

Implicit differentiation of variational quantum algorithms

Shahnawaz Ahmed^{1,2}, Nathan Killoran², and Juan Felipe Carrasquilla Álvarez^{3,4}

¹Department of Microtechnology and Nanoscience, Chalmers University of Technology, 412 96 Gothenburg, Sweden

²Xanadu, Toronto, ON, M5G 2C8, Canada

³Vector Institute for Artificial Intelligence, MaRS Centre, Toronto, ON, Canada M5G 1M1

⁴Department of Physics and Astronomy, University of Waterloo, Waterloo, Ontario, N2L 3G1, Canada

Several quantities important in condensed matter physics, quantum information, and quantum chemistry, as well as quantities required in meta-optimization of machine learning algorithms, can be expressed as gradients of implicitly defined functions of the parameters characterizing the system. Here, we show how to leverage implicit differentiation for gradient computation through variational quantum algorithms and explore applications in condensed matter physics, quantum machine learning, and quantum information. A function defined implicitly as the solution of a quantum algorithm, e.g., a variationally obtained ground- or steady-state, can be automatically differentiated using implicit differentiation while being agnostic to how the solution is computed. We apply this notion to the evaluation of physical quantities in condensed matter physics such as generalized susceptibilities studied through a variational quantum algorithm. Moreover, we develop two additional applications of implicit differentiation — hyperparameter optimization in a quantum machine learning algorithm, and the variational construction of entangled quantum states based on a gradient-based maximization of a geometric measure of entanglement. Our work ties together several types of gradient calculations that can be computed using variational quantum circuits in a general way without relying on tedious analytic derivations, or approximate finite-difference methods.




Figure 1: Automatic differentiation through a black-box algorithm that implicitly defines a solution function $z^*(a)$ to some problem using an optimality condition $f(z, a) = 0$. Solution points (shown with the red stars) can be obtained using an iterative black-box solver without having access to the explicit solution. However, if we have access to local gradients $(\partial_z f, \partial_a f)$ at the solution points, implicit differentiation can compute the gradient of the solution function $\partial_a z^*(a)$ without differentiating through the black box.

1 Introduction

Variational quantum algorithms (VQAs) cast the solution of a computational problem as the classical optimization of the expectation value of an observable with respect to a state generated by powerful parametrized quantum circuits [1, 2]. In light of their potential for enabling computational breakthroughs, VQAs have been developed for a wide range of applications encompassing most of the domains originally targeted by quantum computing, such as finding ground states of molecules, determining steady states of open quantum systems, simulating quantum dynamics, solving linear systems of equations, machine learning, combinatorial optimization, and many more [1, 2].

Variational quantum circuits can express highly-entangled complex quantum states beyond classical models such as deep Boltzmann machines and tensor networks with small bond dimensions [3]. The increased computational power

of quantum states comes at a significant price related to the difficulty of estimating their properties, which are key to all scientific and technological applications of quantum computing [4]. An important family of such quantities can be written as gradients of functions computed on a quantum state. For instance, nuclear forces in quantum chemistry [5, 6, 7], permanent electric dipole [8], static polarizability [8], the static hyperpolarizabilities of various orders [8], generalized susceptibilities [9], fidelity susceptibilities [10], and geometric tensors [11]. Moreover, in the emerging area of quantum machine learning, analogous gradients may prove critical in the development of hyperparameter tuning, meta-learning, and architectural search techniques. Thus, automatizing gradient computations may expand the algorithmic flexibility of VQAs and ease the development of a wider variety of gradient computations by removing the tedious burden of computing and implementing these derivatives by hand.

A natural approach to tackle problems involving gradient computation within a VQA is automatic differentiation (AD). Recently, AD has seen significant development within the quantum information research community in part due to the availability of modern quantum software such as PennyLane [12] and TensorFlow Quantum [13], which provide tools to implement AD for functions evaluated on simulated and physical quantum devices. While it is already straightforward to calculate several important gradients in VQAs, most prominently energy gradients (with respect to the parameters in the circuits [14, 15], and with respect to parameters in the Hamiltonian through the Hellmann-Feynman theorem [5]), gradient computation of expectation values of general operators using standard AD techniques remains generally unavailable. These generic gradients may become computationally difficult since in principle they require keeping track of all intermediate optimization steps in a VQA.

Here we develop an implicit differentiation strategy to automatically evaluate gradients and higher-order derivatives of expectation values of generic operators beyond the Hellman-Feynman theorem. Specifically, this implicit AD approach applies for the case where the operators, states, or gates making up the expectation value are themselves obtained as the solution of a complex op-

timization problem involving a quantum circuit, such as finding the ground state of a Hamiltonian.

Our strategy combines differentiable programming of quantum computers [12] with modular implicit differentiation [16] as a tool to automatically evaluate implicit gradients through variational quantum algorithms. A sketch explaining the idea of implicit differentiation is given in Fig. 1, and Fig. 2 shows how this idea can be extended to VQAs. We present three applications of using this technique: computing generalized susceptibilities, hyperparameter tuning in a quantum machine learning algorithm, and obtaining a Bell state by variationally optimizing a geometric measure of entanglement. Our implementation integrates a classical modular implicit differentiation package (JAXopt [16]) with a quantum AD framework (PennyLane [12]), which allows us to effortlessly add implicit differentiation to VQAs.

The paper is organized as follows. In Sec. 2 we discuss the implicit function theorem and the derivations of gradient formulas for implicitly defined functions. In Sec. 3 we discuss the prototypical example of differentiating through the ground state of a Hamiltonian to demonstrate how implicit differentiation can be leveraged through a ground-state-finding algorithm. Then, in Sec. 4, we show three applications of our approach. First, we show how generalized susceptibilities can be computed through a VQA that obtains the ground state of a parameterized Hamiltonian in Sec. 4.1. Then we show two examples that go beyond differentiating ground-state solutions. In Sec. 4.2 we demonstrate how hyperparameter optimization of a quantum machine learning algorithm is possible using implicit differentiation. Finally, in Sec. 4.3 we generate a maximally entangled Bell state variationally by optimizing an entanglement witness. The data and code for all numerical experiments are available in [17].

