



Title:

Single-component gradient rules for variational quantum algorithms

Author(s):

Thomas Hubregtsen, Frederik Wilde, Shozab Qasim, Jens Eisert

Document type: Postprint

Terms of Use: Copyright applies. A non-exclusive, non-transferable and limited right to use is granted. This document is intended solely for personal, non-commercial use.

Citation:

"Thomas Hubregtsen, Frederik Wilde, Shozab Qasim, Jens Eisert
Quantum Sci. Technol., 7 (3), 035008 ; <https://iopscience.iop.org/article/10.1088/2058-9565/ac6824>"
Archiviert unter <http://dx.doi.org/10.17169/refubium-38822>

Single-component gradient rules for variational quantum algorithms

Thomas Hubregtsen^{1*}, Frederik Wilde^{1*}, Shozab Qasim¹, Jens Eisert^{1,2}

¹ Dahlem Center for Complex Quantum Systems, Freie Universität Berlin, 14195 Berlin, Germany

² Helmholtz-Zentrum Berlin für Materialien und Energie, 14109 Berlin, Germany

Abstract

Many near-term quantum computing algorithms are conceived as variational quantum algorithms, in which parameterized quantum circuits are optimized in a hybrid quantum-classical setup. Examples are variational quantum eigensolvers, quantum approximate optimization algorithms as well as various algorithms in the context of quantum-assisted machine learning. A common bottleneck of any such algorithm is constituted by the optimization of the variational parameters. A popular set of optimization methods work on the estimate of the gradient, obtained by means of circuit evaluations. We will refer to the way in which one can combine these circuit evaluations as gradient rules. This work provides a comprehensive picture of the family of gradient rules that vary parameters of quantum gates individually. The most prominent known members of this family are the parameter shift rule and the finite differences method. To unite this family, we propose a generalized parameter shift rule that expresses all members of the aforementioned family as special cases, and discuss how all of these can be seen as providing access to a linear combination of exact first- and second-order derivatives. We further prove that a parameter shift rule with one non-shifted evaluation and only one shifted circuit evaluation does not exist, and introduce a novel perspective for approaching new gradient rules.

1 Introduction

Quantum computing has enjoyed a continuous growth in attention, increasingly so since quantum computers are now able to perform tasks promised to be intractable for classical devices [1–3]. While these recent experiments impressively showcase the remarkable progress in hardware design and manufacturing, it is unclear on which type of problems near-term quantum computers, without the ability to perform full quantum error correction, can possibly achieve a practical advantage [4]. Still, as more and more small-scale quantum computers become available, the search for a task for which one can reasonably expect a practical quantum advantage is gaining traction.

Various quantum algorithms are being explored as a potential mechanism for achieving a real-world advantage. The automotive industry attempts to leverage the quantum approximate optimization algorithm [5] for finance and logistics [6]. The aerospace industry investigates quantum machine learning techniques for the simulation of fluid dynamics [7]. In the field of quantum chemistry, various readings of the variational quantum eigensolver [8] are being examined to approximate molecular ground-state energies of molecules [9–13]. Variational quantum-assisted machine learning algorithms [14], or instances of quantum neural networks, are also examined for practical advantage over classical neural networks [10]. Closely linked are kernel methods [15, 16] and trainable kernel methods [17]. One can also go a step further and use variational quantum circuits to design other applications of quantum technologies, such as in quantum metrology [18, 19], so that one quantum device is delivering a blueprint for another.

All these algorithms are commonly conceived as hybrid quantum-classical algorithms [20] in which a quantum circuit depends on a set of free parameters which are typically optimized in an iterative loop controlled by a classical optimizer. This for good reason, as any reasonable quantum architecture in the foreseeable future has limited depth, with the gates fully dedicated to the task at hand. The optimization is done by defining an appropriate cost function as the expectation value of some observable with respect to the output state of the parameterized circuit. While this idea appears straightforward, it turned out that achieving suitably efficient and effective classical control of hybrid algorithms is a challenging task, often giving rise a bottleneck for the application

*These authors contributed equally to this work.

of such algorithms. This applies in particular to settings in which the evaluation of expectation values is costly, such as in cold atomic architectures, where a single measurement can take some tens of seconds, so that meaningful variational algorithms can only be implemented using the most sophisticated tools of classical control.

While several optimization methods for cost functions of this kind have been studied, it remains an open task to find the optimal method given a specific problem and algorithm [21]. In recent years, first-order optimization methods, utilizing the gradient of the cost function, have increasingly gained traction in the field of variational quantum algorithms. Most notable methods for computing gradient components are finite difference methods for approximation and the parameter shift rule [22–29] to calculate the components exactly. Alternative methods entail the quantum natural gradient [30], which calculates the gradient in the distribution space as opposed to the parameter space, and simultaneous perturbation stochastic approximation [31], which estimates multiple parameters at once. These approaches can be combined with approaches aimed at maximizing the information gain from the coherence available to speed up estimation of the expectation values [32].

