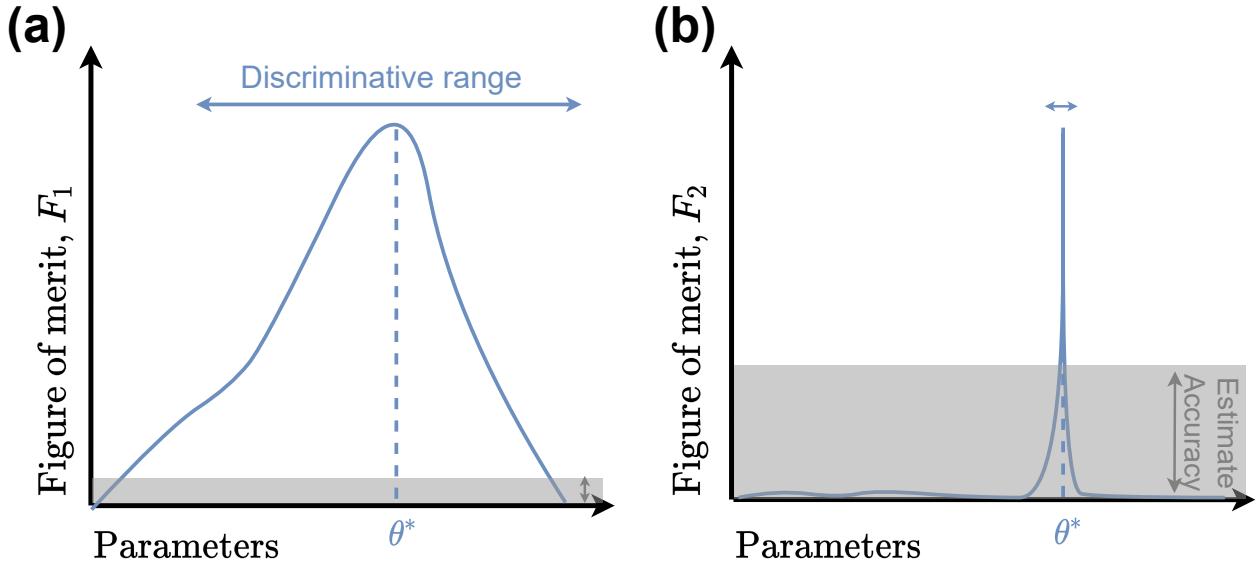


Figure 4.1: Sketch of a good (a) and a bad (b) figure of merit. These two figures are maximized for the same parameter θ^* , and thus are faithful to the same optimization problem admitting solution θ^* . Still, they present different characteristics making them more or less suitable in the context of optimization. Two quantitative criteria are employed when judging a figure of merit. First, it should be experimentally accurate, that is it should be possible to accurately estimate its values based on a realistic number of experimental repetitions. This statistical accuracy is quantified by the variance $\Delta^2[\tilde{F}]$ in the estimate \tilde{F} of the figure F , and is represented by means of a vertical gray arrow and shaded region. Second, a figure of merit should be able to discriminate between parameters, that is, it should adopt different values for varied parameters θ . This ability to discriminate is quantified by the variance $D_\theta(F)$, defined in Equation (4.2), of the figure F over the parameter space and is depicted by means of a horizontal blue arrow. The good figure F_1 (a) adopts different values for different set of parameters, and when estimated over a realistic number of measurement repetitions results in small statistical variance. As such a general optimization algorithm should be able to distinguish between good and bad parameters and is likely to succeed. In contrast, the bad figure F_2 (b) adopts almost everywhere similar values (in other words, values of F_2 concentrate to a narrow interval), which cannot be distinguished due to the large statistical variance. Hence, any optimization starting from parameters not close enough to the solution θ^* is most likely to fail.

4.1 Desirable properties

Three properties are established as desirable when developing a figure of merit F : it should be faithful to the problem at hand, experimentally accurate and discriminative (or equivalently it should not concentrate to experimentally indistinguishable values). These properties are illustrated in Figure 4.1 and are now discussed in greater details.

Faithfulness

A prerequisite for any figure of merit to be adequate for a given problem is faithfulness [203], that is, it should reach its maximum value if and only if the desired solution has been attained. This restriction, however, does not single out a unique figure. For instance consider the expectation values of two Hermitian operators M_1 and M_2 represented in the same basis, and with eigenvalues $\lambda_1 = 1 > \lambda_2 \geq \dots \geq \lambda_d \geq 0$ arranged in decreasing order, as

$$M_1 = \begin{bmatrix} \lambda_1 = 1 & 0 & \dots & 0 \\ 0 & 0 & \dots & 0 \\ \vdots & & \ddots & \vdots \\ 0 & \dots & & 0 \end{bmatrix}, \text{ and } M_2 = \begin{bmatrix} \lambda_1 = 1 & 0 & \dots & 0 \\ 0 & \lambda_2 < 1 & \vdots & 0 \\ \vdots & & \ddots & \vdots \\ 0 & \dots & & 0 \leq \lambda_d \end{bmatrix}. \quad (4.1)$$

Despite their differences, they both reach their maximal expectation value, $\max \langle M_{i=1,2} \rangle = 1$, for the same state. As such, the figures of merit corresponding to the expectation values of M_1 and M_2 are deemed adequate for the same problem of preparing the eigenstate $|\lambda_1\rangle$ (corresponding to the eigenvalue λ_1). This freedom in designing many operators of the form M_2 for the same task, will be exploited when engineering alternative figures of merit. Before that, we need to establish further criteria allowing to compare figures of merit which are faithful to the same problem.

Experimental accuracy

To be of practical interest experimentally, one should be able to *accurately estimate* a figure of merit based on a realistic number of measurement repetitions. As discussed in Section 2.2, accuracy is understood in terms of the statistical variance of the estimation errors. While some (bounded) amount of noise can be tolerated by the optimization routine (in particular with BO and up to the case of single shot binomial noise as demonstrated in Section 3.2), it remains the case that a reduced level of noise is in general beneficial from an optimization perspective. For unbiased estimates, as considered in this thesis, recall that the variance of the estimation errors is equal to the variance of the estimate itself.

With estimation accuracy in mind, and denoting $\tilde{F}_1(\tilde{F}_2)$ the (unbiased) estimates of $F_1(F_2)$, the figure F_1 is deemed better than F_2 if, for similar number of measurement repetitions, it entails a smaller estimation variance, that is, $\Delta^2[\tilde{F}_1] < \Delta^2[\tilde{F}_2]$. Equivalently, this criterion implies that in order to reach the same accuracy $\Delta^2[\tilde{F}_1] = \Delta^2[\tilde{F}_2]$, a smaller experimental effort would be necessary for the figure F_1 . Such variances will be evaluated based on the expressions obtained in Section 2.2.2.

Discriminative

Finally, as these figures are used in the context of optimization we also need to consider their ability to distinguish between different realizations of a quantum dynamics; these different realizations result from the choice of different values of the control parameters $\boldsymbol{\theta}$ of the system under study (that is, different control fields in QOC or different angles of parameterized gates in VQAs). In particular, it seems sensible to require that a figure of merit does not adopt similar values for the vast majority of the parameter space. This ability to discriminate is quantified in terms of the variance $D_{\boldsymbol{\theta}}(F)$ in the values of the figure F over a distribution of parameters $p(\boldsymbol{\theta})$, and is defined as

$$D_{\boldsymbol{\theta}}(F) = \int p(\boldsymbol{\theta}) (F(\boldsymbol{\theta}) - \langle F \rangle_{\boldsymbol{\theta}})^2, \quad (4.2)$$

with $\langle F \rangle_{\boldsymbol{\theta}} = \int p(\boldsymbol{\theta}) F(\boldsymbol{\theta})$ the average of F (over the same distribution). A typical distribution $p(\boldsymbol{\theta})$ would be a uniform distribution over the domain of $\boldsymbol{\theta}$.

Notably, the variance in Equation (4.2) is evaluated for a distribution of parameter values rather than a distribution of probabilistic measurement outcomes, as is the case for the estimation accuracy discussed earlier. Still, this variance $D_{\boldsymbol{\theta}}(F)$ should be understood in light of the estimation accuracy $\Delta^2[\tilde{F}]$: to distinguish between typical values of F one would need to estimate them with an accuracy $\Delta^2[\tilde{F}]$ of the order of $D_{\boldsymbol{\theta}}(F)$, that is, with a number of measurement repetitions scaling as $D_{\boldsymbol{\theta}}(F)^{-1}$. In particular, a vanishing variance $D_{\boldsymbol{\theta}}(F)$ indicates that the values of F will *concentrate* into a very narrow range¹ indistinguishable experimentally, thus impeding the success of any optimization starting from random, or at least not close to the optimal parameters $\boldsymbol{\theta}^*$. With this ability to discriminate in mind, a figure of merit F_1 is deemed better than F_2 if its values are more wildly distributed, that is when $D_{\boldsymbol{\theta}}(F_1) > D_{\boldsymbol{\theta}}(F_2)$, and we now discuss how such variances can be assessed.

It is only in certain situations that the variance in Equation (4.2) can be exactly evaluated. For instance, assuming that random values of the parameters $\boldsymbol{\theta}$ result in a distribution of random

¹Recall that Chebyshev inequality imposes that the probability of a value F to differ from its mean value $\langle F \rangle_{\boldsymbol{\theta}}$ by more than ε is bounded by $Pr(|F - \langle F \rangle_{\boldsymbol{\theta}}| \geq \varepsilon) \leq D_{\boldsymbol{\theta}}(F)/\varepsilon^2$. Hence for a vanishing value of $D_{\boldsymbol{\theta}}$, the deviation ε must also be vanishing for the probability to be significant enough.

states² facilitate the evaluation of Equation (4.2). Given a figure of merit F corresponding to the expectation value of an operator M , the variance $D_{rdm}(F)$ over such random states distribution can be expressed as (derived in Equation (C.10) of the Appendix C)

$$D_{rdm}(F) = \frac{Tr[M^2]}{d(d+1)} - \frac{Tr^2[M]}{d^2(d+1)} \leq \frac{Tr[M^2]}{d(d+1)}, \quad (4.3)$$

which depends only on the eigenvalues of the operator M (entering Equation (4.3) through the traces involved), and on the dimension d of the system (which grows exponentially with the size n of a system).

Already, Equation (4.3) permits to appreciate the contribution of the eigenspectrum of M in the variance $D_{rdm}(F)$, and also to unveil some of the challenges in optimizing figures of merit defined over large Hilbert spaces. For an operator with a single non-null eigenvalue – for instance, a projector $|\psi_{tgt}\rangle\langle\psi_{tgt}|$ onto a target state $|\psi_{tgt}\rangle$, as involved in the fidelity measure – $Tr[M^2]$ has constant value, and the variance $D_{rdm}(F_1)$ scales at best as d^{-2} . This exponential decay of the variance, with respect to the system size n , highlights prominent issues when using the fidelity as a figure of merit [205] for large value of n . A choice of an operator of the form M_2 in Equation (4.1) with a number d of non vanishing eigenvalues would result instead in an improved scaling of (at best) d^{-1} , as $Tr[M^2]$ now scales as d . This is deemed favorable as it delays phenomena of concentration to larger values of n . Still, in both situations it can be seen that the variance in the values of the figure of merit will inevitably vanish for large system sizes, thus preventing the success of any optimization starting with non-informed initial parameters.

The scalings just discussed, however, rely on the assumption of states randomly distributed, which may not be justified. In the following, we will rather consider the true distributions realized by the specific controlled quantum systems that are investigated. In such cases, $D_{\theta}(F)$ does not admit an analytical expression, but, can be approximated by means of Monte Carlo sampling, as the empirical variance

$$D_{emp}(F) = \frac{1}{S} \sum_{s=1}^S (F(\boldsymbol{\theta}_s) - \bar{F})^2, \quad (4.4)$$

of the values of the figures of merit F evaluated for a finite number S of parameters $\boldsymbol{\theta}^s$ sampled from $p(\boldsymbol{\theta})$, and where \bar{F} denotes the empirical mean $\bar{F} = \sum F(\boldsymbol{\theta}^s)/S$.

This ability to discriminate, as quantified by Equation (4.2), is closely related to the phenomenon of Barren plateaus [109], which has been exposed in the context of VQAs. Such Barren plateaus are characterized by vanishing amplitudes of the gradients of the figure of merit. In such situation, a choice of random initial parameters (of the parameterized quantum

²These are pure states which are uniformly distributed over the $d-1$ hypersphere, or more technically, which have a Haar measure. This technical aspect is discussed in greater length in Appendix C.

circuits involved) will almost inevitably corresponds to gradients close to zero, such that any gradient-descent method would fail. Here, rather than considering the variance in the amplitudes of the gradients, we consider instead the variance in the values of the figure of merit as the most relevant quantity to be studied. This is motivated by the fact that our optimizer of choice, BO, and many other optimizers used in QOC do not rely on gradients but only on evaluations of F . Nonetheless, under reasonable assumptions [206], a vanishing variance in the values of the figure of merit and the Barren plateaus phenomenon are equivalent.

Overall we propose three criteria to evaluate the adequateness of a figure of merit. The advantage of choosing these criteria lies in the fact that they can be easily assessed and are general enough to be relevant to any optimization problem, or to any optimization strategy considered. This makes them appealing when comparing and developing new figures of merit as we now proceed to show.

4.2 Engineering and assessment of alternative figures of merit

4.2.1 Preparation of separable states

So far, the discussion about the desirable properties of a figure of merit has remained general. We now apply the quantitative criteria freshly established to compare two alternative figures adequate for the preparation of separable states. For the sake of concreteness we first restrict the discussion to the preparation of the n -qubit target state $|\mathbf{0}\rangle = |0\rangle^{\otimes n}$. Despite the apparent simplicity of the state to be realized, we conclude this section by discussing several practical applications.

State fidelity and local fidelity

As considered in Section 3.1.1 and Section 3.2.3, a canonical choice of figure of merit for problems of quantum state preparation is the state fidelity. For a prepared state ρ and the target $|\mathbf{0}\rangle$, it is defined as

$$\mathcal{F}(\rho) = \text{Tr}[|\mathbf{0}\rangle\langle\mathbf{0}|\rho]. \quad (4.5)$$

Given the separable physical nature of the target state, an alternative to the fidelity [204] can be engineered based on the local probabilities $\mathcal{F}_i(\rho) = \text{Tr}[|0\rangle\langle 0|_i\rho]$ of each part i of the state ρ to match its corresponding target $|0\rangle_i$. More explicitly, the local fidelity \mathcal{F}_{loc} is defined as the

expectation value of the operator

$$F_{loc} = \frac{1}{n} \sum_{i=1}^n |0\rangle\langle 0|_i \otimes \tilde{I}^i, \quad (4.6)$$

with \tilde{I}^i being the identity acting on all qubits but qubit i . By construction, F_{loc} is diagonal in the computational basis, with eigenvalues $\lambda_{\mathbf{b}}$ equal to the number of $|0\rangle$ qubits in the basis state $|\mathbf{b}\rangle$ (divided by n), such that $\mathcal{F}_{loc} \in [0, 1]$. Furthermore \mathcal{F}_{loc} reaches its maximal value only when the target state $|\mathbf{0}\rangle$ has been prepared, that is, is faithful to the state preparation problem and can thus be compared further to the fidelity.

Variance in estimates of the state and local fidelities

Given that the fidelity $\mathcal{F}(\rho)$ is effectively the probability to observe a bitstring outcome $\mathbf{0} \equiv (0 \dots 0)$ (when measuring ρ in the computational basis), an estimate $\tilde{\mathcal{F}}(\rho)$ obtained over N measurement repetitions is the frequency of such $\mathbf{0}$ outcome. This estimate has a variance (dropping the dependency on ρ)

$$\Delta^2[\tilde{\mathcal{F}}] = \frac{\mathcal{F}(1 - \mathcal{F})}{N}. \quad (4.7)$$

Akin to the fidelity, estimating \mathcal{F}_{loc} only requires measurements in the computational basis. In this case, each of the probability \mathcal{F}_i with $(i = 1, \dots, n)$ is estimated based on the measurement outcomes of the corresponding qubit i , and averaged over the n qubits. It is shown in Appendix B.2 that the variance in the estimates of \mathcal{F}_{loc} , based on N measurements, can be bounded by

$$\Delta^2[\mathcal{F}_{loc}] \leq \frac{\mathcal{F}_{loc}(1 - \mathcal{F}_{loc})}{N}, \quad (4.8)$$

with the higher bounds saturated only in cases when the measurement outcomes over each of the n qubits are perfectly correlated. Hence, in general, the variance in Equation (4.8) will exhibit a better scaling than the variance of the state fidelity estimates shown in Equation (4.7). In particular, for uncorrelated measurement outcomes for each of the n qubits – arising for separable states $\rho = \rho^{(1)} \otimes \dots \otimes \rho^{(n)}$ – this variance will scale as $1/(Nn)$, which decreases with the system size n due to the fact that each of the n qubit outcomes, per measurement, contributes independently to the estimate. We now verify numerically these scalings, and also probe more general cases than the independent and perfectly correlated measurement outcomes just discussed.

In Figure 4.2, we report the variances obtained when estimating the two figures given a fixed number of $N = 10\,000$ measurement repetitions, and for varied system sizes in between $n = 2$ to 9 qubits. In any case, the errors between the true value of the figure of merit F and its estimate \tilde{F} are rescaled by a factor $F(1 - F)$, such that the scalings in Equation (4.7) and

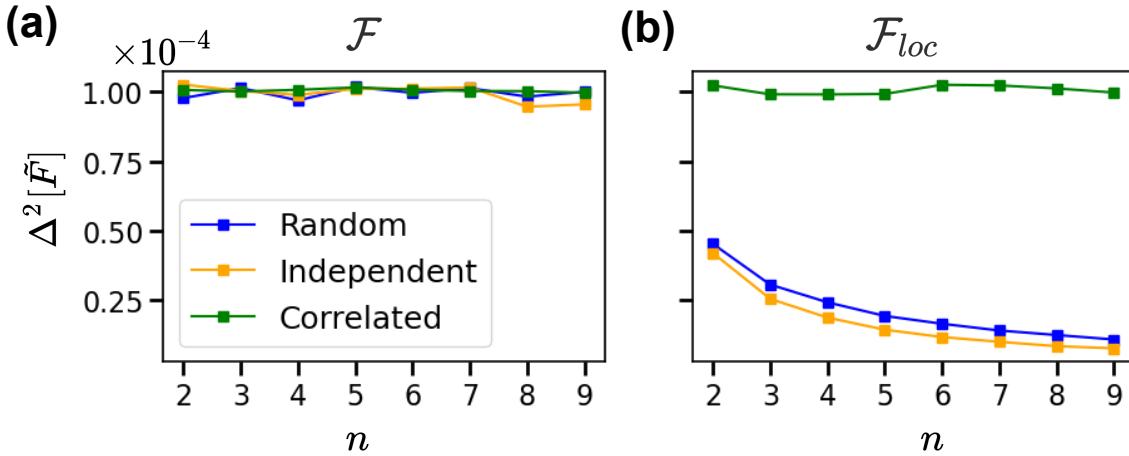


Figure 4.2: Comparison of the variance $\Delta^2[\tilde{F}]$ of the rescaled (see main text) estimation errors for (a) the state fidelity \mathcal{F} defined in Equation (4.5), and for (b) the local fidelity \mathcal{F}_{loc} defined in Equation (4.6). In all cases, the estimates are obtained given $N = 10\,000$ measurement repetitions, and the variances are obtained over a total of $S = 10\,000$ states sampled from 3 different ensembles. In addition to a distribution of random states (blue), we also consider two ensembles of states which result in either independent (orange) or fully correlated (green) qubit measurement outcomes. These ensembles are defined more precisely in the main text. Variances are reported for system sizes ranging from $n = 2$ to 9 qubits.

Equation (4.8) can be decoupled from the underlying values of the figures estimated. For each system size n studied, the variances in these (rescaled) errors are obtained and averaged over 10 000 states which are randomly drawn from 3 different ensembles. We consider the case of random states (in blue), separable states (in orange) in which case each of the n qubit states is randomly generated, and states of the form $|\psi_{corr}\rangle \propto (|0\rangle^{\otimes n} + \alpha|1\rangle^{\otimes n})$ which entail fully correlated measurement outcomes in the computational basis (in green).

In accordance to Equation (4.7)(a), it is verified in Figure 4.2 that the variance of the estimation errors for the fidelity neither depends on the system size n nor on the nature of the states measured, but remains constant with value of $\Delta^2[\tilde{F}] \propto 1/N = 10^{-4}$. However, as can be seen in Figure 4.2(b), the situation differs when considering errors in the local fidelity measure. As expected, for the specific case of fully correlated measurements (green), variances of the errors are found with the same constant value of $1/N = 10^{-4}$ as for the fidelity. This corresponds to the upper bound of Equation (4.8). For separable states the expected scaling of n^{-1} is also verified. Interestingly, for states randomly generated, we numerically observe a similar scaling, inversely proportional to the system size, albeit with a larger constant factor. Already, this increased experimental accuracy points toward a practical advantage when performing optimization based on the local fidelity \mathcal{F}_{loc} (especially for increased values of n). Even more importantly, as we now see, the local fidelity can also avoid issues of concentrations that were mentioned earlier.

Concentrations of the values of the state and local fidelities

We saw that, when evaluated over an ensemble of random states, any figure of merit taken to be the expectation value of an operator (with bounded eigenvalues) will exhibit a variance $D_{rdm}(F)$ scaling at best as d^{-1} , that is, exponentially decreasing with the system size. We now show that, for more structured distributions of states, such exponentially vanishing scaling can be avoided, provided an adequate choice of the figure of merit. For that purpose, we return to the task of preparation of the Mott-insulator state which was studied in Section 3.1.1. In Section 3.2.2 we argued for the merits of an alternative to the state fidelity, especially regarding its scalability to large system sizes, and deferred its study. We now develop further this aspect.

Recall that for the QOC task of preparing the Mott-insulator (MI) state, which is a (separable) Fock state, an alternative to the state fidelity was defined in Equation (3.10) as the expectation value of the operator

$$U = \frac{1}{N} \sum_{i=1}^n |1\rangle\langle 1|_i \otimes \tilde{I}^i, \quad (4.9)$$

with the local projector $|1\rangle\langle 1|_i$ corresponding to the occupation of lattice site i with exactly 1 boson, and the identity \tilde{I}^i acting on all sites but i . This operator corresponds to the local fidelity which was defined in Equation (4.6) for a $|\mathbf{0}\rangle$ target state, but here adapted to the target MI state.

The controlled system considered, which is sketched in Figure 4.3(a) and was detailed in Section 3.2.2, consists of an initial superfluid (SF) state driven by means of a time-dependent control function $\Gamma(t)$. A random realization of such system is obtained by evolving the SF state under a random control function. These random control functions are obtained by uniformly sampling the control parameters $\boldsymbol{\theta}_s$ and reconstructing the time-dependent functions associated. This is repeated $S = 250$ times, and at the end of each control sequence, we record the exact values of the fidelity $\mathcal{F}(\boldsymbol{\theta}_s)$ and the local fidelity $\mathcal{F}_{loc}(\boldsymbol{\theta}_s)$ (called average probability in Section 3.2.2).

In Figure 4.3(b), we report the statistics obtained over these 250 repetitions for different system sizes ranging from $L = 4$ to 10 lattice sizes, with unit filling $B = L$. The empirical average $\langle F \rangle_{\boldsymbol{\theta}}$ and the variances $D_{\boldsymbol{\theta}}(F)$ for both the figures are displayed in the top and bottom panels respectively. For the fidelity \mathcal{F} (blue), it can be seen that both its average value and its variance decay exponentially with the size of the system. That is, even under this new distribution of states, issues of concentration remain present. In sharp contrast, the average values and variances of the alternative figure \mathcal{F}_{loc} (orange) are found to remain constant with increased system size. Provided that this scaling holds for larger systems, this suggests that for large values of L and B any optimization based on the fidelity would systematically fail (or would have required extremely good initial guess in the control parameters), but that optimization

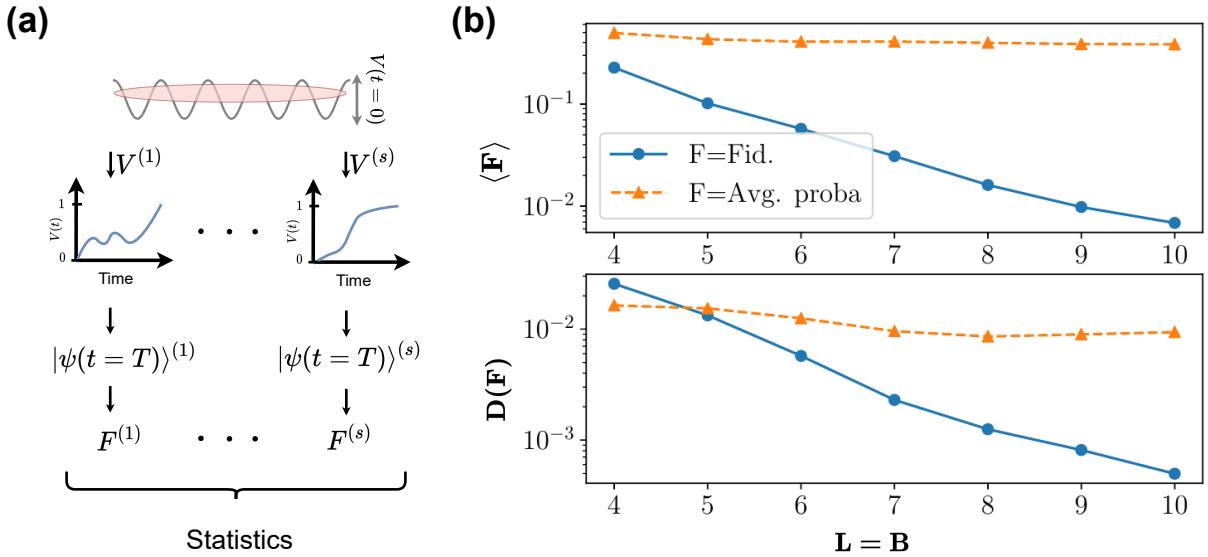


Figure 4.3: Comparison of two figures of merit employed for the SF–MI control problem: the state fidelity and the average probability of unit filling (which effectively corresponds to the local fidelity adapted to the target MI state). (a) The statistics are obtained empirically by drawing $S = 250$ set of random control parameters θ_s from which can be reconstruct the full time–dependent control functions. The initial superfluid state is evolved according to these controls, and for each of the S realizations we evaluate the values of the resulting figures of merit. (b) Plot of the empirical average and variances of the fidelity (blue) and average probability (orange).

with the alternative figure is, in principle, possible.

Finally, we highlight that the lack of concentration of the local fidelity that is reported here for a boson system which is (i) driven over a short period of time (when compared to adiabatic timescales) and with (ii) short–range interactions and hopping, is similar to the findings in [204] obtained for a qubit state evolved by means of a (i) shallow quantum circuit and with (ii) short–range interactions (m –qubit gates acting on adjacent qubits). Given these similarities, it may be expected that the analytical results for quantum circuits presented in [204] could be extended to more general systems, and dynamics.

Concluding remarks

Overall, the local fidelity studied in this section was found to exhibit improved characteristics, both in terms of estimation accuracy and variances of its values, compared to the state fidelity. Resorting to this alternative could thus accelerate, or even enable, tasks of quantum state preparation of product states in many-body quantum experiments, such as the MI state considered earlier and also others [179]. Additionally, even improving on the preparation of the seemingly trivial $|0\rangle^{\otimes n}$ qubit state could find practical applications. For instance, it could be employed in problems of optimal control for the fast initialization and reset of qubit registers.

Also, we remark that the fidelity with respect to the $|0\rangle^{\otimes n}$ state, is a quantity of interest when characterizing the average error over a set of gates by means of randomized benchmarking [207–209]. Given that such measure has been used for problems of quantum gate engineering [45, 56], the proposed alternative could be substituted to the fidelity and may prove advantageous, especially when optimizing over larger system sizes than the 1-qubit quantum gates which were optimized in [45, 56].

4.2.2 Preparation of stabilizer states

As highlighted in the previous example, despite its popularity, the state fidelity may not be the most adequate figure of merit for quantum state preparation purposes. However the alternative figure of merit proposed in Equation (4.6) can only be constructed for separable states, limiting its applicability. We now consider the more challenging task of preparing stabilizer states [15]. This family of quantum states holds a particular appeal in the quantum computing and information toolbox as being entangled but with a well defined structure, compared to random states, facilitating their study and practical utilization. As such, they have often been used to benchmark quantum computing capabilities [210–213], and are central resources in quantum metrology [214] and measurement-based quantum computing [215]. Being able to prepare such states faithfully is thus of great importance and has been approached by means of QOC [81, 84].

Following the methodology employed to develop entanglement witnesses in [216], we devise a suite of alternative figures of merit, that are called witness fidelities, for the preparation of generic stabilizer states, and show their advantage compared to the state fidelity measure. We then discuss how such advantage can even be further increased when considering realization of specific stabilizer states, such as the popular GHZ state.

Witness-inspired figure of merit

Elements of the stabilizer formalism [15, 217] are presented in Appendix D. Here, we only review the aspects necessary for the construction of the new figure of merit. We denote the group generated by n -qubit operators G_i ($i = 1, \dots, l$) – that we restrict to be n -fold tensor products of Pauli matrices and identity, with prefactor ± 1 – as $\langle G_1, \dots, G_l \rangle$. Given a group S generated by n independent and commuting generators³ G_i , it can be shown that there exists a unique (up to global phase) state $|\psi_S\rangle$ which remains invariant under the action of any of the generators G_i , and thus under the action of any of the operators of the group

³The exact conditions that these generators need to satisfy are listed in more details in Appendix D.1

$S = \langle G_1, \dots, G_n \rangle$. That is,

$$|\psi_S\rangle \Leftrightarrow \forall O \in S, O|\psi_S\rangle = |\psi_S\rangle. \quad (4.10)$$

This state $|\psi_S\rangle$ is said to be stabilized by the group S , and we now detail how an alternative to the state fidelity, with respect to $|\psi_S\rangle$, can be constructed.