2 Implicit differentiation

The main tool we use in our work is the implicit function theorem (IFT). Below we state the theorem in its simplest form for an analytic function [18, 19].

Theorem 1 (IFT for analytic functions) *Let $f(z, a)$ be an analytic function of complex*

variables (z, a) in a neighbourhood around the point $(z_0, a_0) \in \mathbb{C}^m \times \mathbb{C}^n$. If $f(z_0, a_0) = 0$ and $\partial_z f \neq 0$ at (z_0, a_0) , then there exists a unique analytic solution function $z^*(a)$ to $f(z, a) = 0$ in a neighbourhood $\mathcal{N}(a_0)$, given by (z_0, a_0) with $z^*(a_0) = z_0$.

Since the solution function $z^*(a)$ is analytic, it can be complex-differentiable in the neighbourhood of points $(z^*(a), a)$ that satisfy

$$f(z^*(a), a) = 0. \quad (1)$$

This allows computing gradients of the solution function by differentiating Eq. (1) w.r.t a at $(z_0 = z^*(a_0), a_0)$

$$\begin{aligned} \partial_a f(z_0, a_0) + \partial_z f(z_0, a_0) \partial_a z^*(a) &= 0 \\ \partial_a z^*(a) &= -(\partial_z f(z_0, a_0))^{-1} \partial_a f(z_0, a_0) \end{aligned} \quad (2)$$

The region $\mathcal{N}(a_0)$ that defines the analytic domain can be lower bounded as shown in [20].

In practice, AD tools can be used to compute the gradients $\partial_a f$ and $\partial_z f$, which are vectors for a multivariate function. In fact f can also be vector-valued such that $\partial_a f$ and $\partial_z f$ are Jacobians. These can be computed column-wise with vector-Jacobian products (VJPs) or row-wise using Jacobian-vector products (JVPs). VJPs—also called reverse-mode autodifferentiation or backpropagation—are efficient for a scalar-valued function, compared to JVPs, where we only have one output but several input parameters. VJPs make the calculation of implicit gradients easier, as we show below.

Mathematically, the VJP of a function $g : \mathbb{R}^n \rightarrow \mathbb{R}^m$ for input parameters $x \in \mathbb{R}^n$ is the operation

$$\text{VJP}(g, x, v) : (x, v) \rightarrow (v^T \partial g(x)), \quad (3)$$

where the vector $v \in \mathbb{R}^m$ has the same dimension as the output of the function $g(x)$ and $\partial g : \mathbb{R}^n \rightarrow \mathbb{R}^m \times \mathbb{R}^n$ is the Jacobian matrix. It is easy to see that setting $v = [1, 1, \dots]$ gives us the gradients of the function for all n parameters,

$$v^T \partial g(x) = \begin{bmatrix} 1, & 1, & \dots \end{bmatrix} \begin{bmatrix} \frac{\partial g_1}{\partial x_1} & \dots & \frac{\partial g_1}{\partial x_n} \\ \vdots & \vdots & \vdots \\ \frac{\partial g_m}{\partial x_1} & \dots & \frac{\partial g_m}{\partial x_n} \end{bmatrix}.$$

We can write a VJP for the implicit function $z^*(a)$ from Eq. (2) as the operation

$$\begin{aligned} \text{VJP}(z^*, a, v) : (a, v) &\rightarrow (v^T \partial_a z^*(a)) \\ &= (-v^T A^{-1} B), \end{aligned} \quad (4)$$

assuming $A = \partial_z f(z_0, a_0)$ and $B = \partial_a f(z_0, a_0)$.

VJPs solve the difficulty of obtaining the inverse of the matrix A (matrix inversion has a cubic scaling). Using VJPs we can avoid the cost of full inversion with a trick that only requires approximately solving the linear system [16]

$$v = A^T u \quad (5)$$

for a vector u that has the same dimension as a . The approximate solution to this linear system using an iterative algorithm such as generalized minimal residual method or conjugate gradient [16] simplifies the inversion as

$$\begin{aligned} [A^T]^{-1} v &= u \\ ([A^T]^{-1} v)^T &= u^T \\ v^T A^{-1} &= u^T. \end{aligned} \quad (6)$$

We can then compute the VJP of the implicit function as

$$\text{VJP}(z^*, a, v) : (a, v) \rightarrow (-u^T B), \quad (7)$$

which gives $\partial_a z^*(a) = -u^T B$.

Implicit differentiation can also be applied to fixed-point equations such as

$$f(z, a) = z, \quad (8)$$

by defining $h(z, a) = z - f(z, a)$, with $h(z, a) = 0$, such that a solution map $z^*(a)$ satisfies

$$h(z^*(a), a) = 0. \quad (9)$$

A similar derivation for such fixed-point equations gives the gradients as

$$\partial_a z^*(a) = [\mathbf{I} - A]^{-1} B. \quad (10)$$

In Eq. (10) the inversion can also be approximated with a similar trick as discussed before, where we now define

$$v = [\mathbf{I} - A^T] u \quad (11)$$

such that we have

$$v + A^T u = u. \quad (12)$$

Eq. (12) now defines a fixed-point function for u that can be solved using any fixed-point solver. Then using Eq. (11) we have

$$u^T = v^T [\mathbb{I} - A]^{-1}, \quad (13)$$

and the VJP for the fixed point solution can be computed as

$$\text{VJP}(z^*, a, v) : (a, v) \rightarrow (-u^T B). \quad (14)$$

The implicit gradient is again $\partial_a z^*(a) = -u^T B$.

Therefore, calculating implicit gradients efficiently requires computing VJPs of a function defining a fixed point equation or optimality condition at the solution point. The inverse term is approximated efficiently with a simple trick that requires solving a linear system of equations or iteratively finding a fixed point. Other approaches for calculating the inverse such as using a Neumann series approximation [21, 22] are also possible.