A first version of the parameter shift rule was first used in the area of control-pulse theory [22]. It entailed two circuit evaluations with a shift of $\pi/2$ for the parameter under evaluation. Similar approaches have been employed in subsequent improvements [23–25]. This work was expanded to hold for any unitary gate $U_G = e^{i\theta G}$ whose Hermitian generator G has two distinct eigenvalues in a circuit $f(\theta) = \langle \psi | U_G^\dagger(\theta) A U_G(\theta) | \psi \rangle$ with observable A [26]. Gates that have generators with more than two unique eigenvalues can be handled by decomposing the gates [27]. A stochastic algorithm for estimating the gradient of any multi-qubit parametric quantum evolution through a mathematically exact formula, without the need for ancillary qubits or the use of Hamiltonian simulation techniques followed [28]. Subsequently, the estimation of higher-order derivatives was proposed, together with a derivation of a formula that allows for gate-independent shifts of opposing angle [29]. We will use the variable γ , to indicate the value of a shift that can be freely chosen, e.g., independent from the characteristics of the gate. This leads to the current centered parameter shift rule (cPSR) to calculate the analytic gradient $f'(\theta)$ using circuit evaluations as

$$g_{\text{cPSR}}(\theta, \gamma) := \frac{r[f(\theta + \gamma) - f(\theta - \gamma)]}{\sin(2r\gamma)}, \quad (1)$$

for any γ that is not an integer multiple of $\pi/(2r)$, where $2r$ is the difference of the eigenvalues of G .

In this work, we unite all known members of the family of methods that use circuit evaluations to calculate each component of the gradient vector individually, by introducing a generalized parameter shift rule. This also yields the variance and bias of all previous methods. Furthermore, we close the discussion on the potential existence of a forward or backward parameter shift rule by proving its non-existence.

We start our work by describing the underlying setting in Section 2. The generalized parameter shift rule is presented in Section 3, and show its relation to the finite differences and parameter shift rule. Section 4 discusses the no-go results for the forward and backward parameter shift rule. We end this work with a discussion of the type of data any of these methods delivers and outlook on future research in Section 5, on how a general picture of gradient rules should be conceived.

2 Setting

Variational algorithms as considered in this work consist of *parameterized quantum circuits*, mapping $\theta \mapsto U(\theta)$, that are defined by the quantum circuit topology and variational parameters $\theta = (\theta_1, \dots, \theta_n)$. Note that in the case of embedding data x , the circuit can be imagined to have an implicit dependence of the data $(\theta, x) \mapsto U(\theta, x)$, where x can be assumed to be constant for the purpose of this work. In order to optimize the circuit we first need to specify a cost function. Cost functions are in general not merely expectation values but functions thereof, but for the purposes of the present work it is sufficient to investigate expectation values as we can use the chain rule.

Definition 1 (Cost function). For variational parameters $\theta = (\theta_1, \dots, \theta_n)$, a parameterized quantum circuit $\theta \mapsto U(\theta)$, an observable A , and a quantum state vector $|\psi\rangle$ reflecting a pure initial state, a cost function $F : \mathbb{R}^n \rightarrow \mathbb{R}$ is defined as

$$F(\theta) := \langle \psi | U(\theta) A U^\dagger(\theta) | \psi \rangle. \quad (2)$$

A single-shot estimator $\hat{F}(\boldsymbol{\theta})$ for the cost function $F(\boldsymbol{\theta})$ corresponds to a single measurement of the Hermitian observable A when the system is described by the state vector $U(\boldsymbol{\theta})|\psi\rangle$. Following these lines, the n -shot estimator $\hat{F}^{(n)}(\boldsymbol{\theta})$ is the sample mean of n i.i.d. samples of the single-shot estimator.

We then complete the variational algorithm by searching for updates to the parameters $\boldsymbol{\theta}$ such that the cost function is minimized. In the entirety of this work, we will focus on optimizing the cost function F by estimating its gradient $\nabla_{\boldsymbol{\theta}}F$. Gradient rules are strategies that help estimate the gradient, leveraging carefully chosen circuit evaluations exploiting symmetries in the problem. These methods dictate shifts for the parameters of the circuit, and how to combine the resulting cost function measurements to form the gradient.