Given that each of the n generators G_i has eigenvalues $\lambda^{G_i} \in \{-1, 1\}$, it follows that $|\psi_S\rangle$ is the only state maximizing simultaneously all their expectation values. Hence, $|\psi_S\rangle$ is also the only state maximizing the expectation value of the operator

$$\mathcal{G} = \sum_{i=1}^n G_i, \quad (4.11)$$

resulting from the sum of each of the n generators. As such, similarly to the projector $F = |\psi_S\rangle\langle\psi_S|$, \mathcal{G} permits the characterization of $|\psi_S\rangle$, and could thus be substituted to F . Additionally, the operator \mathcal{G} presents the advantage of being decomposed into n separable operators, rather than the 2^n required for the projector F .

To be comparable to the fidelity, the alternative figure of merit $\mathcal{F}_W(\rho)$ is defined as the expectation value $\langle F_W \rangle_\rho$ of the operator

$$F_W = \frac{1}{2} + \frac{\mathcal{G}}{2n}, \quad (4.12)$$

with respect to the prepared state ρ . Akin to \mathcal{G} , the expectation value of F_W is maximized only for $|\psi_S\rangle$ (that is, \mathcal{F}_W is faithful to the task of preparing the state $|\psi_S\rangle$), but yields values of \mathcal{F}_W bounded in the range $[0, 1]$, such that it can be compared with \mathcal{F} .

Given the general formula of the variance of an estimate (obtained by importance sampling) established in Equation (2.30) and the decomposition of F_W in Equation (4.12), the variance of an estimate $\tilde{\mathcal{F}}_W$ (based on N measurement repetitions) of \mathcal{F}_W is readily obtained as

$$\Delta^2[\tilde{\mathcal{F}}_W] = \frac{n \sum_{i=1}^n 1}{(2n)^2 N} = \frac{1 - \mathcal{F}_W}{4N}. \quad (4.13)$$

This is to be contrasted to the variance of an estimate of the fidelity, which was found in Section 2.2.3 to be

$$\Delta^2[\tilde{\mathcal{F}}] = \frac{1 - \mathcal{F}}{N}, \quad (4.14)$$

for any stabilizer states.

Hence, at fixed estimation accuracy (and assuming similar values of the two figures), the estimation of \mathcal{F}_W will require 4 times less measurements than for the fidelity. Furthermore, it can be shown (Appendix D.2) that F_W decomposes over $\sim d$ eigenvectors with non-vanishing eigenvalues, resulting in a variance $D_{rdm}(F_W)$, in the values of \mathcal{F}_W over a distribution of random states,

scaling as $(nd)^{-1}$ (with the exact expression derived in Equation (C.16) of the Appendix C.3). Again, this compares favorably with the fidelity, but still signals issues of concentration. As we now see, such advantages, both in terms of estimation accuracy and concentration, can be improved further when considering more practical problems.

Preparation of a GHZ state with shallow quantum circuits

Here, we restrict our focus to the preparation of a n -qubit GHZ state of the form $|\psi_{tgt}\rangle = 1/\sqrt{n}(|0\rangle^{\otimes n} + |1\rangle^{\otimes n})$. Such state is stabilized by the n generators [216]

$$\begin{cases} G_1 = \bigotimes_{j=1}^n X_j \\ G_{2 \leq j \leq n} = Z_{j-1} \otimes Z_j. \end{cases} \quad (4.15)$$

Crucially, as any pair of operators $G_{2 \leq j \leq n}$ commute, their expectation values can be estimated given the same measurement setting, entailing a measurement of all the qubits in the computational basis. Hence, estimating the alternative figure in Equation (4.11) only requires to consider two measurement settings for the case of GHZ states. To highlight this decomposition, \mathcal{G}_{GHZ} from Equation (4.11) is rewritten as the sum

$$\mathcal{G}_{GHZ} = \mathcal{G}_X + \mathcal{G}_Z, \text{ with } \mathcal{G}_X = \bigotimes_{j=1}^n X_j, \text{ and } \mathcal{G}_Z = \sum_{j=1}^{n-1} Z_j Z_{j+1}. \quad (4.16)$$

When estimating \mathcal{G}_{GHZ} , one could distribute uniformly the measurements to be taken for the estimation of the expectation values of \mathcal{G}_X and \mathcal{G}_Z . However, we saw that such allocation may not be optimal, and rather derived optimal allocation of the measurements in Equation (2.31). For the GHZ state, this yields the two (importance) sampling probabilities $Pr(i = \mathcal{G}_X, \mathcal{G}_Z)$, with i specifying the index of the measurement to be taken.

Recall that determining such optimal allocation requires evaluating the expectation values of $\langle \mathcal{G}_Z^2 \rangle_\rho$ and $\langle \mathcal{G}_X^2 \rangle_\rho$. Given that $\mathcal{G}_X^2 = I$, it follows that $\langle \mathcal{G}_X^2 \rangle_\rho = 1$ which does not depend on the exact details of the underlying state ρ measured. Values of $\langle \mathcal{G}_Z^2 \rangle_\rho$, however, depend on the state ρ which is unknown and would vary over the course of an optimization. Still, in such case one can approximate these values by their average obtained over a distribution of random states. This average is evaluated in Appendix C.4 as $\langle \mathcal{G}_Z^2 \rangle_{rdm} = n - 1$. Following Equation (2.31), this yields the distribution

$$Pr(i = \mathcal{G}_X) = \frac{1}{\sqrt{n-1} + 1}, \text{ and } Pr(i = \mathcal{G}_Z) = \frac{\sqrt{n-1}}{\sqrt{n-1} + 1}, \quad (4.17)$$

which allocates (in average) more measurements to the observable \mathcal{G}_Z as it contributes the most

to the overall variance of \mathcal{G} . The variance of the estimates obtained with such allocation, when averaged over random states, can be bounded as (derivations in Appendix C.5)

$$\Delta_{2\text{obs}}^2[\tilde{\mathcal{F}}_W] \leq \frac{1}{4n} + \frac{\sqrt{n-1}}{2n^2} - \frac{1}{4n(d+1)}, \quad (4.18)$$

which, for large n , scales as n^{-1} . This is an improvement of a factor n compared to the general case where $\tilde{\mathcal{F}}_W$ requires n distinct measurement settings (as given in Equation (4.13)). We now proceed to verify such scalings numerically for random states, and also assess if they hold in the case of the preparation of a GHZ states by means of a shallow quantum circuit.

Numerical results for the GHZ state

In this section we present a numerical comparison of the estimation accuracy $\Delta^2[\mathcal{F}]$ and of the variances $D_{\theta}(F)$, between the state fidelity $F = \mathcal{F}$ (blue), and the witness fidelity $F = \mathcal{F}_W$ (orange), both constructed with respect to a n -qubit GHZ target state. These statistics are evaluated over a total of $S = 10\,000$ states, and for different system sizes ranging from $n = 2$ to 9 qubits.

Statistics presented in the first row of Figure 4.4 are based on an ensemble of random states. The values of the variances $D_{\theta}(F)$, reported in Figure 4.4(a), exhibit an exponential decay with the system size n , following closely the analytical scalings (depicted in dashed lines) of d^{-2} and $(nd)^{-1}$ expected for the state fidelity and the witness fidelity respectively.

Variances of the estimates, displayed in Figure 4.4(b), are obtained given a fixed number of $N = 10\,000$ measurements. In the case of the witness fidelity, we report results assuming either that n distinct experimental observables need to be measured separately (orange), which is representative of the estimation of a generic stabilizer state, or that only two measurement settings (green) are required, as is possible for the GHZ state. For the latter case, the allocation of the measurements is performed according to Equation (4.17). As can be seen, the estimates of the witness fidelity are always at least 4 times more accurate than the estimates of the state fidelity. Furthermore when only 2 distinct measurement settings are required, the variance of the corresponding estimates is found to be inversely proportional to the system size n . Again, these numerical results confirm the scalings (depicted in dashed lines) which were provided in Equation (4.14) for the case of the state fidelity, and in Equations (4.13) and (4.18) for the witness fidelity decomposed over n and 2 observables, respectively. Overall, we have verified the correctness of the analytical results established earlier, and we can now proceed to reexamine these statistics for an ensemble of states differing from the random states.

We now consider the n -qubit states prepared by means of a shallow quantum circuit, sketched in Figure 4.4 (bottom left). This circuit is structured as a sequence of $n - 1$ controlled NOT

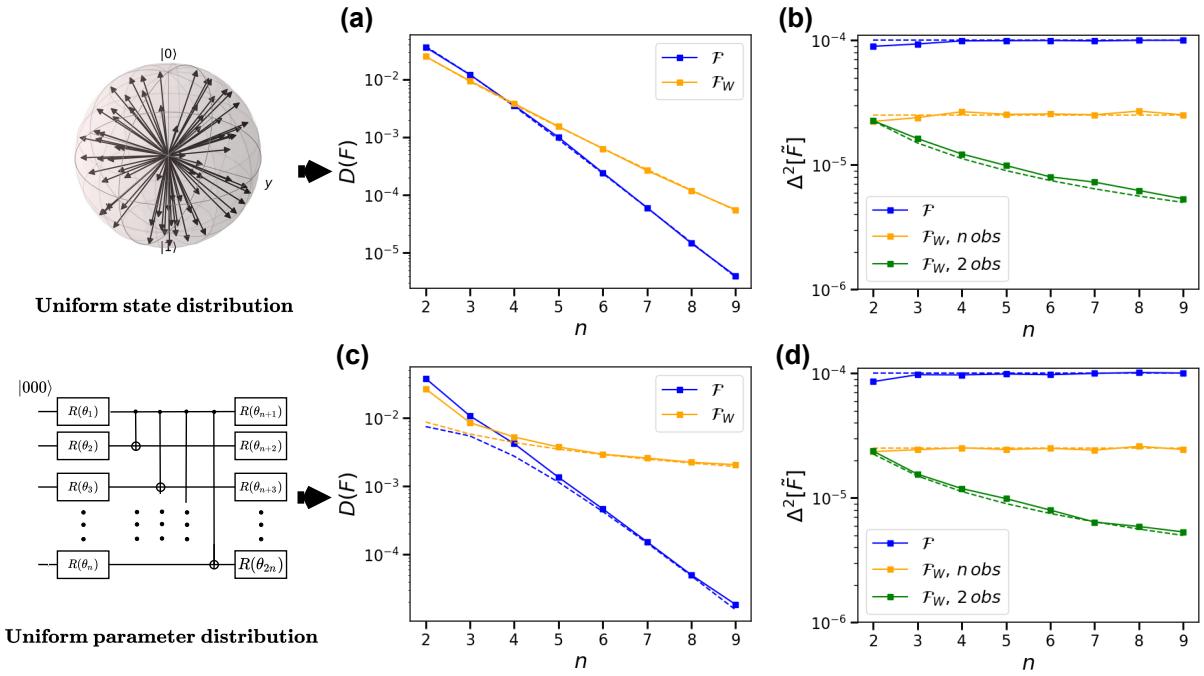


Figure 4.4: Comparison of the state fidelity \mathcal{F} and the alternative witness fidelity \mathcal{F}_W with respect to a target GHZ state. For each figure of merit, we report the variance in its values $D(F)$ (a,c) and the variance $\Delta^2[\tilde{F}]$ in its estimates given $N = 10\,000$ measurements (b,d), for different system sizes ranging from $n = 2$ to 9 qubits. These statistics (depicted with filled symbols and plain lines) are systematically evaluated given an ensemble of $S = 10\,000$ states which are either random states (first row) or states generated by a quantum circuit with random parameters (second row, with the circuit sketched on the left). (a,c) The witness fidelity (orange) concentrates less than the state fidelity (blue), and in particular, does not decay exponentially with the system size n for the ensemble of circuit–output states (c). (b,d) Comparison of the variances in estimates of the state (blue) and witness fidelities. For the witness fidelity \mathcal{F}'_W two strategies are presented, one where n distinct observables are measured (orange), and one where this number is compressed to 2 (green), as is possible for GHZ states. The scalings (dashed lines) either obtained analytically, or fitted to the numerical data are discussed in the main text.

gates with a fixed control qubit ($i = 1$) and targeting each of the remaining qubits ($i > 1$), to which has been prepended and appended $2n$ parameterized single-qubit rotations. Each of these parameterized rotations is defined as

$$R(\boldsymbol{\theta}_i) = R_X(\boldsymbol{\theta}_i^{(1)})R_Z(\boldsymbol{\theta}_i^{(2)})R_X(\boldsymbol{\theta}_i^{(3)}), \quad (4.19)$$

in terms of the Pauli rotations $R_{\sigma \in \{X,Y,Z\}}(\theta) = \exp^{-i\frac{\theta}{2}\sigma}$, and can realize an arbitrary single-qubit rotation for an appropriate choice of the parameters $\boldsymbol{\theta}_i$.

Given an initial Hadamard gate applied to the first qubit of the system, and any of the remaining $2n - 1$ parameterized rotations fixed to be the identity, this circuit is known to realize perfectly the desired GHZ state. However, due to control errors that are typically found in NISQ devices, such idealized circuit is likely to produce a different state when implemented on a real hardware.

Deviations between the realized state and the desired GHZ state could be (at least partially) compensated by means of optimizing the $2n$ parameterized rotations (and thus the $6n$ angles). In such a way, the circuit presented could be used to refine the preparation of a GHZ state on a quantum computing platform. Furthermore, it is representative of some shallow quantum circuits (with a number of sequential gates and a number of parameters to be optimized scaling linearly in the number of qubits) which are often found in VQAs.

The distribution of states corresponding to this circuit is obtained by randomly sampling many of the free rotation angles θ uniformly over the domain $[0, 4\pi]$, and evolving an initial state $|0\rangle^{\otimes n}$ with the resulting circuit. Similar study as the one performed in the case of random states is repeated with this new ensemble of states. Crucially, one can see in Figure 4.4(c) that while the variance in the values of the state fidelity (blue) remains exponentially vanishing in the system size n , this is not the case anymore for the witness fidelity (orange). Instead, for the latter, a scaling of $1/n$ (displayed in dashed orange line) is empirically found to match the numerical data. Consequently, the optimization of such circuits, for large values of n , would only be possible with the witness fidelity. Additionally, scalings in the variances of the estimates shown in Figure 4.4(d) are empirically found similar to the ones obtained for the ensemble of random states in Figure 4.4(b), that is, with enhanced accuracy for the witness fidelity.

Concluding remarks

In conclusion, in light of our criteria for a good figure of merit, we find that the witness fidelity defined in Equation (4.12) provides a practical advantage when compared to the state fidelity. In particular, the values of the latter are found to systematically concentrate, such that, for even moderately large system sizes n , any problem of quantum state preparation (relying on the state fidelity measure) started from random parameters would most likely fail. In contrast, in the case of the witness fidelity, such effect is either retarded or suppressed depending on the specificities of the preparation routine considered. These findings are similar to the ones reported for the local fidelity, which was defined in the previous section for separable states. Additionally, it was shown that the witness fidelity can be more efficiently estimated, especially in the case of GHZ states. Given that the GHZ state is a popular state in quantum information and metrology, requiring accurate preparation, the alternative figure presented could find immediate application. More generally, these results highlight the limits of the state fidelity for optimization purposes, and should encourage the development of substitutes. In addition to these problems of quantum state preparation, we now discuss alternative figures for problems of gate engineering.

4.2.3 Alternative to the process fidelity

Quantum gate engineering aims at the optimization of the dynamics of a quantum system such as to mimic the action of a target unitary, typically taken as an elementary gate central to the development of any quantum computing platform. Such optimization relies on the ability to measure similarities between the target and the realized dynamics. As briefly mentioned earlier, randomized benchmarking (RB) [207–209] provides a scalable path towards the characterization of errors averaged over a set of gates, and has the benefit to decouple such errors from state preparation and measurement errors. However, its application is often limited to the characterization of a specific gate-set, typically Clifford gates, and the interpretation of the errors that RB measures relies on assumptions which do not always hold [218]. A more general figure of merit, commonly adopted in problems of quantum gate engineering, is the process fidelity [171, 219].

In this section we recall the definition of the process fidelity, and derive formulas quantifying the statistical errors incurred when estimating such quantity experimentally. As we will see, for general target unitaries, acting on Hilbert space with dimension d , the number of measurement repetitions required to achieve constant estimate accuracy scales as d^2 . Such scaling quickly limits the use of the process fidelity to the optimization of gates acting on small system sizes. An alternative to the process fidelity is proposed in [170] under the name of 0-fidelity. Amongst its properties making it well-suited for problems of quantum gate engineering, we derive its estimation accuracy, showing a quadratic improvement compared to the process fidelity. Finally, we discuss the use of this new figure of merit in the context of the optimization of a 3-qubit gate with a parameterized quantum circuit.

It is stressed that the author contributed to [170] in regards to the analytical derivations of the variances of the estimates for both the process and 0-fidelity, which form the core of this section. His contribution, however, does not extend to neither the original construction of the 0-fidelity, nor its application to the variational optimization of a 3-qubit gate on an IBM quantum chip. Still, to put our work into perspective, we will review the definition of the 0-fidelity, and also briefly discuss the optimization results achieved.

Process fidelity for quantum gate engineering

Definition. The process fidelity naturally extends the state fidelity, which measures the overlap between states, to quantum processes. This extension is achieved by leveraging the fact that quantum processes can be mapped to quantum states in larger spaces, such that the state fidelity can be employed. More precisely, there exists an one-to-one mapping between any quantum process Λ acting on the Hilbert space \mathcal{H} , with dimension d , and a state ρ_Λ belonging

to the set of operators acting on the extended space $\mathcal{H}' = \mathcal{H} \otimes \mathcal{H}$ which has dimension d^2 . Given an arbitrary basis $\{|i\rangle\}$ of \mathcal{H} , this mapping, known as the Jamiolkowski isomorphism [220], is defined as

$$\rho_\Lambda = (I \otimes \Lambda)|\Phi\rangle\langle\Phi|, \quad (4.20)$$

with the input $|\Phi\rangle = \sum |j\rangle|j\rangle/\sqrt{d} \in \mathcal{H}'$ being a state with maximal bipartite entanglement in between the two copies of \mathcal{H} .

The process fidelity between two processes Λ and Γ is defined as the state fidelity between their corresponding states ρ_Λ and ρ_Γ :

$$\mathcal{F}_{pro}(\Gamma, \Lambda) = \mathcal{F}(\rho_\Gamma, \rho_\Lambda). \quad (4.21)$$

By construction it inherits properties of the state fidelity such as, having bounded values $\mathcal{F}_{pro} \in [0, 1]$, and reaching its maximal value if and only if $\Gamma = \Lambda$. In problems of gate engineering one aims at reducing discrepancies between an implemented process Γ , and a targeted one Λ . For unitary targets of the form $\Lambda(\rho) = U_{tgt}\rho U_{tgt}^\dagger$, the process fidelity can be expressed [219] as the averaged overlap

$$\mathcal{F}_{pro}(\Gamma, \Lambda) = \frac{1}{d^2} \sum_{j=1}^{d^2} \text{Tr}[\Gamma(W_j^\dagger)\Lambda(W_j)], \quad (4.22)$$

between d^2 pairs of operators, $\Gamma(W_j)$ and $\Lambda(W_j)$, resulting from the action of the two processes on each element W_j of an orthonormal basis $\{W_j\}$ of unitary operators (which are scaled here by a factor $1/\sqrt{d}$). In the following, we restrict the discussion to n -qubit systems, and this basis $\{W_j\}$ is taken to be the set of the Pauli observables with unit norm. These observables were defined in Section 2.2.3 as the n -fold tensor products of either Pauli matrices (X, Y and Z) or the identity (I), rescaled by a factor $1/\sqrt{d}$. Given that such operators are Hermitian, we will often substitute W_j for its Hermitian transpose W_j^\dagger , and $\Lambda(W_j)$ for $\Lambda(W_j)^\dagger$.

Estimation by importance sampling. In order to relate the process fidelity to quantities which can be experimentally measured, it is necessary to rearrange Equation (4.22). First, it is required for the inputs entering the implemented channel Γ to be physical states which can be accurately prepared. Recall that the operators W_j have eigendecompositions

$$W_j = \sum_{k=1}^d \lambda_j^k |\lambda_j^k\rangle\langle\lambda_j^k|, \quad (4.23)$$

with eigenvalues $\lambda_j^k = \pm 1/\sqrt{d}$, and eigenvectors $|\lambda_j^k\rangle$ that are tensor products of 1-qubit states (these states are the eigenvectors of the individual Pauli operators, or identities, that form W_j). Since the preparation of $|\lambda_j^k\rangle$ only requires local operations which can be realized faithfully, these states are deemed adequate input states. Second, the operators $\Lambda(W_j)$ to be measured need to

be decomposed in terms of the experimental observables $\{W_i\}$, that is,

$$\Lambda(W_j) = \sum_{i=1}^{d^2} \text{Tr}[\Lambda(W_j)W_i]W_i. \quad (4.24)$$

Inserting Equations (4.23) and (4.24) into Equation (4.22) results in an expression of the process fidelity given as a sum

$$\mathcal{F}_{pro}(\Gamma, \Lambda) = \sum_{i,j=1}^{d^2} \sum_{k=1}^d \alpha(i, j, k) \langle W_i \rangle_{\rho_j^k}, \quad (4.25)$$

with $\begin{cases} \alpha(i, j, k) = \lambda_j^k \text{Tr}[\Lambda(W_j)W_i]/d^2 \\ \rho_j^k = \Gamma(|\lambda_j^k\rangle\langle\lambda_j^k|) \end{cases}$

of elements $\langle W_i \rangle_{\rho_j^k}$ which can be estimated experimentally, weighted by the scalars $\alpha(i, j, k)$. Both these terms are function of a triplet (i, j, k) of integers. Each triplet corresponds to a unique experimental setting comprised of the choice of the observable W_i to be measured, and of the choice of the input eigenstate $|\lambda_j^k\rangle$ to be initially prepared.

The formulas obtained in Section 2.2.2 (regarding the estimation of observables by means of importance sampling) are directly applicable to the estimation of the process fidelity in Equation (4.25), provided that the indices $i \rightarrow (i, j, k)$ and the expectation values $\langle W_i \rangle_{\rho} \rightarrow \langle W_i \rangle_{\rho_j^k}$ are relabelled appropriately. In particular, the probability of choosing one setting to be measured, which was given in Equation (2.29), becomes

$$Pr(i, j, k)^* = \frac{\alpha^2(i, j, k) \langle W_i^2 \rangle_{\rho_j^k}}{\sum_i \alpha^2(i, j, k) \langle W_i^2 \rangle_{\rho_j^k}}. \quad (4.26)$$

In turn, the corresponding variance of an estimate Y (obtained for a single measurement) of the process fidelity \mathcal{F}_{pro} , which was given in Equation (2.30), becomes

$$\Delta^2[Y] = D \sum_{i,j=1}^{d^2} \sum_{k=1}^d \alpha^2(i, j, k) \langle W_i^2 \rangle_{\rho_j^k} - \langle \mathcal{F}_{pro} \rangle^2, \quad (4.27)$$

where D denotes the number of distinct non-zero coefficients $\alpha(i, j, k)$. Equation (4.27) can be simplified, and we now detail the steps of derivation involved.

Evaluating the estimation accuracy. First, remark that for any unitary target of the form

$\Lambda(\rho) = U_{tgt}\rho U_{tgt}^\dagger$, we have

$$\begin{aligned}\Lambda^2(W_j) &= U_{tgt}^\dagger W_j U_{tgt} U_{tgt}^\dagger W_j U_{tgt} \\ &= U_{tgt}^\dagger \frac{I}{d} U_{tgt} = \frac{I}{d},\end{aligned}\tag{4.28}$$

which follows from $W_j^2 = I/d$. Second, note that

$$\begin{aligned}Tr[\Lambda^2(W_j)] &= Tr\left[\left(\sum_{i=1}^d Tr[\Lambda(W_j)W_i]W_i\right)\left(\sum_{i=1}^d Tr[\Lambda(W_j)W_{i'}]W_{i'}\right)\right] \\ &= \sum_{i,i'=1}^{d^2} Tr[\Lambda(W_j)W_i]Tr[\Lambda(W_j)W_{i'}]Tr[W_{i'}W_i] \\ &= \sum_{i=1}^{d^2} Tr^2[\Lambda(W_j)W_i].\end{aligned}\tag{4.29}$$

The first equality in Equation (4.29) results from the decomposition of each operator $\Lambda(W_j)$ in the basis $\{W_{i(i')}\}$. The second line results from the linearity of the trace. Finally, due to the orthonormality of the basis $\{W_i\}$, $Tr[W_iW_{i'}] = \delta_{i,i'}$ permitting the simplification of the sum and yielding the final expression. It follows from Equations (4.28) and (4.29) that

$$\sum_{i=1}^{d^2} Tr^2[\Lambda(W_j)W_i] = 1.\tag{4.30}$$

Given this last result, we can now simplify the expression of the variance of the estimates of the process fidelity, which was given in Equation (4.27).

$$\begin{aligned}\Delta^2[Y] &= \frac{D}{d^6} \sum_{i,j=1}^{d^2} \sum_{k=1}^d Tr^2[\Lambda(S_j)W_i] - \langle \mathcal{F}_{pro} \rangle^2 \\ &= \frac{D}{d^6} \sum_{j=1}^{d^2} \sum_{k=1}^d 1 - \langle \mathcal{F}_{pro} \rangle^2 \\ &= \frac{D}{d^3} - \langle \mathcal{F}_{pro} \rangle^2.\end{aligned}\tag{4.31}$$

The r.h.s. of the first line is obtained by inserting the expression of $\alpha(i,j,k)$ from Equation (4.25), recalling that $(\lambda_j^k)^2 = 1/d$, and that $W_i^2 = I/d$ which entails $\langle W_i^2 \rangle_\rho = 1/d$ for any state ρ . The second line follows from Equation (4.30) which allows us to simplify the summation over the index i . Given that the remaining indices j and k run over d^2 and d terms respectively, the overall sum becomes d^3 . This yields the final expression relating the variance $\Delta^2[Y]$ to the number D of non-zero coefficients $\alpha(i,j,k)$. We now discuss further this number, and the corresponding variances that it entails, for general and specific families of target unitaries.

Scaling for general and Clifford unitaries. As corresponding to the input states to be prepared, the indices j and k always run over the d^2 elements of the basis $\{W_j\}$, and over the d eigenvectors of the operator W_i respectively. However, the number of indices i resulting in non-zero values of $\alpha(i, j, k)$ can vary depending on the exact details of the target channel Λ . Given the definition of $\alpha(i, j, k)$ provided in Equation (4.25), it can be seen that this coefficient becomes null when $\text{Tr}[\Lambda(W_j)W_i] = 0$. Hence, for a fixed choice of indices (j, k) , the number $D_i(j, k)$ of non-null coefficients will always adopt integer values in the interval $[1, d^2]$ resulting in a number D of non-null coefficients in the interval $[d^3, d^5]$. The variance of an estimate of the process fidelity, given in Equation (4.31), can thus be bounded as

$$1 - \langle \mathcal{F}_{\text{pro}} \rangle^2 \leq \Delta^2[Y] \leq d^2 - \langle \mathcal{F}_{\text{pro}} \rangle^2. \quad (4.32)$$

We now discuss these lower and upper bounds.

For general channels Λ such that any of the terms $\text{Tr}[\Lambda(W_j)W_i] \neq 0$, one obtains a number $D_i(j, k) = d^2$ yielding a number $D_{\text{gen}} = d^5$ of non-zero values of $\alpha(i, j, k)$. When inserted in Equation (4.31), it results in a variance of the estimate Y of \mathcal{F}_{pro} given by

$$\Delta_{\text{gen}}^2[Y] = d^2 - \langle \mathcal{F}_{\text{pro}} \rangle^2, \quad (4.33)$$

and scaling as d^2 . Such scaling shows that any attempt to estimate the process fidelity with respect to general unitary targets will quickly require a large number of measurement repetitions. For instance, given a system size of $n = 10$ qubits, aiming at an estimation accuracy of $\Delta_{\text{gen}}^2[\tilde{\mathcal{F}}] = 10^{-4}$ (that is, a standard deviation of 1%) would already necessitate a number $N \approx 10^4 \times 2^{20} \approx 10^{10}$ of measurement repetitions. Nonetheless, for more specific unitary targets a better scaling than d^2 can be achieved. In particular, we now characterize the targets resulting in the lower bound variance seen in Equation (4.32).

Recall the definition of a Clifford gate C as a *normalizer* of the Pauli group, that is, which action on any of the W_j Pauli observables result in another Pauli observable $W_{j'}$ (that is, $C^\dagger W_j C = W_{j'}$). Therefore, given a Clifford target gate $\Lambda(\rho) = C^\dagger \rho C$, the terms $\text{Tr}[\Lambda(W_j)W_i] = \text{Tr}[W_{j'}W_i] = \delta_{i,j'}$ which are non-null only for a single index $i = j'$, per input W_j . Hence, for any of the input states $|\lambda_j^k\rangle\langle\lambda_j^k|$ (eigenstates of W_j) it is only required to measure the operator $W_{j'}$. This translates into a reduced number of experimental settings $D_{\text{Cliff}} = d^3$, and a variance of the estimate Y of the process fidelity given by

$$\Delta_{\text{Cliff}}^2[Y] = 1 - \langle \mathcal{F}_{\text{pro}}^2 \rangle, \quad (4.34)$$

which is bounded in $[0, 1]$.