Now, we can connect implicit differentiation to quantum algorithms, as illustrated in Fig. 2.

3 Implicit differentiation of variational quantum algorithms

We can use implicit differentiation on variational quantum algorithms in several different ways, depending on which elements of the expectation value (e.g., the states or the observable) are parameterized. As a motivating example, consider a variational quantum algorithm whose objective function is defined through an operator $H(a)$, with parameters a , that codifies the solution to a computational problem. Here, $H(a)$ may represent a quantum Hamiltonian, the objective function of a quantum machine learning algorithm, an Ising Hamiltonian encoding a combinatorial optimization problem, etc. Variational quantum algorithms optimize the expected value of $H(a)$ over a quantum state generated by a parameterized quantum circuit, represented by a unitary matrix $U(z)$ endowed with parameters z . The unitary $U(z)$ acts on an input state $|\psi_0\rangle$, typically a product state, as $|\psi_z\rangle = U(z)|\psi_0\rangle$. A quantum computer provides a strategy to evaluate the expectation value of $H(a)$ on $|\psi_z\rangle$,

$$E(z, a) = \langle \psi_z | H(a) | \psi_z \rangle, \quad (15)$$




Figure 2: Implicit differentiation through variational quantum algorithms: A quantum computer generates a variational state $|\psi_z\rangle$ and computes expectation values of quantities such as the energy $E(z, a)$ of a parameterized Hamiltonian $H(a)$. A classical optimization provides a solution map $z^*(a)$ for variational parameters z , minimizing the expectation value and therefore defining an optimality condition. We can then apply implicit differentiation to calculate gradients $\partial_a z^*(a)$ which would otherwise require unrolling all the steps in obtaining $z^*(a)$ using standard automatic differentiation (AD). In a realistic VQA, unrolling the optimization and storing intermediate variables may not be feasible for direct application of AD. However using the higher-order gradients of expectation values computed by the quantum computer at the solution we can compute $\partial_a z^*(a)$ using implicit differentiation.

as well as its derivatives with respect to z . A variational quantum algorithm minimizes $E(z, a)$ with respect to z , which gives optimal parameters

$$z^*(a) = \arg \min_z E(z, a). \quad (16)$$

Concretely, suppose we are interested in approximating the ground state of a quantum many-body Hamiltonian $H(a)$. The algorithm aims at minimizing the expectation value of the Hamiltonian over a variational quantum circuit to find an approximation to its ground-state properties.

The solution $z^*(a)$ corresponds to a minimum of the function $E(z, a)$ which means that

$$f(z^*(a), a) = \partial_z E(z^*(a), a) = 0. \quad (17)$$

The approximate ground-state solution of $H(a)$, $|\psi_{z^*(a)}\rangle$, is therefore implicitly defined by $z^*(a)$ and depends on the Hamiltonian parameters a . We are interested in computing how a change in the Hamiltonian parameter a affects the solution $z^*(a)$, and functions of the solution state.

Beyond the motivating example of a parameterized Hamiltonian, we can also construct implicit gradients for an arbitrary function that

combines variational parameters z and parameters of interest a , e.g., a loss function $\mathcal{L}(z, a)$ defined with a variational state ψ_z , where a represents the strength of regularization. Similarly, we can consider z and a to be variational parameters for two different quantum circuits, as we show in our example in Sec. 4.3 using a geometric definition of an entanglement measure using two variational quantum states. The example of taking gradients through a ground-state minimization task is perhaps most straightforward, and therefore we focus on it in more detail below.

The gradient of an expectation value defines a function $f(z, a) = \partial_a E(z, a)$ along with an optimality or fixed-point condition, i.e., $f(z, a) = 0$. In this case, the function $f(z, a)$ itself is the first-order gradient of the expectation value. Provided we can compute the Hessian of $E(z, a)$ w.r.t. parameters z , and its inverse, implicit differentiation gives a strategy to evaluate $\partial_a z^*(a)$. As we demonstrate below, $\partial_a z^*(a)$ can be used to compute other quantities that depend on the ground state.

We first consider a function that computes the expectation value of an operator A on the ground state,

$$\langle A \rangle = \langle \psi_{z^*(a)} | A | \psi_{z^*(a)} \rangle. \quad (18)$$

If $A = H(a)$, i.e., we are computing the ground-state energy, the Hellman-Feynman theorem [23] can be applied to obtain

$$\partial_a \langle A \rangle = \langle \psi_{z^*(a)} | \partial_a A | \psi_{z^*(a)} \rangle. \quad (19)$$

In quantum-chemistry problems such as molecular geometry optimization [6], the Hellman-Feynman theorem simplifies the gradient computation as we only require expectation values of $\partial_a A$. However, in a more general case, when the operator A is not the Hamiltonian, the computation of $\partial_a \langle A \rangle$ requires implicit gradients

$$\begin{aligned} \partial_a \langle A \rangle &= \partial_a z^*(a) \cdot \partial_z \langle A \rangle \\ &+ \langle \psi_{z^*(a)} | \partial_a A | \psi_{z^*(a)} \rangle. \end{aligned} \quad (20)$$

All the quantities required to compute Eq. (20) can be evaluated on a quantum computer for a large family of quantum circuits. Assuming the existence of a parameter shift-rule [14, 24, 15] for $U(z)$, the term $\partial_z \langle A \rangle$ can be evaluated using a similar strategy used to compute gradients

of the objective function of the variational quantum algorithm. Likewise, the second term, which is analogous to the Hellman-Feynman theorem, is evaluated by computing the expectation value of an operator $\partial_a A$ on the variationally obtained solution $\psi_{z^*(a)}$. Finally, the evaluation of $\partial_a z^*(a)$ (see Eq. 2) requires access to $\partial_a f, \partial_z f$, which are second-order derivatives of the expectation value $E(z, a)$, accessible on a quantum computer for large classes of variational circuits [25].