Definition 2 (Gradient rule). To estimate the i -th component of the gradient at the parameter point $\boldsymbol{\theta} \in \mathbb{R}^n$, a gradient rule g requires the evaluation of the cost function F at shifted parameter points $\boldsymbol{\theta} + \mathbf{s}_1, \dots, \boldsymbol{\theta} + \mathbf{s}_k$ (depending on i) and describes how the resulting measurements form this gradient component, i.e.

$$g(\boldsymbol{\theta}, \mathbf{s}_1, \dots, \mathbf{s}_k) = \frac{\partial F}{\partial \theta_i}. \quad (3)$$

Corresponding to the gradient rule g , we define the estimator \hat{g} where all cost function evaluations $F(\boldsymbol{\theta} + \mathbf{s})$ are replaced by single-shot measurement estimates $\hat{F}(\boldsymbol{\theta} + \mathbf{s})$.

By applying gradient rules multiple times, one can recover higher order derivatives such as the Hessian. In the next section, we will exclusively focus on the family of gradient rules that alter each parameter individually. We will refer to these as single-component gradient rules. As an example, in the case of the original parameter shift rule [24] for the i -th component, we have $\mathbf{s}_1 = \pi/(4r)\mathbf{e}_i$ and $\mathbf{s}_2 = -\pi/(4r)\mathbf{e}_i$, where \mathbf{e}_i is the i -th canonical basis vector.

3 Single-component gradient rules

For the scenario of gradient rules that alter the parameters individually, we can now define a single-component cost function for the i -th component of the parameter vector, defined as

$$f(\theta) := F(\theta_1, \dots, \theta_{i-1}, \theta, \theta_{i+1}, \dots, \theta_n) = \langle \psi | U(\boldsymbol{\theta}) A U(\boldsymbol{\theta})^\dagger | \psi \rangle, \quad (4)$$

where all parameters, except θ_i are fixed. This gives rise to a family of single-component gradient rules giving access to $f' = \partial f / \partial \theta$. To rigorously state the (generalized) parameter shift rule, we define r -gates and the set of circuits parameterized by them.

Definition 3 (r -gate). A parameterized gate $e^{-i\theta G}$ is called an r -gate if its Hermitian generator G has exactly two eigenvalues e_0 and e_1 , such that $e_0 - e_1 = 2r$.

Definition 4 (Set of r -gate parameterized expectation values). Let $q \geq 1$ be the number of qubits in the system at hand. Then, the set \mathcal{F}_r of all r -gate parameterized expectation values on this system is defined as follows. A single-parameter family of expectation values $f : [0, 2\pi/r) \rightarrow \mathbb{R}$ is contained in \mathcal{F}_r if and only if there exists a state vector $|\psi\rangle \in \mathcal{H} := \mathbb{C}^{\otimes q}$, an r -gate $e^{-i\theta G}$, and a Hermitian operator A acting on \mathcal{H} , such that for all $\theta \in [0, 2\pi/r)$

$$f(\theta) = \langle \psi | e^{i\theta G} A e^{-i\theta G} | \psi \rangle. \quad (5)$$

Remark 4.1 (Convention for the spectrum of generators). We assume $e_0 = -e_1 = r$ for the purposes of this work [26]. If this was not the case, we could simply add a phase shift $e^{i\theta(e_0 - e_1)/2}$ to the gate $e^{-i\theta G}$, which would have no effect on the expectation value. With this assumption, we can express the gate as

$$e^{-i\theta G} = \cos(r\theta) - i \frac{G}{r} \sin(r\theta), \quad (6)$$

where we have exploited the fact that $G^2 = r^2 \mathbb{1}$.

In particular, we will be now be taking a closer look at and generalizing the *centered*, *backward* and *forward finite difference rules* (cFD, bFD and fFD respectively), as well as the *centered parameter*

shift rule (cPSR) [22–29]. In the above language, they can be stated as

$$g_{\text{cFD}}(\theta, h) := \frac{f(\theta + h) - f(\theta - h)}{2h} \simeq f'(\theta), \quad (7)$$

$$g_{\text{bFD}}(\theta, h) := \frac{f(\theta) - f(\theta - h)}{h} \simeq f'(\theta), \quad (8)$$

$$g_{\text{hFD}}(\theta, h) := \frac{f(\theta + h) - f(\theta)}{h} \simeq f'(\theta), \quad (9)$$

$$g_{\text{cPSR}}(\theta, \gamma) := \frac{r[f(\theta + \gamma) - f(\theta - \gamma)]}{\sin(2r\gamma)} = f'(\theta). \quad (10)$$

Here, h and γ represent the shift of the parameter, where h is small compared to the period of f while γ can take any value except $k\pi/(2r)$ for any $k \in \mathbb{Z}$. It is important to note that while the first three rules aim at approximating the first derivative, the last rule provides an exact expression on the level of expectation values for all $\gamma \in \mathbb{R}$. In the following definition, we will present the generalized parameter shift rule, which we show in the subsequent proof to express the first- and second-order derivatives. In the subsequent corollaries, we show how the generalized parameter shift rule relates to the gradient rules in this family, as well as their relation to the (higher-order) derivatives. This will also include exact expressions for the bias, whereas previously the asymptotic relation in equations (7-9) could only be approximated.

