

TITLE

Solving nonlinear differential equations with differentiable quantum circuits

AUTHORS

Kyriienko, O; Paine, AE; Elfving, VE

JOURNAL

Physical Review A

DEPOSITED IN ORE

19 May 2021

This version available at

<http://hdl.handle.net/10871/125753>


COPYRIGHT AND REUSE

Open Research Exeter makes this work available in accordance with publisher policies.

A NOTE ON VERSIONS

The version presented here may differ from the published version. If citing, you are advised to consult the published version for pagination, volume/issue and date of publication

Solving nonlinear differential equations with differentiable quantum circuits

Oleksandr Kyriienko ^{1,2}, Annie E. Paine,^{1,2} and Vincent E. Elfving ²

¹*Department of Physics and Astronomy, University of Exeter, Stocker Road, Exeter EX4 4QL, United Kingdom*

²*Qu & Co B.V., PO Box 75872, 1070 AW, Amsterdam, The Netherlands*



(Received 18 December 2020; accepted 13 April 2021; published 17 May 2021)

We propose a quantum algorithm to solve systems of nonlinear differential equations. Using a quantum feature map encoding, we define functions as expectation values of parametrized quantum circuits. We use automatic differentiation to represent function derivatives in an analytical form as differentiable quantum circuits (DQCs), thus avoiding inaccurate finite difference procedures for calculating gradients. We describe a hybrid quantum-classical workflow where DQCs are trained to satisfy differential equations and specified boundary conditions. As a particular example setting, we show how this approach can implement a spectral method for solving differential equations in a high-dimensional feature space. From a technical perspective, we design a Chebyshev quantum feature map that offers a powerful basis set of fitting polynomials and possesses rich expressivity. We simulate the algorithm to solve an instance of Navier-Stokes equations and compute density, temperature, and velocity profiles for the fluid flow in a convergent-divergent nozzle.

DOI: [10.1103/PhysRevA.103.052416](https://doi.org/10.1103/PhysRevA.103.052416)

I. INTRODUCTION

Differential equations are ubiquitous in various fields of science [1]. Their applications range from predicting motion in mechanics and fluid dynamics [2], to describing reactions in chemistry, ecosystem balance in ecology, market dynamics in finance, and disease spreading in epidemiology. In many cases, finding solutions to nonlinear systems of differential equations (DEs) is challenging, and requires advanced numerical techniques [3,4]. The difficulties arise in DEs for many functions and variables (known as “the curse of dimensionality”), models with high degree of nonlinearity, optimal control problems, as well as stiff and numerically unstable systems [5,6]. Classical numerical methods for solving DEs can be divided into local and global methods. Local approaches rely on discretization of the space of variables, with derivatives being approximated with numerical differentiation techniques (finite differencing and Runge-Kutta methods). Often a fine grid for multivariable functions is required to represent a solution qualitatively and quantitatively [3], leading to increasing computational cost. Global methods represent the solution in terms of a suitable basis set [4]. This recasts the problem to finding optimal coefficients for the polynomial approximation (e.g., Fourier or Chebyshev) of the sought function. Finding spectral solutions for complex problems may require ever-increasing basis sets to achieve high accuracy, enlarging the global differential operator, and introduces difficulties when dealing deterministically with boundary conditions.

Quantum computers offer a fundamentally different approach to perform computation, and possess algorithmically superior scaling for certain problems that include amplitude amplification and Abelian hidden subgroup problems [7,8]. For prototypical linear algebra tasks, quantum computers can be used as exponential accelerators for solving linear system

of equations (LSE), as offered by the HHL algorithm [9] and other quantum LSE approaches with improved scaling [10–14]. These developments, among others, have led to the nascent field of quantum machine learning [15–24]. Quantum LSE solvers can be applied to linear differential equations that are rewritten as systems of algebraic equations using a finite differencing scheme (Euler’s method). Several studies have developed generic quantum solvers for linear differential equations, and have shown algorithmic improvements in terms of accessed quantum oracles [25–33]. For the global numerical differentiation methods a spectral solver was proposed in Ref. [34]. The power of HHL-based algorithms relies on *amplitude encoding*, where a function is represented as a quantum state $|u\rangle = \sum_{k=1}^{2^N} u_k |k\rangle$ encoded in complex u_k amplitudes of computational basis states $|k\rangle$ for N qubits. This allows for compressing large grids into a small qubit register, providing exponential memory advantage. However, there are several caveats that should be considered [9]. First, the proposed quantum algorithms solve an equivalent quantum task of finding a state $|u\rangle = \hat{A}^{-1}|b\rangle$ given the matrix \hat{A} acting on predefined (boundary) state $|b\rangle$. Reading out the function values from the quantum state $|u\rangle$ in general requires exponential sampling, representing a so-called data “output problem.” Second, the preparation of a generic input state $|b\rangle$ requires sophisticated techniques such as quantum random access memory (QRAM) [35], thus leading to a so-called data “input problem.” Finally, translating quantum oracles that encode DEs (matrix \hat{A}) can introduce huge computational overheads, making simulation infeasible even for a fault-tolerant quantum computer [36].

Modern quantum processors are not perfect, and hardware available today correspond to noisy intermediate scale quantum (NISQ) devices [37]. They operate for a limited circuit depth, but still enjoy a Hilbert space that increases exponentially with the number of qubits. Recently, such a

setup was shown to provide advantage over classical methods for a tailored problem [38]. However, finding NISQ protocols that can offer advantage for industrially relevant tasks represents an open challenge. One promising direction corresponds to simulation of material science and chemistry, where hybrid quantum-classical approaches are developed for preparing ground states of correlated molecules [39–41]. The paradigmatic hybrid algorithm corresponds to the variational quantum eigensolver (VQE) [42,43], where a quantum computer is used to prepare parametrized quantum states and measure their energy, while classical optimization is used for finding the best variational parameters. This strategy enables quantum calculations with relatively noisy devices, and allows for numerous advances from the experimental [44–52] and theoretical [53–66] perspectives.

The rise of VQE has launched further development of variational quantum algorithms (VQAs) [67–71], targeting generic data analysis and machine learning tasks. Performed by *parametrized quantum circuits*, also dubbed as quantum neural networks (QNNs) [67], VQAs rapidly expand applications to include classification [72–78], regression [73], generative modeling [79–87], clustering [88,89], reinforcement learning [90–95], quantum simulation [96–99], quantum metrology [100–102], and other tasks. The power of VQAs for machine learning (ML) tasks comes from the high expressibility of quantum circuits, where data points are mapped to quantum states. Also, a large role for designing efficient algorithms is played by automatic differentiation and natural gradient techniques [73,103–111], developing deep networks and new QNN architectures [112–115], as well as proposing operationally meaningful cost functions for optimization [70,116], and optimal ansatz search [117–122]. First experiments show promising results for systems of small and increasing size [74,123,124]. Various strategies of error mitigation were proposed that can further improve the performance of algorithms when run on physical devices [125–132]. Finally, considering linear algebra problems, several VQAs were also proposed to solve LSEs [133–138], and ongoing efforts are directed towards improving their workflow. To date, several small scale linear equation solvers have been implemented using a nuclear magnetic resonance-based setup [139,140].

As for *nonlinear* differential equation solvers, available quantum protocols remain scarce and the field is only starting to grow. An intriguing approach was proposed in Ref. [141], where nonlinearity is introduced using quantum nonlinear processing units realized by ancillary quantum registers and controlled-multiqubit operations. The approach uses amplitude encoding, offering memory saving while potentially facing the data input and output problems. Another hybrid quantum-classical approach is presented in Ref. [142], where computational fluid dynamics applications are considered. This relies heavily on classical computation, with the quantum processor being delegated to perform function averaging (using amplitude estimation as subroutine). Using an oracle-based approach, amplitude-encoded states, and two quantum Fourier transforms, the discussed algorithm remains an option more suitable for future fault-tolerant devices.

In the classical computational methods, the advanced methods for solving differential equations are also currently

developed based on machine learning techniques. The solutions are encoded by deep neural networks as universal function approximators [143], advancing on the original steps with shallow artificial neural networks [144]. The approach relies on training physically informed neural networks [145,146], Fourier neural operators [147] or universal differential equation networks [148]. Recently, it has become a part of scientific machine learning ecosystem [149]. Another intriguing direction of contemporary research in differential calculus is representation of solutions using tensor networks [150]. This was also considered as a quantum-inspired and quantum approach to provide speedup for certain problems [151].