As was the case when estimating the state fidelity with respect to a target stabilizer state

(discussed in Section 2.2.3), the estimation of the process fidelity with respect to a target Clifford gate can be achieved with a number of measurements which does not depend on the system size. For instance, achieving the estimate accuracy of $\Delta_{Cliff}^2[\tilde{F}] = 10^{-4}$ mentioned earlier would require a number $N \approx 10^4$ of measurement repetitions for any Clifford unitary. Such results are encouraging⁴ given that Clifford gates are central to many quantum computation schemes. Nonetheless, improving on the scaling of the estimate accuracy for more general target unitaries (as seen in Equation (4.33)) would ease the optimization of gates which do not belong to the Clifford group. We now discuss an alternative to the process fidelity which exhibits improved estimate properties for general targets.

Zero-fidelity for quantum gate engineering

Alternatives to the process fidelity are presented in [170]. These are provided as a hierarchy of k -fidelities \mathcal{F}_k , with k varying from 0 to n , which are each an approximation of the process fidelity. While the case $k = n$ recovers exactly the process fidelity, we rather restrict our attention to the $k = 0$ -fidelity measure \mathcal{F}_0 which presents desirable properties when used in the context of optimization. We briefly review the definition and properties of \mathcal{F}_0 , before comparing it to the process fidelity.

Definition and properties of \mathcal{F}_0 . Consider a set of 4 single-qubit states $\{\rho_i^1\}$ which have a minimal overlap $\sum_{i \neq i'} \text{Tr}[\rho_i^1 \rho_{i'}^1]$. These states, known as symmetric informationally complete (SIC), have analytical solutions [223] which can be visualized as the 4 vertices of a tetrahedron on the Bloch sphere. While these are often employed as optimal measurements in quantum state tomography, as we are interested in processes rather than states, we will employ them as input states to the channels to be characterized. When probing a general n -qubit process, we will consider their n -fold tensor products⁵, which are denoted $\rho_j = \rho_{j_1}^1 \otimes \dots \otimes \rho_{j_N}^1$. Given such a set $\{\rho_j\}$ of d^2 states spanning the full space of n -qubit states, the 0-fidelity is defined as

$$\mathcal{F}_0(\Gamma, \Lambda) = \frac{1}{d^2} \sum_{i,j=1}^{d^2} \text{Tr}[\Gamma(\rho_j) W_i] \text{Tr}[\Lambda(\rho_j) W_i], \quad (4.35)$$

with $\{W_j\}$ denoting, as usual, the orthonormal basis of Pauli observables (Section 2.2.3).

Notably, it can be shown that \mathcal{F}_0 is maximized if and only if the process fidelity \mathcal{F}_{pro} is maximized [170], that is, when both the target and the realized processes wholly agree, $\Lambda = \Gamma$. The 0-fidelity is thus faithful to the original task of gate engineering. Furthermore it adopts

⁴This ability to characterize Clifford gates with an experimental effort which does not depend on the system size has been reported several times, for instance in [166, 221, 222] but with explanations often tied to several additional technical details.

⁵These tensor product states are not the most general states minimizing the overlap $\sum_{i \neq i'} \text{Tr}[\rho_i \rho_{i'}]$, which are only known for specific dimensions [223], but are optimal among the set of separable states [170]

bounded values $\mathcal{F}_0 \in [0, 1]$ which makes it directly comparable to \mathcal{F}_{pro} , to which it can be related further by means of the bounds [224]:

$$1 - \frac{3}{2}(1 - \mathcal{F}_0) \leq \mathcal{F}_{pro} \leq \mathcal{F}_0. \quad (4.36)$$

Finally, as can be seen in Equation (4.35), \mathcal{F}_0 can be decomposed in terms of a total of d^4 unique experimental settings which are indexed by the pair of integers (i, j) both varying in between 1 and d^2 . Given this reduced number of terms appearing in the decomposition of \mathcal{F}_0 (compared to the maximal number of d^5 terms for \mathcal{F}_{pro}), one may except an improved estimation accuracy, a feature which is now verified.

Variance of the estimates. In the following, we adopt the same methodology as used for the process fidelity in order to derive the variance of an estimate of \mathcal{F}_0 obtained by means of importance sampling. For that purpose, Equation (4.35) is first recast as the weighted sum

$$\mathcal{F}_0(\Gamma, \Lambda) = \sum_{i,j=1}^{d^2} \alpha_0(i, j) \langle W_i \rangle_{\Gamma(\rho_j)} \quad (4.37)$$

$$\text{with } \alpha_0(i, j) = \text{Tr}[\Lambda(\rho_j) W_i] / d^2$$

over the expectation values $\langle W_i \rangle_{\Gamma(\rho_j)}$ of the operator W_i with respect to the state $\Gamma(\rho_j)$, which can be directly estimated. It follows that the probability of selecting one setting (i, j) corresponding to a measurement of $\langle W_i \rangle_{\rho_j}$ is given by

$$Pr(i, j)^* = \frac{\alpha_0^2(i, j) \langle W_i^2 \rangle_{\Gamma(\rho_j)}}{\sum_i \alpha_0^2(i, j) \langle W_i^2 \rangle_{\Gamma(\rho_j)}}. \quad (4.38)$$

This results in an estimate Y_0 (obtained given a single measurement) of \mathcal{F}_0 with variance

$$\begin{aligned} \Delta^2[Y_0] &= D_0 \sum_{i,j=1}^{d^2} \alpha_0^2(i, j) \langle W_i^2 \rangle_{\rho_j} - \langle \mathcal{F}_0 \rangle^2 \\ &= \frac{D_0}{d^5} - \langle \mathcal{F}_0 \rangle^2, \end{aligned} \quad (4.39)$$

where we have performed similar derivation steps as in Equation (4.31), and with D_0 denoting the number of distinct non-zero coefficients $\alpha_0(i, j)$.

Given that the indices i and j run over the d^2 input states ρ_j and the d^2 experimental observables W_i respectively, we can bound the number of coefficients as $D_0 \leq d^4$. This yields a variance

$$\Delta^2[Y_0] \leq d - \langle \mathcal{F}_0 \rangle^2. \quad (4.40)$$

Compared to the variance of the process fidelity given Equation (4.32), it can be observed a

quadratic improvement in the scaling of (the higher bound of) the variance with respect to the dimension d of the system. Despite this improvement, it remains the case that the variance of the estimate Y_0 grows with d . Still, \mathcal{F}_0 offers a practical advantage compared to \mathcal{F}_{pro} : at given system size and estimation accuracy, the former will necessitate d times less measurements. Returning to our example of a targeted accuracy of $\Delta_{gen}^2[\tilde{\mathcal{F}}_0] = 10^{-4}$ for a system size of 10 qubits, this entails a number $N \approx 10^4 \times 2^{10} \approx 10^7$ (which was 10^{10} for the process fidelity) of measurements.

Overall, we have seen that, akin to the case of the state fidelity, an alternative to the process fidelity can be constructed. This alternative ensures faithfulness to the same target gate, and has the benefit of being easier (in general) to estimate. Additionally, as based on the state fidelity, the process fidelity would display similar problems of concentration in its values, which may be improved when resorting to the 0-fidelity. We leave the study of the concentration in the values of \mathcal{F}_0 to future work, and now rather briefly discuss the application of the 0-fidelity in the context of optimization.

Variational optimization of a quantum gate based on the 0-fidelity

The 0-fidelity is used in [170] to optimize the implementation of a controlled NOT gate⁶ acting on next-neighbour qubits on an IBM quantum chip. The (3-qubit) quantum circuit to be optimized is composed of a fixed sequence of gates – entailing the perfect realization of the target gate in the case of flawless operation of the chip – to which is appended and prepended parameterized single-qubit rotations (as defined earlier in Equation (4.19)). In the same spirit as the circuit discussed for the realization of a GHZ state in Figure 4.4(c,d), such additional free rotations could, if properly tuned, compensate for intrinsic control errors as found in current NISQ devices.

This task of quantum gate engineering involves a total of $P = 18$ parameters to be optimized with the aim of minimizing discrepancies between the target gate and the realized process. This optimization was performed based on the 0-fidelity measure and with BO. Over 10 repeated optimization runs, the optimized circuits were found to perform systematically better than the idealized implementation of the gate (that is, with the 6 free rotations fixed to be the identity). Such improvement was assessed in terms of the process fidelity, rather than the approximate 0-fidelity that was employed during optimization, and it was found an average of 12(17)% of absolute(relative) improvement after optimization.

⁶Note that for such gate belonging to the Clifford group, the process fidelity could have been preferred to the 0-fidelity as per the scaling in Equation (4.34). Still, given that the amplitudes of the 0-fidelity estimation errors do not depend on the target gate, the optimization results presented in [170] can be equally representative of any other problem of quantum gate engineering defined over similar system sizes.

4.3 Concluding remarks

In addition to the choice of the right optimizer which was studied in Chapter 3, here we have investigated an aspect of optimization often ignored, namely the choice of the right figure of merit. As was shown in several examples, it is often the case that alternative figures (compared to the figures of merit commonly employed) can be discovered and could reveal beneficial. Deciding which figure of merit is appropriate requires clear guidelines. For that purpose, we have started this chapter by establishing criteria for adequate figures of merit. These were taken general enough to be relevant to many problems of optimizations, and such that they can be easily assessed. For instance, ensuring that bounded estimates error can be obtained, is required for any optimization performed on experimental data. Quickly characterizing the magnitudes of such errors was made possible by means of the general statistical methodology for estimation presented in Section 2.2. This was extended in Section 4.2.3 to the case when varied input states need to be taken into consideration (as needed, for instance, to characterize quantum gates). Additionally, ensuring that the values of a figure of merit do not concentrate to an (experimentally) indistinguishable range is also a prerequisite for any successful optimization starting from random parameters. The scaling of such concentrations can be studied over small system sizes and extrapolated to larger sizes. Taken together, these criteria allow assessing the feasibility of an optimization problem, and can be used to contrast different figures of merit.

Based on these criteria, we compared several figures of merit relevant to problems of state preparation and gate engineering. In any case, we found improved characteristics for the alternatives which were studied. These results, should encourage the discovery of new figures exceeding the examples studied. Ultimately, the figures suggested should be judged on full-fledged optimizations, rather than on the study of their properties, as presented here. For that purpose, we discussed several practical applications along the way. These could be explored to benchmark and validate further the alternatives proposed. For instance, the preparation of large-scale GHZ state [81, 165, 225] can provide an ideal testbed for assessing the benefits of the witness fidelity discussed in Section 4.2.2. Similarly, the preparation of many-body product states in ultracold-atoms platforms [179], or the optimization of gate sets by means of randomized benchmarking [45, 56] could be improved by employing the local fidelity presented in Section 4.2.1. Given that such tasks are of great interest for quantum computing, simulation and metrology purposes, the proposed figures of merit could find immediate application.

It was discussed at the end of Section 4.2.3 a successful optimization of a 3-qubit gate using BO and the alternative 0-fidelity. This optimization was performed on a NISQ device, and enabled the improvement of a controlled NOT gate (over non adjacent qubits) compared to its idealized implementation. This permits to appreciate the advantage that can be obtained by optimizing the dynamics of a quantum system directly based on experimental outcomes,

rather than relying on a flawed model. Still, to fully reap the benefits of such approach, and to improve on the implementation of a general quantum circuit, one would need to optimize not only a single gate, as was detailed here, but many of such gates. Such situation, involving many optimizations to be performed jointly, calls for even more principled optimization methodologies and is discussed in the next chapter.

Chapter 5

From one to many optimizations

So far we have been concerned with individual tasks of optimization, such as the realization of a Mott insulating state in Section 3.1, the creation of a single target GHZ state in Section 3.2 or the optimization of a controlled NOT gate discussed in Section 4.2.3. In practice, however, optimization tasks often come in families. Consider for instance the following problems:

1. Resolve the electronic ground-state energies of an molecule for different nuclear separations [20].
2. Resolve the ground-state energies of an Ising model over varied model parameters [25].
3. Maximize the state transfer between 2 levels in a 3-level system given varied values of the energies detunings which could originate, for instance, from manufacturing imperfections [54].
4. Learn to optimally control continuous families of 2-qubit entangling gates [226–228].

All these examples share the common feature that many optimizations have to be performed all together. Undertaking these extended problems requires a principled and efficient optimization methodology.

A naive approach towards the optimization of a family of problems would consist of performing each of the many optimization tasks independently. For instance, in situation Item 1, one would start by discretizing the values of the nuclear separation and would perform separate optimizations for each of the configurations. This, however, becomes quickly prohibitively demanding for a fine discretization, or for configuration space with dimension greater than one (for instance, with more than one inter-atomic distance to be varied). In all the examples mentioned above, Items 1 to 4, it is important to recognize that the many tasks to be solved belong to a same family, and share similarities. Hence, one would hope to leverage information gained from one task to another. For that purpose, we present two frameworks intended to efficiently address the situation where optimizations over many *related* problems need to be performed, one for the case of VQAs and one for QOC problems.

In Section 5.1 we leverage the model-based approach of BO (presented in Section 2.1.3) to reuse

measurements obtained over the course of one optimization in order to accelerate any other optimization. The scheme that is presented in [229] lends itself both to parallel or sequential optimization of VQAs. When applied to families of variational quantum eigensolver (VQE) problems, it is found to provide one, or two, order-of-magnitude improvement compared to independent optimizations. Additionally, in Section 5.2 we propose to lift the task of optimal control of a single target operation (for instance, a controlled NOT gate) to the optimal control of continuous families of targets (for instance, controlled $X(\alpha)$ rotations with a variable rotation angle α). For that purpose, we develop a framework based on neural networks producing controls conditioned on the details of the target to be realized. This framework is applied to the control of continuous families of 1- to 3-qubit gates and achieves low control errors and improved operational times. In contrast to the rest of this thesis, the training of this last framework assumes exact numerical simulations. Nonetheless, we discuss how it could also be trained directly based on experimental data. In all cases, the schemes presented permit us to solve families of optimization problems with considerable reduced effort when compared to existing approaches.

The framework of BO with information sharing, that is presented in the next section, is the fruit of a collaborative effort [229]. The author of this thesis contributed to the conceptual idea and to the implementation of the optimization methodology which permits its applicability to general problems of VQAs. These aspects form the core of Section 5.1. However, the author's contribution does not extend neither to the choice of the problems which are addressed nor to the choice of the circuit ansatzes which are used. Still, to exemplify the applicability of the framework, we give an overview of the results which were made possible by its adoption. Finally, we note that a strategy exhibiting some similarities with the framework presented was also proposed in [230], but with details only briefly sketched, and applied to sequential optimization which is in contrast to the parallel approach that we now describe.

5.1 Variational quantum algorithms with information sharing

In this section, we consider the problem of finding the (many) ground states, and energies, of a family of n -qubit Hamiltonians

$$H(\alpha) = \sum_{i \in \mathcal{I}} w_i(\alpha) P_i \quad (5.1)$$

decomposed over the same subset $\{P_i\}_{i \in \mathcal{I}} \subset \{X, Y, Z, I\}^{\otimes n}$ of tensored Pauli operators (and identity), and with weights $\mathbf{w}(\alpha) = \{w_i(\alpha)\}_{i \in \mathcal{I}}$ depending on a problem index α . This problem index α can be a scalar or a vector, adopting either discrete or continuous values which would

typically correspond to the parameters of a physical model. In the spirit of VQAs (Section 1.2), these ground states preparations are attempted with a parameterized quantum circuit $U(\boldsymbol{\theta})$, with variational parameters $\boldsymbol{\theta}$, applied to an initial state $|\psi_0\rangle$. We stress that both the circuit structure and the initial state are kept fixed independently of α . Hence, each problem to be solved is understood as the optimization of the figure of merits

$$F_\alpha(\boldsymbol{\theta}) = \langle H(\alpha) \rangle_{\boldsymbol{\theta}}, \quad (5.2)$$

where we have used the compact notation $\langle M \rangle_{\boldsymbol{\theta}} = \text{Tr}[M|\psi(\boldsymbol{\theta})\rangle\langle\psi(\boldsymbol{\theta})|]$ to denote the expectation value of the operator M with respect to the output state $|\psi(\boldsymbol{\theta})\rangle = U(\boldsymbol{\theta})|\psi_0\rangle$ of the circuit.

This general formulation encompasses many problems of interests found in the field of quantum chemistry and condensed-matter physics, where ground states often need to be obtained over varied configurations of a model. For instance, understanding bond dissociation of a molecule requires to resolve its potential energy for varied inter-atomic distances. Similarly, studying a condensed-matter model typically consists of understanding the properties of its ground (and sometimes excited) states over varied configurations of the model parameters.

As alluded previously, a straightforward but inefficient approach would involve choosing discretized values α_j , over the domain of the parameters α under study, and performing independent optimizations for each of the corresponding figures of merit F_{α_j} in Equation (5.2). However, the overhead incurred by these repeated optimizations quickly becomes excessive. For instance, it was reported in [20], an experimental time of more than 6 days to apply the VQE algorithm to find the ground state energies of a small molecule $He - H^+$ given a single atomic separation α which was discretized over 90 values. Going to larger molecules, and adding a second inter-atomic distance to be varied over (and assuming a similar discretization), would incur more than a year of quantum computations. Experimental progress have been achieved since this first demonstration of VQE [20], still these figures highlight the urgent necessity for more frugal optimization frameworks.

Bayesian optimization with information sharing

As discussed in Section 2.1.3, a key distinction of BO compared to more traditional optimizers, lies in its ability to build a surrogate model of the function to be maximized based on the available body of observations. In particular, recall that the predictive distribution $p(f(\boldsymbol{\theta})|\mathbf{y})$, which was given in Equation (2.8) and used at each BO step, is conditioned on the full dataset of observations \mathbf{y} . Given an additional evaluation y^* of the figure of merit at arbitrary parameters $\boldsymbol{\theta}^*$, this dataset is readily extended to $\mathbf{y} \rightarrow \mathbf{y}' = \mathbf{y} \cup y^*$. In principle, this would result in an improved predictive distribution $p(f(\boldsymbol{\theta})|\mathbf{y}')$. Such feature is in stark contrast to non model-based optimizers. Indeed, these require evaluations of the figure of merit (or its gradients)

according to a predefined strategy, and thus cannot benefit from additional data if obtained for arbitrary parameter values. We build on this specific feature of BO to allow for measurement information obtained over the course of distinct optimizations to be shared. In the following, this strategy is referred as information sharing (IS), and is now detailed.

We first consider the case where two optimization tasks, indexed by α_j with $j = 1$ or 2 , are performed in parallel. This is illustrated in Figure 5.1(a), where we denote $\mathcal{B}(\alpha_j)$ the two Bayesian optimizers aiming at minimizing the figures of merit F_{α_j} . At every iteration of the algorithm, each optimizer requests the evaluation of its corresponding figure of merit for a distinct set of parameters $\boldsymbol{\theta}_j$. For instance $\mathcal{B}(\alpha_1)$ will suggest a new set of circuit parameters $\boldsymbol{\theta}_1$ to be implemented in order to evaluate $F_{\alpha_1}(\boldsymbol{\theta}_1)$. In practice, this evaluation entails measuring the expectation values $\langle P_i \rangle_{\boldsymbol{\theta}_1}$ of the Pauli operators appearing in Equation (5.1) given the circuit parameters $\boldsymbol{\theta}_1$. In turn, these expectation values are used to recompose the figure $F_{\alpha_1}(\boldsymbol{\theta}_1) = \sum_i w_i(\alpha_1) \langle P_i \rangle_{\boldsymbol{\theta}_1}$ given the set of weights $\mathbf{w}(\alpha_1)$. Crucially, the same Pauli measurements can also be used to evaluate $F_{\alpha_2}(\boldsymbol{\theta}_1) = \sum_i w_i(\alpha_2) \langle P_i \rangle_{\boldsymbol{\theta}_1}$, provided that the Pauli measurements are now weighted with $\mathbf{w}(\alpha_2)$. This measurements sharing is depicted in Figure 5.1(a) as a dashed diagonal arrow. A similar arguments applies to the evaluation of $F_{\alpha_1}(\boldsymbol{\theta}_2)$ given measurement data originally intended for $F_{\alpha_2}(\boldsymbol{\theta}_2)$. Hence, any set of Pauli measurements obtained can be systematically shared in between the two optimizers.

This scheme, described so far for the case of two optimizations ran in a parallel, can be extended to an arbitrary number of parallel, or sequential, optimizations. In the case of parallel optimizations over M different problems, each iteration involves Pauli measurements for the M distinct set of parameters (suggested by each of the M optimizers). These measurements enable the evaluation of a total of M^2 figures $F_{\alpha_j}(\boldsymbol{\theta}_{j'})$ with $(j, j' = 1, \dots, M)$, and each optimizer effectively benefits from M new observations per iteration. In the sequential case, any new optimization would start with a body of observations corresponding to all the parameters that have been evaluated over the course of the previous optimizations. A hybrid strategy alternating parallel and sequential BOs with IS could even be envisioned. A first batch of parallel optimizations over coarse-grained values of α would be performed. This would be followed by a second (or more) batch of parallel optimizations over a new discretization of α chosen based on the previous results. In all these cases, IS ensures that all the measurements available at a certain point in time can be maximally exploited.

While, in principle, it should always be beneficial to share as much information as possible, that is to broadcast measurement data obtained for the task α_j to any of the optimizers $\mathcal{B}(\alpha_{j'})$, it is also the case that the computational effort in updating the model in BO grows cubically with the number of observations available (this is discussed in Section 2.1.3). In situations where this computational burden becomes excessive, it may be beneficial to share only a curated set of measurements. Different sharing strategies are depicted in Figure 5.1(b). The cases of inde-

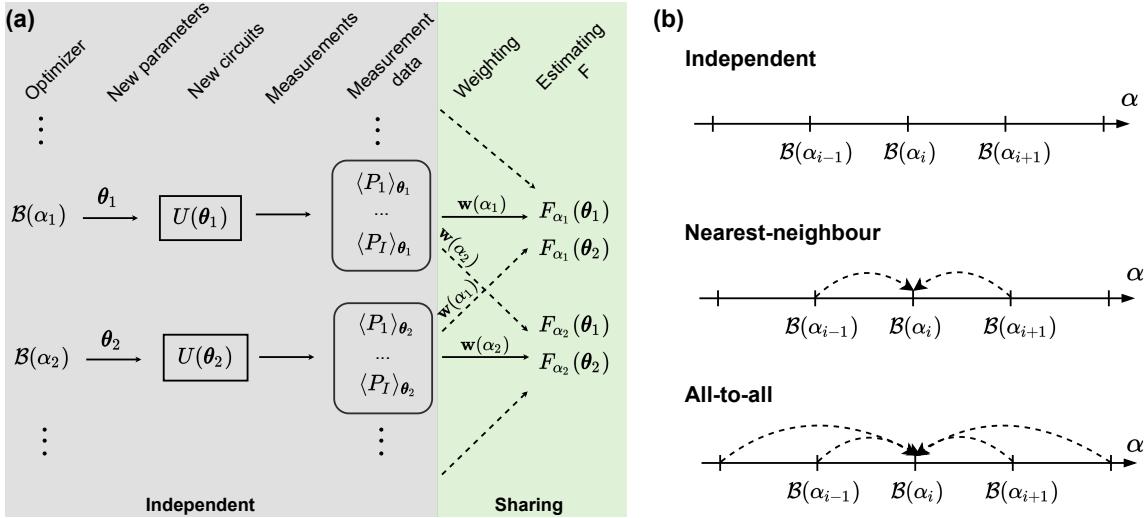


Figure 5.1: Adapted from [229]. (a) Schematic of one iteration of BO with information sharing (IS). In situation when several problems, indexed by α_j , need to be optimized concurrently, information (in terms of measurement results) can be shared in between the optimizers in order to increase their individual performances. Distinct optimizers $\mathcal{B}(\alpha_j)$ optimize for distinct problems with corresponding figures of merit F_{α_j} , defined in Equation (5.2). Each optimizer $\mathcal{B}(\alpha_j)$ requests the evaluation of its figure of merit for a new set of circuit parameters θ_j . These evaluations require measurements of the Pauli operators $\langle P_i \rangle_{\theta_j}$ which, when weighted by the coefficients $w(\alpha_j)$, allow to assess $F_{\alpha_j}(\theta_j)$ (horizontal plain arrows). When weighted with $w(\alpha_{j' \neq j})$, the same measurements also permit the assessment of any figure $F_{\alpha_{j'}}(\theta_j)$ and thus provide additional information that can be used by the optimizer $\mathcal{B}(\alpha_{j'})$ (this information sharing is depicted by diagonal dashed arrows). This forms the basis of the IS scheme, enabling to maximally utilize each measurement data. (b) Different rules can be used when sharing measurements in-between the different optimizers. Here, are depicted (in dashed arrows) the additional measurements received by the optimizer $\mathcal{B}(\alpha_i)$. In the extreme case of independent optimization (first row) no measurement is shared in-between the optimizers. In the opposite limit of all-to-all sharing (last row) all the measurements are systematically shared in between the optimizers. Intermediate strategies, such as the nearest-neighbour sharing (middle row) allow to restrict measurement sharing to neighbouring optimization tasks.

pendent optimization (first row), where no measurement is shared, and of all-to-all sharing (last row), where all the measurements are shared, have both already been discussed. Additionally, one can limit the measurements to be shared only between optimizers corresponding to values of α close in parameter space. The case of nearest-neighbour sharing is depicted in the middle row of Figure 5.1(b) for a one-dimension parameter α . Overall, it allows to balance the amount of information to be shared and the computational burden entailed. In the following we present numerical results which permit to appreciate the differences between these strategies.

Numerical results for a spin model with information sharing

To demonstrate the effectiveness of IS, we will discuss results obtained in [229] for the ground state preparations of a quantum spin chain with Hamiltonians

$$H(\alpha_X, \alpha_Z) = \sum_{\langle i,j \rangle} Z_i Z_j - \sum_i (\alpha_X X_i + \alpha_Z Z_i) \quad (5.3)$$

indexed by the parameters α_X and α_Y which are dimensionless transverse and longitudinal fields respectively. Values of these parameters are limited to $\alpha_X = \alpha_Z = h \in [0, 0.9]$ for which the ground states of Equation (5.3) remain in the same phase. These values are discretized over 15 uniformly spaced points $h_{0 \leq i \leq 14} = i \times \delta$ with $\delta = 0.9/14$, thus forming a family of 15 problems to be solved. The system is taken to be comprised of four spins, with open boundary conditions. In experimental realizations, the success of any VQA is sensitive to different sources of experimental noise, such as sampling effects due to the finite number of measurement repetitions, and device noise which may fluctuate in time. To decouple such effects to the analysis of the IS scheme, the results presented are obtained based on exact numerical evaluations of the expectation values of Equation (5.3). Further technical details, such as the choice of the ansatz, can be found in [229] and are not discussed here.

The performance of four different flavours of the IS scheme are reported in Figure 5.2. In addition to the "Independent" (Figure 5.2(a)), the "Nearest-neighbour" (Figure 5.2(c)) and the "All-to-all" (Figure 5.2(d)) sharing strategies already discussed, we also include an "Independent+Random" setting (Figure 5.2(b)) which will allow us to probe further the importance of data acquisition in BO. This latter setting corresponds to independent runs of BO for each of the 15 problems, and, with each optimizer receiving two additional evaluations of the figure of merit taken at random parameters, per iteration. Results are systematically reported in terms of the difference between the minimum energies E^* found after optimization and the true ground state energies E_{exact} . Each configuration is repeated 100 times, with the empirical means and histograms of the results plotted for each of the 15 values of h . Furthermore, to facilitate the comparison between the different strategies, the total number of exact evaluations performed and the number of evaluations effectively available per Bayesian optimizer are re-

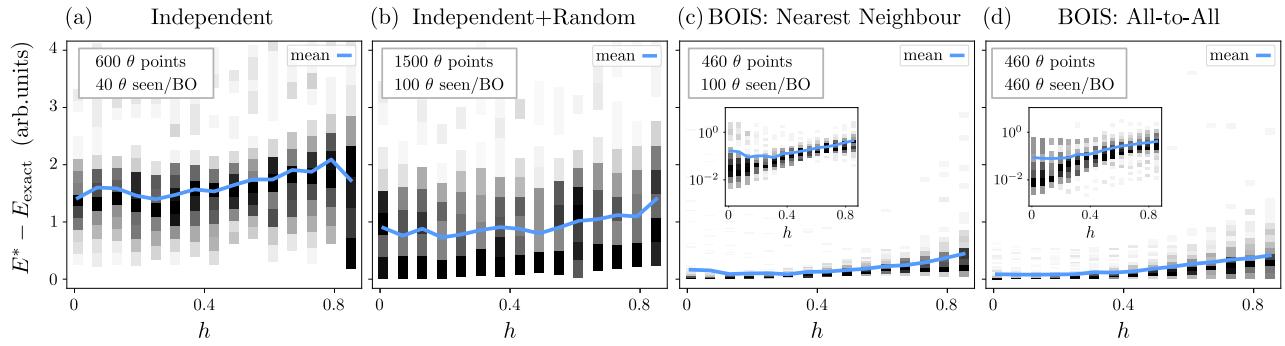


Figure 5.2: From [229]. Comparison between 4 different information sharing strategies (a-d) for BO of the ground states of the Hamiltonian in Equation (5.3) over 15 discretized values of $h = \alpha_X = \alpha_Z \in [0, 0.9]$ (this effectively corresponds to 15 optimization tasks to be solved). In all cases, histograms (gray scales) and means (blue line) of the differences $E^* - E_{exact}$ between optimized and exact ground states are obtained over 100 repetitions. The total number of evaluations and the number seen per optimizer are provided in boxes. (a) Optimizers do not share any information. (b) Similar as before, except than each optimizer now receives 2 additional random evaluations per iteration. (c) Measurements are shared in between neighbor optimizers. (f) All the measurements are shared in between all the optimizers.

ported (in boxes). Each strategy is ran for a total of 30 iterations, following an initial stage where the figure of merit is evaluated for 10 random parameters.