Definition 5 (Generalized parameter shift rule). The generalized parameter shift rule (gPSR) is defined as

$$g_{\text{gPSR}}(\theta, \gamma_1, \gamma_2) := r [f(\theta + \gamma_1) - f(\theta - \gamma_2)], \quad (11)$$

for $\gamma_1, \gamma_2 \in \mathbb{R}$, where $f \in \mathcal{F}_r$ is an r -gate parameterized expectation value (see Definition 4).

Theorem 1 (Deriving rules from the generalized parameter shift rule). Given two shifts $\gamma_1, \gamma_2 \in \mathbb{R}$, the gPSR as defined above relates to the first and second derivative of f as

$$g_{\text{gPSR}}(\theta, \gamma_1, \gamma_2) = \frac{\sin(2r\gamma_1) + \sin(2r\gamma_2)}{2} f'(\theta) - \frac{\cos(2r\gamma_1) - \cos(2r\gamma_2)}{4r} f''(\theta). \quad (12)$$

Proof. We first need to decompose the r -gates in γ and θ -dependent gates as

$$e^{-i(\theta+\gamma)G} = e^{-i\theta G} e^{-i\gamma G}. \quad (13)$$

The gate depending on θ can be absorbed to simplify the equation, by defining $|\phi(\theta)\rangle := e^{-i\theta G} |\psi\rangle$. In contrast, the gate depending on γ needs to be decomposed using Euler's identity from equation (6), as done previously in Ref. [27], but here with unequal shifts γ_1 and γ_2 , in order to get

$$g_{\text{gPSR}}(\theta, \gamma_1, \gamma_2) = r \left\{ \langle \phi(\theta) | \left[\cos(r\gamma_1) \mathbb{1} + i \frac{a}{r} \sin(r\gamma_1) G \right] A \left[\cos(r\gamma_1) \mathbb{1} - i \frac{a}{r} \sin(r\gamma_1) G \right] | \phi(\theta) \rangle \right. \\ \left. - \langle \phi(\theta) | \left[\cos(r\gamma_2) \mathbb{1} + i \frac{a}{r} \sin(r\gamma_2) G \right] A \left[\cos(r\gamma_2) \mathbb{1} + i \frac{a}{r} \sin(r\gamma_2) G \right] | \phi(\theta) \rangle \right\}. \quad (14)$$

We can now expand and reorganize this equation into six terms. Note that $\langle \phi(\theta) | A | \phi(\theta) \rangle = f(\theta)$ and $\langle \phi(\theta) | GAG | \phi(\theta) \rangle / r^2 = f(\theta + \pi/(2r))$ together can be combined into the second derivative using the double-angle formula $\cos^2(x) = (\cos(2x) + 1)/2$. The remaining terms involving $\langle \phi(\theta) | [iG, A] | \phi(\theta) \rangle$ form the first derivative. This gives

$$g_{\text{gPSR}}(\theta, \gamma_1, \gamma_2) = r \cos^2(r\gamma_1) \langle \phi(\theta) | A | \phi(\theta) \rangle - r \cos^2(r\gamma_2) \langle \phi(\theta) | A | \phi(\theta) \rangle \\ + \frac{1}{r} \sin^2(r\gamma_1) \langle \phi | GAG | \phi(\theta) \rangle - \frac{1}{r} \sin^2(r\gamma_2) \langle \phi(\theta) | GAG | \phi(\theta) \rangle \\ + \cos(r\gamma_1) \sin(r\gamma_1) \langle \phi(\theta) | [iG, A] | \phi(\theta) \rangle + \cos(r\gamma_2) \sin(r\gamma_2) \langle \phi(\theta) | [iG, A] | \phi(\theta) \rangle \\ = \frac{\sin(2r\gamma_1) + \sin(2r\gamma_2)}{2} f'(\theta) - \frac{\cos(2r\gamma_1) - \cos(2r\gamma_2)}{4r} f''(\theta). \quad (15)$$

□

Corollary 1.1 (Uses of the generalized PSR). The generalized PSR, when multiplied with a pre-factor and for the listed values for shift angles γ_1 and γ_2 , can express any higher-order derivative.