Our implicit differentiation formula can therefore extend the capabilities of automatic differentiation tools such as PennyLane to compute gradients of a variationally-obtained solution state without relying on tedious analytic computations or numerical approximations. Further, the computation does not require keeping track how the solution is computed or unrolling all the steps in the optimization, which might be infeasible in a variational quantum computation.

4 Examples

In the following examples we show how implicit differentiation can be applied to calculate quantities such as generalized susceptibility, perform hyperparameter optimization in a quantum machine learning algorithm, and to generate multipartite entanglement using variational quantum circuits.

4.1 Generalized susceptibility calculation

The response of an observable A upon the introduction of a perturbation in a Hamiltonian $H(a)$ can be quantified through a generalized susceptibility. These types of linear response functions are written in terms of the gradient of the expectation value of the operator A measured on the ground state $\psi_{\text{GS}}(a)$,

$$\partial_a \langle A \rangle = \partial_a \langle \psi_{\text{GS}}(a) | A | \psi_{\text{GS}}(a) \rangle, \quad (21)$$

where $\psi_{\text{GS}}(a)$ represents the ground state. Here a represents, e.g., the strength of a field B in the Hamiltonian, and A is the operator whose response to this perturbation we aim to compute. Examples of generalized susceptibilities important in condensed matter physics include the uniform magnetic susceptibility, or transport properties such as the superfluid stiffness, all of which




Figure 3: Susceptibility computation for a $N = 5$ spin system. The exact computation finds the ground state of a parameterized Hamiltonian $H(a)$ and unrolls the eigendecomposition to compute gradients (susceptibility) $\partial_a \langle A(a) \rangle$ using AD. In the variational setting, we can approximate the ground state by minimizing the expectation value of $H(a)$ using a variational state $|\psi_z\rangle$. The variational state is generated by the two-design ansatz with L layers. Once the ground-state solution $z^*(a)$ is found, implicit differentiation can compute the susceptibility efficiently without requiring to keep track of how the variational ground state is obtained.

are central to the characterization of strongly correlated phases of matter and their corresponding phase transitions [9].

The computation of susceptibilities and fidelity susceptibilities are typically carried out analytically for sufficiently tractable problems [26, 27, 28], using finite-difference methods through exact diagonalization and density matrix renormalization group [26, 29], through a costly computation of excited states using exact or variational techniques [30], or through quantum Monte Carlo simulations for Hamiltonians which do not have a sign problem [9]. Alternatively, we can compute susceptibilities in a variational quantum algorithm using implicit differentiation. Here, we consider a parameterized Hamiltonian for a linear spin-chain model in an external magnetic field given by

$$H(a) = - \sum_i \sigma_i^z \sigma_{i+1}^z - \gamma \sum_i \sigma_i^x - a \sum_i \sigma_i^z.$$

We aim at computing the response of the average magnetization $\partial_a \langle A \rangle$ to a change in a magnetic field along the z direction, i.e., $A = \frac{1}{N} \sum_i \sigma_i^z$ and $B = \sum_i \sigma_i^z$. Here $\sigma_i^{x,y,z}$ represent the Pauli spin operators at the site i . A small field $\delta \sum_i \sigma_i^z$ is added to the Hamiltonian for numerical stability near $a \approx 0$.

In a variational setting, a parameterized ground state $|\psi_{z^*(a)}\rangle$ can be obtained by optimizing the state $|\psi_z\rangle$ with respect to its variational parameters z . With the variationally-obtained ground state, by applying implicit differentiation we can compute the susceptibility using Eq. (20), going beyond the Hellman-Feynmann theorem.

In our experiments, the variational ground state is obtained by running a standard VQE optimization. Our variational circuit is a simplified two-design ansatz consisting of Pauli-Y rotation gates and entangling controlled-Z gates. We can consider multiple layers for a more expressive ansatz. The parameters of the circuit are optimized with respect to the energy $E(z, a)$ using gradient descent.

In Fig. 3 we show how implicit differentiation through a variationally-optimized ground state compares against an exact calculation of the susceptibility $\partial_a \langle A \rangle$. The exact ground state $|\psi_{\text{GS}(a)}\rangle$ is obtained through eigendecomposition of $H(a)$, where we evaluate $\partial_a \langle A \rangle$ by differentiating through the eigendecomposition [31, 32]. The gradient computation through eigendecomposition is implemented with the AD tool Jax [33] using the formulation in [31] (specifically Eq. (4.60)). For a sufficiently expressive ansatz with $L = 5$ layers, we can closely match the susceptibility computed by eigendecomposition. However for $L = 4$ layers, the gradient computation significantly deviates from the exact values.

Compared to the eigendecomposition approach, which scales exponentially with system size, the variational approach may scale well in practice for larger system sizes in future quantum devices beyond systems that can be simulated classically. Provided that the variational algorithms find accurate representations of the ground state and the availability of parameter-shift rules, we anticipate that variational quantum algorithms paired with implicit differentiation will enable the computation of susceptibilities and other response functions in such systems. While we have derived the susceptibility formula in Eq. (20) motivated by their use in a quantum algorithm, we emphasize that Eq. (20) can be applied to the simulation of response functions in classical variational techniques such as in neural network quantum states [34] and density matrix renormalization group calculations [35].




Figure 4: Comparing results of hyperparameter optimization in a quantum classifier against classical methods. The validation score is shown in the inset for each model. In case of the quantum classifier, the decision boundary obtained without hyperparameter optimization has a worse performance compared to optimized hyperparameters. The hyperparameter optimization is faster since we can use a gradient-based optimization by taking implicit gradients of the validation loss.

4.2 Hyperparameter optimization in a quantum classifier

We now consider an example motivated by the applications of implicit differentiation to hyperparameter optimization in machine learning [16, 36]. Hyperparameters play an important role in model selection and controlling the complexity of machine learning models [37]. Our demonstration uses the data-reuploading classifier proposed in [38], where a single-qubit quantum system is shown to be a universal classifier similar to neural networks with at least one hidden layer. In addition to training the parameters of the classifier, we consider adding hyperparameters and tuning them through gradient descent, using so-called hyper-gradients [39] calculated by implicit differentiation.