Starting with a comparison of Figure 5.2(a) and Figure 5.2(b), one can see that the additional evaluations, obtained at random parameter values for the "Independent+Random" strategy, translates into improved optimization performances. This highlights the ability of the model approach of BO to benefit from the inclusion of additional data, even when not curated. However, as can be seen in Figure 5.2(b), this improvement remains limited, and the optimized energies are still far from the true minimum energies. In the case of the "Nearest-neighbour" sharing strategy, displayed in Figure 5.2(c), each Bayesian optimizer $\mathcal{B}(h)$ shares measurements only with its two closest neighbors, that is with $\mathcal{B}(h \pm \delta)$. As these additional measurements originate from similar problems, they are deemed more informative than the ones obtained for random parameters. Indeed, despite having access to exactly the same number of 100 evaluations per Bayesian optimizer, the IS scheme used in Figure 5.2(c) is found to significantly outperform the strategy displayed in Figure 5.2(b).

Finally, the full potential of the IS scheme can be appreciated in Figure 5.2(d) which displays performances of BO in the case where all measurements are shared. The substantial improvement compared to the independent optimizations (which results are displayed in Figure 5.2(a)), and also compared to any other setting, highlights the appeal of reusing data as much as is possibly permitted by the IS strategy. Only in the case where $h \approx 0.9$, for which the ground states are the most entangled [229], that the optimizers are found to not always fully converge. Nonetheless, already with only 30 iterations per optimizer, several runs of BO already convergence fully. In any case, the IS strategy largely outperforms independent optimizations, that

is, optimizations as commonly performed.

Further results and outlook

In addition to these numerical results for the spin model of Equation (5.3), IS is also studied experimentally in [229] on IBM quantum chips for the VQE of the two-atom molecules H_2 and LiH , and the three-atom linear chain H_3 . In the latter case, the two inter-atomic distances (in between the first and second, and second and third atoms) are varied, thus entailing a two-dimension ground energy surface to be resolved. In all cases, convergence is reached in 10 – 50 evaluations per problem. This is to be contrasted to typical numbers of evaluations in the 100s or 1000s with non BO and non IS optimizations. For instance, a single iteration of gradient descent may already involve 10s or more of such evaluations, which are needed in order to estimate each element of the gradient vector (Section 2.1.1). Furthermore, for the linear chain of H_3 , which is the largest configuration space considered (with two-dimension parameters α), only 10 evaluations per problem were required. This hints to the potential good scalability of the IS strategy for increased dimension of α : for larger configuration spaces many more optimizations need to be performed, but, at the same time many more measurements are collected and re-used. In conclusion, we have introduced IS as a pragmatic strategy to make maximum use of quantum measurements in the common situation where many related VQAs have to be optimized jointly. The scheme presented here is prone to further extensions that we now discuss.

First, while presented in tandem with BO, the IS scheme could equally be applied to any model-based optimization strategy [231, 232], which may prove more practical than BO in case of very large number of parameters and observations. Second, in the example of the spin chain discussed, parameters of the spin system were chosen such as no phase transition was crossed. In the case of a phase transition, it is expected that the optimal circuit parameters will change drastically in between the different phases. As such, one may expect that measurements obtained for close-to-optimal parameters in one phase may not be informative to problems pertaining to another phase. Still, the inclusion of additional data, even when taken at random parameters, improves the success of BO in general (as was seen in Figure 5.2(b)). Furthermore the explorative nature of BO ensures that measurements are collected in different regions of the parameter space, and are thus likely to benefit any of the optimization tasks. As such, we would still expect some significant advantage offered by IS, and it would instructive to investigate further this situation.

Finally, it should be stressed that the sharing strategy relies on the choice of a fixed ansatz U , and on families of Hamiltonians in Equation (5.1) decomposed over the same subset of Pauli operators. Exceeding these limitations would broaden further the applicability of IS, and could

be achieved when combining it to the latest proposals in the field of VQAs. For instance, resorting to classical shadows [168] could enable the estimation of the expectation values of broader families of operators, for which IS could be employed. Additionally, one could envision a scenario where a core circuit, shared by any problems of the same family, is optimized by means of IS, and later quickly specialized to each of the individual problems. This last stage could be achieved, for instance, by appending additional ad-hoc layers [233, 234]. Such are examples of the possibilities offered by the sharing scheme.

In any case, the methodology proposed here addresses a central challenge in quantum technology, namely data reusability. Quantum resources are, and will remain in a foreseeable future, scarce and expensive. Still, in most situations quantum measurements cannot be recycled. Here we have taken a step towards more systematic reuse of measurements, and have shown how this can provide concrete benefits for the optimization of many VQAs. Note, however, that the approach presented here relies on the ability to communicate measurements, but does not attempt to learn how optimal circuit parameters relate to the details α of the different problems. Learning directly such parameter–problem dependency, rather than treating the various tasks independently, could also provide a distinct path towards more efficient optimization over many problems. This aspect is studied further in the following section, in the context of the optimal control of many quantum operations.

5.2 Optimal control of continuous families of gates

To date, QOC has been successfully applied to a wide variety of quantum platforms, each time focusing on the improvement of a single, or a few, specific operations. As quantum technologies mature it becomes equally important to also enlarge the range of operations which can be accurately performed by the same platform. For the sake of concreteness, this is now further discussed in the context of – but not limited to – quantum computing platforms.

Current quantum computers implement a restricted set of elementary operations, typically single-qubit gates and fixed two-qubit entangling gates, which have been carefully optimized for. This small set of operations already permits the execution of arbitrary quantum circuits [235], but requires such circuits to be decomposed in terms of often lengthy sequences of elementary gates. Given the current lack of error correction, errors accrue at each step of the decomposition, quickly rendering the outcomes of the computation noisy and impractical. Augmenting the set of elementary gates which can be composed would lead to shorter decompositions, and thus would result in more faithful computations. Already, it was shown that the inclusion of specific continuous families of two-qubit gates entails significant gate-count reductions in popular quantum algorithms [226–228]. More generally, the ability to optimize a broad range of operations has the potential to substantially increase the utility of current quantum hardware.

While one would have hoped the QOC methodology to accompany and to facilitate the optimization of large set of operations, it remains the case that QOC only considers the realization of a single target operation at a time, but is not naturally suited for the optimization of families of related targets. A straightforward path towards learning to control more operations, is to repeat independent QOC optimizations many times. However, this approach fails in recognizing that similar targets could be achieved with similar controls, and quickly becomes impractical for large, and potentially continuous, families of operations. In cases when the optimal controls are expected to vary smoothly with changes in the target, one would aim at exploiting such regularities to devise more effective control frameworks.

In the following, we lift the original scope of QOC from the control of a single quantum gate to the control of a continuous family of gates. This new objective is formalized in Section 5.2.1, where it is contrasted to a traditional problem of QOC. Then, in order to efficiently address this extended problem, we propose a framework aiming at learning how to adapt the controls as a function of the target to be realized. This is achieved by introducing a neural network modeling this target-control dependency, and we detail its training. Next, in Section 5.2.2, we exemplify the benefit of this approach in the context of quantum gate engineering for families of 1- to 3-qubit quantum gates. In particular, it is verified that the framework proposed enables the accurate (analog) control of continuous families of multi-qubit gates, with the benefit of reducing the overall implementation times compared to the time that would have been entailed by a sequential gate decomposition. In Section 5.2.3, we conclude by pointing towards several extensions of the framework. In particular, we detail how it could be trained directly based on experimental data, as opposed to the exact numerical simulations which are assumed here.

5.2.1 Methodology

As was presented in Section 1.1, the central objective in a traditional QOC problem of quantum gate engineering is the identification of the time-dependent controls $\mathbf{f}(t)$ that induce a propagator $U(t = T)$ mimicking the action of a target gate U^{tgt} at time T . Given a measure of similarity $S(U^{tgt}, U(T))$ between the target and realized gates, such aim is formulated in terms of the figure of merit $F(\mathbf{f}(t)) = S(U^{tgt}, U(T))$ that has to be maximized. A common example would be the task of realizing a two-qubit controlled NOT gate $U^{tgt} = |0\rangle\langle 0| \otimes I + |1\rangle\langle 1| \otimes X$.

Instead, a QOC problem of continuous families of target gates aims at the accurate realization of the family of gates $U_\alpha^{tgt} = |0\rangle\langle 0| \otimes I - i|1\rangle\langle 1| \otimes \exp(i\alpha X)$ for a continuous angle α . In such situation, the problem to be solved becomes the identification of a continuum of controls $\{\mathbf{f}_\alpha(t)\}$, indexed by the angle α , such that any of the propagators $U_\alpha(t)$ induced by $\mathbf{f}_\alpha(t)$ approximates

its corresponding target U_α^{tgt} , as accurately as possible, at $t = T$. Given the individual figures

$$F_\alpha(\mathbf{f}_\alpha(t)) = S(U_\alpha^{tgt}, U_\alpha(T)) \quad (5.4)$$

to be maximized for each target angle α , the (global) figure of merit for the continuous problem becomes the average

$$F(\{\mathbf{f}_\alpha(t)\}) = \langle F_\alpha(\mathbf{f}_\alpha(t)) \rangle_\alpha, \quad (5.5)$$

taken over α .

Such continuous problem could be addressed in terms of (i) many control tasks to be solved separately over a finite number of values $\{\alpha_i\}$, combined with an additional step of (ii) interpolation for any angle $\alpha' \notin \{\alpha_i\}$. Already, the choice of an appropriate discretization of the values of α and of the interpolation scheme to be employed is not straightforward. On one hand, given that the most adequate interpolation scheme is unknown, any coarse grained discretization of the values of α would likely result in rough approximations of the interpolated controls. Too fine discretization, on the other hand, quickly entails an excessive effort of solving many independent tasks which becomes prohibitive for dimension of α greater than one. Even more fundamentally, there is no guarantee for the controls \mathbf{f}_{α_1} and \mathbf{f}_{α_2} , found as optimal for similar values of α_1 and α_2 , to be themselves similar. In general, many solutions of the same control task could exist, making any attempt to interpolate in between \mathbf{f}_{α_1} and \mathbf{f}_{α_2} inconsistent. In order to avoid such issues, it is highly desirable to require that the optimized controls change smoothly with α .

This requirement can be realized in terms of a continuous parametrization of the dependence of $\mathbf{f}_\alpha(t)$ with respect to α . Resorting to an adequate parametrization could ensure well behaved changes in the controls over the whole family of targets, thus avoiding problems of erratic interpolations. Furthermore, as we will see, this approach permits a distribution of the training effort over many more values of α_i compared to what is permitted by the naive approach described in the previous paragraph. Ultimately, the continuum of control solutions $\{\mathbf{f}_\alpha(t)\}$ will depend on the choice of the parametrization adopted, and there would always be a compromise between the requirement of smoothness in α and the aim of finding good solutions. Since neural networks¹ (NN) provide a convenient way to approximate continuous functions up to any desired precision [240], these are deemed ideally suited for our purpose. In the following we discuss more precisely the structure of the NNs considered and their training.

¹Reviewing the working details of neural networks is out of the scope of this thesis, as many resources are now available. For instance, the reader is referred to [236, 237] for a comprehensive treatment of neural networks, or to [238, 239] for a more pedagogical introduction geared towards physicists.

Controls conditioned on target details by means of neural networks

An appropriate NN for our purpose takes parameters α and time t as input, and produces controls $\mathbf{f}_\alpha(t)$ as output. Hence, these output control values depend both on the time and on the details α of the target U_α^{tgt} to be realized. Such NN is sketched in Figure 5.3(a) for the case of 2–dimension target parameters α (that is, a total of $2 + 1 = 3$ inputs) and a single control function $f_\alpha(t)$ to be learnt (that is, a single output). This can be adapted to arbitrary problem dimensions by changing the dimensions of the inputs and outputs of the NN.

A fixed choice of NN weights² ϕ_{NN} entirely specifies a continuum of functions $\{\mathbf{f}_\alpha(t)\}$, with any individual set of functions $\mathbf{f}_\alpha(t)$ corresponding to a fixed input α , but varied input $t \in [0, T]$ of the NN. In turn, $\mathbf{f}_\alpha(t)$ permits the construction of the system Hamiltonian $H_\alpha(t)$. Given this (time–dependent) Hamiltonian, the system can be evolved such as to obtain its propagator $U_\alpha = U_\alpha(t = T)$ at final time $t = T$. Finally, comparing this propagator to its corresponding target U_α^{tgt} yields the figure of merit F_α (Equation (5.4)). When numerically performed, the step of evolution (dashed lines in Figure 5.3(a)), requires to solve numerically an ordinary differential equation (ODE), typically the Schrödinger equation in Equation (1.1), governing the dynamics of the system.

This evaluation of the figure of merit F_α discussed for a single α can be repeated for any value of α . Overall, the optimal control problem of continuous targets amounts to training the NN weights to produce controls maximizing the figure of merit F in Equation (5.5), which is the average of F_α over α .

Training the framework by gradient descent

The optimization of the NN (that is, the training of the framework) could be achieved with a variety of techniques, but gradient-descent training has the advantage of simplicity and scalability to high-dimension parameter space. As depicted in Figure 5.3(b), evaluations of the gradients $\nabla_{\phi_{NN}} F$ of the figure of merit with respect to the weights of the NN, require to obtain and combine gradients (illustrated with green arrows) over each of the computational steps involved. These include the evolution of the system and the NN processing. To facilitate such tedious task, we resort to auto–differentiation tools (discussed in Section 2.1.1), which allow us to automatically keep track and to compose gradients. While automatic differentiation over NNs is readily available in any machine–learning framework, the ability to efficiently differentiate over numerical ODE solver remains less widely spread. For that specific purpose, we resort to techniques presented recently in [135], with implementation in [241], built on top of the popular machine learning library Pytorch [127]. Equipped with this machinery, it is only

²These weights are the free parameters of the NN, which can be trained.

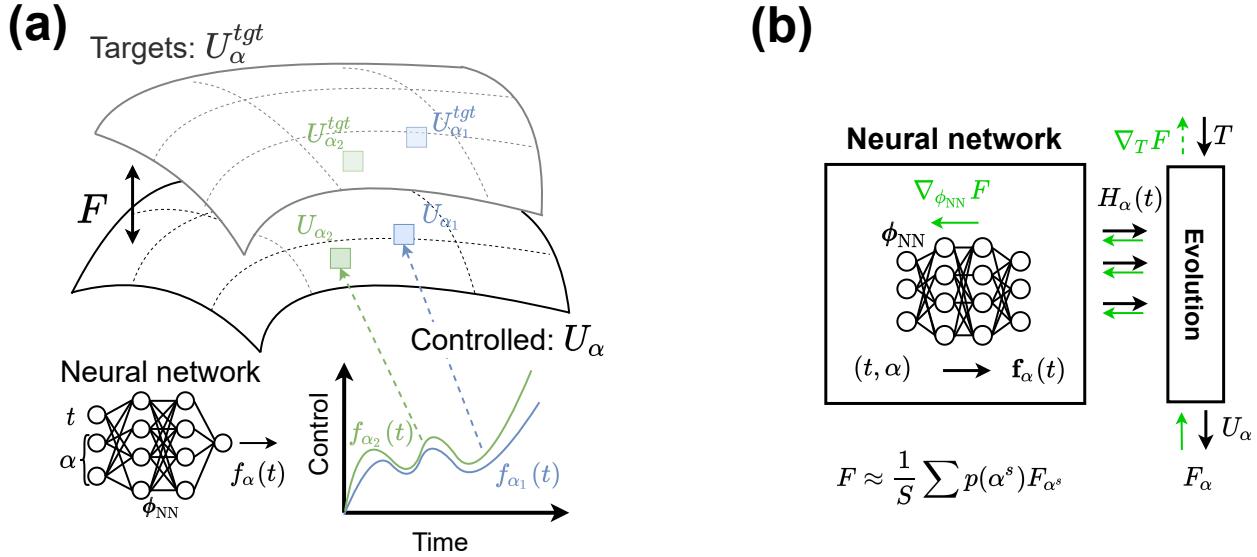


Figure 5.3: (a) Quantum optimal control of a continuous family of targets U_α^{tgt} , indexed by the target variable α (which can be a scalar or a vector). We propose to model the controls $\mathbf{f}_\alpha(t)$, which depends both on the time t and on the variable α , by a neural network (NN). This NN effectively parameterizes a continuous family of controlled operations U_α where each element corresponds to the propagator obtained by evolving the system in time (dashed arrows) under the controls produced by the NN. Accurate control is realized when the similarity F between the controlled and target families is maximized. For that purpose, the weights ϕ_{NN} of the NN are trained by gradient descent. (b) Given a fixed value of α but variable t , the NN produces the controls $\mathbf{f}_\alpha(t)$, permitting the construction of the time-dependent Hamiltonian $H_\alpha(t)$ used to evolve the system. This evolution is performed by numerically solving the Schrödinger equation, yielding the propagator U_α realized by the controlled system at time T . Finally, comparison of this propagator U_α to its corresponding target U_α^{tgt} permits the evaluation of the individual figure of merit F_α . Gradients of this figure can be back-propagated (green arrows) to the weights ϕ_{NN} of the NN (and optionally to the time T , if treated as variable). At each iteration, several values of α are randomly sampled, and the evaluation of F_α (and its gradients) are averaged accordingly (as seen in Equation (5.6)).

required to define how the system is propagated and how the figure of merit is evaluated, but, it is left to the underlying numerical libraries to deal with the intricacies of gradient computation.

Then, each step of optimization corresponds to an update of the weights ϕ_{NN} of the NN performed by gradient descent. Even though the evaluation of the averaged figure of merit F , and its gradients, would always be based on an average over a finite number n_B of values of α such as

$$F(\{\mathbf{f}_\alpha(t)\}) \approx \frac{1}{n_B} \sum_{s=1}^{n_B} F_{\alpha^s}(\mathbf{f}_{\alpha^s}(t)), \quad (5.6)$$

rather than the proper integral in Equation (5.5), the outputs of the neural network are still continuous in α , and systematically choosing different random sampling points α^s at each iteration minimizes artifacts resulting from finite sampling.

Such training of the NN performed with values of α^s that are resampled at each step of the optimization, is in contrast with the approach of fully solving control problems for specific values of α^s (and subsequently interpolating the control solutions) that was discussed before. At fixed numerical effort, the former will be trained on much more diverse values of α than the latter. In addition to the general methodology that has been presented, we now discuss more explicit details of the training and of the choice of the NNs employed.

Choice the hyperparameters

It is known that the quality of the solutions found by a NN can be affected by many design choices commonly referred as hyperparameters [242]. In particular, the choice of the size of the NN, the way its weights are initialized, and the learning rates employed during its optimization are often considered as of particular importance.

As the optimal hyperparameters values are not known beforehand, they need to be identified by means of trial-and-error. For instance, it is common to randomly sample [243] many hyperparameters configurations and to select a posteriori the ones resulting in optimal training. As each of these random configurations is evaluated independently, this approach has the benefit of being straightforward to parallelize and we will follow this methodology. However, in order to reduce the overall computational effort, we base this search on only partial training of the NNs, consisting on a reduced number of training steps in the range of 10 to 50 iterations. This already permits the identification of the most promising hyperparameters values. Once identified, these configurations are reused to perform a complete training of the framework. Explicit details of the range of hyperparameters which are explored are provided in the next section.

5.2.2 Results for problems of quantum gate engineering

In the following we study the control of a n -qubit system evolving under the prototypical Hamiltonian

$$H_n(t) = \sum_{i < j=1}^n f_{XX}^{(i,j)}(t) X_i X_j + \sum_{i=1}^n [f_Y^{(i)}(t) Y_i + f_Z^{(i)}(t) Z_i], \quad (5.7)$$

with tuneable XX interactions between each pair of qubits, local controls over Y and Z Pauli terms and amplitudes systematically bounded in $f \in [-1, 1]$. Such a choice of Hamiltonian allows for the controllability of the system [244], that is, ensures that any unitary U^{tgt} could be realized provided sufficient time T . Still, finding how to realize an arbitrary unitary, without decomposing it in terms of elementary gates, remains non-trivial and require resorting to the QOC machinery. For the system sizes of $n = 1, 2$ and 3 qubits that we investigate, this will necessitate learning a total of $C = 2, 5$ and to 9 time-dependent (and target-dependent) control functions respectively.

The measure of similarity S between realized and target unitaries is taken to be the process fidelity³, which was defined for general processes in Equation (4.25). In the case of unitaries U and V it simplifies to

$$S(U, V) = \mathcal{F}_{pro}(U, V) = \left| \frac{1}{2^n} \text{Tr}[U^\dagger V] \right|^2. \quad (5.8)$$

This measure is used when evaluating the individual figures of merit F_α in Equation (5.4) both during the training and testing of the framework. Global and individual errors are defined as $E = 1 - F$ and $E_\alpha = 1 - F_\alpha$ respectively, with low values indicating accurate control, and with the minimal value of $E = 0$ reached when any of the unitaries belonging to the continuous family of targets is exactly realized.

For all the control problems that we consider, we perform a search of good hyperparameters over the same domain which is now detailed. The NNs considered are taken to contain in between 4 and 10 layers with, in between 100 to 300 nodes per layer. The inner layers are systematically chosen to be ReLU layers [245], which are known to limit issues of gradient vanishing when training NNs. However the output layer is taken to be a sigmoid layer. Compared to the values ranging from 0 to infinity produced by a ReLU function, a sigmoid function ensures bounded output in $[0, 1]$ which can further be shifted and rescaled in order to produce control amplitudes bounded in the desired range $[-1, 1]$. Optimization of the weights of the NNs is performed using Adam [121], with learning rates l in the range $[10^{-4}, 10^{-2}]$, which is a common choice when training NNs. Additionally, we find the amplitudes of the initial weights of the NN to be critical to the successful training of the framework. In particular, we notice that the default distribution [246] of the initial weights results in small initial control amplitudes,

³The 0-fidelity defined in Section 4.2.3 could also have been used. However, as its benefit was shown for experimental situations, rather than the numerical evolution of the system employed here, we do not resort to this alternative.

and which are often found to quickly converge to 0. To avoid such effect, we rescale the initial weights by a factor $\beta \in [1.8, 2.2]$ such that the initial control values produced by the NNs are different enough from 0. With this explicit formulation of the problems to be tackled, and of the training to be performed, we can now proceed to discuss the results obtained with the NN control framework.

Optimal controls for arbitrary single-qubit rotations

To gain insights in the working of the framework we start by applying it to the Hamiltonian in Equation (5.7), in the case when $n = 1$ qubit, and with the aim of controlling the manifold of single-qubit rotations defined as

$$U_{1q}(\alpha = [\alpha_1, \alpha_2, \alpha_3]) = R_Z(\alpha_1)R_Y(\alpha_2)R_Z(\alpha_3), \quad (5.9)$$

in terms of the Pauli rotations $R_{\sigma \in \{X, Y, Z\}}(\alpha) = e^{-i\alpha\sigma/2}$, and with target parameters α taken in the domain $\chi = [0, \pi]^3$. Despite the choice of target representation appearing in Equation (5.9), we aim at realizing any of such rotations in a single analog control sequence, that is, without resorting to a decomposition in terms of atomic Pauli rotations. The framework is trained for $M = 400$ iterations, and each iteration involves averaging the process fidelity in Equation (5.8) over $n_B = 128$ control problems (corresponding to values of α randomly sampled).

To assess the success of the NN, after its training we apply it to the control of new target rotations. To this intent, new values of α are sampled and fed to the NN which, in turn, produces the control functions to be employed. As an example, these controls are plotted in Figure 5.4(a) for 4 control tasks. Given these controls, the system is propagated in time such that errors between the realized and target gates can be evaluated. These errors are also reported in Figure 5.4(a). As can be seen, in any of these four examples the NN succeeds in producing controls tailored to the individual targets resulting in low errors $E_\alpha \leq 10^{-4}$.

A more thorough assessment of the performances of the framework requires repeating such study over a larger ensemble of targets. For 250 rotations, we find an average (and standard deviation denoted in square brackets) error of $E = 2[3] \times 10^{-4}$. With this more exhaustive study, we can conclude that despite having been trained only on a finite number of tasks, the NN is able to accurately control any of the rotations belonging to the continuous family of gates U_{1q} .

To visualize further the dependency in between the controls and the targets we report as a 3-d plot in Figure 5.4(b) the control functions obtained for varied $\alpha_3 \in [0, \pi]$ and fixed $\alpha_1 = \alpha_2 = 3\pi/4$. As can be seen, both the control values $f_Y(t)$ and $f_Z(t)$ change smoothly over the whole range of targets and times considered. Similar results are found for varied α_1 and α_2 .

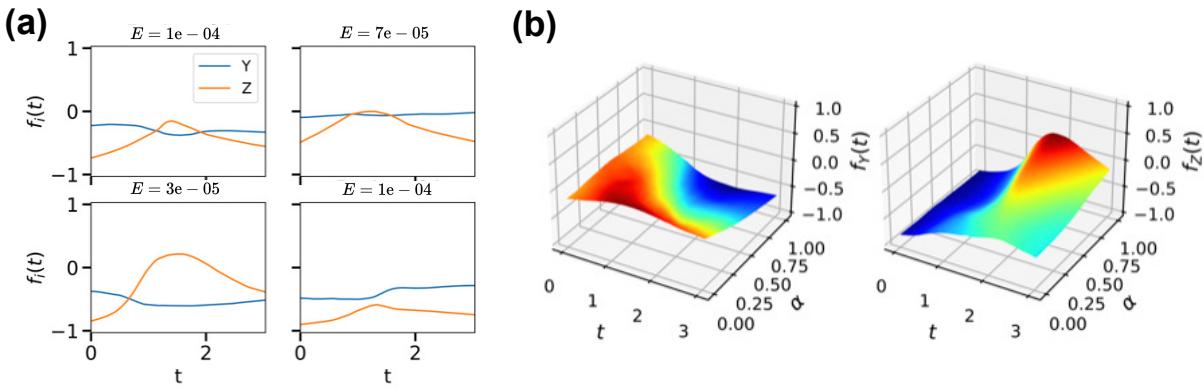


Figure 5.4: Problem of optimal control of the continuous family of single-qubit rotations $U_{1q}(\alpha_1, \alpha_2, \alpha_3)$ defined in Equation (5.9), with the Hamiltonian in Equation (5.7). After training, the framework is tested on new control problems (corresponding to parameters α_i randomly sampled). (a) Results for 4 test problems. The values of the amplitudes $f_Y(t)$ (blue) and $f_Z(t)$ (orange) produced by the NN are plotted, and the corresponding errors E_α are reported for each of these tasks. In all cases, the NN has learnt to produce tailored control resulting in low errors. (b) To visualize the dependency between the controls learnt and the target details, $\alpha_1 = \alpha_2 = 3\pi/4$ are fixed and $\alpha_3 = \alpha\pi$ is varied in $\alpha \in [0, 1]$. The amplitudes of the two control fields f_Y (left panel), and f_Z (right panel) vary continuously with the time t and target variable α .

but are not displayed here. This ability to smoothly adapt the controls to the specific targets is key to the accurate control of, not only a finite set of rotations, but rather of the whole continuous family of targets. In addition to this ability to learn how to control continuous targets in fixed time, we now show how the framework can also be extended to learn the final time to be adopted, or even to learn such final time taylored to each of the individual rotations.

Optimization of both the controls and times for arbitrary single-qubit rotations

In the previous example the final time, $T = \pi$, was taken large enough to ensure the accurate realization of any of the target rotations. Defining a minimum final time T_{min} ensuring that any of the control targets can be realized accurately is in general non-trivial, and it is common to assess it empirically by means of repeated optimizations for varied values of T . Here, this extra effort can be avoided by leveraging the fact that the derivatives $\nabla_T F$ can also be evaluated (Appendix.C [135]). Hence, by treating the final time T as a variable, one can optimize jointly the weights of the NN and the time, to maximize a new figure of merit

$$F'(\{\mathbf{f}_\alpha(t)\}, T) = F(\{\mathbf{f}_\alpha(t)\}) - \lambda_T T, \quad (5.10)$$

which consists of the original figure of merit F , and an additional term $-\lambda_T T$, with $\lambda_T > 0$, which penalizes long control durations T . Values of λ_T can be varied to put more or less emphasize on the objective of accurate control or on the objective of fast operations. We find

the value of $\lambda_T = 10^{-2}$ to be a good compromise, and use this value in the following.

Extending the framework even further, it is also possible to learn target-dependent final times T_α . This is achieved by introducing a second NN, with input α and output T_α . This time T_α is used when evolving the system (for the control of U_α^{tgt}) and now depends on the target to be controlled. The two NNs, one producing the control values and the other one the final times, can be trained in parallel to maximize the figure of merit

$$F''(\{\mathbf{f}_\alpha(t)\}, T) = F(\{\mathbf{f}_\alpha(t)\}) - \lambda_T \langle T_\alpha \rangle_\alpha, \quad (5.11)$$

which is similar to Equation (5.10), with the exception that now the final times T_α varying with α need to be averaged.