We propose to solve nonlinear differential equations (DEs) of a general form in a radically different way, where we construct solutions using *differentiable quantum circuits* (or equivalently, *derivative* quantum circuits) that we refer to as DQCs in the following [152]. These can be viewed as quantum neural network circuits that are designed to deal with functions and their derivatives using automatic differentiation rules. The overview of our approach is outlined in Fig. 1. It relies on defining the DQC structure, choosing the optimization strategy, and training of the circuit parameters that allows for reproducing the solution function. The required elements and workflow are detailed in the next section.

Our approach targets near-term quantum devices, where hybrid quantum-classical computing can work with noisy operation and reasonably shallow circuit depth. We alleviate the data input problem, as we do not rely on amplitude encoding, but use a latent space representation defined in the high-dimensional Hilbert space. This brings advantage in finding solutions of differential equations, as we use the large expressive power of quantum feature maps and parametrized quantum circuits. For efficient readout, we project the high-dimensional function representation to its scalar values by measuring the expectation of an observable (cost operator), and separate evaluation of quantum circuits, offering a NISQ-friendly implementation. We note that our approach also can work with quantum kernels [153], an advantageous choice for the loss function that allows for efficient multipoint training. However, this approach is more suitable for future highly coherent quantum devices, and will be considered in forthcoming studies.

II. DQC-BASED DIFFERENTIAL EQUATIONS SOLVER: GENERAL OVERVIEW

We start with an overview description of the differential equation solver based on derivative circuits. To solve DEs, we prepare trial solutions of the differential equation(s) as quantum circuits parametrized by a variable $x \in \mathbb{R}$ (or a collection of v variables, $\mathbf{x} \in \mathbb{R}^v$). As the discussion is generalized straightforwardly to the case of v variables, for brevity we henceforth use the simplified single variable notation x . We use a *quantum feature map* circuit $\hat{U}_\varphi(x)$ to encode the pre-defined *nonlinear* function of variables $\varphi(x)$ to amplitudes of the quantum state $\hat{U}_\varphi(x)|\emptyset\rangle$ prepared from some initial state $|\emptyset\rangle$. A quantum feature map represents a latent space encoding that, unlike amplitude encoding, does not require access to each amplitude and is controlled by classical gate parameters,




FIG. 1. DQC optimization workflow. The problem is set as a system of differential equations for functions u , variables x , and specified boundary conditions. The derivative quantum circuit is constructed by choosing the encoding circuit and optimization schedule. The solution is optimized by evaluating the function and derivative circuits at the defined grid of points X (which may be multidimensional), and using these values to calculate the loss function derivative. The variational angles are updated in the hybrid quantum-classical loop until a specified exit condition is reached.

mapping real parameter x to the corresponding variable value. Sometimes this is also called a quantum *embedding* [154,155], referring to the way data is embedded in the circuit. This follows the steps of quantum circuit learning (QCL) [73], shown to work well for solving problems of regression and classification. Next, we add a variational quantum circuit \hat{U}_θ parametrized by a vector of variational parameters θ that can be adjusted in a quantum-classical optimization loop. The resulting state $|f_{\phi,\theta}(x)\rangle = \hat{U}_\theta \hat{U}_\phi(x)|\emptyset\rangle$ for optimal angles contains the x -dependent amplitudes sculptured to represent the sought function. Finally, the real valued function can be read out as an expectation value of a predefined Hermitian cost operator \hat{C} , such that the trial function that aims to approximate the solution $u(x)$ reads [Fig. 2(a)]

$$f(x) = \langle f_{\phi,\theta}(x) | \hat{C} | f_{\phi,\theta}(x) \rangle. \tag{1}$$

One crucial step of our algorithm is the differentiation of the quantum feature map circuit, $d\hat{U}_\phi(x)/dx = \sum_j \hat{U}_{d\phi,j}(x)$, where the action differential can be represented as a sum of modified circuits $\hat{U}_{d\phi,j}$. This allows function derivatives to be represented using the product derivative rule. In the case of quantum feature map generated by strings of Pauli matrices

or any involutory matrix we can use the *parameter shift rule* [73,103] such that the function derivative is expressed as a sum of expectations [Fig. 2(b)]

$$df(x)/dx = \frac{1}{2} \sum_j \left(\langle f_{d\phi,j,\theta}^+(x) | \hat{C} | f_{d\phi,j,\theta}^+(x) \rangle - \langle f_{d\phi,j,\theta}^-(x) | \hat{C} | f_{d\phi,j,\theta}^-(x) \rangle \right), \tag{2}$$

with $|f_{d\phi,j,\theta}^\pm(x)\rangle$ defined through the parameter shifting, and index j running through the individual quantum operations used in the feature map encoding. Applying the parameter shift rule once again we can get the second-order derivative $d^2f(x)/dx^2$ with four shifted terms for each generator [104,156]. Importantly, as described we use the automatic differentiation (AD) technique to perform quantum circuit differentiation. AD allows the function derivative to be represented by the exact analytical formula using a set of simple computational rules, as opposed to numerical differentiation. Since automatic differentiation provides an analytical derivative of the circuit at any value of variable x , our scheme does not take on the accumulated error from approximating




FIG. 2. Differentiable quantum circuits. (a) Quantum circuit used for encoding the value of a function at a specific value of the variable $x = x_i$. The circuit consists of a feature map \hat{U}_φ that encodes the x -dependence, followed by variational ansatz \hat{U}_θ , and an observable-based readout for the set of operators \hat{C}_ℓ . The measurement result is classically post-processed to provide a quantum function representation $f(x)$ as a sum of expectations, where coefficients α_ℓ can be optimized in a quantum-classical hybrid loop. To compose the loss function circuit measurements for different points of optimization grid \mathbb{X} are required. (b) Derivative of the sought function $f(x)$ evaluated at specific point $x = x_i$ is estimated as a sum of expectations for derivative quantum circuits. The full structure follows from the feature map differentiation described in the text and shown for example in Fig. 3.

the derivatives. Notably, typical schemes for quantum ODE solvers involve numerical differentiation using Euler's method and finite difference schemes that suffer from approximation errors, and often require a fine discretization grid. This problem is alleviated in our approach.

Our goal is to construct and define the conditions for the quantum circuit to represent the solution of differential equations, generally written as

$$F[\{d^m f_n/dx^m\}_{m,n}, \{f_n(x)\}_n, x] = 0, \quad (3)$$

where the functional $F[\cdot]$ is provided by the problem for function derivatives of different order m and function/variable polynomials of varying degree n . For simplicity we refer to f as a function, and show that the same analysis holds for a vector of functions. The condition above demands that derivatives and nonlinear functions give a net zero contribution. Thus we can rewrite the task as an optimization problem and use the loss function $\mathcal{L}_\theta[d_x f, f, x]$. This corresponds to minimization of $F[x]|_{x \rightarrow x_i}$ at points in the set $\mathbb{X} = \{x_i\}_{i=1}^M$, and additionally ensuring that the boundary conditions are satisfied. Once the optimal angles

$$\theta_{\text{opt}} = \underset{\theta}{\text{argmin}}(\mathcal{L}_\theta[d_x f, f, x]) \quad (4)$$

are found, we can reproduce the solution from Eq. (1) as a function $f(x)|_{\theta \rightarrow \theta_{\text{opt}}} \approx u(x)$.

Summarizing the high level description of the proposed approach we note that we (1) use quantum feature map encoding, thus overcoming the complexity of amplitude encoding for preparing the solution at the boundary; (2) perform automatic differentiation of the quantum feature map circuit, allowing us to represent function derivatives without the imprecision error characteristic to numerical differentiation (finite differencing); (3) search for the suitable solution in the exponential space of fitting polynomials, thus resembling the spectral and finite element methods with improved scaling; and (4) avoid the data readout problem, as the solution is encoded in the observable of an operator, such that expectation can be routinely calculated. For the latter point, it differs from amplitude

encoding $|u\rangle$ in HHL and related methods, where getting the full solution from amplitudes is exponentially costly and requires tomographic measurements. Finally, we note that our goal is to construct circuits that can work for quantum processors with limited computational power, meaning the gate depth (number of operations to performed in series) is limited to a certain amount. This largely defines the training procedure, where we rely on the classical optimization loop. Once deep circuits can be implemented, we can also exploit parallel training strategies for the quantum circuit and quantum state encoding, coming closer to the ideal quantum operation regime.