A set of inputs (α_i, β_i) , each with class labels $c_i \in \{0, 1\}$, forms our data set. The inputs are normalized and loaded into a quantum circuit as the parameters $(\alpha_i, \beta_i, 0)$ of a single-qubit unitary given by $U(x_i)$ ($x_i \equiv (\alpha_i, \beta_i, 0)$) followed by unitaries $U(z_l)$ where z_l represents model parameters for a layer l of quantum gates. The unitaries are simply rotations on the Bloch sphere with three parameters and we set the last parameter to zero. After applying L layers of such

unitaries on some initial quantum state $|\psi_0\rangle$ we obtain the variational state with trainable parameters $z \equiv \{z_L, \dots, z_0\}$ as

$$|\psi(z)\rangle = U(z_L)U(x_i) \dots U(z_0)U(x_i)|\psi_0\rangle. \quad (22)$$

Therefore we interleave unitaries that load the data and unitaries that act as learnable weights. A measurement of the operator σ_z on the output state gives a single-shot outcome representing our prediction for the class label c_i . We take a simpler approach than [38] for the measurements and trainable loss function. By repeating the measurement, we can obtain expectation values of the measurement operator. After normalizing the expectation values to the range $[0, 1]$ such that they can be interpreted as probabilities p_i , we define the binary cross-entropy loss function on the training data as

$$\mathcal{L}_{\text{bincross}}(z) = \frac{1}{N} \sum_i c_i \log p_i + (1 - c_i) \log(1 - p_i). \quad (23)$$

Now we can introduce additional hyperparameters in the loss by considering an L2 (ridge regularization [40]) penalty on the parameters of each layer to define the training loss

$$\mathcal{L}_T(z, a) = \mathcal{L}_{\text{bincross}}(z) + \sum_l a_l \|z_l\|_2^2, \quad (24)$$

where $a = \{a_l\}$ are the hyperparameters for the parameters z_l in each layer of the variational quantum circuit. We optimize the parameters z to minimize the training loss as

$$z^*(a) = \arg \min_z \mathcal{L}_T(z, a). \quad (25)$$

A validation loss is constructed using a validation set of the data not used for training

$$L_V(a) = L_{\text{bincross}}(z^*(a)). \quad (26)$$

We can minimize the validation loss with respect to the hyperparameters a using gradient descent and the learned parameters $z^*(a)$. The gradient computation $\partial_a L_V(a)$ requires differentiating the function $z^*(a)$, which we evaluate through implicit differentiation. We consider a single-qubit data-reuploading classifier with $L = 5$ layers and five hyperparameters, corresponding to each a_l controlling the strength of the L2 regularization in Eq. (24). In Fig. 5, we show how the validation loss and the hyperparameters change with




Figure 5: Hyperparameter optimization using implicit gradients. (a) The validation loss as a function of gradient-based updates of the hyperparameters. We compute the gradients of five hyperparameters (one for each layer in the circuit) using implicit differentiation of the validation loss in a quantum classifier. (b) The value of the five hyperparameters as we optimize them by minimizing the validation loss. The implicitly computed gradients allow a meta optimization of the hyperparameters going beyond a costly grid search in the hyperparameter space.

each gradient-descent update using implicit differentiation to compute the so-called hypergradients. We use the stochastic gradient-descent algorithm for both parameter and hyperparameter optimization with a learning rate of 0.01. All the hyperparameters are set to the same numerical value at the start of their optimization. As the validation loss improves, we find that the optimal hyperparameters differ significantly and converge to different values. The effect of the hyperparameter optimization on the classifier performance can be further seen by considering the decision boundaries in Fig. 4 where we also show a comparison to classical machine learning models such as Gaussian processes or simple neural networks. We can see that the decision boundary of the quantum classifier with an unoptimized set of hyperparameters seems to give a worse performance on a validation set compared to the case where hyperparameters are optimized.

The behaviour of the data-reuploading classifier with and without hyperparameter optimization




Figure 6: Variational generation of entangled quantum states using implicit differentiation. The distance of a variational quantum state $|\psi_a\rangle$ to the closest separable state $|\psi_{z^*(a)}\rangle$ can define a geometric measure of entanglement $E(\psi_a)$. This measure is dependent on the variational parameters a . The closest separable state can be obtained by maximizing $|\langle \psi_z | \psi_a \rangle|^2$ computed using a SWAP test. The optimization implicitly defines the closest separable state $z^*(a)$ and therefore the measure of entanglement, $E(a)$. The gradient of the entanglement measure $\partial_a E(a)$ can be computed using implicit differentiation to maximize the entanglement using gradient-based optimization.

tion shows that tuning the hyperparameters leads to a better learning of the decision boundary. In most hyperparameter tuning techniques, a grid search is performed for each of the hyperparameters. However, with implicit differentiation, it is possible to optimize millions of hyperparameters faster [21].

4.3 Variational optimization of a geometric measure of quantum entanglement

Quantum entanglement is an important resource for quantum communication, quantum cryptography, and plays a pivotal role in the possible speedup provided by quantum algorithms [41]. An entangled quantum state can be defined as a state that cannot be written as a separable state, i.e., as $|\psi\rangle = |\psi_1\rangle \otimes |\psi_2\rangle \otimes |\psi_3\rangle \otimes \dots \otimes |\psi_N\rangle$. A measure of entanglement for an arbitrary quantum state is any positive real-valued function of the state that cannot increase under local operations on parts of the state and classical communication. There exists several examples of quantum entanglement measures, e.g., concurrence, logarithmic negativity, von Neumann entropy [42]. A number of such measures become simple for bipartite quantum systems, but quantifying multipartite entanglement is generally difficult. We use the geometric measure of quantum entanglement de-

defined in [43] as,

$$E(\psi) = 1 - \max_{|\psi'\rangle} |\langle\psi'|\psi\rangle|^2, \quad (27)$$

where the optimization is over the set of separable states $|\psi'\rangle$. This geometric measure of entanglement can be intuitively understood as the distance between an arbitrary quantum state and the nearest separable quantum state found by solving Eq. (27). We now demonstrate how a variational algorithm can generate entangled quantum states by optimizing this geometric measure using implicit differentiation.