Given that the controls produced by the (original) NN should change depending on the control duration T (or T_α), we extend it to also take T (or T_α) as an input. Note that this additional input is only required when the control times vary during the course of the optimization, that is, this refinement is not needed when T is fixed (as was the case in the previous set of results presented). In the following, we revisit the control problem of realizing continuous single-qubit rotations U_{1q} , but now aim at learning both the optimal controls and the optimal times which maximize Equation (5.11).

As before, after 400 steps of training the framework is tested on 250 new target rotations. The average of the testing errors is found to be $E = 4[5] \times 10^{-4}$ and the averaged control times, which now depend on the targeted rotations, is found to be $\bar{T} = 1.66[0.47]$. That is, both the control functions and times were successfully learnt, resulting in a substantial decrease in the average time taken per control task (roughly half of the fixed time that was considered earlier) at the cost of only a modest increase in the errors.

To dissect further these results, we restrict our attention to a 2-dimension subset of rotations, corresponding to a fixed value of $\alpha_3 = 3\pi/4$, which allows for visual examination. Values of α_1 and α_2 are discretized over a 75×75 grid. In Fig. 5.5(a) we plot as a heat-map the testing errors found for each of the $75 \times 75 = 5625$ corresponding rotations. As can be seen, most of the errors are of the order of 10^{-4} or less, except for a few outliers lying at the edge of the domain with errors around 10^{-3} . The control times produced by the framework are plotted in Fig. 5.5 (b). A minimum value of $T \approx 3\pi/8$ is obtained for the target rotation $R_Z(3\pi/4)$, which corresponds to $\alpha_1 = \alpha_2 = 0$ (bottom left of the heat-map). Such rotation can indeed be realized with constant values of $f_Z(t) = 1$ and $f_Y(t) = 0$ applied for a duration of $T = 3\pi/8$. For non-null values of α_1 and α_2 , these times are found to be increasing with values of α_i as would have been expected.

In Figure 5.5(c) we report the relative improvement in time achieved by the controlled rotations compared to the time T_α^{XYZ} that would have been entailed by the Pauli decomposition

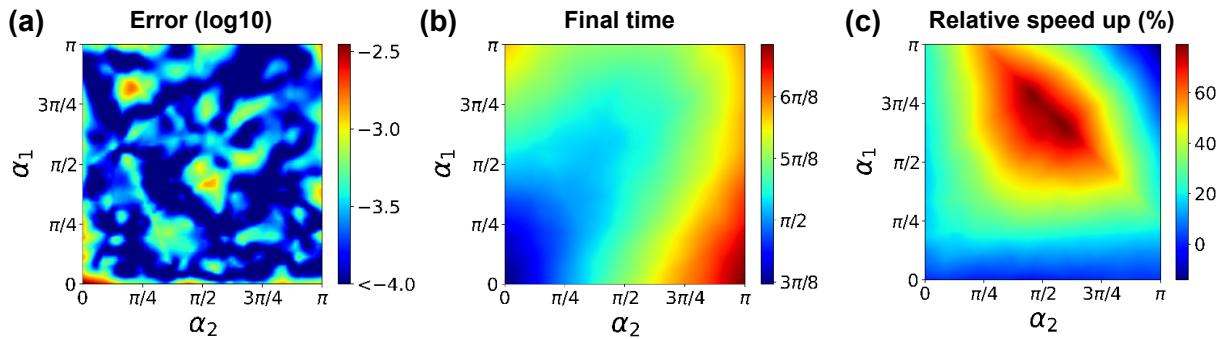


Figure 5.5: Similar problem as depicted in Figure 5.4, except that now both the controls and final times are learnt jointly to minimize Equation (5.11). For the sake of visualization only a 2-d subset of the 3-d family of target rotations is shown, with $\alpha_3 = 3\pi/4$, and values of $\alpha_1, \alpha_2 \in [0, \pi]$ discretized over a grid of 75×75 regularly spaced points. For each of these 5625 targets, a heat-map indicates the values of the individual control errors (a, in log scale), the control times predicted by the framework (b), and the relative speed up (c) entailed by the controls learnt compared to a decomposition of the target gates in terms of Y and Z Pauli rotations as appearing in Equation (5.9).

appearing in Equation (5.9). $T_\alpha^{ZY\bar{Z}}$ is assessed assuming that each of the rotations $R_{\sigma \in \{Y, Z\}}(\pm\theta)$ is performed in time $\theta/2$, which is the fastest that can be achieved under the Hamiltonian in Equation (5.7) with control amplitudes f bounded in the range $[-1, 1]$ considered here. With T_α being the controlled times learnt by the NN, the speed up reported in Figure 5.5(c) is defined as $(T_\alpha^{ZY\bar{Z}} - T_\alpha)/T_\alpha$. In almost all cases, the dynamics learnt by the NN is found to be the fastest, up to a speed up of 80%. Such results highlight the advantage of analog control compared to a decomposition in terms of a limited set of elementary gates.

Control of continuous families of 2- and 3-qubit gates

Having presented the foundations and extensions of the framework for single-qubit gates, we now apply it to larger control problems, with $n = 2$ and 3 qubits. A summary of the results obtained is presented in Table 5.1 detailing the families of targets learnt, the average testing errors achieved and the hyperparameters which were used. For all these problems, the time T is taken to be trainable but shared among different targets belonging to the same family, and is also reported. As before, the training of the NN consists of $M = 400$ iterations, with $n_B = 128$ targets sampled at each iteration, and the statistics reported are systematically evaluated over 250 new targets randomly generated.

For the system sizes of $n = 2$ qubits, a total of 4 families of targets are studied. These include the controlled-Z gate $CZ(\alpha_1)$ and the unitary $e^{-i\alpha_1 ZZ}$ which are both central to popular quantum algorithms [9, 21]. Additionally, we also investigate the control of higher-dimension families of targets such as the unitaries $e^{-i(\alpha_1 XX + \alpha_2 YY + \alpha_3 ZZ)}$. This family already encompasses all possible 2-qubit rotations up to local rotations [247], and is representative of unitaries commonly

Table 5.1: Control of the Hamiltonian in Eq. 5.7, to realize many continuous families of $n = 2$ – and 3–qubits gates (corresponding to $C = 5$ and 9 control functions to be learnt respectively). For each problem the family of target U_α^{tgt} and the domain χ of the target parameters α are reported. After 400 steps of training, the average E (and standard deviation in square brackets) of the errors is evaluated over 250 randomly sampled elements of the target families. Additionally the times T learnt (in unit of π), and the details of the hyperparameters used (with the number n_L of layers, the number n_U of units, and the learning rate l) are reported.

U_α^{tgt}	χ	n	E	T/π	$n_L \times n_U$	$l(10^{-3})$
$\text{CZ}(\alpha_1) = 0\rangle\langle 0 \otimes I + 1\rangle\langle 1 \otimes R_Z(\alpha_1)$	$[0, \pi]$	2	0.01[1]%	1.0	6×150	1
$e^{-i\alpha_1 ZZ}$	$[0, \frac{\pi}{2}]$	2	0.00[0]%	1.0	6×150	1
$e^{-i(\alpha_1 XX + \alpha_2 YY + \alpha_3 ZZ)}$	$[0, \frac{\pi}{2}]^3$	2	0.04[4]%	1.7	8×300	3
$CU_{1q}(\alpha_1, \alpha_2, \alpha_3)$	$[0, \pi]^3$	2	0.03[4]%	1.0	8×300	1
$e^{-i\alpha_1 ZZZ}$	$[0, \frac{\pi}{2}]$	3	0.01[0]%	1.0	10×200	4
$e^{-i(\alpha_1 XXX + \alpha_2 YYY + \alpha_3 ZZZ)}$	$[0, \frac{\pi}{2}]^3$	3	0.30[25]%	2.0	10×300	2
$CCU_{1q}(\alpha_1, \alpha_2, \alpha_3)$	$[0, \frac{\pi}{2}]^3$	3	0.17[8]%	1.5	10×300	4

employed in digital quantum simulation [248–251]. These unitaries are typically implemented by successive applications of their individual components $e^{-i(\alpha\sigma\sigma)}$, which themselves need to be decomposed further in terms of elementary gates. Finally, we also investigate the case of the controlled-unitary $CU_{1q}(\alpha) = |1\rangle\langle 1| \otimes U_{1q}(\alpha) + |0\rangle\langle 0| \otimes I$, with the single-qubit arbitrary rotation U_{1q} defined earlier.

As seen in Table 5.1, the framework was successfully trained for any of these problems, resulting in averaged testing errors E systematically smaller than 5×10^{-4} , and with the largest values of $E = 3$ and 4×10^{-4} obtained for the 3–dimension families of targets studied. Additionally, we can compare the times T learnt by the framework to the times T_{dec} entailed by gate decomposition. The circuit decompositions permitting to assess T_{dec} are obtained using Qiskit [252] routines, assuming a set of elementary gates composed of single–qubit Y and Z rotations and 2–qubit XX rotations (as permitted by the Hamiltonian in Equation (5.7)). Given such methodology, we find the NN controlled gates to be 1.25 to 2.25 faster, in average, than their corresponding decompositions.

Going further, we also investigate the control of unitaries generated by more than two–qubit interactions [253–255]. These can benefit the simulation of exotic phase of matter, the preparation of topological states and the realization of quantum error correction protocols. For system sizes of $n = 3$ qubits a total of 3 families of targets are studied, which results can be seen in the last 3 lines of Table 5.1.

For gates generated by an effective αZZZ interaction term, low control error of $E = 1 \times 10^{-4}$ is achieved by the framework. For the more general unitaries of the form $e^{-i(\alpha_1 XXX + \alpha_2 YYY + \alpha_3 ZZZ)}$ and the control–control–unitaries $CCU_{1q}(\alpha) = |11\rangle\langle 11| \otimes U_{1q}(\alpha) + (I - |11\rangle\langle 11|) \otimes I$, we note an increase in the testing errors up to $E = 3 \times 10^{-3}$ and $E = 1.7 \times 10^{-3}$ respectively. For

the fixed number of training steps adopted here, these increased errors are to be expected as they correspond to the most challenging problems that are addressed, necessitating to learn 9 time-dependent functions in order to control 3-dimension families of targets. Nonetheless, we discuss in the next section how these errors can be further reduced.

As before, we also compare the times T learnt by the framework and the times T_{dec} entailed by gate decomposition. For these families of 3-qubit gates, the controlled gates are found to be at least twice faster than their decompositions, and up to 10 times faster for the case of the control-control-unitaries CCU_{1q} . It should be highlighted that the decomposition used to assess T_{dec} , based on Qiskit routines, may not be optimal. However, we leave the identification of potentially better gate decompositions to future work.

Overall, given the test errors systematically found in the sub-percent regime for the 7 families of gates studied, and the reduced control times which were reported, we can see that fast and accurate control over (multidimensional) continuous families of targets can indeed be achieved by means of the framework that was presented. We now discuss how its training can be further refined.

Improving the training of the NN controller

In all the cases studied, the largest errors are found to be close to the edge of the domain χ (as can already be seen in Fig. 5.5(a)), that is, where less neighbouring training points are to be found. A simple strategy to avoid such boundary effects consists in training the NN on a larger domain $\chi' \supset \chi$ but restrict its use to its initial domain χ after training.

Furthermore, by inspecting (a posteriori) the controls learnt by the NN framework, one can notice patterns which can be exploited. For instance, for the control task of the CCU_{1q} family of gates we observe that (i) the amplitudes of the Z terms of the two control qubits are fixed to zero, (ii) the amplitudes of the Y terms of the control qubits adopt identical values, and (iii) the amplitudes of the XX interaction terms in between each of the two pairs of control-target qubits are identical. Pruning the redundant degrees of freedom (corresponding to observations i-iii) permits the reduction of the effective number of control functions that has to be learnt, and ease the training of the framework. For the example of the CCU_{1q} family discussed, this results in a reduced number of 5 independent control functions to be learnt (after pruning), compared to the 9 originally considered.

With this more advanced training strategy – taking into consideration both the use of a larger parameter domain χ' during training, and the pruning of the redundant control degrees of freedom – we achieve average test errors of $E = 9[5] \times 10^{-4}$ and $E = 6[2] \times 10^{-4}$ for the families of $e^{-i(\alpha_1 XXX + \alpha_2 YYY + \alpha_3 ZZZ)}$ and CCU_{1q} unitaries respectively, that is, a three-fold improvement

compared to the results reported in Table 5.1.

Even further improvement is expected with increased training effort, or faster identification of optimal hyperparameters. In all the examples detailed above, the training was limited to 400 iterations and to a maximal wall-time of 1 day (with CPU-only resources). Already, the use of GPU resources will permit to perform many more training steps in similar time, which is likely to yield reduced errors. Also, as noted earlier, among the hyperparameters studied we identified the initialization of the weights of the NN as affecting the most the training performances. Default weights were rescaled by a factor $\beta > 1$ to avoid collapse of control values to 0. More principled approaches towards weights initialization such as [256] could result in accelerated convergence. More generally, any of the recent proposals aiming at accelerating the training of NNs, for instance [257, 258], can be beneficial to the framework. All these will directly improve on the training of the framework, and in turn, could permit us to explore its applicability to even larger problems than the ones considered here.

5.2.3 Conclusion and outlook

In this section, we have extended the original scope of QOC, from the control of a single targeted operation to the control of continuous families of such operations. To this intent, we have proposed to model the dependency between targets to be realized and controls to be applied by means of a NN. Its training was demonstrated on families of popular quantum computing gates, resulting in low errors (systematically assessed on new control targets which were not seen during training) and improved operational speed. Ultimately, this ability to optimize the control of families of targets, can broaden the range of operations that could be performed accurately by a same device, and thus increase the utility of current quantum hardware.

We note that QOC schemes already relying on the use of NNs [31, 54, 55, 133, 134, 259], could directly benefit from the ideas introduced here. In particular, extending the inputs of the NNs involved (in these cited frameworks) to include the target parameters α would avoid inefficient optimizations over many targets treated independently. To overcome such inefficiencies, a heuristic strategy was proposed in [54, 55] which consists in reusing the control parameters optimized for one target as initial parameters for the optimization of the another target. Even with such strategy, it is reported in [55] a maximum wall-time of 5 hours per target rotation (which are 2-qubit gates parameterized by a single rotation angle α). For the 31 discrete values of the angle α that the authors consider, this results in more than 6 days of computations. Introducing a second(third) rotation angle with similar discretization would have resulted in 200 days(17 years) of computations. These figures assume a sequential treatment of each optimization problem (as was done in [55]) but could be, at least partially,

parallelized to reduce the overall wall-time. Nonetheless, these show the limits of independent optimization, even when supplemented with an improved initialization strategy, and should be contrasted with the maximum wall-time of 24 hours that we allowed for any of the families of gates which were studied.

For the sake of concreteness, we limited ourselves to demonstrations in the idealized situation specified by the system Hamiltonian in Equation (5.7). Still, the applicability of the framework presented does not limit to such cases. Resorting to auto-differentiation tools, as done here, ensures that, as long as the system dynamics can be numerically simulated, the framework can be trained. For instance, the inclusion of experimental limitations, such as finite ramp times or distortion of the controls, or the treatment of large system sizes by means of approximate evolution could all equally be coped with. Even the case where the dynamics of the system is described by a stochastic ordinary differential equation is amenable to such framework [260], thus extending its application to the control of some quantum systems with active feedback [261, 262].

Other interesting lines of research include the training of the framework directly based on experimental outcomes, thus avoiding the necessary steps of precise Hamiltonian characterisation and numerical simulations assumed here. For such purpose two options are conceivable. First, this could simply be achieved by resorting to non-gradient optimization. While BO could be used in principle, for problems of optimization over a large number of NN weights more local optimization strategies, such as explored in [263], may be favoured. Another option would be to restrict the time-dependency of the controls and the dynamics of the system to situations where gradients can be estimated experimentally (see for instance [137]). With the identification of such parametrization – which would typically involve a piece-wise constant parametrization of the control fields, as seen in Equation (1.6) – the NN displayed in Figure 5.3(a,b) would be adjusted to only take parameters α (but not the time) as input, and to output the values of the control parameters. For instance, in the case of a single control field parameterized with 10 piece-wise constant values, the output of the NN would now consist of 10 output nodes, corresponding to the 10 parameters needed to reconstruct the control field.

Taking a step further, we envision the task of learning how to accurately control a quantum experiment in varied *contexts*. Here, this context was in one to one correspondence with the target operation to be realized, but, could be more inclusive. For instance, the inputs α of the NN could incorporate details of the system under control (such as its energy splittings [54] or its size [76]) and also details of its environment (such as heating rates [264] or information about nearby operations resulting in cross-talk [64]). Provided that the effects of varying these context variables can be simulated and that the optimal controls are expected to vary continuously with these variables, one could learn, with the framework developed here, to accurately operate a quantum system in very broad context.

5.3 Concluding remarks

In this chapter, we have addressed the demanding situation where many optimizations need to be performed concurrently. This setup naturally arises in applications of VQAs to quantum chemistry and condensed-matter models, and also in the extended problems of QOC that were introduced. As such, the ability to efficiently optimize over many related tasks was identified as an important aspect, and potential roadblock, in the scalability of quantum technologies. Still, to date families of related optimization problems, have been predominantly approached independently, quickly limiting the reach of current methodologies. In contrast, we have devised two frameworks aiming at exploiting commonalities in between the different problems to be solved in order to accelerate the corresponding optimizations.

Given that quantum data is expensive and limited, it is of great interest to reuse it as much as possible. For that purpose, we have developed a framework based on BO permitting the reusability of measurement data and illustrated its benefit for families of VQAs. A second approach towards many optimizations was proposed in the context of QOC over families of targets. In cases where control solutions are expected to change continuously with details of the underlying targets, we show that it was more efficient to directly learn such control–target functional dependency. Adopting a typical machine–learning perspective, we modeled the unknown dependency by means of a NN, and demonstrated its training for tasks of quantum gate engineering, where low control errors over families of target gates were achieved. When direct comparison was possible, we found the frameworks to significantly outperform existing methodologies. While each framework was developed with a specific application in mind, the general underlying principles – of either sharing measurement data or learning unknown relationship between optimal parameters and details of a problem – which were laid down in this chapter could inspire many more applications, exceeding the examples tackled.

As novel, the methods proposed here are prone to further refinements, and several extensions were discussed along the way. For instance, combining the scheme of information sharing with the recently developed technique of classical shadows could greatly extend our ability to reuse quantum measurement data. For the tasks of QOC of continuous targets by means of a NN controller, it will be particularly appealing to validate its applicability in situations where optimizations are performed directly based on experimental data. Overall, we believe that the conceptual developments and concrete implementations detailed in this chapter could advance our ability to efficiently utilize quantum devices at larger scale than was previously possible.

Conclusion

In this thesis, we have sought to develop new methodologies for the efficient optimization of quantum dynamics based on experimental data. This broad objective was undertaken from three distinct, but complementary, perspectives.

In Chapter 3, we established BO as an appealing alternative to traditional optimization frameworks. Its merit was validated based on thorough numerical comparisons and experiments performed on a NISQ device. Furthermore, it was demonstrated that the original framework can be tailored to the peculiarities of quantum physics experiments, in order to enhance its ability to converge in minimal experimental effort. This advantage offered by BO comes at the expense of an increased computational effort, and we identified its application to complex experimental setups (such as the ones found in ultracold-atoms quantum simulations) as particularly promising. More generally, given the availability of cheap resources for classical computation, compared to the scarcity of near-term quantum resources, we expect that the methods presented will accompany and enable technological developments on a wide range of quantum physics platforms.

Complementary to the choice of an adequate optimizer, we argued in Chapter 4 for the consideration of alternative figures of merit. For that purpose, we established quantitative criteria enabling their study. These were formulated to be general enough such as to indicate the potential success, or failure, of any corresponding optimization. Based on these criteria, it was shown in distinct examples that new figures can be discovered and prove superior to the ones commonly employed. Taken together, these findings should encourage the design of new figures of merit as a compelling path toward more efficient optimization.

In Chapter 5, we tackled the demanding situation where many optimization problems need to be solved simultaneously. As was exemplified for the problems of VQAs and QOC that were investigated, we identified such situations (of simultaneous optimisation) as of particular interest to the practical use of near-term quantum computers. Still, with a handful number of exceptions, principled methodologies that allow families of optimisation problems to be efficiently solved are lacking. Rather than performing each optimization independently, as would typically be done, we devised two frameworks aiming to exploit commonalities in between the different tasks to be solved. When direct comparison was possible, these frameworks were

shown to substantially outperform existing approaches which quickly necessitate a prohibitive optimization effort.

Overall, we have addressed several key aspects affecting the success of optimizations based on experimental data. To this intent, we have developed novel methods inspired from recent advances in the field of ML, and principally probabilistic ML. In particular, we saw fruitful applications of BO and Gaussian processes to the realm of quantum technology. Given that the field of probabilistic ML is moving at fast pace and has remained, as yet, mostly unexplored, we expect the results found in this thesis to pave the way to further research. Ultimately, the methodologies presented should be employed – and would reveal their true potential – on large experimental realizations, exceeding the numerical simulations and the small-scale demonstrations on NISQ devices that were possible within the scope of this thesis.

Bibliography

- [1] IBM, “The ibm quantum experience,” <http://www.research.ibm.com/quantum>.
- [2] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. S. L. Brando, D. A. Buell, B. Burkett, Y. Chen, Z. Chen, B. Chiaro, R. Collins, W. Courtney, A. Dunsworth, E. Farhi, B. Foxen, A. Fowler, C. Gidney, M. Giustina, R. Graff, K. Guerin, S. Habegger, M. P. Harrigan, M. J. Hartmann, A. Ho, M. Hoffmann, T. Huang, T. S. Humble, S. V. Isakov, E. Jeffrey, Z. Jiang, D. Kafri, K. Kechedzhi, J. Kelly, P. V. Klimov, S. Knysh, A. Korotkov, F. Kostritsa, D. Landhuis, M. Lindmark, E. Lucero, D. Lyakh, S. Mandrà, J. R. McClean, M. McEwen, A. Megrant, X. Mi, K. Michelsen, M. Mohseni, J. Mutus, O. Naaman, M. Neeley, C. Neill, M. Y. Niu, E. Ostby, A. Petukhov, J. C. Platt, C. Quintana, E. G. Rieffel, P. Roushan, N. C. Rubin, D. Sank, K. J. Satzinger, V. Smelyanskiy, K. J. Sung, M. D. Trevithick, A. Vainsencher, B. Villalonga, T. White, Z. J. Yao, P. Yeh, A. Zalcman, H. Neven, and J. M. Martinis, “Quantum supremacy using a programmable superconducting processor,” *Nature* **574**, 505–510 (2019).
- [3] H.-S. Zhong, H. Wang, Y.-H. Deng, M.-C. Chen, L.-C. Peng, Y.-H. Luo, J. Qin, D. Wu, X. Ding, Y. Hu, P. Hu, X.-Y. Yang, W.-J. Zhang, H. Li, Y. Li, X. Jiang, L. Gan, G. Yang, L. You, Z. Wang, L. Li, N.-L. Liu, C.-Y. Lu, and J.-W. Pan, “Quantum computational advantage using photons,” *Science* **370**, 1460–1463 (2020).
- [4] Collaboration, “Enhanced sensitivity of the ligo gravitational wave detector by using squeezed states of light,” *Nature Photonics* **7**, 613–619 (2013).
- [5] B. J. Bloom, T. L. Nicholson, J. R. Williams, S. L. Campbell, M. Bishof, X. Zhang, W. Zhang, S. L. Bromley, and J. Ye, “An optical lattice clock with accuracy and stability at the 10^{-18} level,” *Nature* **506**, 71–75 (2014).
- [6] C. J. Kennedy, E. Oelker, J. M. Robinson, T. Bothwell, D. Kedar, W. R. Milner, G. E. Marti, A. Derevianko, and J. Ye, “Precision metrology meets cosmology: Improved constraints on ultralight dark matter from atom-cavity frequency comparisons,” *Phys. Rev. Lett.* **125**, 201302 (2020).

- [7] C. W. Chou, D. B. Hume, T. Rosenband, and D. J. Wineland, “Optical clocks and relativity,” *Science* **329**, 1630–1633 (2010).
- [8] S. Lloyd, “Universal quantum simulators,” *Science* **273**, 1073–1078 (1996).
- [9] P. Shor, “Algorithms for quantum computation: discrete logarithms and factoring,” in *Proceedings 35th Annual Symposium on Foundations of Computer Science* (1994) pp. 124–134.
- [10] A. M. Childs, R. Cleve, E. Deotto, E. Farhi, S. Gutmann, and D. A. Spielman, “Exponential algorithmic speedup by a quantum walk,” in *Proceedings of the Thirty-Fifth Annual ACM Symposium on Theory of Computing*, STOC ’03 (Association for Computing Machinery, New York, NY, USA, 2003) p. 59–68.
- [11] A. Y. Kitaev, “Quantum measurements and the abelian stabilizer problem,” arXiv preprint (1995).
- [12] L. K. Grover, “A fast quantum mechanical algorithm for database search,” in *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing* (1996) pp. 212–219.
- [13] A. W. Harrow, A. Hassidim, and S. Lloyd, “Quantum algorithm for linear systems of equations,” *Phys. Rev. Lett.* **103**, 150502 (2009).
- [14] J. Preskill, “Quantum Computing in the NISQ era and beyond,” *Quantum* **2**, 79 (2018).
- [15] D. Gottesman, “Stabilizer codes and quantum error correction,” arXiv preprint (1997).
- [16] D. A. Lidar and T. A. Brun, *Quantum error correction* (Cambridge university press, 2013).
- [17] J. Werschnik and E. K. U. Gross, “Quantum optimal control theory,” *Journal of Physics B: Atomic, Molecular and Optical Physics* **40**, R175–R211 (2007).
- [18] S. J. Glaser, U. Boscain, T. Calarco, C. P. Koch, W. Köckenberger, R. Kosloff, I. Kuprov, B. Luy, S. Schirmer, T. Schulte-Herbrüggen, D. Sugny, and F. K. Wilhelm, “Training schrödinger’s cat: quantum optimal control,” *The European Physical Journal D* **69**, 279 (2015).
- [19] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, “The theory of variational hybrid quantum-classical algorithms,” *New Journal of Physics* **18**, 023023 (2016).
- [20] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O’Brien, “A variational eigenvalue solver on a photonic quantum processor,” *Nature Communications* **5**, 4213 EP – (2014), article.

- [21] E. Farhi, J. Goldstone, and S. Gutmann, “A quantum approximate optimization algorithm,” arXiv preprint (2014).
- [22] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. J. Coles, “Variational Quantum Algorithms,” arXiv preprint (2020).
- [23] K. Bharti, A. Cervera-Lierta, T. H. Kyaw, T. Haug, S. Alperin-Lea, A. Anand, M. Degroote, H. Heimonen, J. S. Kottmann, T. Menke, *et al.*, “Noisy intermediate-scale quantum (nisq) algorithms,” arXiv preprint (2021).
- [24] G. A. Quantum, Collaborators, F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, S. Boixo, M. Broughton, B. B. Buckley, D. A. Buell, B. Burkett, N. Bushnell, Y. Chen, Z. Chen, B. Chiaro, R. Collins, W. Courtney, S. Demura, A. Dunsworth, E. Farhi, A. Fowler, B. Foxen, C. Gidney, M. Giustina, R. Graff, S. Habegger, M. P. Harrigan, A. Ho, S. Hong, T. Huang, W. J. Huggins, L. Ioffe, S. V. Isakov, E. Jeffrey, Z. Jiang, C. Jones, D. Kafri, K. Kechedzhi, J. Kelly, S. Kim, P. V. Klimov, A. Korotkov, F. Kostritsa, D. Landhuis, P. Laptev, M. Lindmark, E. Lucero, O. Martin, J. M. Martinis, J. R. McClean, M. McEwen, A. Megrant, X. Mi, M. Mohseni, W. Mruczkiewicz, J. Mutus, O. Naaman, M. Neeley, C. Neill, H. Neven, M. Y. Niu, T. E. O’Brien, E. Ostby, A. Petukhov, H. Putterman, C. Quintana, P. Roushan, N. C. Rubin, D. Sank, K. J. Satzinger, V. Smelyanskiy, D. Strain, K. J. Sung, M. Szalay, T. Y. Takeshita, A. Vainsencher, T. White, N. Wiebe, Z. J. Yao, P. Yeh, and A. Zalcman, “Hartree-fock on a superconducting qubit quantum computer,” *Science* **369**, 1084–1089 (2020).
- [25] G. Pagano, A. Bapat, P. Becker, K. S. Collins, A. De, P. W. Hess, H. B. Kaplan, A. Kyprianidis, W. L. Tan, C. Baldwin, L. T. Brady, A. Deshpande, F. Liu, S. Jordan, A. V. Gorshkov, and C. Monroe, “Quantum approximate optimization of the long-range ising model with a trapped-ion quantum simulator,” *Proceedings of the National Academy of Sciences* **117**, 25396–25401 (2020).
- [26] M. I. Jordan and T. M. Mitchell, “Machine learning: Trends, perspectives, and prospects,” *Science* **349**, 255–260 (2015).
- [27] Y. LeCun, Y. Bengio, and G. Hinton, “Deep learning,” *Nature* **521**, 436–444 (2015).
- [28] J. Carrasquilla and R. G. Melko, “Machine learning phases of matter,” *Nature Physics* **13**, 431–434 (2017).
- [29] G. Carleo and M. Troyer, “Solving the quantum many-body problem with artificial neural networks,” *Science* **355**, 602–606 (2017).
- [30] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, and G. Carleo, “Neural-network quantum state tomography,” *Nature Physics* **14**, 447–450 (2018).