III. METHODS

Below we discuss the set of tools that are required to build a differentiable circuit as a solution of differential equations. This corresponds to the main ingredients of the circuit with Sec. III A describing quantum feature maps and their derivatives; Sec. III B variational quantum circuits (ansatze); Sec. III C cost functions; and Sec. III D loss functions for the optimization loop. Additionally we detail proposed boundary handling techniques in Sec. III E, possible strategies for regularization in Sec. III F, encoding multiple functions in Sec. III G, and summarize a complete optimization schedule in Sec. III H.

A. Quantum feature maps and their derivatives

A quantum feature map is a unitary circuit $\hat{U}_\varphi(x)$ that is parametrized by the variable x and typically nonlinear function $\varphi(x)$. Acting on the state it maps x to $\hat{U}_\varphi(x)|\emptyset\rangle$ such that the x dependence is translated into quantum state amplitudes [67]. This is also referred to as a latent space mapping. There exist different ways of feature map encoding. Below we describe some possible choices and propose a distinct *Chebyshev quantum feature map* that allows the approximation of highly nonlinear functions. We also describe the procedure




FIG. 3. Product feature map and its derivative. (a) Quantum feature map of a product type, where single qubit rotations (here chosen as \hat{R}_y) act at each qubit individually and are parametrized by a function of variable x . Specifically, the expectation value of the circuit is shown, with thin pink and green blocks depicting the variational ansatz and the cost function measurement respectively. For the nonlinear feature encoding the nonlinear function $\varphi(x)$ is used as an angle of rotation. The product feature map can be further generalized to several layers, and different functions $\{\varphi\}$. Several feature maps can be concatenated to represent a multivariable function. (b) shows the derivative quantum circuit for the product feature map. Differentiation over variable x follows the chain rule, with the expectation value of the derivative written as a sum of separate expectations with shifted phases, repeated for each x -dependent rotation. (c) Generalized product feature map, which we refer to as *evolution-enhanced feature map*, where the layer of rotation is followed by the unitary evolution generated by Hamiltonian \hat{H} . For complicated multiqubit Hamiltonians the encoded state comprises of exponentially many unique x -dependent amplitudes. Time interval τ can be set variationally, or annealed from zero to finite value during the optimization procedure.

of feature map differentiation, the crucial step of constructing quantum circuits for solutions of differential equations.

1. Product feature maps

First, we consider a product feature map that uses qubit rotations, but has nonlinear dependence on the encoded variable x . In the simplest case, this corresponds to a single layer of rotations written in the form

$$\hat{U}_\varphi(x) = \bigotimes_{j=1}^{N'} \hat{R}_{\alpha,j}(\varphi[x]), \tag{5}$$

where $N' \leq N$ is the number of qubits used for the encoding. $\hat{R}_{\alpha,j}(\varphi) := \exp(-i\frac{\varphi}{2}\hat{P}_{\alpha,j})$ is a Pauli rotation operator for Pauli matrices $\hat{P}_{\alpha,j} = \hat{X}_j, \hat{Y}_j, \text{ or } \hat{Z}_j$ ($\alpha = x, y, z$, respectively) acting on qubit j with phase φ . This represents the feature map choice used in quantum circuit learning [73], and several similar encodings were discussed in Refs. [86,157], and reviewed in Ref. [67]. The next step is to assign a nonlinear function for rotation, with a popular choice being $\varphi(x) = \arcsin x$ and $\alpha = y$ such that only real amplitudes are generated. The corresponding circuit is shown in Fig. 3(a). The unitary operator Eq. (5) is then rewritten as

$$\hat{U}_\varphi(x) = \bigotimes_{j=1}^N \exp\left(-i\frac{\arcsin x}{2}\hat{Y}_j\right), \tag{6}$$

leading to amplitudes that depend on the encoded variables as $\cos[(\arcsin x)/2]$ and $\sin[(\arcsin x)/2]$. Acting on the initial state $|\emptyset\rangle$ this feature map encodes the variable as an N -th degree polynomial formed by $\{1, x, \sqrt{1-x^2}\}$ and their products [73]. The redundancy from many qubits thus forms a basis set for function fitting [157].

The product feature map can be generalized to several layers of rotations $\ell = 1, 2, \dots, L$, various nonlinear functions

φ_ℓ and specific subsets of qubits \mathbb{N} , written as

$$\hat{U}_\varphi(x) = \prod_{\ell=1}^L \bigotimes_{j \in \mathbb{N}_\ell} \hat{R}_{\alpha,j}^{(\ell)}(\varphi_\ell[x]). \tag{7}$$

Next, we show how the quantum feature map can be differentiated. Let us use the example in Eq. (5) considering $\alpha = y$ rotations and a full layer. The derivative for a unitary operator generated by an involutory matrix (length-1 Pauli string in this case) can be written as

$$\begin{aligned} \frac{d}{dx} \hat{U}_\varphi(x) &= \frac{1}{2} \left(\frac{d}{dx} \varphi(x) \right) \sum_{j'=1}^N \bigotimes_{j=1}^N (-i\hat{Y}_{j'}\delta_{j,j'}) \hat{R}_{y,j}(\varphi[x]) \\ &= \frac{1}{2} \left(\frac{d}{dx} \varphi(x) \right) \sum_{j'=1}^N \bigotimes_{j=1}^N \hat{R}_{y,j}(\varphi[x] + \pi\delta_{j,j'}), \end{aligned} \tag{8}$$

where Euler's formula can be used to rewrite the derivative into the form of a sum of unitaries, where x -dependent rotations are shifted one-by-one. Next, the formula can be generalized to the expectation value of any observable $\langle \hat{C} \rangle$ for the encoded state, following the steps of a standard parameter shift rule [73,103]. This reads

$$\frac{d}{dx} \langle \emptyset | \hat{U}_\varphi(x)^\dagger \hat{C} \hat{U}_\varphi(x) | \emptyset \rangle = \frac{1}{4} \left(\frac{d}{dx} \varphi(x) \right) (\langle \hat{C} \rangle^+ - \langle \hat{C} \rangle^-), \tag{9}$$

where $\langle \hat{C} \rangle^+$ and $\langle \hat{C} \rangle^-$ are the sums of shifted unitaries

$$\langle \hat{C} \rangle^\pm = \sum_{j'=1}^N \bigotimes_{j=1}^N \langle \emptyset | \hat{R}_{y,j}^\dagger(\varphi[x] \pm \frac{\pi}{2}\delta_{j,j'}) \hat{C} \hat{R}_{y,j}(\varphi[x] \pm \frac{\pi}{2}\delta_{j,j'}) | \emptyset \rangle. \tag{10}$$

The corresponding derivative quantum circuits (9) are shown in Fig. 3(b), where differentiation of the cost function for

feature map is performed using the chain rule (highlighted rotations). A similar strategy can be applied for generic multilayer feature maps and a different choice of nonlinear map $\varphi(x)$. Finally, in the cases where the generator of the feature map (encoding Hamiltonian \hat{H}) is not an involutory matrix we can rewrite them as a sum of unitary operators, and measure the derivative as a sum of overlap measurements using the SWAP test [104,158,159].