Proof. This statement immediately follows from equation (15) using shift angles $\gamma_1, \gamma_2 = \gamma$, resulting in

$$f'(\theta) = \frac{g_{\text{gPSR}}(\theta, \gamma, \gamma)}{\sin(2r\gamma)} \quad (16)$$

for the first derivative and

$$f''(\theta) = 2 g_{\text{gPSR}}\left(\theta, \frac{\pi}{2r}, 0\right) \quad (17)$$

for the second derivative. Note that equation (16) is the standard cPSR. For arbitrary $n \in \mathbb{N}$, higher order derivatives of f satisfy

$$f^{(n+2)} = -\frac{1}{4r^2} f^{(n)}, \quad (18)$$

where $f^{(n)}$ refers to the n -th derivative of f . This is an important insight, in that it shows that all expressions involving derivatives of arbitrary order can be written as a linear combination of first and second derivatives of the cost function f , which in turn is the data that any of the above rule exactly delivers on the level of expectation values. All higher-order derivatives have been previously identified as following from applying the centered PSR multiple times [29]. \square

Corollary 1.2 (Further uses of the generalized PSR). *The generalized PSR, multiplied by a known pre-factor and with shift angles $\gamma_1 = \gamma_2$ is equivalent to the centered finite differences method, and provides an exact expression for the bias and variance of the centered finite differences.*

Proof. By setting $\gamma_1 = \gamma_2 = h$ we have

$$g_{\text{gPSR}}(\theta, h, h) = 2rh g_{\text{cFD}}(\theta, h), \quad (19)$$

for the centered finite differences, while the bias of the centered finite differences becomes

$$\text{Bias}(\hat{g}_{\text{cFD}}(\theta, h)) = \mathbb{E}[\hat{g}_{\text{cFD}}(\theta, h)] - f'(\theta) = \frac{g_{\text{gPSR}}(\theta, h, h)}{2rh} - f'(\theta) \quad (20)$$

$$= \left[\frac{\sin(2rh)}{2rh} - 1 \right] f'(\theta). \quad (21)$$

The variance was already known to be

$$\text{Var}(\hat{g}_{\text{cFD}}(\theta, h)) = \frac{\sigma_1^2(\theta + h) + \sigma_1^2(\theta - h)}{4h^2}, \quad (22)$$

where σ_1^2 is the one-shot variance [29]. \square

Corollary 1.3 (Backward and forward finite differences from the generalized PSR). *The generalized PSR, multiplied by a known pre-factor and with shift angles 0 and h , is equivalent to both the backward and forward finite differences, and provides an exact expression for the bias and variance of these finite difference methods.*

Proof. For the backward FD we set $\gamma_1 = 0$ and $\gamma_2 = h$, to get

$$g_{\text{gPSR}}(\theta, 0, h) = h g_{\text{bFD}}(\theta, h), \quad (23)$$

$$\begin{aligned} \text{Bias}(\hat{g}_{\text{bFD}}(\theta, h)) &= \mathbb{E}[\hat{g}_{\text{bFD}}(\theta, h)] - f'(\theta) = \frac{g_{\text{gPSR}}(\theta, 0, h)}{h} - f'(\theta) \\ &= \left[\frac{\sin(2rh)}{2h} - 1 \right] f'(\theta) + \frac{\cos(2rh) - 1}{4rh} f''(\theta), \end{aligned} \quad (24)$$

$$\text{Var}(\hat{g}_{\text{bFD}}(\theta, h)) = \frac{\sigma_1^2(\theta) + \sigma_1^2(\theta - h)}{4h^2}. \quad (25)$$

By choosing $\gamma_1 = h$ and $\gamma_2 = 0$ the same relation for bias and variance is obtained for the forward FD method. \square

The here presented gPSR unites all known forms of the finite differences and the centered PSR to the analytic gradient, thereby uniting all known methods in the family of single-component gradient rules and settling the discussion on trade-off between these approximate and exact methods. For the remaining two unknown forms, the backward and forward PSR, we provide a no-go theorem in the the following section.

4 No-go result for a forward and backward parameter shift rule

It is reasonable to expect near-term quantum devices to allow for wider and deeper circuits in the pursuit of achieving a practical quantum advantage, thereby resulting in the need to optimize a large number of parameters in the variational algorithm. Evaluating the gradient for n parameters requires $2n$ cost function evaluations per gradient vector evaluation with the parameter shift rule. As a hypothetical backward or forward parameter shift rule, in the spirit of the backward and forward finite differences, would require only $n + 1$ evaluations of the cost function, it is desirable to find such a rule. Here, we show that it is impossible for such a rule to exist when one considers the set of all cost functions expressed through variational quantum circuits.

Theorem 2 (Impossibility of a backward and forward parameter shift rules). *For all $0 < \gamma < \pi/r$ and $r \in \mathbb{R}$ there exists no function $g : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ with the property that*

$$g[f(\theta), f(\theta + \gamma)] = f'(\theta) \quad (26)$$

for all circuit expectation values $f \in \mathcal{F}_r$ (see Definition 4) and all $\theta \in [0, \pi/r)$.