- [31] M. Bukov, A. G. R. Day, D. Sels, P. Weinberg, A. Polkovnikov, and P. Mehta, “Reinforcement learning in different phases of quantum control,” *Phys. Rev. X* **8**, 031086 (2018).
- [32] T. Fösel, P. Tighineanu, T. Weiss, and F. Marquardt, “Reinforcement learning with neural networks for quantum feedback,” *Phys. Rev. X* **8**, 031084 (2018).
- [33] K. P. Murphy, *Machine learning: a probabilistic perspective* (MIT press, 2012).
- [34] Z. Ghahramani, “Probabilistic machine learning and artificial intelligence,” *Nature* **521**, 452–459 (2015).
- [35] C. E. Wieman, D. E. Pritchard, and D. J. Wineland, “Atom cooling, trapping, and quantum manipulation,” *Rev. Mod. Phys.* **71**, S253–S262 (1999).
- [36] A. Zeilinger, “Experiment and the foundations of quantum physics,” *Rev. Mod. Phys.* **71**, S288–S297 (1999).
- [37] Y. Makhlin, G. Schön, and A. Shnirman, “Quantum-state engineering with josephson-junction devices,” *Rev. Mod. Phys.* **73**, 357–400 (2001).
- [38] D. Leibfried, R. Blatt, C. Monroe, and D. Wineland, “Quantum dynamics of single trapped ions,” *Rev. Mod. Phys.* **75**, 281–324 (2003).
- [39] L. M. K. Vandersypen and I. L. Chuang, “Nmr techniques for quantum control and computation,” *Rev. Mod. Phys.* **76**, 1037–1069 (2005).
- [40] R. S. Judson and H. Rabitz, “Teaching lasers to control molecules,” *Phys. Rev. Lett.* **68**, 1500–1503 (1992).
- [41] R. Gordon and S. Rice, “Active control of the dynamic of atoms and molecules,” *Annual review of physical chemistry* **48**, 601–41 (1997).
- [42] C. J. Bardeen, V. V. Yakovlev, J. A. Squier, and K. R. Wilson, “Quantum control of population transfer in green fluorescent protein by using chirped femtosecond pulses,” *Journal of the American Chemical Society* **120**, 13023–13027 (1998).
- [43] A. Acín, I. Bloch, H. Buhrman, T. Calarco, C. Eichler, J. Eisert, D. Esteve, N. Gisin, S. J. Glaser, F. Jelezko, S. Kuhr, M. Lewenstein, M. F. Riedel, P. O. Schmidt, R. Thew, A. Wallraff, I. Walmsley, and F. K. Wilhelm, “The quantum technologies roadmap: a european community view,” *New Journal of Physics* **20**, 080201 (2018).
- [44] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition* (Cambridge University Press, 2010).

- [45] 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 J. M. Martinis, “Optimal quantum control using randomized benchmarking,” *Phys. Rev. Lett.* **112**, 240504 (2014).
- [46] S. Machnes, E. Assémat, D. Tannor, and F. K. Wilhelm, “Tunable, flexible, and efficient optimization of control pulses for practical qubits,” *Phys. Rev. Lett.* **120**, 150401 (2018).
- [47] K. Singer, U. Poschinger, M. Murphy, P. Ivanov, F. Ziesel, T. Calarco, and F. Schmidt-Kaler, “Colloquium: Trapped ions as quantum bits: Essential numerical tools,” *Rev. Mod. Phys.* **82**, 2609–2632 (2010).
- [48] T. Choi, S. Debnath, T. A. Manning, C. Figgatt, Z.-X. Gong, L.-M. Duan, and C. Monroe, “Optimal quantum control of multimode couplings between trapped ion qubits for scalable entanglement,” *Phys. Rev. Lett.* **112**, 190502 (2014).
- [49] C. Figgatt, A. Ostrander, N. M. Linke, K. A. Landsman, D. Zhu, D. Maslov, and C. Monroe, “Parallel entangling operations on a universal ion-trap quantum computer,” *Nature* **572**, 368–372 (2019).
- [50] M. M. Müller, M. Murphy, S. Montangero, T. Calarco, P. Grangier, and A. Browaeys, “Implementation of an experimentally feasible controlled-phase gate on two blockaded rydberg atoms,” *Phys. Rev. A* **89**, 032334 (2014).
- [51] D. Petrosyan, F. Motzoi, M. Saffman, and K. Mølmer, “High-fidelity rydberg quantum gate via a two-atom dark state,” *Phys. Rev. A* **96**, 042306 (2017).
- [52] C.-Y. Guo, L.-L. Yan, S. Zhang, S.-L. Su, and W. Li, “Optimized geometric quantum computation with a mesoscopic ensemble of rydberg atoms,” *Phys. Rev. A* **102**, 042607 (2020).
- [53] I. N. Hincks, C. E. Granade, T. W. Borneman, and D. G. Cory, “Controlling quantum devices with nonlinear hardware,” *Phys. Rev. Applied* **4**, 024012 (2015).
- [54] R. Porotti, D. Tamascelli, M. Restelli, and E. Prati, “Coherent transport of quantum states by deep reinforcement learning,” *Communications Physics* **2**, 61 (2019).
- [55] M. Y. Niu, S. Boixo, V. N. Smelyanskiy, and H. Neven, “Universal quantum control through deep reinforcement learning,” *npj Quantum Information* **5**, 33 (2019).
- [56] 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**, 14 (2021).

- [57] R. W. Heeres, P. Reinhold, N. Ofek, L. Frunzio, L. Jiang, M. H. Devoret, and R. J. Schoelkopf, “Implementing a universal gate set on a logical qubit encoded in an oscillator,” *Nature Communications* **8**, 94 (2017).
- [58] K. S. Chou, J. Z. Blumoff, C. S. Wang, P. C. Reinhold, C. J. Axline, Y. Y. Gao, L. Frunzio, M. H. Devoret, L. Jiang, and R. J. Schoelkopf, “Deterministic teleportation of a quantum gate between two logical qubits,” *Nature* **561**, 368–373 (2018).
- [59] S. Touzard, A. Grimm, Z. Leghtas, S. O. Mundhada, P. Reinhold, C. Axline, M. Reagor, K. Chou, J. Blumoff, K. M. Sliwa, S. Shankar, L. Frunzio, R. J. Schoelkopf, M. Mirrahimi, and M. H. Devoret, “Coherent oscillations inside a quantum manifold stabilized by dissipation,” *Phys. Rev. X* **8**, 021005 (2018).
- [60] X. Wu, S. L. Tomarken, N. A. Petersson, L. A. Martinez, Y. J. Rosen, and J. L. DuBois, “High-fidelity software-defined quantum logic on a superconducting qudit,” *Phys. Rev. Lett.* **125**, 170502 (2020).
- [61] M. Abdelhafez, B. Baker, A. Gyenis, P. Mundada, A. A. Houck, D. Schuster, and J. Koch, “Universal gates for protected superconducting qubits using optimal control,” *Phys. Rev. A* **101**, 022321 (2020).
- [62] N. Grzesiak, R. Blümel, K. Wright, K. M. Beck, N. C. Pisenti, M. Li, V. Chaplin, J. M. Amini, S. Debnath, J.-S. Chen, and Y. Nam, “Efficient arbitrary simultaneously entangling gates on a trapped-ion quantum computer,” *Nature Communications* **11**, 2963 (2020).
- [63] T. Olsacher, L. Postler, P. Schindler, T. Monz, P. Zoller, and L. M. Sieberer, “Scalable and parallel tweezer gates for quantum computing with long ion strings,” *PRX Quantum* **1**, 020316 (2020).
- [64] A. Winick, J. J. Wallman, and J. Emerson, “Simulating and mitigating crosstalk,” *Phys. Rev. Lett.* **126**, 230502 (2021).
- [65] I. M. Georgescu, S. Ashhab, and F. Nori, “Quantum simulation,” *Rev. Mod. Phys.* **86**, 153–185 (2014).
- [66] M. Lewenstein, A. Sanpera, and V. Ahufinger, *Ultracold Atoms in Optical Lattices: Simulating quantum many-body systems* (Oxford University Press, 2012).
- [67] I. Bloch, J. Dalibard, and S. Nascimbène, “Quantum simulations with ultracold quantum gases,” *Nature Physics* **8**, 267–276 (2012).
- [68] C. Gross and I. Bloch, “Quantum simulations with ultracold atoms in optical lattices,” *Science* **357**, 995–1001 (2017).

- [69] R. Blatt and C. F. Roos, “Quantum simulations with trapped ions,” *Nature Physics* **8**, 277–284 (2012).
- [70] J. Zhang, G. Pagano, P. W. Hess, A. Kyprianidis, P. Becker, H. Kaplan, A. V. Gorshkov, Z.-X. Gong, and C. Monroe, “Observation of a many-body dynamical phase transition with a 53-qubit quantum simulator,” *Nature* **551**, 601–604 (2017).
- [71] S. Rosi, A. Bernard, N. Fabbri, L. Fallani, C. Fort, M. Inguscio, T. Calarco, and S. Montangero, “Fast closed-loop optimal control of ultracold atoms in an optical lattice,” *Phys. Rev. A* **88**, 021601 (2013).
- [72] U. Hohenester, P. K. Rekdal, A. Borzì, and J. Schmiedmayer, “Optimal quantum control of bose-einstein condensates in magnetic microtraps,” *Phys. Rev. A* **75**, 023602 (2007).
- [73] S. E. Sklarz and D. J. Tannor, “Loading a bose-einstein condensate onto an optical lattice: An application of optimal control theory to the nonlinear schrödinger equation,” *Phys. Rev. A* **66**, 053619 (2002).
- [74] A. D. Tranter, H. J. Slatyer, M. R. Hush, A. C. Leung, J. L. Everett, K. V. Paul, P. Vernaz-Gris, P. K. Lam, B. C. Buchler, and G. T. Campbell, “Multiparameter optimisation of a magneto-optical trap using deep learning,” *Nature Communications* **9**, 4360 (2018).
- [75] G. Jäger, D. M. Reich, M. H. Goerz, C. P. Koch, and U. Hohenester, “Optimal quantum control of bose-einstein condensates in magnetic microtraps: Comparison of gradient-ascent-pulse-engineering and krotov optimization schemes,” *Phys. Rev. A* **90**, 033628 (2014).
- [76] S. van Frank, M. Bonneau, J. Schmiedmayer, S. Hild, C. Gross, M. Cheneau, I. Bloch, T. Pichler, A. Negretti, T. Calarco, and S. Montangero, “Optimal control of complex atomic quantum systems,” *Scientific Reports* **6**, 34187 (2016).
- [77] V. Giovannetti, S. Lloyd, and L. Maccone, “Quantum metrology,” *Phys. Rev. Lett.* **96**, 010401 (2006).
- [78] C. L. Degen, F. Reinhard, and P. Cappellaro, “Quantum sensing,” *Rev. Mod. Phys.* **89**, 035002 (2017).
- [79] V. Giovannetti, S. Lloyd, and L. Maccone, “Quantum-enhanced measurements: Beating the standard quantum limit,” *Science* **306**, 1330–1336 (2004).
- [80] G. Tóth and I. Apellaniz, “Quantum metrology from a quantum information science perspective,” *Journal of Physics A: Mathematical and Theoretical* **47**, 424006 (2014).

- [81] A. Omran, H. Levine, A. Keesling, G. Semeghini, T. T. Wang, S. Ebadi, H. Bernien, A. S. Zibrov, H. Pichler, S. Choi, J. Cui, M. Rossignolo, P. Rembold, S. Montangero, T. Calarco, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin, “Generation and manipulation of schrödinger cat states in rydberg atom arrays,” *Science* **365**, 570–574 (2019).
- [82] V. Bergholm, W. Wieczorek, T. Schulte-Herbrüggen, and M. Keyl, “Optimal control of hybrid optomechanical systems for generating non-classical states of mechanical motion,” *Quantum Science and Technology* **4**, 034001 (2019).
- [83] D. Basilewitsch, C. P. Koch, and D. M. Reich, “Quantum optimal control for mixed state squeezing in cavity optomechanics,” *Advanced Quantum Technologies* **2**, 1800110 (2019).
- [84] R. Mukherjee, H. Xie, and F. Mintert, “Bayesian optimal control of greenberger-horne-zeilinger states in rydberg lattices,” *Phys. Rev. Lett.* **125**, 203603 (2020).
- [85] A. Larrouy, S. Patsch, R. Richaud, J.-M. Raimond, M. Brune, C. P. Koch, and S. Gleyzes, “Fast navigation in a large hilbert space using quantum optimal control,” *Phys. Rev. X* **10**, 021058 (2020).
- [86] Q.-S. Tan, M. Zhang, Y. Chen, J.-Q. Liao, and J. Liu, “Generation and storage of spin squeezing via learning-assisted optimal control,” *Phys. Rev. A* **103**, 032601 (2021).
- [87] J. Yang, S. Pang, and A. N. Jordan, “Quantum parameter estimation with the landau-zener transition,” *Phys. Rev. A* **96**, 020301 (2017).
- [88] J. Liu and H. Yuan, “Control-enhanced multiparameter quantum estimation,” *Phys. Rev. A* **96**, 042114 (2017).
- [89] F. Poggiali, P. Cappellaro, and N. Fabbri, “Optimal control for one-qubit quantum sensing,” *Phys. Rev. X* **8**, 021059 (2018).
- [90] J. Liu and H. Yuan, “Quantum parameter estimation with optimal control,” *Phys. Rev. A* **96**, 012117 (2017).
- [91] Z. Hou, R.-J. Wang, J.-F. Tang, H. Yuan, G.-Y. Xiang, C.-F. Li, and G.-C. Guo, “Control-enhanced sequential scheme for general quantum parameter estimation at the heisenberg limit,” *Phys. Rev. Lett.* **123**, 040501 (2019).
- [92] H. Xu, J. Li, L. Liu, Y. Wang, H. Yuan, and X. Wang, “Generalizable control for quantum parameter estimation through reinforcement learning,” *npj Quantum Information* **5**, 82 (2019).

- [93] P. Titum, K. Schultz, A. Seif, G. Quiroz, and B. D. Clader, “Optimal control for quantum detectors,” *npj Quantum Information* **7**, 53 (2021).
- [94] T. Caneva, T. Calarco, and S. Montangero, “Chopped random-basis quantum optimization,” *Phys. Rev. A* **84**, 022326 (2011).
- [95] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbrüggen, and S. J. Glaser, “Optimal control of coupled spin dynamics: design of nmr pulse sequences by gradient ascent algorithms,” *Journal of Magnetic Resonance* **172**, 296–305 (2005).
- [96] D. M. Reich, M. Ndong, and C. P. Koch, “Monotonically convergent optimization in quantum control using krotov’s method,” *The Journal of Chemical Physics* **136**, 104103 (2012).
- [97] U. Boscain, M. Sigalotti, and D. Sugny, “Introduction to the foundations of quantum optimal control,” arXiv preprint (2020).
- [98] H. Rabitz, R. de Vivie-Riedle, M. Motzkus, and K. Kompa, “Whither the future of controlling quantum phenomena?” *Science* **288**, 824–828 (2000).
- [99] P. B. Wigley, P. J. Everitt, A. van den Hengel, J. W. Bastian, M. A. Sooriyabandara, G. D. McDonald, K. S. Hardman, C. D. Quinlivan, P. Manju, C. C. N. Kuhn, I. R. Petersen, A. N. Luiten, J. J. Hope, N. P. Robins, and M. R. Hush, “Fast machine-learning online optimization of ultra-cold-atom experiments,” *Scientific Reports* **6**, 25890 EP – (2016), article.
- [100] L. S. Pontryagin, *Mathematical theory of optimal processes* (Routledge, 2018).
- [101] V. Bergholm, J. Izaac, M. Schuld, C. Gogolin, M. S. Alam, S. Ahmed, J. M. Arrazola, C. Blank, A. Delgado, S. Jahangiri, *et al.*, “Pennylane: Automatic differentiation of hybrid quantum-classical computations,” arXiv preprint (2018).
- [102] X.-Z. Luo, J.-G. Liu, P. Zhang, and L. Wang, “Yao. jl: Extensible, efficient framework for quantum algorithm design,” *Quantum* **4**, 341 (2020).
- [103] M. Broughton, G. Verdon, T. McCourt, A. J. Martinez, J. H. Yoo, S. V. Isakov, P. Massey, M. Y. Niu, R. Halavati, E. Peters, *et al.*, “Tensorflow quantum: A software framework for quantum machine learning,” arXiv preprint (2020).
- [104] P. J. J. O’Malley, R. Babbush, I. D. Kivlichan, J. Romero, J. R. McClean, R. Barends, J. Kelly, P. Roushan, A. Tranter, N. Ding, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. G. Fowler, E. Jeffrey, E. Lucero, A. Megrant, J. Y. Mutus, M. Neeley, C. Neill, C. Quintana, D. Sank, A. Vainsencher, J. Wenner, T. C. White, P. V. Coveney, P. J. Love, H. Neven, A. Aspuru-Guzik, and J. M. Martinis, “Scalable quantum simulation of molecular energies,” *Phys. Rev. X* **6**, 031007 (2016).

- [105] K. Sharma, S. Khatri, M. Cerezo, and P. J. Coles, “Noise resilience of variational quantum compiling,” *New Journal of Physics* **22**, 043006 (2020).
- [106] M. E. S. Morales, T. Tlyachev, and J. Biamonte, “Variational learning of grover’s quantum search algorithm,” *Phys. Rev. A* **98**, 062333 (2018).
- [107] E. Anschuetz, J. Olson, A. Aspuru-Guzik, and Y. Cao, “Variational quantum factoring,” in *International Workshop on Quantum Technology and Optimization Problems* (Springer, 2019) pp. 74–85.
- [108] C. Cîrstoiu, Z. Holmes, J. Iosue, L. Cincio, P. J. Coles, and A. Sornborger, “Variational fast forwarding for quantum simulation beyond the coherence time,” *npj Quantum Information* **6**, 82 (2020).
- [109] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, “Barren plateaus in quantum neural network training landscapes,” *Nature Communications* **9**, 4812 (2018).
- [110] K. J. Sung, J. Yao, M. P. Harrigan, N. Rubin, Z. Jiang, L. Lin, R. Babbush, and J. R. McClean, “Using models to improve optimizers for variational quantum algorithms,” *Quantum Science and Technology* **5**, 044008 (2020).
- [111] J. F. Gonthier, M. D. Radin, C. Buda, E. J. Doskocil, C. M. Abuan, and J. Romero, “Identifying challenges towards practical quantum advantage through resource estimation: the measurement roadblock in the variational quantum eigensolver,” *arXiv preprint* (2020).
- [112] L. Zhou, S.-T. Wang, S. Choi, H. Pichler, and M. D. Lukin, “Quantum approximate optimization algorithm: Performance, mechanism, and implementation on near-term devices,” *Phys. Rev. X* **10**, 021067 (2020).
- [113] M. Willsch, D. Willsch, F. Jin, H. De Raedt, and K. Michielsen, “Benchmarking the quantum approximate optimization algorithm,” *Quantum Information Processing* **19**, 197 (2020).
- [114] O. R. Meitei, B. T. Gard, G. S. Barron, D. P. Pappas, S. E. Economou, E. Barnes, and N. J. Mayhall, “Gate-free state preparation for fast variational quantum eigensolver simulations: ctrl-vqe,” *arXiv preprint* (2020).
- [115] M. Schuld, V. Bergholm, C. Gogolin, J. Izaac, and N. Killoran, “Evaluating analytic gradients on quantum hardware,” *Phys. Rev. A* **99**, 032331 (2019).
- [116] M. Sarovar, T. Proctor, K. Rudinger, K. Young, E. Nielsen, and R. Blume-Kohout, “Detecting crosstalk errors in quantum information processors,” *Quantum* **4**, 321 (2020).

- [117] E. Zahedinejad, S. Schirmer, and B. C. Sanders, “Evolutionary algorithms for hard quantum control,” *Phys. Rev. A* **90**, 032310 (2014).
- [118] S. Machnes, U. Sander, S. J. Glaser, P. de Fouquières, A. Gruslys, S. Schirmer, and T. Schulte-Herbrüggen, “Comparing, optimizing, and benchmarking quantum-control algorithms in a unifying programming framework,” *Phys. Rev. A* **84**, 022305 (2011).
- [119] J. Stokes, J. Izaac, N. Killoran, and G. Carleo, “Quantum Natural Gradient,” *Quantum* **4**, 269 (2020).
- [120] D. Goodwin and I. Kuprov, “Auxiliary matrix formalism for interaction representation transformations, optimal control, and spin relaxation theories,” *The Journal of chemical physics* **143**, 084113 (2015).
- [121] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” arXiv preprint (2014).
- [122] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, and D. Amodei, “Language models are few-shot learners,” in *Advances in Neural Information Processing Systems*, Vol. 33, edited by H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (Curran Associates, Inc., 2020) pp. 1877–1901.
- [123] U. Schollwöck, “The density-matrix renormalization group,” *Rev. Mod. Phys.* **77**, 259–315 (2005).
- [124] M. Bartholomew-Biggs, S. Brown, B. Christianson, and L. Dixon, “Automatic differentiation of algorithms,” *Journal of Computational and Applied Mathematics* **124**, 171–190 (2000), *numerical Analysis 2000. Vol. IV: Optimization and Nonlinear Equations*.
- [125] A. G. Baydin, B. A. Pearlmutter, A. A. Radul, and J. M. Siskind, “Automatic differentiation in machine learning: a survey,” *Journal of machine learning research* **18** (2018).
- [126] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng, “TensorFlow: Large-scale machine learning on heterogeneous systems,” (2015), software available from tensorflow.org.

- [127] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, “Pytorch: An imperative style, high-performance deep learning library,” in *Advances in Neural Information Processing Systems 32*, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett (Curran Associates, Inc., 2019) pp. 8024–8035.
- [128] H.-J. Liao, J.-G. Liu, L. Wang, and T. Xiang, “Differentiable programming tensor networks,” *Phys. Rev. X* **9**, 031041 (2019).
- [129] J.-G. Liu, L. Wang, and P. Zhang, “Tropical tensor network for ground states of spin glasses,” arXiv preprint (2020).
- [130] H. Xie, J.-G. Liu, and L. Wang, “Automatic differentiation of dominant eigensolver and its applications in quantum physics,” *Phys. Rev. B* **101**, 245139 (2020).
- [131] N. Leung, M. Abdelhafez, J. Koch, and D. Schuster, “Speedup for quantum optimal control from automatic differentiation based on graphics processing units,” *Phys. Rev. A* **95**, 042318 (2017).
- [132] M. Abdelhafez, D. I. Schuster, and J. Koch, “Gradient-based optimal control of open quantum systems using quantum trajectories and automatic differentiation,” *Phys. Rev. A* **99**, 052327 (2019).
- [133] F. Schäfer, M. Kloc, C. Bruder, and N. Lörch, “A differentiable programming method for quantum control,” *Machine Learning: Science and Technology* **1**, 035009 (2020).
- [134] L. Coopmans, D. Luo, G. Kells, B. K. Clark, and J. Carrasquilla, “Protocol discovery for the quantum control of majoranas by differentiable programming and natural evolution strategies,” *PRX Quantum* **2**, 020332 (2021).
- [135] R. T. Q. Chen, Y. Rubanova, J. Bettencourt, and D. K. Duvenaud, “Neural ordinary differential equations,” in *Advances in Neural Information Processing Systems*, Vol. 31, edited by S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Curran Associates, Inc., 2018).
- [136] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii, “Quantum circuit learning,” *Phys. Rev. A* **98**, 032309 (2018).
- [137] L. Banchi and G. E. Crooks, “Measuring Analytic Gradients of General Quantum Evolution with the Stochastic Parameter Shift Rule,” *Quantum* **5**, 386 (2021).
- [138] A. F. Izmaylov, R. A. Lang, and T.-C. Yen, “Analytic gradients in variational quantum algorithms: Algebraic extensions of the parameter-shift rule to general unitary transformations,” arXiv preprint (2021).

- [139] D. Wierichs, J. Izaac, C. Wang, and C. Y.-Y. Lin, “General parameter-shift rules for quantum gradients,” arXiv preprint (2021).
- [140] J. C. Spall, “Multivariate stochastic approximation using a simultaneous perturbation gradient approximation,” IEEE Transactions on Automatic Control **37**, 332–341 (1992).
- [141] J. C. Spall, “Implementation of the simultaneous perturbation algorithm for stochastic optimization,” IEEE Transactions on aerospace and electronic systems **34**, 817–823 (1998).
- [142] C. Ferrie and O. Moussa, “Robust and efficient in situ quantum control,” Phys. Rev. A **91**, 052306 (2015).
- [143] R. J. Chapman, C. Ferrie, and A. Peruzzo, “Experimental demonstration of self-guided quantum tomography,” Phys. Rev. Lett. **117**, 040402 (2016).
- [144] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, “Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets,” Nature **549**, 242–246 (2017).
- [145] P. Doria, T. Calarco, and S. Montangero, “Optimal control technique for many-body quantum dynamics,” Phys. Rev. Lett. **106**, 190501 (2011).
- [146] N. Rach, M. M. Müller, T. Calarco, and S. Montangero, “Dressing the chopped-random-basis optimization: A bandwidth-limited access to the trap-free landscape,” Phys. Rev. A **92**, 062343 (2015).
- [147] D. Beasley, D. R. Bull, and R. R. Martin, “An overview of genetic algorithms: Part 1, fundamentals,” University computing **15**, 56–69 (1993).
- [148] N. B. Lovett, C. Crosnier, M. Perarnau-Llobet, and B. C. Sanders, “Differential evolution for many-particle adaptive quantum metrology,” Phys. Rev. Lett. **110**, 220501 (2013).
- [149] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, I. Antonoglou, H. King, D. Kumaran, D. Wierstra, S. Legg, and D. Hassabis, “Human-level control through deep reinforcement learning,” Nature **518**, 529–533 (2015).
- [150] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton, Y. Chen, T. Lillicrap, F. Hui, L. Sifre, G. van den Driessche, T. Graepel, and D. Hassabis, “Mastering the game of go without human knowledge,” Nature **550**, 354–359 (2017).
- [151] R. S. Sutton and A. G. Barto, *Introduction to Reinforcement Learning*, 1st ed. (MIT Press, Cambridge, MA, USA, 1998).

- [152] A. A. Melnikov, H. Poulsen Nautrup, M. Krenn, V. Dunjko, M. Tiersch, A. Zeilinger, and H. J. Briegel, “Active learning machine learns to create new quantum experiments,” *Proceedings of the National Academy of Sciences* **115**, 1221–1226 (2018).
- [153] H. P. Nautrup, N. Delfosse, V. Dunjko, H. J. Briegel, and N. Friis, “Optimizing Quantum Error Correction Codes with Reinforcement Learning,” *Quantum* **3**, 215 (2019).
- [154] P. Andreasson, J. Johansson, S. Liljestrand, and M. Granath, “Quantum error correction for the toric code using deep reinforcement learning,” *Quantum* **3**, 183 (2019).
- [155] Y. Baum, M. Amico, S. Howell, M. Hush, M. Liuzzi, P. Mundada, T. Merkh, A. R. Carvalho, and M. J. Biercuk, “Experimental deep reinforcement learning for error-robust gateset design on a superconducting quantum computer,” arXiv preprint (2021).
- [156] S. Khairy, R. Shaydulin, L. Cincio, Y. Alexeev, and P. Balaprakash, “Learning to optimize variational quantum circuits to solve combinatorial problems,” *Proceedings of the AAAI Conference on Artificial Intelligence* **34**, 2367–2375 (2020).
- [157] X. B. Peng, M. Andrychowicz, W. Zaremba, and P. Abbeel, “Sim-to-real transfer of robotic control with dynamics randomization,” in *2018 IEEE international conference on robotics and automation (ICRA)* (IEEE, 2018) pp. 3803–3810.
- [158] E. Brochu, V. M. Cora, and N. De Freitas, “A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning,” arXiv preprint (2010).
- [159] J. Snoek, H. Larochelle, and R. P. Adams, “Practical bayesian optimization of machine learning algorithms,” arXiv preprint (2012).
- [160] P. I. Frazier, “A tutorial on bayesian optimization,” arXiv preprint (2018).
- [161] J. Snoek, H. Larochelle, and R. P. Adams, “Practical bayesian optimization of machine learning algorithms,” in *Advances in neural information processing systems* (2012) pp. 2951–2959.
- [162] R. Calandra, A. Seyfarth, J. Peters, and M. P. Deisenroth, “Bayesian optimization for learning gaits under uncertainty,” *Annals of Mathematics and Artificial Intelligence* **76**, 5–23 (2016).
- [163] R.-R. Griffiths and J. M. Hernández-Lobato, “Constrained bayesian optimization for automatic chemical design,” arXiv preprint (2017).
- [164] C. E. Rasmussen, “Gaussian processes in machine learning,” in *Summer school on machine learning* (Springer, 2003) pp. 63–71.