2. Chebyshev feature maps

Next, we consider a distinct choice of nonlinear quantum feature map that we name the *Chebyshev feature map*. Belonging to the product feature map family, it drastically changes the basis set for function representation. As a building block we use a single qubit rotation $\hat{R}_{y,j}(\varphi[x])$, but with nonlinearity introduced as $\varphi(x) = 2n \arccos x$, $n = 0, 1, 2, \dots$, such that the encoding circuit reads

$$\hat{U}_\varphi(x) = \bigotimes_{j=1}^N \hat{R}_{y,j}(2n[j] \arccos x). \quad (11)$$

Here we consider that the coefficient $n[j]$ may in general depend on the qubit position j . Note that the seemingly small change of factor two multiplication for $\varphi(x) = 2 \arccos x$ (as compared to previously considered product map with $\varphi(x) = \arccos x$) goes a surprisingly long way. Namely, let us expand the rotation using Euler's formula, getting

$$\begin{aligned} \hat{R}_{y,j}(\varphi[x]) &= \exp\left(-i \frac{2n \arccos(x)}{2} \hat{Y}_j\right) \\ &= \cos(n \arccos(x)) \mathbb{1}_j - i \sin(n \arccos(x)) \hat{Y}_j \\ &= T_n(x) \mathbb{1}_j + \sqrt{1-x^2} U_{n-1}(x) \hat{X}_j \hat{Z}_j. \end{aligned} \quad (12)$$

The resulting decomposition (12) corresponds to a unitary operation with matrix elements defined by degree- n Chebyshev polynomials of first and second kind, denoted as $T_n(x)$ and $U_n(x)$, respectively [160]. The low degree Chebyshev polynomials of the first kind are written explicitly as $T_0(x) = 1$, $T_1(x) = x$, $T_2(x) = 2x^2 - 1$, $T_3(x) = 4x^3 - 3x$, and higher degrees can be deduced using the recursion relation

$$T_{n+1}(x) = 2xT_n(x) - T_{n-1}(x). \quad (13)$$

Similarly, we can write second-kind Chebyshev polynomials as $U_0(x) = 1$, $U_1(x) = 2x$, $U_{n+1}(x) = 2xU_n(x) - U_{n-1}(x)$. The crucial properties of Chebyshev polynomials are their chaining properties, nesting properties, and simple differentiation rules. The chaining properties for polynomials of the first and second kind read as $2T_m(x)T_n(x) = T_{m+n}(x) + T_{|m-n|}(x)$ and $U_m(x)U_n(x) = \sum_{k=0}^n U_{m-n+2k}(x)$, respectively. Derivatives can be obtained as $dT_n(x)/dx = nU_{n-1}(x)$. Nesting corresponds to the relation $T_n(T_m(x)) \equiv T_{nm}(x)$. Finally, polynomials of different kinds can be converted between as $U_n(x) = 2 \sum_{j \text{ even}}^n T_j(x)$ when n is even, and $U_n(x) = 2 \sum_{j \text{ odd}}^n [T_j(x) - 1]$ when n is odd. Finally, we note that Chebyshev polynomials represent oscillating functions defined in the region $x = (-1, 1)$, and their derivatives diverge at the boundaries of this interval.

The power of the described representation can be inferred from approximation theory. It states that any smooth function




FIG. 4. Fitting with Chebyshev polynomials. As an example we show a fit of the function $y(x) = \exp(-1/25x^2)\cos(20x)$ using different number n Chebyshev polynomials up to degree $n - 1$. We note here that the target function is oscillatory and is decaying exponentially. The thick solid curve denotes the desired function $y(x)$ and the dashed curves are the obtained fitting curves (see legend). Their fit improves with n . The basis set with $n = 2$ is unable to represent the function faithfully (only the linear fit is available). Increased basis set size with $n = 16$ represents the desired function in some regions of x , and $n = 64$ successfully provides an accurate fit.

can be represented as $f(x) = \sum_{n=0}^{\infty} A_n T_n(x)$, $|x| \leq 1$. An example of fitting a set of Chebyshev polynomials of the first kind to a function is shown in Fig. 4. As expected the more basis functions we use, the more accurate a fit we can obtain. As shown in the Remez algorithm [161,162], Chebyshev polynomials form an optimal set of basis functions in the sense of the uniform L_∞ norm. This is why they are at the foundation of spectral algorithms for solving ODEs [34], and also give an edge in quantum simulation [10,163].

In the present study, we consider two types of Chebyshev quantum feature maps. The first version corresponds to a *sparse Chebyshev feature map* defined as

$$\hat{U}_\varphi(x) = \bigotimes_{j=1}^N \hat{R}_{y,j}(2 \arccos x), \quad (14)$$

where the encoded degree is homogeneous and equal to one. Here we make use of the chaining properties $T_n(x)$ and $U_n(x)$, noting that once we create states with Chebyshev polynomials as prefactors, the basis set will grow further by concatenating elements. Henceforth we drop the *sparse* distinction and simply refer to (14) as *Chebyshev feature map*. Importantly, the unitary operators in Eq. (14) have nonlinear dependence on variable x , leading to a harmonic feature mapping.

The second version we consider corresponds to a *Chebyshev tower feature map* defined as

$$\hat{U}_\varphi(x) = \bigotimes_{j=1}^N \hat{R}_{y,j}(2j \arccos x), \quad (15)$$

where the encoded degree grows with the number of qubits, creating a tower-like structure of polynomials with increasing $n = j$. Again, as polynomials chain together and morph between two kinds and their degrees, the basis set is largely enriched. This is the choice we exploit when large expressibility is needed without increasing system size and number of rotations. Eq. (15) allows the representation of generic functions, and can be improved further by using layers of rotations as in Eq. (7).

3. Evolution-enhanced feature maps

Product feature maps induce nonlinear mappings between variable(s) x and quantum states described by tensor products of separate single-qubit wave functions. These states are limited to the subspace of product states. To utilize the power of the entire Hilbert space of the system, approaching the amplitude encoding case, we need to populate independently distinct amplitudes, including the subspace of entangled states. To make the described feature maps even more expressive, we suggest enhancing product feature maps (and specifically the layered Chebyshev map) with additional entangling layers represented by Hamiltonian evolution. Namely, after the set of single qubit rotations we consider another unitary $\exp(-i\hat{H}\tau)$ which acts for time τ and is generated by the Hamiltonian \hat{H} . The sketch of the circuit is shown in Fig. 3(c). By choosing \hat{H} as a complex many-body Hamiltonian we ensure that exponentially many amplitudes are generated. It is known that the quantum simulation of dynamics leads to a volume-like increase of entanglement. One important choice is when \hat{H} corresponds to a hard problem from NP-hard complexity class, as proposed in Ref. [74]. Then using two layers of rotations plus evolution the embedding becomes difficult to simulate classically, but can be implemented as a unitary evolution on a quantum computer. Additionally, we envisage an evolution layer that is parameter-dependent, $\tau(x)$. In this case the *evolution-enhanced feature map* can also be seen through the prism of the recently proposed Fourier feature maps [155]. This class of quantum feature maps is based on the evolution operator $\exp(-i\hat{H}_{\text{data}}x)$, which is applied for some set of qubits. The Fourier map lets functions be encoded as Fourier series defined by the differences of the eigenvalues of \hat{H}_{data} . As it involves unitary operators with phases being linear functions of x , this is a fully harmonic mapping. The evolution-enhanced feature map then joins the Chebyshev and Fourier basis sets, encoded in the full Hilbert space for complex \hat{H} .

4. Digital quantum feature maps

Another possibility for encoding the data using a feature map is to transform a data instance into a computational basis state, which we refer to as a *digital quantum feature map*. This relates x , written in binary form, to the corresponding state $|x\rangle$ in binary representation. The feature map circuit $\hat{U}_\varphi(x)$ used to encode the binary variable x reads $\hat{U}_\varphi(x) = \bigotimes_{j=1}^N \exp(-i\frac{\pi}{2}x_j\hat{X}_j)$, where $\{x_j\}$ denote binary values for the parameter x in j th digit. The differentiation of the digital feature map then relies on the product rule for N rotations, and also includes the binary derivative of the variable from the product rule.

Another possibility is converting the variable into a decimal representation as $x_{\text{int}} = \sum_j x_j \cdot 2^j$. For the reverse procedure, we can identify each binary digit x_j by the remainder of the repeated division $x_j = \text{mod}(x_{\text{int}}, 2^j)$. We can thus rewrite $\hat{U}_\varphi(x)$ as a function of x_{int} , and learn how to differentiate circuits with this feature map with respect to $x = x_{\text{int}}$.