Proof. Let $\zeta \in \mathbb{R}$ and $0 < \gamma < \pi/r$ be arbitrary but fixed. We pick an $f \in \mathcal{F}_r$ along with some r -gate F where

$$f(\theta) = \langle \psi | e^{i\theta G} A e^{-i\theta G} | \psi \rangle, \quad (27)$$

and the conditions

$$c_{\perp} |\zeta_{\perp}\rangle := e^{-i\gamma F} |\zeta\rangle - \langle \zeta | e^{-i\gamma F} | \zeta \rangle |\zeta\rangle \neq 0, \quad (28)$$

$$\langle \zeta | [G - F, A] | \zeta \rangle \neq 0, \quad (29)$$

are fulfilled using the short-hand $|\zeta\rangle := e^{-i\zeta G} |\psi\rangle$, where we have imposed $|\zeta_{\perp}\rangle$ to have unit norm. For most choices of G , F , A , and $|\psi\rangle$ these conditions will be fulfilled. However, it suffices to show that one such choice exists: For instance, choose $G = X$, $F = (Y + Z)/\sqrt{2}$, $A = Y$, and $|\psi\rangle = e^{i\zeta X} |0\rangle$. Now we can define a second element of \mathcal{F}_r as

$$\tilde{f}(\theta) := \langle \zeta | e^{i(\theta-\zeta)F} B e^{-i(\theta-\zeta)F} | \zeta \rangle, \quad (30)$$

with the observable

$$B := A - \frac{\langle \zeta | e^{i\gamma G} A e^{-i\gamma G} | \zeta \rangle - \langle \zeta | e^{i\gamma F} A e^{-i\gamma F} | \zeta \rangle}{|c_{\perp}|^2} |\zeta_{\perp}\rangle \langle \zeta_{\perp}|. \quad (31)$$

One can now verify that with this choice, we have

$$\tilde{f}(\zeta) = \langle \zeta | A | \zeta \rangle = f(\zeta), \quad (32)$$

$$\tilde{f}(\zeta + \gamma) = \langle \zeta | e^{i\gamma G} A e^{-i\gamma G} | \zeta \rangle = f(\zeta + \gamma). \quad (33)$$

Meanwhile, the derivatives at the point ζ are given by

$$f'(\zeta) = \langle \zeta | [iG, A] | \zeta \rangle, \quad (34)$$

$$\tilde{f}'(\zeta) = \langle \zeta | [iF, A] | \zeta \rangle, \quad (35)$$

which are unequal by the assumption in Eq. (29). Therefore, the function g discussed here, as described in the statement of the theorem, would receive the same inputs for f and \tilde{f} at $\theta = \zeta$ and $\theta = \zeta + \gamma$, while having to return different outputs. Hence, such a function g cannot exist. \square

5 Discussion

Variational quantum algorithms hold great potential for practical quantum applications. These algorithms are limited by their scaling in terms of number of variational parameters n , due to the increase in optimization difficulty. Gradient methods have gained significant attention over the past years as key method for optimization, and come in various forms. In our work, we focused on the single-component gradient methods, including multiple variations of the finite differences (FD) and

the parameter shift rule (PSR). The key lesson to be learned from our work, is that any of the above methods deliver linear combinations of higher-order derivatives, which can be reduced back to a linear combination of the first and second derivatives, of the cost function

$$g(\theta, \gamma) = a(\gamma)f'(\theta) + b(\gamma)f''(\theta) \quad (36)$$

with known functional dependence of the coefficients. The centered parameter shift rule (cPSR) solely targets the first derivative directly - the other methods however deliver instances of such linear combinations. We have formalized this, and proposed the generalized parameter shift rule (gPSR). This gPSR relates, in an exact manner, to all versions of the FD method, as well as the cPSR. Furthermore, it allowed us to calculate the bias and variance for all previously mentioned methods. To fully cover the family of single-component gradient rules, we have also presented a proof that neither a backward, nor forward parameter shift rule can exist.

The notion that each method delivers a linear combination of first- and second-order derivatives, points at an opportunity for future work to explore how these linear combination can be exploited, when addressing the task of classical control of variational quantum algorithms. Traditional second order optimization methods are known to outperform first order methods, in that they allow for the assignment of better directions of updates, as well as better step sizes; allowing for jumps directly to the minimum of the local quadratic approximation [33]. Here, the situation is slightly different and somewhat unusual from the perspective of second order optimization - in that first and second derivatives in a fixed combination are natively provided by the data, and most resource efficient approaches should use those native data at hand when addressing a control problem.