In Fig. 6 we show how two different variational quantum circuits can be used to create an entangled state $|\psi_a\rangle$ with variational parameters a . One circuit consists of only local (single-qubit) unitaries specified by parameters z , which creates separable states $|\psi_z\rangle$. The other circuit contains general operations that could possibly create entangled states $|\psi_a\rangle$. A SWAP test, its improved version [44], or any choice of randomized measurement overlap estimator [45, 46, 4], can provide a measure of the distinguishability between the two states [47, 48], and we can obtain a separable quantum state $|\psi_{z^*(a)}\rangle$ that is close to a fixed state $|\psi_a\rangle$ by the following optimization

$$z^*(a) = \arg \max_z |\langle\psi_z|\psi_a\rangle|^2. \quad (28)$$

The optimization problem defines an implicit function $z^*(a)$ that gives a measure of the entanglement of the state $E(|\psi_a\rangle) = E(a)$. Practically, we solve Eq. (28) by minimizing $-|\langle\psi_z|\psi_a\rangle|^2$. Now, using implicit differentiation we can obtain the gradients of the entanglement measure $\partial_a E(a)$ and maximize it. In Fig. 7, we show the results of this procedure by minimizing the loss function

$$\mathcal{L}(a) = -\log E(a) \quad (29)$$

The gradient of the loss function can be computed without the need for backpropagation through the quantum algorithm that generates the nearest separable state. We only need the value of the overlap and local gradients at $z^*(a)$. We optimize both Eq. (28) and Eq. (29) with stochastic gradient-descent with a learning rate of 0.001. The results of this bi-level optimization is shown in Fig. 7 where we can see that the entanglement measure increases as the loss function (negative log of the entanglement) decreases. As a result




Figure 7: Optimization of a geometric entanglement measure using gradient descent. The loss and entanglement measure after each update of the variational parameters a . Using implicit differentiation, we compute the gradient of the loss function and minimize it to maximize the entanglement measure.

we variationally obtained one of the Bell states starting from a random initial state, without ever defining it, as shown in Fig. 8. We can similarly extend this approach to multiple qubits to generate entangled quantum states variationally using gradient-based optimization. The algorithm is sensitive to the choice of the loss function, optimization, and hyperparameters and here we demonstrate one choice that allowed the creation of Bell states. In future investigations, we can determine the best choice of loss function and hyperparameters for a robust generation of entangled quantum states.

5 Conclusion

In this work, we have presented a common approach to applying implicit differentiation to quantum algorithms that encompasses a range of use-cases—from calculating ground-state property gradients using a variational quantum algorithm to hyperparameter optimization in quantum machine learning. We show how leveraging modern automatic-differentiation tools working in tandem with quantum software can simplify the application of implicit differentiation to wide variety of problems.

Our work shows how we can go beyond the Hellman-Feynmann theorem to compute derivatives of the expectation of arbitrary operators and not just the total energy in a parameterized Hamiltonian. As a demonstration, we take a simple spin-chain model to show how a variational




Figure 8: The final state $|\psi_a\rangle$ (bottom) generated by our variational approach starting from a random set of initial parameters (top). We use a variational quantum circuit that can create entanglement but learn the parameters using implicit differentiation and gradient descent; see Fig. 6. The geometric measure of entanglement uses another variational quantum circuit that only consists of single qubit operations, i.e., can only create separable states, such that computing an overlap between the separable state and a (possibly) entangled state gives a measure of entanglement. Implicit differentiation allows us to compute gradients through this geometric measure of entanglement. Optimizing the entanglement using gradient descent, we arrive at a maximally entangled Bell state $\frac{|00\rangle + |11\rangle}{\sqrt{2}}$ without ever explicitly defining the Bell state.

quantum algorithm can be used to obtain ground-state gradients by a parameterized quantum circuit. Such calculations may be easily performed on emerging NISQ devices.

Finally, we demonstrated two powerful applications of implicit differentiation: hyperparameter optimization of quantum machine learning algorithms, and variationally generating entangled quantum states using a geometric measure of entanglement. We demonstrate how hyperparameter optimization with implicit differentiation can not only be applied in the classical machine learning setting but also in quantum algorithms. We have shown that optimizing regularization strength using implicit differentiation could lead to better performance where implicit differentiation allowed a meta-optimization on the hyperparameters to avoid a costly grid search. The idea of entanglement generation using a bi-level optimization in a variational setting could possibly allow creating large multipartite entanglement beyond specific classes such as GHZ or W states.

Implicit differentiation is an idea as old as calculus itself with the first notable application by Fermat to compute tangents in the folium of Descartes curve [49]. As it was with the idea of applying gradient descent in both classical and quantum machine learning, new breakthroughs came with efficient and algorithmic ways to compute gradients by automatic differentiation. Recent demonstrations of implicit differentiation for inverse design in steady-state quantum systems [50] show how powerful this idea is with potential applications to quantum chemistry, control, and quantum machine learning. Our work opens up avenues to explore ideas such as neural quantum ODEs and compute interesting physical quantities using quantum computers and implicit differentiation.

The practical implementations of implicit differentiation on quantum hardware would require access to gradient computations as well as fast solutions to the optimality condition. Future works could investigate in detail how feasible such computations are when scaled to larger systems. It is also important to analyze how errors affect the gradient computation and its scaling. During the preparation of this manuscript, we became aware of another work that focuses on using implicit differentiation to compute the fidelity susceptibility

and analyzes some of these issues [51]. We hope that our work together with similar demonstrations of implicit differentiation [50] can find many more applications in the emerging area of quantum computation and quantum machine learning.