Digital quantum feature maps offer a potentially powerful technique when dealing with functions of discrete variables. Paired with expressive variational ansatzes they may offer

function encoding with memory savings of the amplitude encoding, while avoiding the input problem. The details of digital feature map training will be discussed in future works.

B. Variational quantum circuits

To construct the solution of differential equations as a quantum circuit we need to manipulate the latent space basis function and bring both derivatives and function to the required form. This is achieved through the variational circuit \hat{U}_θ , typically referred to as a *variational quantum ansatz*. Below we detail architectures employed.

1. Hardware efficient ansatz

As a first choice of variational circuit \hat{U}_θ we consider layers of parametrized rotations, followed by layers of CNOT operations. This is known as a hardware efficient ansatz (HEA), which was originally proposed for VQE for chemistry applications [45]. The structure of a HEA quantum circuit corresponds to concatenated layers of single qubit rotations and global entangling layers for all N qubits, shown schematically in Fig. 5(a). Rotations are arranged in a \hat{R}_z - \hat{R}_x - \hat{R}_z sequence parametrized by independent angles θ such that arbitrary single-qubit operations can be reproduced. The entangling layer is chosen as a network of CNOTs. We here consider specifically a linear quantum device connectivity, but the entangling layer can be generalized to other connectivities. The block of rotations plus CNOTs is then repeated for a depth of d times. As the number of layers d grows, the circuit's expressive power increases. Expressive power is the ability to represent arbitrary N -qubit unitary gates. However, this comes with an increased number of controlled parameters which can complicate the search for an optimal θ_{opt} for the solution, known as a problem in trainability. This can change depending on the variational ansatz type and cost function choice [70].

2. Alternating blocks ansatz

A second option is to use the alternating blocks ansatz (ABA), where instead of global entangling layers separate subblocks are used, interleaved into a checkerboard form [Fig. 5(b)]. Each subblock has a hardware efficient form shown in Fig. 5(a) for the specified depth b . For the first layer the width of the subblock (number of active qubits) is equal to N_b such that $\lceil N/N_b \rceil$ blocks are used (and is smaller than N_b if N/N_b is not an integer). The next layer consists of the same subblocks, but is shifted by $\lfloor N_b/2 \rfloor$, where subblocks at the ends are adjusted to span the remaining qubits. The described checkerboard-like structure is repeated for d_{layers} . The motivation behind ABA is that we want to entangle qubits locally first, and gradually form a correlated state by interleaving subblocks. This helps to improve trainability of the circuit together while maintaining high expressibility [70,164].

Here, several points are due. First, we note that the choice of ansatz is sensitive to the choice of the cost function operator (see the next subsection), as dictated by the symmetry. Namely, as a consequence of cost function choice we need to choose noncommuting generators \hat{G}_j for the variational ansatz such that $[\hat{C}, \hat{G}_j] \neq 0$. This ensures that the solution space can be spanned. Also, we can account for the symmetries,




FIG. 5. Variational quantum circuits. (a) A variational ansatz in the hardware efficient form. It consists of a parametrized rotation layer forming \hat{R}_z - \hat{R}_x - \hat{R}_z pattern, such that an arbitrary single qubit rotation can be implemented. Variational angles θ are set for each rotation individually. The rotation layer is then followed by an entangling layer chosen as CNOT operations between nearest neighbours. The blocks of “rotations-plus-entangler” are repeated d times to form the full variational circuit \hat{U}_θ . (b) Alternating blocks ansatz. The variational circuit consists of blocks of width N_b qubits ($N_b/2$ for boundary qubits). Blocks are chosen in the hardware efficient form shown in (a) with depth of b . The blocks are placed in a checkerboard pattern, and repeated d_{layers} times. The goal is to entangle qubits locally, while avoiding global entangling operations that can result in vanishing gradients during θ optimization.

reducing the Hilbert space for the solution search. In many cases generators can be chosen such that only real amplitudes are generated, see for instance [70]. We can also use adaptive strategies or a genetic search [54,55,165].

Second, to search for the optimal circuit parameters we use stochastic gradient descent, and specifically its adaptive version represented by Adam [166]. For this, the gradients of the variational circuit $\nabla_\theta \hat{U}_\theta$ are measured using the automatic differentiation approach, as performed in previous studies [73,81,103]. The gradient is also weighted with a learning rate, commonly referred to as α . Choosing an ansatz parametrized by single-qubit rotations allows the application of the parameter shift rule, while overlap measurement opens up options for more general strategies [158].

Finally, we note that choosing the optimal ansatz is an open challenge for improving DQC training. In comparison to chemistry problems where the structure may be inspired by physical reasoning [39,40], this is not the case for differential equation solvers and further studies in this direction are needed.

C. Cost function

To read out information we choose a Hermitian cost operator \hat{C} to measure as an observable. This allows relating the expectation of \hat{C} , parametrized by nonlinear variable dependence and variational angles, to the scalar function $f(x)$ as $\langle f_{\varphi,\theta}(x) | \hat{C} | f_{\varphi,\theta}(x) \rangle$. In general many possible choices of cost operators are available. The simplest example corresponds to the magnetization of a single qubit j , $\langle \hat{Z}_j \rangle$. Note that this choice allows representing functions in range $[-1, 1]$, and requires rescaling for other intervals. Other choices include total magnetization in the system $\hat{C} = \sum_j \hat{Z}_j$ with equal or randomized weights.

Additionally, we can choose the cost as a quantum Hamiltonian that has a provably complex spectrum, and for instance belongs to ergodic phase. This can be written as an Ising Hamiltonian with additional transverse and longitudinal magnetic fields,

$$\hat{C} = \sum_j J_{j,j+1} \hat{Z}_j \hat{Z}_{j+1} + h_j^z \hat{Z}_j + h_j^x \hat{X}_j, \quad (16)$$

where the Ising couplings $J_{j,j+1}$ and $h_j^{z,x}$ can be inhomogeneous. We note that for cost functions with several noncommuting groups of observables we can use a Hamiltonian averaging procedure, where term-by-term measurement is performed. Moving on from nearest-neighbour Hamiltonians, we can also exploit spin-glass type cost functions of the form

$$\hat{C} = \sum_{i<j} J_{i,j} \hat{Z}_i \hat{Z}_j + \sum_j h_j^z \hat{Z}_j. \quad (17)$$

These are known to include NP-hard problem instances, and they allow for high expressibility of the circuit describing the DE's solution. Finally, a generic cost function may contain a large set of Pauli strings, similar to some instances in quantum chemistry.

Together with measuring individual cost functions, we also consider functions being represented by a classically weighted sum of observables. This reads

$$\hat{C} = \sum_\ell \alpha_\ell \hat{C}_\ell, \quad (18)$$

where $\alpha_\ell \in \mathbb{R}$ are weighting coefficients, and \hat{C}_ℓ are cost functions that can be chosen from the pool of operators described above. Importantly, we consider the coefficients α_ℓ to be tunable, such that the gradient descent (represented by Adam in our case) can adjust the cost to have optimal form. This procedure further harnesses the strength of the proposed hybrid quantum-classical workflow.

D. Loss function

To solve the system of differential equations we need to provide a way to quantify how well the DQC-suggested trial function matches the conditions to represent the solution of the problem being considered. The classical optimizer can then update the variational parameters to reduce this “distance.” This distance corresponds to the difference between a differential equation (all terms collected on one side) and zero. The difference is evaluated at a set of points. We also need to check if the solution matches initial and boundary

conditions. This can be recast as an optimization problem for a loss function of derivatives and functions evaluated at a grid of points.