This line of reasoning can also be extended to the multi-component gradient rules, where multiple parameters are updated simultaneously. A natural first step would be to explore the existence of an exact multi-parameter shift rule, aimed at improving the number of circuit evaluations per chosen parameter. General rules of this type would give rise to scalar products with gradients and Hessians of the cost function. These can be exploited in suitable second order methods. Approaches of multi-parameter optimization would also open up the possibility of exploiting further structure in the gradient, inspired from the domain of signal processing - this includes resorting to sparse and low-rank structures [34]. However, in line with results from this work, we believe one should look beyond this, and aim to exploit the fact that multiple cross-derivatives between parameters can be accessed as a linear combination; just as we showed for the single-component setting. For instance, as higher-order derivatives can be expressed in terms of first and second derivatives, the cost function can also be expanded in *all* parameters simultaneously. Here the l -th derivative, represented by a rank- l tensor, can be computed at a cost significantly lower than n^l , where n is the number of parameters. Moreover, the series has finite length.

In particular, for a setting with two parameterized gates, the circuit evaluation of the expectation value of $f(\theta_1 + \gamma_1, \theta_2 + \gamma_2)$ can not only express first- and second-order derivatives for each parameter of $(\theta_1, \theta_2) \mapsto f(\theta_1, \theta_2)$, but also mixed (higher-order) derivatives. Given this access, one could envision a gradient rule that yields a linear combination of a subset of these derivatives, and an optimization method that works directly on the outcome of this gradient rule. One example of such an approach would be the following set of circuit evaluations, which yield a linear combination of the Hessian elements as

$$\begin{aligned} & f(\theta_1 + \gamma_1, \theta_2 + \gamma_2) + f(\theta_1 - \gamma_1, \theta_2 - \gamma_2) - 2 \sin^2(r\gamma_1) \sin^2(r\gamma_2) f\left(\theta_1 + \frac{\pi}{2}, \theta_2 + \frac{\pi}{2}\right) \\ &= 2 \sin(2r\gamma_1) \sin(2r\gamma_2) \frac{\partial}{\partial \theta_1} \frac{\partial}{\partial \theta_2} f(\theta_1, \theta_2) + \frac{1}{4} \cos^2(r\gamma_2) (5 - 3 \cos(2r\gamma_1)) \frac{\partial^2}{\partial \theta_1^2} f(\theta_1, \theta_2) \\ &+ \frac{1}{4} \cos^2(r\gamma_1) (5 - 3 \cos(2r\gamma_2)) \frac{\partial^2}{\partial \theta_2^2} f(\theta_1, \theta_2). \end{aligned} \quad (37)$$

When using this as input for the optimization, no classical access to the individual components of the Hessian are required; thereby reducing the required number of evaluations of the quantum circuit. Such a setup could give rise to pseudo-Newton methods that make use of this linear combination, directly in its optimization process. A comprehensive solution would extend this methodology, and economically estimate the full gradient over all parameters, based on suitable control patterns over the variational parameters. It is our hope that the present work stimulates such research endeavours in near-term quantum computing.

Acknowledgements

We would like to thank Johannes Jakob Meyer, Paul K. Faehrmann, and Ryan Sweke for helpful discussions. This work has been supported by the BMWi (PlanQK) and the BMBF (Hybrid and FermiQP). It has also received funding from the DFG under Germany’s Excellence Strategy – MATH+: The Berlin Mathematics Research Center, EXC-2046/1 – project ID: 390685689, and the CRC 183, as well as the European Union’s Horizon 2020 research and innovation programme under grant agreement No. 817482 (PASQuanS).