- [165] R. Mukherjee, F. Sauvage, H. Xie, R. Löw, and F. Mintert, “Preparation of ordered states in ultra-cold gases using bayesian optimization,” *New Journal of Physics* **22**, 075001 (2020).
- [166] S. T. Flammia and Y.-K. Liu, “Direct fidelity estimation from few pauli measurements,” *Phys. Rev. Lett.* **106**, 230501 (2011).
- [167] M. P. da Silva, O. Landon-Cardinal, and D. Poulin, “Practical characterization of quantum devices without tomography,” *Phys. Rev. Lett.* **107**, 210404 (2011).
- [168] H.-Y. Huang, R. Kueng, and J. Preskill, “Predicting many properties of a quantum system from very few measurements,” *Nature Physics* **16**, 1050–1057 (2020).
- [169] S. Pallister, N. Linden, and A. Montanaro, “Optimal verification of entangled states with local measurements,” *Phys. Rev. Lett.* **120**, 170502 (2018).
- [170] S. Greenaway, F. Sauvage, K. E. Khosla, and F. Mintert, “Efficient assessment of process fidelity,” *Phys. Rev. Research* **3**, 033031 (2021).
- [171] A. Gilchrist, N. K. Langford, and M. A. Nielsen, “Distance measures to compare real and ideal quantum processes,” *Phys. Rev. A* **71**, 062310 (2005).
- [172] M. Hein, J. Eisert, and H. J. Briegel, “Multiparty entanglement in graph states,” *Phys. Rev. A* **69**, 062311 (2004).
- [173] W. Dür, G. Vidal, and J. I. Cirac, “Three qubits can be entangled in two inequivalent ways,” *Phys. Rev. A* **62**, 062314 (2000).
- [174] F. Sauvage and F. Mintert, “Optimal quantum control with poor statistics,” *PRX Quantum* **1**, 020322 (2020).
- [175] F. Schäfer, T. Fukuhara, S. Sugawa, Y. Takasu, and Y. Takahashi, “Tools for quantum simulation with ultracold atoms in optical lattices,” *Nature Reviews Physics* **2**, 411–425 (2020).
- [176] M. Greiner, O. Mandel, T. Esslinger, T. W. Hänsch, and I. Bloch, “Quantum phase transition from a superfluid to a mott insulator in a gas of ultracold atoms,” *Nature* **415**, 39–44 (2002).
- [177] A. S. Sørensen, E. Altman, M. Gullans, J. V. Porto, M. D. Lukin, and E. Demler, “Adiabatic preparation of many-body states in optical lattices,” *Phys. Rev. A* **81**, 061603 (2010).
- [178] T. Cantat-Moltrecht, R. Cortiñas, B. Ravon, P. Méhaignerie, S. Haroche, J. M. Raimond, M. Favier, M. Brune, and C. Sayrin, “Long-lived circular rydberg states of laser-cooled rubidium atoms in a cryostat,” *Phys. Rev. Research* **2**, 022032 (2020).

[179] H. Bernien, S. Schwartz, A. Keesling, H. Levine, A. Omran, H. Pichler, S. Choi, A. S. Zibrov, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin, “Probing many-body dynamics on a 51-atom quantum simulator,” *Nature* **551**, 579–584 (2017).

[180] M. P. A. Fisher, P. B. Weichman, G. Grinstein, and D. S. Fisher, “Boson localization and the superfluid-insulator transition,” *Phys. Rev. B* **40**, 546–570 (1989).

[181] D. Jaksch, C. Bruder, J. I. Cirac, C. W. Gardiner, and P. Zoller, “Cold bosonic atoms in optical lattices,” *Phys. Rev. Lett.* **81**, 3108–3111 (1998).

[182] T. Albash and D. A. Lidar, “Adiabatic quantum computation,” *Rev. Mod. Phys.* **90**, 015002 (2018).

[183] A. Walther, F. Ziesel, T. Ruster, S. T. Dawkins, K. Ott, M. Hettrich, K. Singer, F. Schmidt-Kaler, and U. Poschinger, “Controlling fast transport of cold trapped ions,” *Phys. Rev. Lett.* **109**, 080501 (2012).

[184] S. Simsek and F. Mintert, “Quantum control with a multi-dimensional Gaussian quantum invariant,” *Quantum* **5**, 409 (2021).

[185] T. Caneva, M. Murphy, T. Calarco, R. Fazio, S. Montangero, V. Giovannetti, and G. Santoro, “Optimal control at the quantum speed limit,” *Physical review letters* **103**, 240501 (2009).

[186] S. Deffner and S. Campbell, “Quantum speed limits: from heisenberg’s uncertainty principle to optimal quantum control,” *Journal of Physics A: Mathematical and Theoretical* **50**, 453001 (2017).

[187] J. H. M. Jensen, F. S. Møller, J. J. Sørensen, and J. F. Sherson, “Achieving fast high-fidelity control of many-body dynamics,” *arXiv preprint* (2020).

[188] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, İ. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors, “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python,” *Nature Methods* **17**, 261–272 (2020).

[189] T. G. authors, “Gpyopt: A bayesian optimization framework in python,” <http://github.com/SheffieldML/GPyOpt> (2016).

[190] W. S. Bakr, J. I. Gillen, A. Peng, S. Fölling, and M. Greiner, “A quantum gas microscope for detecting single atoms in a hubbard-regime optical lattice,” *Nature* **462**, 74–77 (2009).

[191] B. M. Henson, D. K. Shin, K. F. Thomas, J. A. Ross, M. R. Hush, S. S. Hodgman, and A. G. Truscott, “Approaching the adiabatic timescale with machine learning,” *Proceedings of the National Academy of Sciences* **115**, 13216–13221 (2018).

[192] D. Zhu, N. M. Linke, M. Benedetti, K. A. Landsman, N. H. Nguyen, C. H. Alderete, A. Perdomo-Ortiz, N. Korda, A. Garfoot, C. Brecque, L. Egan, O. Perdomo, and C. Monroe, “Training of quantum circuits on a hybrid quantum computer,” *Science Advances* **5** (2019).

[193] H. Moon, D. T. Lennon, J. Kirkpatrick, N. M. van Esbroeck, L. C. Camenzind, L. Yu, F. Vignneau, 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**, 4161 (2020).

[194] T. L. Nguyen, *Study of dipole-dipole interaction between Rydberg atoms : toward quantum simulation with Rydberg atoms*, Ph.D. thesis, Paris 6 (2016), thèse de doctorat dirigée par Brune, Michel Physique quantique.

[195] C. C. N. Kuhn, G. D. McDonald, K. S. Hardman, S. Bennetts, P. J. Everitt, P. A. Altin, J. E. Debs, J. D. Close, and N. P. Robins, “A bose-condensed, simultaneous dual-species mach–zehnder atom interferometer,” *New Journal of Physics* **16**, 073035 (2014).

[196] M. Endres, H. Bernien, A. Keesling, H. Levine, E. R. Anschuetz, A. Krajenbrink, C. Senko, V. Vuletic, M. Greiner, and M. D. Lukin, “Atom-by-atom assembly of defect-free one-dimensional cold atom arrays,” *Science* **354**, 1024–1027 (2016).

[197] V. Havlíček, A. D. Córcoles, K. Temme, A. W. Harrow, A. Kandala, J. M. Chow, and J. M. Gambetta, “Supervised learning with quantum-enhanced feature spaces,” *Nature* **567**, 209–212 (2019).

[198] R. Moriconi, M. P. Deisenroth, and K. S. Sesh Kumar, “High-dimensional bayesian optimization using low-dimensional feature spaces,” *Machine Learning* **109**, 1925–1943 (2020).

[199] Z. Wang, F. Hutter, M. Zoghi, D. Matheson, and N. De Freitas, “Bayesian optimization in a billion dimensions via random embeddings,” *J. Artif. Int. Res.* **55**, 361–387 (2016).

[200] J. R. Gardner, G. Pleiss, D. Bindel, K. Q. Weinberger, and A. G. Wilson, “Gpytorch: Blackbox matrix-matrix gaussian process inference with gpu acceleration,” in *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, NIPS’18 (Curran Associates Inc., Red Hook, NY, USA, 2018) p. 7587–7597.

[201] J. Hensman, N. Fusi, and N. D. Lawrence, “Gaussian processes for big data,” arXiv preprint (2013).

- [202] B. Lakshminarayanan, A. Pritzel, and C. Blundell, “Simple and scalable predictive uncertainty estimation using deep ensembles,” arXiv preprint (2016).
- [203] S. Khatri, R. LaRose, A. Poremba, L. Cincio, A. T. Sornborger, and P. J. Coles, “Quantum-assisted quantum compiling,” *Quantum* **3**, 140 (2019).
- [204] M. Cerezo, A. Sone, T. Volkoff, L. Cincio, and P. J. Coles, “Cost function dependent barren plateaus in shallow parametrized quantum circuits,” *Nature Communications* **12**, 1791 (2021).
- [205] B. T. Kiani, G. De Palma, M. Marvian, Z.-W. Liu, and S. Lloyd, “Quantum earth mover’s distance: A new approach to learning quantum data,” arXiv preprint (2021).
- [206] A. Arrasmith, M. Cerezo, P. Czarnik, L. Cincio, and P. J. Coles, “Effect of barren plateaus on gradient-free optimization,” arXiv preprint (2020).
- [207] J. Emerson, R. Alicki, and K. Życzkowski, “Scalable noise estimation with random unitary operators,” *Journal of Optics B: Quantum and Semiclassical Optics* **7**, S347–S352 (2005).
- [208] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland, “Randomized benchmarking of quantum gates,” *Phys. Rev. A* **77**, 012307 (2008).
- [209] E. Magesan, J. M. Gambetta, and J. Emerson, “Scalable and robust randomized benchmarking of quantum processes,” *Phys. Rev. Lett.* **106**, 180504 (2011).
- [210] T. Monz, P. Schindler, J. T. Barreiro, M. Chwalla, D. Nigg, W. A. Coish, M. Harlander, W. Hänsel, M. Hennrich, and R. Blatt, “14-qubit entanglement: Creation and coherence,” *Phys. Rev. Lett.* **106**, 130506 (2011).
- [211] H.-S. Zhong, Y. Li, W. Li, L.-C. Peng, Z.-E. Su, Y. Hu, Y.-M. He, X. Ding, W. Zhang, H. Li, L. Zhang, Z. Wang, L. You, X.-L. Wang, X. Jiang, L. Li, Y.-A. Chen, N.-L. Liu, C.-Y. Lu, and J.-W. Pan, “12-photon entanglement and scalable scattershot boson sampling with optimal entangled-photon pairs from parametric down-conversion,” *Phys. Rev. Lett.* **121**, 250505 (2018).
- [212] G. J. Mooney, C. D. Hill, and L. C. L. Hollenberg, “Entanglement in a 20-qubit superconducting quantum computer,” *Scientific Reports* **9**, 13465 (2019).
- [213] G. J. Mooney, G. A. L. White, C. D. Hill, and L. C. L. Hollenberg, “Generation and verification of 27-qubit Greenberger-Horne-Zeilinger states in a superconducting quantum computer,” arXiv preprint (2021).

- [214] N. Shettell and D. Markham, “Graph states as a resource for quantum metrology,” *Phys. Rev. Lett.* **124**, 110502 (2020).
- [215] R. Raussendorf and H. J. Briegel, “A one-way quantum computer,” *Phys. Rev. Lett.* **86**, 5188–5191 (2001).
- [216] G. Tóth and O. Gühne, “Entanglement detection in the stabilizer formalism,” *Phys. Rev. A* **72**, 022340 (2005).
- [217] M. A. Nielsen and I. Chuang, “Quantum computation and quantum information,” (2002).
- [218] T. Proctor, K. Rudinger, K. Young, M. Sarovar, and R. Blume-Kohout, “What randomized benchmarking actually measures,” *Phys. Rev. Lett.* **119**, 130502 (2017).
- [219] M. A. Nielsen, “A simple formula for the average gate fidelity of a quantum dynamical operation,” *Physics Letters A* **303**, 249–252 (2002).
- [220] A. Jamiołkowski, “Linear transformations which preserve trace and positive semidefiniteness of operators,” *Reports on Mathematical Physics* **3**, 275–278 (1972).
- [221] D. Lu, H. Li, D.-A. Trottier, J. Li, A. Brodutch, A. P. Krismanich, A. Ghavami, G. I. Dmitrienko, G. Long, J. Baugh, and R. Laflamme, “Experimental estimation of average fidelity of a clifford gate on a 7-qubit quantum processor,” *Phys. Rev. Lett.* **114**, 140505 (2015).
- [222] A. Erhard, J. J. Wallman, L. Postler, M. Meth, R. Stricker, E. A. Martinez, P. Schindler, T. Monz, J. Emerson, and R. Blatt, “Characterizing large-scale quantum computers via cycle benchmarking,” *Nature Communications* **10**, 5347 (2019).
- [223] J. M. Renes, R. Blume-Kohout, A. J. Scott, and C. M. Caves, “Symmetric informationally complete quantum measurements,” *Journal of Mathematical Physics* **45**, 2171–2180 (2004), <https://doi.org/10.1063/1.1737053> .
- [224] K. Mayer, Private Communication (2021).
- [225] Y. Zhao, R. Zhang, W. Chen, X.-B. Wang, and J. Hu, “Creation of greenberger-horne-zeilinger states with thousands of atoms by entanglement amplification,” *npj Quantum Information* **7**, 24 (2021).
- [226] B. Foxen, C. Neill, A. Dunsworth, P. Roushan, B. Chiaro, A. Megrant, J. Kelly, Z. Chen, K. Satzinger, R. Barends, F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, S. Boixo, D. Buell, B. Burkett, Y. Chen, R. Collins, E. Farhi, A. Fowler, C. Gidney, M. Giustina, R. Graff, M. Harrigan, T. Huang, S. V. Isakov, E. Jeffrey, Z. Jiang, D. Kafri, K. Kechedzhi, P. Klimov, A. Korotkov, F. Kostritsa, D. Landhuis, E. Lucero, J. McClean, M. McEwen, X. Mi, M. Mohseni, J. Y. Mutus, O. Naaman, M. Neeley, M. Niu,

A. Petukhov, C. Quintana, N. Rubin, D. Sank, V. Smelyanskiy, A. Vainsencher, T. C. White, Z. Yao, P. Yeh, A. Zalcman, H. Neven, and J. M. Martinis (Google AI Quantum), “Demonstrating a continuous set of two-qubit gates for near-term quantum algorithms,” *Phys. Rev. Lett.* **125**, 120504 (2020).

[227] N. Lacroix, C. Hellings, C. K. Andersen, A. Di Paolo, A. Remm, S. Lazar, S. Krinner, G. J. Norris, M. Gabureac, J. Heinsoo, A. Blais, C. Eichler, and A. Wallraff, “Improving the performance of deep quantum optimization algorithms with continuous gate sets,” *PRX Quantum* **1**, 110304 (2020).

[228] D. M. Abrams, N. Didier, B. R. Johnson, M. P. d. Silva, and C. A. Ryan, “Implementation of xy entangling gates with a single calibrated pulse,” *Nature Electronics* **3**, 744–750 (2020).

[229] C. N. Self, K. E. Khosla, A. W. R. Smith, F. Sauvage, P. D. Haynes, J. Knolle, F. Mintert, and M. S. Kim, “Variational quantum algorithm with information sharing,” *npj Quantum Information* **7**, 116 (2021).

[230] C. Kokail, C. Maier, R. van Bijnen, T. Brydges, M. K. Joshi, P. Jurcevic, C. A. Muschik, P. Silvi, R. Blatt, C. F. Roos, and P. Zoller, “Self-verifying variational quantum simulation of lattice models,” *Nature* **569**, 355–360 (2019).

[231] J. Snoek, O. Rippel, K. Swersky, R. Kiros, N. Satish, N. Sundaram, M. Patwary, M. Prabhat, and R. Adams, “Scalable bayesian optimization using deep neural networks,” in *International conference on machine learning* (PMLR, 2015) pp. 2171–2180.

[232] K. J. Sung, J. Yao, M. P. Harrigan, N. C. Rubin, Z. Jiang, L. Lin, R. Babbush, and J. R. McClean, “Using models to improve optimizers for variational quantum algorithms,” *Quantum Science and Technology* **5**, 044008 (2020).

[233] H. R. Grimsley, S. E. Economou, E. Barnes, and N. J. Mayhall, “An adaptive variational algorithm for exact molecular simulations on a quantum computer,” *Nature Communications* **10**, 3007 (2019).

[234] M. Ostaszewski, E. Grant, and M. Benedetti, “Structure optimization for parameterized quantum circuits,” *Quantum* **5**, 391 (2021).

[235] A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. A. Smolin, and H. Weinfurter, “Elementary gates for quantum computation,” *Phys. Rev. A* **52**, 3457–3467 (1995).

[236] C. M. Bishop *et al.*, *Neural networks for pattern recognition* (Oxford university press, 1995).

[237] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, 2016).

[238] P. Mehta, M. Bukov, C.-H. Wang, A. G. Day, C. Richardson, C. K. Fisher, and D. J. Schwab, “A high-bias, low-variance introduction to machine learning for physicists,” *Physics Reports* **810**, 1–124 (2019), a high-bias, low-variance introduction to Machine Learning for physicists.

[239] F. Marquardt, “Machine Learning and Quantum Devices,” *SciPost Phys. Lect. Notes* , 29 (2021).

[240] K. Hornik, “Approximation capabilities of multilayer feedforward networks,” *Neural Networks* **4**, 251–257 (1991).

[241] R. T. Q. Chen, “Pytorch implementation of differentiable ode solvers,” <https://github.com/rtqichen/torchdiffeq> (2018).

[242] M. Claesen and B. De Moor, “Hyperparameter search in machine learning,” in *Proc. of the 11th Metaheuristics International Conference* (2015) pp. 1–5.

[243] J. Bergstra and Y. Bengio, “Random search for hyper-parameter optimization.” *Journal of machine learning research* **13** (2012).

[244] S. G. Schirmer, H. Fu, and A. I. Solomon, “Complete controllability of quantum systems,” *Physical Review A* **63**, 063410 (2001).

[245] V. Nair and G. E. Hinton, “Rectified linear units improve restricted boltzmann machines,” in *Proceedings of the 27th International Conference on International Conference on Machine Learning*, ICML’10 (Omnipress, Madison, WI, USA, 2010) p. 807–814.

[246] pytorch, “Update weight initialisations to current best practices,” <https://github.com/weinbe58/QuSpin> (since 2019).

[247] B. Kraus and J. I. Cirac, “Optimal creation of entanglement using a two-qubit gate,” *Phys. Rev. A* **63**, 062309 (2001).

[248] B. P. Lanyon, C. Hempel, D. Nigg, M. Müller, R. Gerritsma, F. Zähringer, P. Schindler, J. T. Barreiro, M. Rambach, G. Kirchmair, M. Hennrich, P. Zoller, R. Blatt, and C. F. Roos, “Universal digital quantum simulation with trapped ions,” *Science* **334**, 57–61 (2011).

[249] Y. Salathé, M. Mondal, M. Oppliger, J. Heinsoo, P. Kurpiers, A. Potočnik, A. Mezzacapo, U. Las Heras, L. Lamata, E. Solano, S. Filipp, and A. Wallraff, “Digital quantum simulation of spin models with circuit quantum electrodynamics,” *Phys. Rev. X* **5**, 021027 (2015).

[250] R. Barends, L. Lamata, J. Kelly, L. García-Álvarez, A. G. Fowler, A. Megrant, E. Jeffrey, T. C. White, D. Sank, J. Y. Mutus, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, I.-C. Hoi, C. Neill, P. J. J. O’Malley, C. Quintana, P. Roushan, A. Vainsencher, J. Wenner, E. Solano, and J. M. Martinis, “Digital quantum simulation of fermionic models with a superconducting circuit,” *Nature Communications* **6**, 7654 (2015).

[251] A. Smith, M. S. Kim, F. Pollmann, and J. Knolle, “Simulating quantum many-body dynamics on a current digital quantum computer,” *npj Quantum Information* **5**, 106 (2019).

[252] I. collaboration, “Qiskit: An open-source framework for quantum computing,” (2021).

[253] C. H. Tseng, S. Somaroo, Y. Sharf, E. Knill, R. Laflamme, T. F. Havel, and D. G. Cory, “Quantum simulation of a three-body-interaction hamiltonian on an nmr quantum computer,” *Phys. Rev. A* **61**, 012302 (1999).

[254] W. Liu, W. Feng, W. Ren, D.-W. Wang, and H. Wang, “Synthesizing three-body interaction of spin chirality with superconducting qubits,” *Applied Physics Letters* **116**, 114001 (2020).

[255] F. Petziol, M. Sameti, S. Carretta, S. Wimberger, and F. Mintert, “Quantum simulation of three-body interactions in weakly driven quantum systems,” *Phys. Rev. Lett.* **126**, 250504 (2021).

[256] K. He, X. Zhang, S. Ren, and J. Sun, “Delving deep into rectifiers: Surpassing human-level performance on imagenet classification,” in *Proceedings of the IEEE international conference on computer vision* (2015) pp. 1026–1034.

[257] S. Ioffe, “Batch renormalization: Towards reducing minibatch dependence in batch-normalized models,” in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS’17 (Curran Associates Inc., Red Hook, NY, USA, 2017) p. 1942–1950.

[258] J. R. Sashank, K. Satyen, and K. Sanjiv, “On the convergence of adam and beyond,” in *International Conference on Learning Representations*, Vol. 5 (2018) p. 7.

[259] M. August and J. M. Hernández-Lobato, “Taking gradients through experiments: Lstms and memory proximal policy optimization for black-box quantum control,” in *International Conference on High Performance Computing* (Springer, 2018) pp. 591–613.

[260] C. Rackauckas, Y. Ma, J. Martensen, C. Warner, K. Zubov, R. Supekar, D. Skinner, A. Ramadhan, and A. Edelman, “Universal differential equations for scientific machine learning,” arXiv preprint (2020).

- [261] F. Schäfer, P. Sekatski, M. Koppenhöfer, C. Bruder, and M. Kloc, “Control of stochastic quantum dynamics by differentiable programming,” *Machine Learning: Science and Technology* **2**, 035004 (2021).
- [262] R. Porotti, A. Essig, B. Huard, and F. Marquardt, “Deep reinforcement learning for quantum state preparation with weak nonlinear measurements,” arXiv preprint (2021).
- [263] T. Salimans, J. Ho, X. Chen, S. Sidor, and I. Sutskever, “Evolution strategies as a scalable alternative to reinforcement learning,” arXiv preprint (2017).
- [264] K. P. Horn, F. Reiter, Y. Lin, D. Leibfried, and C. P. Koch, “Quantum optimal control of the dissipative production of a maximally entangled state,” *New Journal of Physics* **20**, 123010 (2018).
- [265] GPy, “GPy: A gaussian process framework in python,” <http://github.com/SheffieldML/GPy> (since 2012).
- [266] B. Collins and P. Śniady, “Integration with respect to the haar measure on unitary, orthogonal and symplectic group,” *Communications in Mathematical Physics* **264**, 773–795 (2006).
- [267] J. Johansson, P. Nation, and F. Nori, “Qutip 2: A python framework for the dynamics of open quantum systems,” *Computer Physics Communications* **184**, 1234–1240 (2013).
- [268] Z. Puchała and J. Miszczak, “Symbolic integration with respect to the haar measure on the unitary groups,” *Bulletin of the Polish Academy of Sciences: Technical Sciences* **65**, 21–27 (2017).
- [269] G. Tóth and O. Gühne, “Detecting genuine multipartite entanglement with two local measurements,” *Phys. Rev. Lett.* **94**, 060501 (2005).

Chapter A

Bayesian Optimization

In this appendix, we review the technical elements of Bayesian optimization (BO) that were glossed over in the main text. First, we provide a formal definition of the Gaussian processes (Appendix A.1) which are employed as surrogate models in BO. Then, we review the choice of the kernel and mean functions (Appendix A.2), and of the hyperparameters (Appendix A.3). We end this appendix by providing the full expression of the predictive distribution (Appendix A.4), which is used at each step of BO.

A.1 Gaussian processes

A random process (sometimes called stochastic process or even random function) extends the concept of a probability distribution defined over a finite set of random variables, to a probability distribution over an infinite collection of random variables. Any function f can be understood as the continuum limit of a set of discrete random variables $\{f(\boldsymbol{\theta}_i)\}$, and since we are interested in functions F of continuous control parameters $\boldsymbol{\theta}$, random processes are the appropriate mathematical structure for the probabilistic modelling of such functions. In particular, Gaussian processes extend Gaussian distributions over finitely many random variables to infinitely many (that is, to functions).

A random process such that any finite subset of random variables $[f(\boldsymbol{\theta}_1), \dots, f(\boldsymbol{\theta}_N)]$ follows a Gaussian distribution is a Gaussian Process. Recall that a (multivariate) Gaussian distribution, over a finite set of random variables $X = [X_1, \dots, X_N]$, is defined as

$$p(X = \mathbf{X}) = \mathcal{N}(\mathbf{m}, C) = \frac{\exp\left[-(\mathbf{X} - \mathbf{m})^T C^{-1}(\mathbf{X} - \mathbf{m})\right]}{\sqrt{(2\pi)^D |C|}}, \quad (\text{A.1})$$

which is entirely specified by a mean vector \mathbf{m} and a (symmetric positive semidefinite) covariance matrix C . In contrast, Gaussian processes are specified in terms of a mean function m

and a (symmetric positive semidefinite) kernel function k

$$\begin{aligned} m(\boldsymbol{\theta}) &= \langle f(\boldsymbol{\theta}) \rangle, \\ k(\boldsymbol{\theta}, \boldsymbol{\theta}') &= \langle f(\boldsymbol{\theta}) f(\boldsymbol{\theta}') \rangle - \langle f(\boldsymbol{\theta}) \rangle \langle f(\boldsymbol{\theta}') \rangle, \end{aligned} \quad (\text{A.2})$$

which define respectively the mean value (denoted with symbols $\langle \cdot \rangle$) of any of the (infinitely many) $f(\boldsymbol{\theta})$ random variables, and the covariance in between any pair of $f(\boldsymbol{\theta})$ and $f(\boldsymbol{\theta}')$ variables.

More explicitly, if f follows a Gaussian process distribution (with mean m and kernel k) then, for an *arbitrary* set of parameters $[\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_N]$ the vector of random variables $[f(\boldsymbol{\theta}_1), \dots, f(\boldsymbol{\theta}_N)]$ follows the (multivariate) Gaussian distribution:

$$p([f(\boldsymbol{\theta}_1), \dots, f(\boldsymbol{\theta}_N)] = \mathbf{f}) = \mathcal{N}(\mathbf{m}, C), \quad (\text{A.3})$$

with mean vector $\mathbf{m} = [m(\boldsymbol{\theta}_1), \dots, m(\boldsymbol{\theta}_N)]$ and covariance matrix C with entries $C_{ij} = k(\boldsymbol{\theta}_i, \boldsymbol{\theta}_j)$.

Hence, the full specification of a Gaussian process corresponds to a particular choice of the functions m and k , which defines the global properties of f , such as its periodicity or smoothness.

A.2 Choice of the mean and kernel functions

Without particular knowledge of the underlying function F to be modeled, it is advised [164] to choose the mean function to be the zero function $m(\boldsymbol{\theta}) = 0$, and rather, to standardize the observed values of F , thus ensuring that their empirical mean is also 0. Additionally, a common choice for the kernel is the *Matérn 5/2* function that only assigns non-null probabilities to functions f which are (at least) twice differentiable [164]. This kernel is defined as:

$$k_{5/2}(\boldsymbol{\theta}, \boldsymbol{\theta}') = k_{5/2}(x = |\boldsymbol{\theta} - \boldsymbol{\theta}'|) = V \left(1 + \frac{x}{l} + \frac{x^2}{3l^2} \right) e^{-x/l}, \quad (\text{A.4})$$

which depends on the distance $x = |\boldsymbol{\theta} - \boldsymbol{\theta}'|$ scaled by a correlation length l , and on a constant scalar variance V .

Essentially, the variance $V = k_{5/2}(x = 0) = \langle f^2(\boldsymbol{\theta}) \rangle$ specifies to which extend any variable $f(\boldsymbol{\theta})$ is expected to deviate from its mean value (here 0 with our choice of mean function m). The length-scale l , appearing in the exponentially decaying term, scales the distance between the parameters $\boldsymbol{\theta}$ and $\boldsymbol{\theta}'$, and is indicative of the extent of the correlation between the corresponding variables $f(\boldsymbol{\theta})$ and $f(\boldsymbol{\theta}')$. In particular, the exponential decay in the distance x ensures that function values $f(\boldsymbol{\theta})$ and $f(\boldsymbol{\theta}')$ are highly correlated when close in parameter space ($x/l \lesssim 1$),

but that their correlations quickly fade away over larger distances. The values of V and l are considered as hyperparameters of the model that can be fitted to the data (this is discussed in the next section).

Other choices of kernel could be adopted. For instance, another *Matérn* function (similar to Equation (A.4) but with a different degree of the polynomial function depending on x) would permit to favour smoother or more irregular functions f [164]. Still, we find the choice of the kernel function in Equation (A.4) to be adequate for the problems that we have considered. We note that the *Matérn* 5/2 function is frequently used in other applications of BO, and is often considered as a standard choice of kernel [189].

A.3 Fitting the hyperparameters

Both the values of the variance V and of the correlation length l entering the kernel function in Equation (A.4), and the variance σ_N^2 of the Gaussian noise contained in the observations (which is discussed in details in Section 2.1.3) are considered as hyperparameters of the model. In general, it is not possible to have a precise idea of the values of these hyperparameters in advance, but, they can be fitted at any stage of the optimization to the available observations $\mathbf{y} = [y_1, \dots, y_D]$. These hyperparameters are typically chosen [164] to maximize the marginal likelihood (discussed in Equation (2.7)) $p(\mathbf{y}) = \int d\mathbf{f} p(\mathbf{f})p(\mathbf{y}|\mathbf{f})$, with the probability $p(\mathbf{f})$ depending (implicitly) on the values of the hyperparameters. Such optimization is implemented in any Gaussian processes library (for instance [265]) and can be performed automatically. In summary, the prior distribution is entirely defined by the choice of the functions k and m , discussed in the previous section, and of the hyperparameters V , and l , discussed in this section.