We can write a loss function parametrized by variational parameters θ in the general form

$$\mathcal{L}_\theta[d_x f, f, x] = \mathcal{L}_\theta^{(\text{diff})}[d_x f, f, x] + \mathcal{L}_\theta^{(\text{boundary})}[f, x], \quad (19)$$

where we split the loss contribution from matching the differentials $\mathcal{L}_\theta^{(\text{diff})}$ and the loss contribution from satisfying the boundary conditions $\mathcal{L}_\theta^{(\text{boundary})}$. The differential loss is defined as

$$\mathcal{L}_\theta^{(\text{diff})}[d_x f, f, x] = \frac{1}{M} \sum_{i=1}^M L(F[d_x f(x_i), f(x_i), x_i], 0) \quad (20)$$

with $L(a, b)$ being a function describing how the distance between the two arguments a and b is being measured. The loss is estimated on a grid of M points, and is normalized by the grid size. Functional F corresponds to the differential equation written in the form $F[d_x u, u, x] = 0$. It can be evaluated by combining values of f and $d_x f$ at the training grid points. We stress that the functional includes information about all differential equations when we deal with the system, such that contributions from all equations are accounted for. The boundary loss contribution reads

$$\mathcal{L}_\theta^{(\text{boundary})}[f, x] = \eta L(f(x_0), u_0), \quad (21)$$

which includes the distance between the function value at the boundary x_0 and given boundary value u_0 . We note that x_0 can be an initial point or a set of boundary points. We also introduce η as a *boundary pinning* coefficient that controls the weight of the boundary term in the optimization procedure. In particular, larger $\eta > 1$ may be used to ensure the boundary is prioritised and represented to higher precision.

We have considered several choices of the loss defined by three distance definitions L . The first loss type corresponds to the *mean square error* (MSE) introduced as

$$L(a, b) = (a - b)^2. \quad (22)$$

While being simple, we find the choice (22) intuitive and performing sufficiently well in numerical simulations. Additionally we also consider *mean absolute error* (MAE) loss defined with distance $L(a, b) = |a - b|$. Finally, several more complex metrics can be used, including variants of Kullback-Leibler (KL) divergence and Jensen-Shannon divergence. Being the loss functions routinely used in statistical modeling, we expect them to perform well for specific systems.

The choice of loss functions dictates how the optimizer perceives the distance between vectors and therefore affects the convergence. MSE places a greater emphasis on larger distances and smaller weight on small distances, strongly discouraging terms with large L . Both MAE and KL do not place such an emphasis and may have slower convergence. However, once close to the optimal solution they can achieve higher accuracy than MSE.

E. Boundary handling

As our goal is to construct a quantum circuit that satisfies a system of differential equations, together with matched derivatives we need to ensure that an initial value or boundary value problem is solved. Generally this corresponds to fixing the function value at a required initial point or a collection of boundary points, thus resembling the quantum circuit learning tasks considered in Ref. [73]. At the same time, there are several ways how the DQC-based function $f_\theta(x)$ can be constructed, leading to varying performance and specific pros/cons when solving particular problems.

Information about the boundary can be included as part of the loss function [Eq. (19)]. For the MSE loss function type the boundary part (21) can be written in the form

$$\mathcal{L}_\theta^{(\text{boundary})}[f, x] = \eta (f_\theta(x_0) - u_0)^2, \quad (23)$$

where x_0 represents the set of boundary points (or an initial point), and u_0 is a vector of boundary values, and η is a pinning coefficient as described previously.

1. Pinned boundary handling

The first option is to include the information about the boundary in the expectation of the cost function. This corresponds to simply choosing a cost operator \hat{C} , and representing the solution in the form

$$f(x) = \langle f_{\varphi, \theta}(x) | \hat{C} | f_{\varphi, \theta}(x) \rangle. \quad (24)$$

The initial value u_0 is then matched via the boundary term in the loss function. The strength of the *pinned boundary* handling is in equivalent treatment of boundary and derivative terms, both being encoded in \hat{C} . At the same time, the weakness corresponds to the necessity of adjusting the boundary value starting from the one represented by initial θ_{init} , typically generated randomly. This can be adjusted by shifting $f(x)$ by a constant-times-identity term added to the cost operator, $\hat{C} = \alpha_0 \mathbb{1} + \sum_{j=1}^M \alpha_j \hat{C}_j$, where α_0 is set such that for $\theta_{\text{init}} \sim \text{random}[0, 2\pi]$ the function $\langle f_{\varphi, \theta_{\text{init}}}(x) | \hat{C} | f_{\varphi, \theta_{\text{init}}}(x) \rangle$ typically lies close to u_0 value when evaluated at $x = x_0$.

2. Floating boundary handling

The second choice of the boundary handler corresponds to iteratively shifting the estimated solution based on the boundary or initial point. For this method the boundary information does not require a separate boundary loss term nor is it encoded in the expectation of the cost function. Instead it is set iteratively within the parametrization of the function. As the function is parametrized to match a specific boundary, information about the boundary is still contained within the function and its derivatives. Therefore boundary information is still present within the loss function despite there not being a separate boundary loss term. We parametrize the function as

$$f(x) = f_b + \langle f_{\varphi, \theta}(x) | \hat{C} | f_{\varphi, \theta}(x) \rangle, \quad (25)$$

with $f_b \in \mathbb{R}$ being a parameter adjusted after each iteration step as

$$f_b = u_0 - \langle f_{\varphi, \theta}(x_0) | \hat{C} | f_{\varphi, \theta}(x_0) \rangle. \quad (26)$$

This effectively allows the solver to find a function $\langle f_{\varphi,\theta}(x) | \hat{C} | f_{\varphi,\theta}(x) \rangle$ which solves the differential equation shifted to any position, then being shifted to the desired initial condition as shown in Eq. (25). This method of boundary handling guarantees exact matching to initial values given and does not require a separate boundary term in the loss function, thus the derivative loss term does not have to compete with the boundary loss. Furthermore, as we allow the cost function to match to the solution shifted by any amount, this simplifies the choice for optimal angles and removes the dependence on initial θ_{init} . However this method does require knowledge of an exact initial value which can be an issue in specific situations. This technique can be generalized for multivariable problems which have an initial condition as a function of a subset of the independent variables. Evaluation of the derivatives of the initial condition function is required for encoding the partial derivatives of the represented function.

3. Optimized boundary handling

Finally, we also propose a boundary handling technique that relies on a classical shift of the solution, but defined by the gradient descent procedure on par with variational angles optimization. This removes the need to include boundary information in the cost expectation, but information still needs to be included in the loss function. Namely, we seek for the solution in the form

$$f(x) = f_c + \langle f_{\varphi,\theta}(x) | \hat{C} | f_{\varphi,\theta}(x) \rangle, \quad (27)$$

where $f_c \in \mathbb{R}$ is a variational parameter alongside the quantum ansatz angles and updated accordingly via the classical optimizer. Therefore the gradients for f_c have to be calculated additionally when using this boundary handler. One strength of the described method is that, due to the classical shift, even if the random initial angles start such that $\langle f_{\varphi,\theta_{\text{init}}}(x_0) | \hat{C} | f_{\varphi,\theta_{\text{init}}}(x_0) \rangle$ is far from the initial value u_0 , the optimizer can quickly and easily update f_c to rectify this. A weakness however is that the boundary and differential terms in the loss may compete against one another.

F. Regularization

Given that our goal is to find a variational spectral representation of the differential equations solution using large basis sets, the optimization procedure benefits from having a good initial guess, or “pre-trained” DQCs. We can achieve this by introducing a regularization procedure [167,168], also helping the optimizer to avoid getting trapped in local minima. Variants of the regularization procedure include: (1) feeding-in prior information about the potential solution; (2) biasing the DQC-based solution into a specific form; and (3) searching for a solution in a region close to the boundary values, and feeding-in points from the first training into next sessions. The input for procedures 1 and 2 consists of regularization points for the variable(s) $\{x_{\text{reg}}\}_{r=1}^R$, together with corresponding function values $\{u_{\text{reg}}\}_{r=1}^R$ for R points. Similarly, we can consider regularization based on the derivative values. We employ the simplest strategy where an additional contribution to the loss function comes from the regularization points, $\mathcal{L}_{\theta}^{(\text{reg})}[f, x]$. This loss is defined such that the DQC-based function matches the regularization values at corresponding grid points. This

has a form analogous to the boundary loss contribution. Using MSE loss as example, the regularization contribution reads

$$\mathcal{L}_{\theta}^{(\text{reg})}[f, x] = \sum_{r=1}^R \zeta(n_j) (f_{\theta}(x_{\text{reg},r}) - u_{\text{reg},r})^2, \quad (28)$$

where n_j denotes the iteration step. $\zeta(n_j)$ is introduced as an iteration step-dependent regularization weight, and thus denoting an *optimization schedule*. In general, we require higher emphasis on the regularization-based training at initial stages, which shall diminish to zero at higher iteration numbers. This leads to the prior information being used at first, setting a rough solution or preferred function behavior, followed by precise derivative loss optimization at later training stages. One possible choice of an optimization schedule corresponds to linearly decreasing regularization weight, $\zeta(n_j) = 1 - n_j/n_{\text{iter}}$, where n_j is current iteration number and n_{iter} is the maximum iteration number. This strategy works for small learning rates and large number of iterations, such that the optimizer has sufficient “time” to adjust to the constantly changing loss landscape. Another choice corresponds to a *reverse sigmoid optimization schedule*, where a smooth drop of regularization weight is performed at pre-defined training stages. We parametrize this schedule as

$$\zeta(n_j) = 1 - \tanh\left(\frac{n_j - n_{\text{drop}}}{\delta_j n_{\text{iter}}}\right), \quad (29)$$

where n_{drop} denotes the iteration step number at which regularization weight drops, and δ_j assigns the transition rate. This allows the DQC to initially focus almost entirely on the regularization optimization, later switching the focus towards the gradient optimization.