Acknowledgements

SA acknowledges discussions with Roeland Wiersema and Josh Izaac. JFCA acknowledges discussions with Federico Becca and Sandro Sorella. This work was supported by Mitacs through the Mitacs Accelerate program. JFCA acknowledges support from the Natural Sciences and Engineering Research Council (NSERC), the Shared Hierarchical Academic Research Computing Network (SHARCNET), Compute Canada, and the Canadian Institute for Advanced Research (CIFAR) AI chair program. Resources used in preparing this research were provided, in part, by the Province of Ontario, the Government of Canada through CIFAR, and companies sponsoring the Vector Institute www.vectorinstitute.ai/#partners

References

- [1] M. Cerezo, Andrew Arrasmith, Ryan Babush, Simon C. Benjamin, Suguru Endo, Keisuke Fujii, Jarrod R. McClean, Kosuke Mitarai, Xiao Yuan, Lukasz Cincio, and Patrick J. Coles. Variational quantum algorithms. *Nat. Rev. Phys.*, 3(9):625–644, 2021.
- [2] Kishor Bharti, Alba Cervera-Lierta, Thi Ha Kyaw, Tobias Haug, Sumner Alperin-Lea, Abhinav Anand, Matthias Degroote, Hermann Heimonen, Jakob S. Kottmann, Tim Menke, Wai-Keong Mok, Sukin Sim, Leong-Chuan Kwek, and Alán Aspuru-Guzik. Noisy intermediate-scale quantum algorithms. *Rev. Mod. Phys.*, 94:015004, 2022.
- [3] Yuxuan Du, Min-Hsiu Hsieh, Tongliang Liu, and Dacheng Tao. Expressive power of parametrized quantum circuits. *Phys. Rev. Research*, 2:033125, 2020.
- [4] Andreas Elben, Steven T. Flammia, Hsin-Yuan Huang, Richard Kueng, John Preskill, Benoît Vermersch, and Peter Zoller. The randomized measurement toolbox. (arXiv:2203.11374), 2022.
- [5] Thomas E. O’Brien, Bruno Senjean, Ramiro Sagastizabal, Xavier Bonet-Monroig, Alicja Dutkiewicz, Francesco Buda, Leonardo DiCarlo, and Lucas Visscher. Calculating energy derivatives for quantum chemistry on a quantum computer. *Npj Quantum Inf.*, 5(1):1–12, 2019.
- [6] P. Pulay. Ab initio calculation of force constants and equilibrium geometries in polyatomic molecules. *Mol. Phys.*, 17(2):197–204, 1969.
- [7] Kosuke Mitarai, Yuya O. Nakagawa, and Wataru Mizukami. Theory of analytical energy derivatives for the variational quantum eigensolver. *Phys. Rev. Research*, 2:013129, 2020.
- [8] Ivan Kassal and Alán Aspuru-Guzik. Quantum algorithm for molecular properties and geometry optimization. *J. Chem. Phys.*, 131(22):224102, 2009.
- [9] Anders W. Sandvik. Computational Studies of Quantum Spin Systems. *AIP Conf. Proc.*, 1297(1):135–338, 2010.
- [10] Wen-Long You, Ying-Wai Li, and Shi-Jian Gu. Fidelity, dynamic structure factor, and susceptibility in critical phenomena. *Phys. Rev. E*, 76:022101, 2007.
- [11] Lorenzo Campos Venuti and Paolo Zanardi. Quantum critical scaling of the geometric tensors. *Phys. Rev. Lett.*, 99:095701, 2007.
- [12] Ville Bergholm, Josh Izaac, Maria Schuld, Christian Gogolin, Shahnawaz Ahmed, Vishnu Ajith, M. Sohaib Alam, et al. PennyLane: Automatic differentiation of hybrid quantum-classical computations. (arXiv:1811.04968), 2022.
- [13] Michael Broughton, Guillaume Verdon, Trevor McCourt, Antonio J. Martinez, Jae Hyeon Yoo, Sergei V. Isakov, Philip Massey, Ramin Halavati, Murphy Yuezhen Niu, Alexander Zlokapa, Evan Peters, Owen Lockwood, Andrea Skolik, Sofiene Jerbi, Vedran Dunjko, Martin Leib, Michael Streif, David Von Dollen, Hongxiang Chen, Shuxiang Cao, Roeland Wiersema, Hsin-Yuan Huang, Jarrod R. McClean, Ryan Babush, Sergio Boixo, Dave Bacon, Alan K. Ho, Hartmut Neven, and Masoud Mohseni. TensorFlow Quantum: A Software Framework for Quantum Machine Learning. (arXiv:2003.02989), 2021.