A.4 Predictive distribution

Given a choice of mean function m , kernel function k and corresponding hyperparameters we can evaluate the predictive distribution central to BO. The sought-after predictive distribution $p(f(\boldsymbol{\theta}^*)|\mathbf{y})$ permits to make (probabilistic) predictions about the values that can adopt the unknown function F at arbitrary parameters $\boldsymbol{\theta}^*$. Given the assumptions detailed in Section 2.1.3, this distribution admits the analytical expression [164]

$$p(f(\boldsymbol{\theta}^*)|\mathbf{y}) = \mathcal{N}(\mu_f(\boldsymbol{\theta}^*), \sigma_f^2(\boldsymbol{\theta}^*)),$$

with,
$$\begin{cases} \mu_f(\boldsymbol{\theta}^*) = \mathbf{k}_*(K + \sigma_n^2 I)^{-1} \mathbf{y} \\ \sigma_f^2(\boldsymbol{\theta}^*) = V - \mathbf{k}_*(K + \sigma_n^2 I)^{-1} \mathbf{k}_*^T \end{cases}$$
 (A.5)

where the row vector \mathbf{k}_* has entries $\mathbf{k}_*^{(j)} = k(\boldsymbol{\theta}^*, \boldsymbol{\theta}_j)$, and the covariance matrix K has entries $K_{i,j} = k(\boldsymbol{\theta}_i, \boldsymbol{\theta}_j)$. Note that, both the mean and standard deviation of this predictive distribution depend implicitly on $\boldsymbol{\theta}^*$ as the elements of \mathbf{k}_* are function of $\boldsymbol{\theta}^*$. Evaluation of this (Gaussian) predictive distribution is illustrated in Figure 2.2. The most computationally demanding part when evaluating values of the mean and variance in Equation (A.5) originates from the inversion of the matrix $K + \sigma_N^2 I$ which scales as D^3 , where D is the number of observations.

Chapter B

Estimation

In this appendix, we derive several expressions related to the estimation of the figures of merit. These include the optimal importance probability distribution (Appendix B.1), which was discussed in Section 2.2.2, and a higher bound on the variance of the local fidelity estimates (Appendix B.2), which was provided in Section 4.2.1.

B.1 Optimal importance probability distribution

Here, we derive the optimal importance distribution which was provided in Equation (2.31). Recall from Equation (2.28), that the variance $\Delta^2[Y_\rho]$, of an unbiased estimate Y_ρ (obtained with importance sampling) of the expectation value $\langle M \rangle_\rho$ of the operator M with respect to the state ρ , can be expressed as

$$\Delta^2[Y_\rho] = \sum_{i=1}^D \frac{\alpha^2(i)}{Pr(i)} \langle W_i^2 \rangle_\rho - \langle M \rangle_\rho^2. \quad (\text{B.1})$$

We aim at finding the importance probability distribution $Pr(i)$ which minimizes Equation (B.1). This corresponds to the optimization problem

$$Pr(i)^* = \arg \min_{\sum_i Pr(i)=1} \Delta^2[Y_\rho], \quad (\text{B.2})$$

under the additional constraints that for all i , $Pr(i) \geq 0$. This can be recast as the minimization of the Lagrangian function

$$\mathcal{L}(Pr(i), \lambda) = \Delta^2[Y_\rho] - \lambda \left(\sum_i Pr(i) - 1 \right), \quad (\text{B.3})$$

with respect to λ and each of the probabilities $Pr(i)$. Equating the derivatives $\partial\mathcal{L}/\partial Pr(i)$ and $\partial\mathcal{L}/\partial\lambda$ to zero, results in the set of conditions

$$\begin{cases} -\frac{\alpha^2(i)}{Pr(i)^2} \langle W_i^2 \rangle_\rho - \lambda = 0, \\ \sum_i Pr(i) - 1 = 0, \end{cases} \quad (\text{B.4})$$

to be simultaneously satisfied. It follows from the first line of Equation (B.4), that

$$Pr(i) = \pm \frac{|\alpha(i)\sqrt{\langle W_i^2 \rangle_\rho}|}{\sqrt{\lambda}}. \quad (\text{B.5})$$

Inserting these equalities into the second line of Equation (B.4) yields

$$\sqrt{\lambda} = \pm \sum_i |\alpha(i)\sqrt{\langle W_i^2 \rangle_\rho}|. \quad (\text{B.6})$$

Finally, given that each probability is greater or equal than 0, we obtain the sought-after optimal distribution

$$Pr(i)^* = \frac{|\alpha_i\sqrt{\langle W_i^2 \rangle_\rho}|}{\sum_i |\alpha_i\sqrt{\langle W_i^2 \rangle_\rho}|}. \quad (\text{B.7})$$

B.2 Estimation of the local fidelity

Here, we derive the bounds (Equation (4.8)) on the variance of estimates of the local fidelity, which was introduced in Section 4.2.1. Akin to the treatment provided in Section 4.2.1, these are derived assuming the target state to be the n -qubit $|0\rangle^{\otimes n}$, but the same bounds can be obtained for any separable target state.

First, recall the definition of the local fidelity as the average

$$\mathcal{F}_{loc}(\rho) = \frac{1}{n} \sum_{i=1}^n \mathcal{F}_i(\rho), \quad (\text{B.8})$$

of the probabilities

$$\mathcal{F}_i(\rho) = Tr[|0\rangle\langle 0|_i\rho], \quad (\text{B.9})$$

of measuring the qubit i of ρ in state $|0\rangle$.

The estimates \tilde{p}_i of each of the individual probabilities \mathcal{F}_i are the frequencies n_i/N of having observed n_i times qubit i in state $|0\rangle$ over N measurements. Each of these estimates has a variance $\Delta^2[\tilde{p}_i] = \mathcal{F}_i(1 - \mathcal{F}_i)/N$. An estimate $\tilde{\mathcal{F}}_{loc}$ of \mathcal{F}_{loc} is obtained by averaging the individual estimates over each of the qubits, that is, $\tilde{\mathcal{F}}_{loc} = 1/n \sum \tilde{p}_i$. It can be shown that the

variance in such an estimate is bounded by:

$$\begin{aligned}
\Delta^2[\tilde{\mathcal{F}}_{loc}] &= \Delta^2\left[\frac{1}{n} \sum \tilde{p}_i\right] \leq \frac{1}{n^2} \left(n \sum_{i=1}^n \Delta^2[\tilde{p}_i]\right) \\
&= \frac{1}{n} \sum_{i=1}^n \frac{\mathcal{F}_i(1 - \mathcal{F}_i)}{N} \\
&\leq \frac{1}{n} \sum_{i=1}^n \frac{\mathcal{F}_{loc}(1 - \mathcal{F}_{loc})}{N} \\
&\leq \frac{\mathcal{F}_{loc}(1 - \mathcal{F}_{loc})}{N}.
\end{aligned} \tag{B.10}$$

The first inequality is obtained in case of maximal correlation (that is, unit correlation) between each of the random variables \tilde{p}_i and is saturated when each $\Delta^2[\tilde{p}_i]$ has the same value. Inserting the expressions of the variance $\Delta^2[\tilde{p}_i]$ of the individual estimates \tilde{p}_i yields the equality in the second line. The new inequality in the third line arises from the maximization of the previous expression, with respect to the terms \mathcal{F}_i and under the constraints that $n\mathcal{F}_{loc} = \sum \mathcal{F}_i$. The maximal bound is saturated for $\mathcal{F}_i = \mathcal{F}_{loc}$ for all qubits i . Finally, the sum over the constant terms simplifies, yielding the inequality provided in Equation (4.8). It follows from these derivation steps, that the higher bound of the final inequality is saturated in case when (i) all the local probabilities \mathcal{F}_i are equal to \mathcal{F}_{loc} , and (ii) the measurement outcomes of each qubit are fully correlated. These two conditions are simultaneously satisfied when measuring states of the form $|\psi_{corr}\rangle \propto (|0\rangle^{\otimes n} + \alpha|1\rangle^{\otimes n})$. This is verified numerically in Figure 4.2.

Chapter C

Random unitaries and states

In this appendix, we first recall the definition of the Haar measure (Appendix C.1) over random unitaries and states. Then we review some of its properties (Appendix C.2) which permit us to derive analytical expressions of the average and variance of the expectation values of operators over an ensemble of random states. We then proceed to apply these formulas to the case of the witness fidelity which was studied in Section 4.2.2 (Appendices C.3 and C.4).

C.1 Random unitaries and states under the Haar measure

The normalized Haar measure $d\mu(U)$, over elements U of the unitary group $U(d)$, is defined [266] as the unique measure invariant by (right or left) action of any of the element of the group, such that

$$d\mu(VU) = d\mu(UV) = d\mu(U), \quad (\text{C.1})$$

for any $U, V \in U(d)$. Normalization further imposes that

$$\int d\mu(U) = 1, \quad (\text{C.2})$$

with integration, if not otherwise stated, always performed over $U(d)$. We call the ensemble of unitaries admitting such measure the random unitaries. Numerical sampling from this ensemble is implemented in many statistical packages, such as Scipy [188], or specialized quantum computing library such as Qutip [267].

We extend the definition of random unitaries to random states, which are defined as the distribution of states resulting from the action of random unitaries onto an arbitrary, but fixed, initial state $|\psi_0\rangle$, such that the pure state $|\psi\rangle = U|\psi_0\rangle$ has measure (dropping the ket notation)

$$d\mu(\psi) = d\mu(U), \quad (\text{C.3})$$

which, similarly to $d\mu(U)$, is invariant under action of any element of $U(d)$.

C.2 Statistics of expectation values under the Haar measure

The first two moments of $U(d)$ under the Haar measure admit simple analytical expressions (see [268] for instance), that we now recall. For a fixed basis $\{|0\rangle, \dots, |d-1\rangle\}$, the matrix elements of U are denoted $U_{i,j}$ (with conjugate $U_{i,j}^*$) such that $U = \sum U_{ij}|i\rangle\langle j|$ and $U^\dagger = \sum U_{ji}^*|i\rangle\langle j|$. The first moment under the Haar measure is expressed as

$$\int d\mu(U) U_{ij} U_{i'j'}^* = \frac{\delta_{ij}^{i'j'}}{d}, \quad (\text{C.4})$$

with $\delta_{\mathbf{i}}^{\mathbf{j}} = 1$ only if $\mathbf{i} = \mathbf{j}$ and 0 otherwise. The second moment is given by

$$\int d\mu(U) U_{ij} U_{kl} U_{i'j'}^* U_{k'l'}^* = \frac{\delta_{ijkl}^{i'j'k'l'} + \delta_{ijkl}^{k'l'i'j'}}{d^2 - 1} - \frac{\delta_{ijkl}^{k'j'i'l'} + \delta_{ijkl}^{i'l'k'j'}}{d(d^2 - 1)}. \quad (\text{C.5})$$

Equipped with Equations (C.4) and (C.5), one can simplify the expressions of the average and variances of the expectation values of an operator M . As usual, the expectation value of an operator M with respect to a state ψ is denoted $\langle M \rangle_\psi = \text{Tr}[\psi M] = \langle \psi | M | \psi \rangle$, and we denote $\langle M \rangle_{rdm}$ its average over random states, which is defined as

$$\langle M \rangle_{rdm} = \int d\mu(\psi) \langle M \rangle_\psi. \quad (\text{C.6})$$

Furthermore, the variance of $\langle M \rangle_\psi$ is denoted $D_{rdm}(M)$, which is defined as $D_{rdm}(M) = \int d\mu(\psi) [\langle M \rangle_\psi - \langle M \rangle_{rdm}]^2$, but which can also be expressed as

$$D_{rdm}(M) = \int d\mu(\psi) [\langle M \rangle_\psi^2 - \langle M \rangle_{rdm}^2]. \quad (\text{C.7})$$

We now proceed to solve Equations (C.6) and (C.7).

To facilitate the derivations, it is convenient to represent the random unitaries U in the eigenbasis of M , and we denote their entries $U_{ij} = \langle \lambda_i | U | \lambda_j \rangle$, given $M = \sum_i \lambda_i | \lambda_i \rangle \langle \lambda_i |$. The average

of $\langle M \rangle_\psi$ over a random states distribution can then be expressed as

$$\begin{aligned}
\langle M \rangle_{rdm} &= \int d\mu(\psi) \langle \psi | M | \psi \rangle = \int d\mu(U) \langle \lambda_1 | U M U^\dagger | \lambda_1 \rangle \\
&= \sum_{i=1}^d \lambda_i \int d\mu(U) \langle \lambda_1 | U | \lambda_j \rangle \langle \lambda_j | U^\dagger | \lambda_1 \rangle \\
&= \sum_{i=1}^d \lambda_i \int d\mu(U) U_{1j} U_{1,j}^* \\
&= \sum_{i=1}^d \frac{\lambda_i}{d} = \frac{\text{Tr}[M]}{d},
\end{aligned} \tag{C.8}$$

where the eigendecomposition of M is used to obtain the second line, and Equation (C.4) is used to simplify the third line. Similarly,

$$\begin{aligned}
\int d\mu(\psi) \langle M \rangle_\psi^2 &= \int d\mu(\psi) (\langle \psi | F | \psi \rangle)^2 \\
&= \int d\mu(U) \langle \lambda_1 | U M U^\dagger | \lambda_1 \rangle \langle \lambda_1 | U M U^\dagger | \lambda_1 \rangle \\
&= \sum_{i,j=1}^d \lambda_i \lambda_j \int d\mu(U) \langle \lambda_1 | U | \lambda_i \rangle \langle \lambda_i | U^\dagger | \lambda_1 \rangle \langle \lambda_1 | U | \lambda_j \rangle \langle \lambda_j | U^\dagger | \lambda_1 \rangle \\
&= \sum_{i,j=1}^d \lambda_i \lambda_j \int d\mu(U) U_{1j} U_{1k} U_{1,j}^* U_{1,k}^* \\
&= \sum_{i,j=1}^d \lambda_i \lambda_j \left[\frac{1 + \delta_i^j}{d^2 - 1} - \frac{1 + \delta_i^j}{d(d^2 - 1)} \right] \\
&= \frac{\sum_i \lambda_i^2 + \sum_{i,j} \lambda_i \lambda_j}{d(d+1)} = \frac{\text{Tr}[M^2] + \text{Tr}^2[M]}{d(d+1)},
\end{aligned} \tag{C.9}$$

where the eigendecomposition of M is used to obtain the third line, and where Equation (C.5) is used to obtain the fifth one. Finally, inserting the two previous expressions into Equation (C.7) yields the variance $D_{rdm}(F)$ of $\langle M \rangle_\psi$ over random states:

$$\begin{aligned}
D_{rdm}(M) &= \int d\mu(\psi) \langle M \rangle_\psi^2 - \langle M \rangle_{rdm} \\
&= \frac{\text{Tr}[M^2]}{d(d+1)} - \frac{\text{Tr}^2[M]}{d^2(d+1)},
\end{aligned} \tag{C.10}$$

which is reported in Equation (4.3) and enables the comparison of the variances of varied figures of merit (provided that they are constructed as expectation values of an operator with known eigenvalues) in Chapter 4.

C.3 Average and variance of the witness fidelity values

Here, we detail the derivations of the average and variance (over random states) of values of the witness fidelity \mathcal{F}_W , that was introduced in Section 4.2.2 as an alternative of the fidelity for stabilizer states. \mathcal{F}_W was defined as the expectation value of the operator

$$F_W = 1/2 + \mathcal{G}/2n, \quad (\text{C.11})$$

with n the number of qubits, and \mathcal{G} the sum of the n generators of a stabilizer group. Notably, the operator \mathcal{G} has $n+1$ eigenvalues $\lambda_j^{\mathcal{G}} = n-2j$ ($j = 0, \dots, n$), with corresponding eigenspaces of dimension $\binom{n}{j}$ (this is detailed in Appendix D.2). Given these details, we can proceed to evaluate its average and variance.

To apply the formulas of the average and variance, established in Equations (C.8) and (C.10) respectively, to \mathcal{G} we need to evaluate the traces of the operators \mathcal{G} and \mathcal{G}^2 . For that purpose, we first recall formulas of sums with binomial coefficients. In particular,

$$\sum_{j=0}^n \binom{n}{j} = 2^n, \quad \sum_{k=0}^n \binom{n}{j} j = n2^{n-1} \quad \text{and,} \quad \sum_{j=0}^n \binom{n}{j} j^2 = (n^2 + n)2^{n-2}. \quad (\text{C.12})$$

It follows that

$$Tr[\mathcal{G}] = \sum_{j=0}^n \binom{n}{j} (n-2j) = (n2^n - 2n2^{n-1}) = 0, \quad (\text{C.13})$$

and that

$$\begin{aligned} Tr[\mathcal{G}^2] &= \sum_{j=0}^n \binom{n}{j} (n-2j)^2 = \sum_{j=0}^n \binom{n}{j} (n^2 - 4nj + 4j^2) \\ &= n^2 2^n - 4n^2 2^{n-1} + 4(n^2 + n)2^{n-2} = nd, \end{aligned} \quad (\text{C.14})$$

where $d = 2^n$.

It follows from Equations (C.8) and (C.13) that the average of the operator \mathcal{G} is $\langle \mathcal{G} \rangle_{rdm} = 0$, such that the average of the witness fidelity \mathcal{F}_W (which is shifted by $1/2$) is equal to

$$\langle F_W \rangle_{rdm} = \frac{1}{2}. \quad (\text{C.15})$$

It follows from Equations (C.10), (C.13) and (C.14) that the variance of the operator \mathcal{G} is $D_{rdm}(\mathcal{G}) = n/(d+1)$, and that the variance of the witness fidelity \mathcal{F}_W (which is rescaled by a

factor $1/4n^2$ due to the $1/2n$ scaling in the definition of F_W) is equal to

$$D_{rdm}(F_W) = \frac{1}{4n(d+1)}, \quad (\text{C.16})$$

which is reported in Section 4.2.2.

C.4 Average of $\langle \mathcal{G}_Z^2 \rangle$

When estimating the witness fidelity for the GHZ state, it is only necessary to perform measurements in two distinct basis (Section 4.2.2). In order to evaluate the importance sampling probabilities (that is, the probabilities to take a measurement in one of the two specific basis), which was given in Equation (2.31), it is necessary to evaluate the statistics $\langle \mathcal{G}_Z^2 \rangle_\psi$ of the operator $\mathcal{G}_Z = \sum Z_{j-1}Z_j$. As these depend on the unknown state ψ , it needs to be approximated by, for instance, the average $\langle \mathcal{G}_Z^2 \rangle_{rdm}$ of $\langle \mathcal{G}_Z^2 \rangle_\psi$ taken over random states, and we now proceed to evaluate this quantity.

\mathcal{G}_Z can be eigendecomposed over n distinct eigensubspaces with eigenvalues $\lambda_j = n - 1 - 2j$ ($j = 0, \dots, n-1$) and with dimension $2\binom{n}{j}$ (this is detailed in Appendix D.3). It follows from this eigendecomposition and Equation (C.8), that the average $\langle \mathcal{G}_Z^2 \rangle_{rdm} = \int d\mu(\psi) \langle \mathcal{G}_Z^2 \rangle_\psi$ can be expressed as:

$$\begin{aligned} \langle \mathcal{G}_Z^2 \rangle_{rdm} &= \frac{Tr[\mathcal{G}_Z^2]}{d} = \frac{1}{d} \sum_{j=0}^{n-1} 2\binom{n}{j} (n-1-2j)^2 \\ &= \frac{2}{d} \sum_{j=0}^{n-1} \binom{n}{j} [(n-1)^2 - 4(n-1)j + 4j^2] \\ &= \frac{2}{d} [(n-1)^2 2^{(n-1)} - 4(n-1)^2 2^{n-2} + 4((n-1)^2 + (n-1)) 2^{n-3}] \\ &= n-1, \end{aligned} \quad (\text{C.17})$$

which is used in Section 4.2.2 to derive the importance probabilities reported in Equation (4.17).

C.5 Average variance in the estimates of \mathcal{F}_W for the GHZ state

Finally we can also evaluate the variance in an estimate Y_ρ of \mathcal{G} obtained with importance sampling over the two measurement settings (corresponding to measurements of the operators \mathcal{G}_X and \mathcal{G}_Z defined in Equation (4.16)), and averaged over random states. For that purpose,

recall that $\langle \mathcal{G}_X^2 \rangle_\psi = 1$ and that $\mathcal{G} = \mathcal{G}_X + \mathcal{G}_Z$. This permits to simplify the general variance, given in Equation (2.32), to

$$\begin{aligned}\Delta^2[Y_\psi] &= \left(1 + \sqrt{\langle \mathcal{G}_Z^2 \rangle_\psi}\right)^2 - \langle \mathcal{G} \rangle_\psi^2 \\ &= 1 + 2\sqrt{\langle \mathcal{G}_Z^2 \rangle_\psi} + \langle \mathcal{G}_Z^2 \rangle_\psi - \langle \mathcal{G} \rangle_\psi^2.\end{aligned}\tag{C.18}$$

While we already know how to average over the term $\langle \mathcal{G}_Z^2 \rangle_\psi$ (see Equation (C.17)) we also need to perform this average over $\sqrt{\langle \mathcal{G}_Z^2 \rangle_\psi}$, which does not seem to admit analytical solution. Still, Jensen's inequality permits to bound this average as

$$\int d\mu(\psi) \sqrt{\langle \mathcal{G}_Z^2 \rangle_\psi} \leq \sqrt{\int d\mu(\psi) \langle \mathcal{G}_Z^2 \rangle_\psi} = \sqrt{\langle \mathcal{G}_Z^2 \rangle_{rdm}}\tag{C.19}$$

It follows that the average variance $\Delta_{rdm}^2[Y_\psi] = \int d\mu(\psi) \Delta^2[Y_\psi]$ can be bounded as

$$\Delta_{rdm}^2[Y_\psi] \leq 1 + 2\sqrt{\langle \mathcal{G}_Z^2 \rangle_{rdm}} + \langle \mathcal{G}_Z^2 \rangle_{rdm} - \int d\mu(\psi) \langle \mathcal{G} \rangle_\psi^2,\tag{C.20}$$

where each term in the r.h.s. can be evaluated. Already $\langle \mathcal{G}_Z^2 \rangle_{rdm}$ follows from Equation (C.17), and it only remains to assess $\int d\mu(\psi) \langle \mathcal{G} \rangle_\psi^2$. Given the expressions of the trace of \mathcal{G} (Equation (C.13)) and \mathcal{G}^2 (Equation (C.14)), Equation (C.9) yields

$$\int d\mu(\psi) \langle \mathcal{G} \rangle_\psi^2 = \frac{Tr[\mathcal{G}^2] + Tr^2[\mathcal{G}]}{d(d+1)} = \frac{n}{d+1}.\tag{C.21}$$

It follows that

$$\Delta_{rdm}^2[Y_\psi] \leq n + 2\sqrt{n-1} - \frac{n}{d+1}.\tag{C.22}$$

Finally, recalling that one needs to rescale \mathcal{G} by $1/2n$ to obtain F_W , we obtain the higher bound for the variance (averaged over random states) of an estimate $\tilde{\mathcal{F}}_W$ of \mathcal{F}_W

$$\Delta_{rdm}^2[\tilde{\mathcal{F}}_W] \leq \frac{1}{4n} + \frac{\sqrt{n-1}}{2n^2} - \frac{1}{4n(d+1)},\tag{C.23}$$

which was reported in Equation (4.18) and plotted in Figure 4.4.

Chapter D

Stabilizer states

In this appendix, we start (Appendix D.1) by reviewing elements of the stabilizer formalism [15, 217], before detailing the eigendecompositions of the operators \mathcal{G} (Appendix D.2) and \mathcal{G}_Z (Appendix D.3), which permit the construction and the study of the witness fidelity developed in Section 4.2.2.

D.1 Elements of the stabilizer formalism

As a starting point, consider the set of 1-qubit matrices, formed by the Pauli matrices (X, Y and Z) and the identity (I) with prefactor ± 1 or $\pm i$,

$$\mathcal{P}_1 = \{\pm I, \pm iI, \pm X, \pm iX, \pm Y, \pm iY, \pm Z, \pm iZ\}. \quad (\text{D.1})$$

This set forms a group (under matrix multiplication), called the Pauli group. It is extended to n -qubit systems as

$$\mathcal{P}_n = \{P_1 \otimes \dots \otimes P_n \mid P_i \in \mathcal{P}_1\} \quad (\text{D.2})$$

containing any distinct n -fold tensor product of elements of \mathcal{P}_1 .

A state which remains invariant when acted on by any element of a group of operators is said to be stabilized by such group, that is,

$$|\psi\rangle \text{ is stabilized by } S \Leftrightarrow \forall O \in S, O|\psi\rangle = |\psi\rangle, \quad (\text{D.3})$$

with S a group of operators O . Similarly, the subspace formed by all the states that are stabilized by a group S , is called the stabilized subspace and denoted V_S . Finally, given a set of l operators G_i , we denote the group generated by such operators (that is, the group which contains all possible products of G_i) $\langle G_1, \dots, G_l \rangle$.

Having established the notations and terminology needed, we can proceed to recall one of the main aspect of the stabilizer formalism. Given $S = \langle G_1, \dots, G_l \rangle$, it can be shown [217] that

if the l generators $G_i \in \mathcal{P}_n$ (i) are independent (ii) commute and if (iii) $-I \notin S$, then the vector space V_S stabilized by S has dimension $\dim(V_S) = 2^{n-l}$. In particular, for n generators G_i satisfying (i-iii), there exists a single state (up to a global phase) which is stabilized by the group S . In such case, we denote this state $|\psi_S\rangle$. Finally, we note that condition (iii) entails that each generator G_i (and each element of the group S) is a Hermitian operator, such that its expectation value can be estimated experimentally.

These properties of a stabilized state $|\psi_S\rangle$ permitted us, in Section 4.2.2, to identify the sum of its n generators G_i ,

$$\mathcal{G} = \sum_{i=1}^n G_i, \quad (\text{D.4})$$

as an observable which expectation value is maximized only for $|\psi_S\rangle$, and, thus which can be used to construct a figure of merit alternative to the fidelity.

D.2 Eigendecomposition of \mathcal{G}

In this section, we study further the properties of the operator \mathcal{G} (corresponding to a stabilizer state $|\psi_S\rangle$) and follow the derivations presented in [269]. For that purpose, we first define the set of states $\{|\psi_S^{\mathbf{i}}\rangle\}$, indexed by a n -bit string \mathbf{i} , as the states stabilized by the groups $S^{\mathbf{i}} = \langle(-1)^{\mathbf{i}_1}G_1, \dots, (-1)^{\mathbf{i}_n}G_n\rangle$. Note that the state $|\psi_S^0\rangle$ is the original stabilizer state $|\psi_S\rangle$. We now show how these states can be used to decompose \mathcal{G} .

By construction, the 2^n (distinct) states $|\psi_S^{\mathbf{i}}\rangle$ are eigenvectors of any of the operators G_l , with eigenvectors 1 (or -1) if $\mathbf{i}_l = 0$ (or 1). Hence, one can always eigendecompose G_l as

$$G_l = \sum_{\mathbf{i}} \lambda_{\mathbf{i}}^{G_l} |\psi_S^{\mathbf{i}}\rangle \langle \psi_S^{\mathbf{i}}|, \text{ with } \lambda_{\mathbf{i}}^{G_l} = 1 - 2\mathbf{i}_l. \quad (\text{D.5})$$

Given that \mathcal{G} is the sum of the n operators G_l , it admits the decomposition

$$\mathcal{G} = \sum_{l=1}^n G_l = \sum_{\mathbf{i}} \lambda_{\mathbf{i}}^{\mathcal{G}} |\psi_S^{\mathbf{i}}\rangle \langle \psi_S^{\mathbf{i}}|, \text{ with } \lambda_{\mathbf{i}}^{\mathcal{G}} = n - 2 \sum_l \mathbf{i}_l. \quad (\text{D.6})$$

That is, \mathcal{G} has $n+1$ distinct eigenvalues $\lambda_{\mathbf{i}}^{\mathcal{G}} = n - 2j \in [-n, n]$ ($j = 0, \dots, n$) and corresponding eigensubspaces E_j with dimension $\dim(E_j) = \binom{n}{j}$. It directly follows that the operator $F_W = \frac{1}{2} + \frac{\mathcal{G}}{2n}$, used to define the witness fidelity, has $n+1$ distinct eigenvalues $\lambda_j^G = 1 - j/n \in [0, 1]$ and corresponding eigensubspaces with dimension $\binom{n}{j}$.

D.3 Eigendecomposition of \mathcal{G}_Z

Recall the definition of the operator \mathcal{G}_Z (given in Equation (4.16)) as the sum of the $n - 1$ operators $Z_{j-1}Z_j$ ($j = 2, \dots, n$)

$$\mathcal{G}_Z = \sum_{i=2}^n Z_{i-1}Z_i. \quad (\text{D.7})$$

Given that the operators $G_{j-1} = Z_{j-1}Z_j$ are $n - 1$ (out of n) generators stabilizing the GHZ state, it can be shown (following the same derivations as in the previous section) that

$$\mathcal{G}_Z = \sum_{l=1}^{n-1} G_l = \sum_{\mathbf{i}} \lambda_{\mathbf{i}}^{\mathcal{G}} |\psi_S^{\mathbf{i}}\rangle\langle\psi_S^{\mathbf{i}}|, \text{ with } \lambda_{\mathbf{i}}^{\mathcal{G}} = n - 1 - 2 \sum_{l=1}^{n-1} \mathbf{i}_l. \quad (\text{D.8})$$

That is, \mathcal{G}_Z has n distinct eigenvalues $\lambda_{\mathbf{i}}^{\mathcal{G}} = n - 1 - 2j \in [1 - n, n - 1]$ ($j = 0, \dots, n - 1$) and corresponding eigensubspaces E_j with dimension $\dim(E_j) = 2 \binom{n}{j}$.