G. Multifunction encoding

When solving a *system of differential equations* we need to decide how multiple functions should be encoded simultaneously. There are several ways that this can be achieved. The first approach we consider is to use the same quantum register for all functions, thus compressing information about the function vector using the same feature map circuit and variational ansatz. The functions are then defined through the choice of different cost operators at the readout stage. This method is resource-frugal, and is suitable for certain systems. However, the choice of suitable cost operators becomes complicated, as in some cases shared register encoding may exhibit competition between the optimization of different functions. This concerns the question of expressibility of the set of cost operators, and may potentially be solved using a weighted sum of operators with optimized weights. A second option is to use separate quantum registers, and correspondingly a different feature map and variational ansatz for each function. This removes the issue of choosing the cost operators, and avoids the related direct competition due to independent parametrization. While requiring more resources we note that function and derivative evaluation can be done in parallel. However, we note that as a combined loss function is considered care must be taken when timing the circuit runs.

H. DQC-based solver: the workflow

Finally, using the elements and strategies described above, we present a workflow for constructing the differential equation solver based on derivative quantum circuits. This is summarized in Fig. 1 showing the flowchart. We start by specifying the input for the solver. This comprises the problem in hand, specified as a set of nonlinear differential equations of various types, together with their respective boundary conditions. Additionally, a set of regularization points may be added to ensure the optimized solution is chosen in the desired qualitative form. Next, we set up the schedule for derivative quantum circuit optimization and choose the quantum circuit composition. For this we choose (a) the type of quantum feature map; (b) the ansatz of variational quantum circuit, including its depth; (c) the cost function type, also choosing if variational weights are considered; (d) the type of the loss function; and (e) the strategy to match the boundary terms and derivatives. We also need to specify the classical optimizer for variational angles and weights (with associated hyperparameters), including the number of iterations and exit conditions. Finally, we specify whether the loss function uses a specific optimization schedule, where it is changing during the training.

Once the DQC structure and optimization schedule are defined, we need to specify a set of points \mathbb{X} for each equation variable. This can be a regular equidistant grid, Chebyshev grid, or a randomly drawn grid. The variational parameters are set to initial values $\theta \leftarrow \theta_{\text{init}}$ (e.g., as random angles). The expectation value over variational quantum state $|u_{\varphi, \theta}(x_i)\rangle$ for the cost function is estimated using the quantum hardware for the chosen point x_i . Then a potential solution at this point is constructed, accounting for the boundary handling procedure. Next, the derivative quantum circuits are constructed and their expectation value is estimated for the specified cost function, at point x_i . Repeating the procedure for all x_i in \mathbb{X} we collect function values and derivatives, and compose the loss function for the entire grid and system of equations (forming required polynomials and cross-terms by classical post-processing). Regularization points may be also added, biasing the solution to take specific values at these points. The goal of the loss function is to assign a “score” to how well the potential solution (parametrized by the variational angles θ) satisfies the differential equation, matching derivative terms and the function polynomial to minimize the loss. With the aim to increase the score (and decrease the loss function), we compute the gradient of the loss function with respect to variational parameters θ . Using the gradient descent procedure (or in principle any other classical optimization procedure) we update the variational angles from iteration $n_j = 1$ into the next one $n_j + 1$, $\theta^{(n_j+1)} \leftarrow \theta^{(n_j)} - \alpha \nabla_{\theta} \mathcal{L}$ (with α being here a “learning” rate), and repeat the steps outlined before until we reach the exit condition. The exit condition may be chosen as (1) the maximal number of iterations n_{iter} reached; (2) loss function value is smaller than pre-specified value; and (3) loss gradient is smaller than a certain value. Once we exit the classical loop, the solution is chosen as a circuit with angles θ_{opt} that minimize the loss. Finally, we extract the full solution by sampling the cost function for optimal angles $\langle u_{\varphi, \theta}(x) | \hat{C} | u_{\varphi, \theta}(x) \rangle$. Notably, this can be done for any

point x , as DQC constructs the solution valid also beyond (and between) the points at which loss is evaluated originally.

IV. RESULTS

A. Differential equation example

Now let us see how the algorithm performs in practice. For this, we choose a differential equation with a known analytical solution, and compare it to the one obtained by the derivative quantum circuit. We choose a single ODE for the initial value problem which reads

$$\frac{du}{dx} + \lambda u(\kappa + \tan(\lambda x)) = 0, \quad u(0) = u_0, \quad (30)$$

where λ and κ are real parameters, and u_0 sets the value of the function u at $x = 0$. Equation (30) has a solution in the form of a damped oscillating function,

$$u(x) = \exp(-\kappa \lambda x) \cos(\lambda x) + \text{const}, \quad (31)$$

where “const” is determined by the initial condition. While the problem is fairly simple being a single ODE, reproducing the damped oscillating solution requires a rich basis of fitting functions, that needs to include both oscillatory and increasing/decreasing functions. As λ and κ grow the function starts to oscillate and decay even more rapidly, and the solution becomes harder to express.

To show how the proposed method works, we use derivative quantum circuits to solve Eq. (30) using optimization of differentiable quantum feature maps. Specifically, we choose two cases with parameters $\lambda = 8$ and 20 , and fixed $\kappa = 0.1$, $u_0 = 1$. These problem parameters are chosen to make DQC construction challenging, with $\lambda = 20$ being a complex case as the resulting solution is highly nonlinear and oscillatory. We consider an equidistant optimization grid of 20 points, starting from $x = 0$, with maximal time of 0.9 (dimensionless units are used). This is chosen such that the region with diverging derivative of the nonlinear feature function $d\varphi(x)/dx$ is avoided, and we note that x can be rescaled to match required boundaries. To find the solution we use a quantum register with $N = 6$ qubits, and the cost function is chosen as total magnetization in the Z direction, $\hat{C} = \sum_{j=1}^N \hat{Z}_j$. For the variational circuit, we choose the standard hardware efficient ansatz described in the Methods section, setting the depth to $d = 5$. To search for optimal angles θ_{opt} we perform adaptive stochastic gradient descent using ADAM [166] with automatic differentiation enabled by analytical derivatives. Specifically, we code the workflow using YAO.JL package [81,169,170] for Julia, which allows fast and efficient implementation. A full quantum state simulator is used in a noiseless setting. In this example, we use the *floating boundary* handling.