- [14] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii. Quantum circuit learning. *Phys. Rev. A*, 98:032309, 2018.
- [15] David Wierichs, Josh Izaac, Cody Wang, and Cedric Yen-Yu Lin. General parameter-shift rules for quantum gradients. *Quantum*, 6:677, 2022.
- [16] Mathieu Blondel, Quentin Berthet, Marco Cuturi, Roy Frostig, Stephan Hoyer, Felipe Llinares-López, Fabian Pedregosa, and Jean-Philippe Vert. Efficient and Modular Implicit Differentiation. (arXiv:2105.15183), 2021.
- [17] Code and data for the article "Implicit differentiation of variational quantum algorithms" is available at <https://github.com/quantshah/quantum-implicit-differentiation>.
- [18] Steven George Krantz and Harold R Parks. *The implicit function theorem: history, theory, and applications*. Springer Science & Business Media, 2002.
- [19] Lars Hormander. *An introduction to complex analysis in several variables*. Elsevier, 1973.
- [20] Hung-Chieh Chang, Wei He, and Nagabhushana Prabhu. The analytic domain in the implicit function theorem. *JIPAM. J. Inequal. Pure Appl. Math*, 4(1), 2003.
- [21] Jonathan Lorraine, Paul Vicol, and David Duvenaud. Optimizing Millions of Hyperparameters by Implicit Differentiation. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, pages 1540–1552. PMLR, 2020.
- [22] Aravind Rajeswaran, Sham M. Kakade, Chelsea Finn, and Sergey Levine. Meta-learning with implicit gradients. *Advances in Neural Information Processing Systems*, 32, 2019.
- [23] R. P. Feynman. Forces in molecules. *Phys. Rev.*, 56:340–343, 1939.
- [24] Maria Schuld, Ville Bergholm, Christian Gogolin, Josh Izaac, and Nathan Killoran. Evaluating analytic gradients on quantum hardware. *Phys. Rev. A*, 99(3):032331, 2019.
- [25] Andrea Mari, Thomas R. Bromley, and Nathan Killoran. Estimating the gradient and higher-order derivatives on quantum hardware. *Phys. Rev. A*, 103:012405, 2021.
- [26] Gergely Barcza, Florian Gebhard, Thorben Linneweber, and Örs Legeza. Ground-state properties of the symmetric single-impurity anderson model on a ring from density-matrix renormalization group, hartree-fock, and gutzwiller theory. *Phys. Rev. B*, 99:165130, 2019.
- [27] Min-Fong Yang. Ground-state fidelity in one-dimensional gapless models. *Phys. Rev. B*, 76:180403, 2007.
- [28] A. N. Kirillov and N. Yu. Reshetikhin. Exact solution of the Heisenberg XXZ model of spin s. *J. Sov. Math.*, 35(4):2627–2643, 1986.
- [29] Juan Carrasquilla, Salvatore R. Manmana, and Marcos Rigol. Scaling of the gap, fidelity susceptibility, and bloch oscillations across the superfluid-to-mott-insulator transition in the one-dimensional bose-hubbard model. *Phys. Rev. A*, 87:043606, 2013.
- [30] Tôru Sakai and Hiroki Nakano. Quantum Spin Fluid Behaviors of the Kagome- and Triangular-Lattice Antiferromagnets. *J. Phys. Conf. Ser.*, 969(1):012127, 2018.
- [31] Christoph Boeddeker, Patrick Hanebrink, Lukas Drude, Jahn Heymann, and Reinhold Haeb-Umbach. On the Computation of Complex-valued Gradients with Application to Statistically Optimum Beamforming. (arXiv:1701.00392), 2017.
- [32] Hao Xie, Jin-Guo Liu, and Lei Wang. Automatic differentiation of dominant eigensolver and its applications in quantum physics. *Phys. Rev. B*, 101:245139, 2020.
- [33] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018.
- [34] Giuseppe Carleo and Matthias Troyer. Solving the quantum many-body problem with artificial neural networks. *Science*, 355(6325):602–606, 2017.
- [35] Steven R. White. Density matrix formulation for quantum renormalization groups. *Phys. Rev. Lett.*, 69:2863–2866, 1992.
- [36] Kaiyi Ji. Bilevel Optimization for Machine Learning: Algorithm Design and Convergence Analysis. (arXiv:2108.00330), 2021.
- [37] Jake Lever, Martin Krzywinski, and Naomi

- Altman. Points of Significance: Regularization. *Nat. Methods*, 13(10):803–804, 2016.
- [38] Adrián Pérez-Salinas, Alba Cervera-Liarta, Elies Gil-Fuster, and José I. Latorre. Data re-uploading for a universal quantum classifier. *Quantum*, 4:226, 2020.
- [39] Dougal Maclaurin, David Duvenaud, and Ryan Adams. Gradient-based Hyperparameter Optimization through Reversible Learning. In *Proceedings of the 32nd International Conference on Machine Learning*, pages 2113–2122. PMLR, 2015.
- [40] Christopher M. Bishop. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag, Berlin, Heidelberg, 2006.
- [41] Richard Jozsa and Noah Linden. On the role of entanglement in quantum-computational speed-up. *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 459(2036):2011–2032, 2003.
- [42] Ryszard Horodecki, Paweł Horodecki, Michał Horodecki, and Karol Horodecki. Quantum entanglement. *Rev. Mod. Phys.*, 81:865–942, 2009.
- [43] Tzu-Chieh Wei and Paul M. Goldbart. Geometric measure of entanglement and applications to bipartite and multipartite quantum states. *Phys. Rev. A*, 68:042307, 2003.
- [44] Lukasz Cincio, Yiğit Subaşı, Andrew T. Sornborger, and Patrick J. Coles. Learning the quantum algorithm for state overlap. *New Journal of Physics*, 20(11):113022, November 2018.
- [45] Andreas Elben, Benoît Vermersch, Rick van Bijnen, Christian Kokail, Tiff Brydges, Christine Maier, Manoj K. Joshi, Rainer Blatt, Christian F. Roos, and Peter Zoller. Cross-platform verification of intermediate scale quantum devices. *Phys. Rev. Lett.*, 124:010504, Jan 2020.
- [46] Leonardo Guerini, Roeland Wiersema, Juan Felipe Carrasquilla, and Leandro Aolita. Quasiprobabilistic state-overlap estimator for NISQ devices, December 2021.
- [47] Adriano Barenco, Andre Berthiaume, David Deutsch, Artur Ekert, Richard Jozsa, and Chiara Macchiavello. Stabilization of quantum computations by symmetrization. *SIAM Journal on Computing*, 26(5):1541–1557, 1997.
- [48] Harry Buhrman, Richard Cleve, John Watrous, and Ronald de Wolf. Quantum fingerprinting. *Phys. Rev. Lett.*, 87:167902, 2001.
- [49] Jaume Paradís, Josep Pla, and Pelegrí Viader. Fermat and the quadrature of the folium of descartes. *The American Mathematical Monthly*, 111(3):216–229, 2004.
- [50] Rodrigo A. Vargas-Hernández, Ricky T. Q. Chen, Kenneth A. Jung, and Paul Brumer. Fully differentiable optimization protocols for non-equilibrium steady states. *New J. Phys.*, 23(12):123006, 2021.
- [51] Olivia Di Matteo and R. M. Woloshyn. Quantum computing fidelity susceptibility using automatic differentiation. (arXiv:2207.06526), 2022.