References

- [1] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. Brandao, D. A. Buell, *et al.*, “Quantum supremacy using a programmable superconducting processor,” *Nature*, vol. 574, no. 7779, pp. 505–510, 2019.
- [2] 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, *et al.*, “Quantum computational advantage using photons,” *Science*, vol. 370, no. 6523, pp. 1460–1463, 2020.
- [3] A. Acin, I. Bloch, H. Buhrman, T. Calarco, C. Eichler, J. Eisert, J. 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 European quantum technologies roadmap,” *New J. Phys.*, vol. 20, p. 080 201, 2018.
- [4] J. Preskill, “Quantum computing in the NISQ era and beyond,” *Quantum*, vol. 2, p. 79, 2018.
- [5] E. Farhi, J. Goldstone, and S. Gutmann, “A quantum approximate optimization algorithm,” *arXiv:1411.4028*, 2014.
- [6] M. Harrigan, K. Sung, M. Neeley, *et al.*, “Quantum approximate optimization of non-planar graph problems on a planar superconducting processor,” *Nature Phys.*, vol. 17, pp. 332–336, 2021.
- [7] O. Kyriienko, A. E. Paine, and V. E. Elfving, “Solving nonlinear differential equations with differentiable quantum circuits,” *arxiv:2011.10395*, 2020.
- [8] 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 Comm.*, vol. 5, no. 1, pp. 1–7, 2014.
- [9] 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*, vol. 549, p. 242, 2017.
- [10] A. Abbas, D. Sutter, C. Zoufal, A. Lucchi, A. Figalli, and S. Woerner, “The power of quantum neural networks,” *arXiv:2011.00027*, 2020.
- [11] C. Hempel, C. Maier, J. Romero, J. McClean, T. Monz, H. Shen, P. Jurcevic, B. P. Lanyon, P. Love, R. Babbush, A. Aspuru-Guzik, R. Blatt, and C. F. Roos, “Quantum chemistry calculations on a trapped-ion quantum simulator,” *Phys. Rev. X*, vol. 8, p. 031 022, 2018.
- [12] G.-L. R. Anselmetti, D. Wierichs, C. Gogolin, and R. M. Parrish, “Local, expressive, quantum-number-preserving VQE ansatzes for fermionic systems,” *arXiv:2104.05695*,
- [13] F. Arute *et al.*, “Hartree-Fock on a superconducting qubit quantum computer,” *Science*, vol. 369, pp. 1084–1089, 2020.
- [14] 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:2101.08448*, 2021.
- [15] M. Schuld, “Quantum machine learning models are kernel methods,” *arXiv:2101.11020*, 2021.
- [16] 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*, vol. 567, no. 7747, pp. 209–212, 2019.
- [17] T. Hubregtsen, D. Wierichs, E. Gil-Fuster, P.-J. H. Derks, P. K. Faehrmann, and J. J. Meyer, “Training quantum embedding kernels on near-term quantum computers,” *arXiv:2105.02276*, 2021.

- [18] B. Koczor, S. Endo, T. Jones, Y. Matsuzaki, and S. C. Benjamin, “Variational-state quantum metrology,” *arXiv:1908.08904*, 2019.
- [19] J. J. Meyer, J. Borregaard, and J. Eisert, “A variational toolbox for quantum multi-parameter estimation,” *npjqi*, 2021.
- [20] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, “The theory of variational hybrid quantum-classical algorithms,” *New J. Phys.*, vol. 18, no. 2, p. 023 023, 2016.
- [21] M. Benedetti, E. Lloyd, S. Sack, and M. Fiorentini, “Parameterized quantum circuits as machine learning models,” *Quantum Science and Technology*, vol. 4, no. 4, p. 043 001, 2019.
- [22] J. Li, X. Yang, X. Peng, and C.-P. Sun, “Hybrid quantum-classical approach to quantum optimal control,” *Phys. Rev. Lett.*, vol. 118, p. 150 503, 2017.
- [23] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii, “Quantum circuit learning,” *Phys. Rev. A*, vol. 98, p. 032 309, 2018.
- [24] M. Schuld, A. Bocharov, K. M. Svore, and N. Wiebe, “Circuit-centric quantum classifiers,” *Phys. Rev. A*, vol. 101, p. 032 308, 2020.
- [25] J.-G. Liu and L. Wang, “Differentiable learning of quantum circuit Born machines,” *Phys. Rev. A*, vol. 98, p. 062 324, 2018.
- [26] M. Schuld, V. Bergholm, C. Gogolin, J. Izaac, and N. Killoran, “Evaluating analytic gradients on quantum hardware,” *Phys. Rev. A*, vol. 99, p. 032 331, 2019.
- [27] G. E. Crooks, “Gradients of parameterized quantum gates using the parameter-shift rule and gate decomposition,” *arXiv:1905.13311*, 2019.
- [28] L. Banchi and G. E. Crooks, “Measuring analytic gradients of general quantum evolution with the stochastic parameter shift rule,” *Quantum*, vol. 5, p. 386, 2021.
- [29] A. Mari, T. R. Bromley, and N. Killoran, “Estimating the gradient and higher-order derivatives on quantum hardware,” *arXiv:2008.06517*, 2020.
- [30] J. Stokes, J. Izaac, N. Killoran, and G. Carleo, “Quantum natural gradient,” *Quantum*, vol. 4, p. 269, 2020.
- [31] J. C. Spall, “An overview of the simultaneous perturbation method for efficient optimization,” *Johns Hopkins apl technical digest*, vol. 19, no. 4, pp. 482–492, 1998.
- [32] G. Wang, D. E. Koh, P. D. Johnson, and Y. Cao, “Minimizing estimation runtime on noisy quantum computers,” *PRX Quantum*, vol. 2, p. 010 346, 2021.
- [33] N. Agarwal, B. Bullins, and E. Hazan, “Second-order stochastic optimization for machine learning in linear time,” *J. Mach. Learn. Res.*, vol. 18, p. 4148, 2017.
- [34] E. J. Candes and M. Wakin, “An introduction to compressive sampling,” *IEEE Signal Process. Mag.*, vol. 25, pp. 21–30, 2008.