We search for the circuit-based solution using three different feature maps described in the Methods section, and compare their performance. These correspond to the *product feature map* [Eq. (6)], the sparse version of the *Chebyshev feature map* [Eq. (14)], and the *tower Chebyshev feature map*, as defined in Eq. (15). We label results for these feature maps as Prod, Chev, and ChevT, respectively. To assess the performance we use several metrics. The first metric is the *full loss* (denoted as L_F). It refers to the loss calculated from the differential equations and any boundary or regularization




FIG. 6. DQC-based solution for the single ODE example. (a) Results from the circuit trained to solve Eq. (30) for $u_0 = 1$, $\lambda = 8$, and $\kappa = 0.1$. We present DQC solutions $f_{\varphi,\theta}(x)$ obtained using three different quantum feature maps $\varphi(x)$ labeled as Prod, Cheb, and ChebT. Corresponding DQC solutions are shown by dashed curves and labeled for clarity (see legend). Analytical solution $u(x)$ is shown by the thick solid curve (label 1). (b) Same as in (a), but for $\lambda = 20$ example. [(d) and (e)] Full loss L_F (dashed curves) and quality of solution L_Q (solid curves) are shown as functions of iteration number n_j for optimization results displayed in (a) and (b), respectively. (c) DQC-based solution of Eq. (30) shown for four different ansatz depths $d = 3, 6, 12$, and 24 , with the analytical solution $u(x)$ presented by the solid curve. The case of $d = 3$ is labeled for clarity. (f) Full loss (dotted curves) and quality of solution (solid curves) shown as a function of iteration number for the solution in (c).

terms, in this case $\mathcal{L}_\theta[d_x f, f, x] = \mathcal{L}_\theta^{(\text{diff})}[d_x f, f, x] + \mathcal{L}_\theta^{(\text{boundary})}[f, x] + \mathcal{L}_\theta^{(\text{reg})}[f, x]$. The second metric corresponds to differential loss (L_D), being a part of the full loss excluding regularization contribution. Finally, the third metric is the quality of solution (L_Q). The quality of the solution is the distance of the current DQC-based solution from the known true solution. This is calculated by evaluating the DQC-based solution and true solution at a set of points and using the MSE loss type, being equal to $L_Q = (1/M) \sum_{i=1}^M [f(x_i) - u(x_i)]^2$. Quality of solution gives us a useful way to compare how two different training setups perform, especially if they are training to solve the same differential equations.

The results of DQC training are shown in Fig. 6. In Figs. 6(a) and 6(b), we show the solutions of Eq. (31) for $\lambda = 8$ and 20 , respectively, where solid curve (with label 1) represents the analytical solution $u(x)$ in Eq. (31). The dashed curves (curves 2, 3, 4) represent the final DQC-based solutions sampled from the cost function at approximation points ($n_{\text{iter}} = 250$ is used). In Figs. 6(d) and 6(e), we show the relevant training metrics as a function of iteration number, where solid curves denote the quality of solution (curves 2, 4, and 6) and dashed curves represent the full loss (curves 1, 3, and 5). We observe that for $\lambda = 8$ both Chebyshev feature maps converge closely to the true solution [Figs. 6(a) and 6(d)]. The more expressible Chebyshev tower feature map (curves 5, 6) takes longer to converge but reaches a solution closer to the

true solution. The less powerful product feature map fails to converge with the loss quickly plateauing (curves 1, 2). For $\lambda = 20$ the true solution is more oscillatory and has stronger damping, making the solution harder to represent [Figs. 6(b) and 6(e)]. The product feature map still fails to converge (curves 1, 2), but now also failing to converge is the simpler Chebyshev feature map (curves 3, 4). The full loss for both cases plateaus rapidly. The more expressible ChebT feature map continues to perform well (curves 5, 6). This supports the hypothesis that choosing a feature map expressible enough for the problem is important, and more simulations with more qubits offers a way to increase the power drastically.

Next, we compare the effect of ansatz depths for the variational circuit \hat{U}_θ . We use $d = 3, 6, 12, 24$, $\lambda = 20$ and the Chebyshev tower feature map, and the rest of the training setup remains the same as previously considered. These results are presented in Figs. 6(c) and 6(f). We observe that for lower depths the solver is slower to converge and does not reach as high accuracy as it does for higher depths. As depth increases more layers of parametrized gates are included in the variational ansatz and so the number of variational angle parameters increase. This causes an increase in the number of gate operations needed in each iteration and how many parameters the classical optimizer needs to update, raising the time taken per iteration. As the depth of the ansatz continues to increase eventually the problem of barren plateaus could

be encountered [171]. Then vanishing gradients would cause the solver to struggle to improve the parameters, however at $d = 24$ we had not yet ran into this with over 400 variational parameters. We also note that the alternating blocks ansatz is designed in the way that vanishing gradients can be avoided for certain conditions [164].

B. Differential equation with highly nontrivial dynamics

To highlight the importance of powerful feature maps for solving differential equations, we provide another example with a rapidly oscillation nonperiodic solution. We consider an initial value problem

$$\frac{du}{dx} - 4u + 6u^2 - \sin(50x) - u \cos(25x) + 1/2 = 0, \quad u(0) = u_0, \quad (32)$$

for the function $u(x)$. This differential equation has a solution that is hard to represent since it involves both oscillatoric terms with high frequency and nonoscillatoric terms representing growth and decay. Therefore a feature map with high expressibility is vital for this problem. To illustrate this point, we solve Eq. (32) using DQCs with both the regular Chebyshev tower feature and an *evolution-enhanced* version of that feature map. As shown in Fig. 6, the Chebyshev tower feature map was the most expressible of the three considered for that example. In this example the evolution-enhanced feature map we consider consists of an application of the Chebyshev tower feature map followed by the unitary evolution $\exp(-i\hat{H}\tau)$ with Hamiltonian \hat{H} . Specifically, we have chosen a nearest-neighbor Ising Hamiltonian $\hat{H} = -J \sum_j \hat{Z}_j \hat{Z}_{j+1} + h \sum_j \hat{X}_j$. The evolution time τ is fixed, and we consider couplings J and magnetic field h to be drawn uniformly from $(0,1)$. Note that the type of Hamiltonian is not crucial, and may depend on hardware implementation.

For the training regime we consider $u_0 = 0.75$, 100 training points, and set evolution time $\tau = 2$. We check that results remain consistent for different J and h generated randomly. We use Adam optimizer with learning rate 0.01 and 200 iterations. The rest of the training setup is the same as considered for problem Eq. (30) in Sec. IV A with a six-qubit register, cost choice of total magnetization in the Z direction, and a hardware efficient variational ansatz with depth $d = 5$.

The results are shown in Fig. 7. First, the circuit is trained with the Chebyshev tower feature map. Being expressive in itself, it manages to come close to the true solution but misses some details for getting a precise solution [Fig. 7(a), red dashed curve labeled as 2]. The angles this training results in are then passed as initial angles for training with the evolution-enhanced feature map. The circuit is now further trained for another 200 iterations to improve upon the solution obtained from the initial training. As can be seen, after this second stage of training, we are able to capture the intricate details of the solution [Fig. 7(a), blue dashed curve labeled by 3], and both the full loss and quality of solution are improved upon [Fig. 7(b)]. The expressive power of the evolution-enhanced map can be increased for more complex \hat{H} and larger τ , though at the expense of trainability. The full discussion of the evolution-enhanced feature map will be the subject of separate work.




FIG. 7. DQC-based solution for highly nontrivial dynamics example. (a) Results from the circuit trained to solve initial value problem in Eq. (32) with $u_0 = 0.75$. We show solutions obtained using two different feature maps: Chebyshev tower feature map (ChebT) and *evolution-enhanced* Chebyshev tower feature map (EvEn). Corresponding DQC solutions are shown by dashed curves (labeled by 2 and 3, respectively). The true solution is shown by a solid curve (labeled by 1). (b) Full loss L_f (dashed curves 1 and 3) and quality of solution L_Q (solid curves 2 and 4) are shown as functions of iteration number n_j for optimization results in (a).

We note that by adding a quantum evolution we increase expressibility of the circuit. This helps to represent highly oscillatoric functions. At the same time, DQCs in this case are more difficult to train, as keeping track of the gradients for oscillatoric function requires a finer training grid and has a slower convergence. We envisage that $\tau[n_j]$ can be set as a *flowing* parameter of the feature map that increases from zero as iteration number n_j grows (similar to annealing procedure as in adiabatic quantum computing). This will ensure easy trainability at the start, followed by adjusting θ 's for a larger fitting function set providing extra precision. Another approach is to take $\tau(x)$, adding variable-dependent amplitudes directly through the evolution. While in this case parameter shift rule is not applicable, the circuit differentiation can be done numerically [172], with wave-function overlaps measured with the SWAP test [158].

C. Strongly coupled equations

Building up on the single ODE example, we proceed to consider a *system* of differential equations, taking two strongly coupled differential equations as an example. This describes the evolution of competing modes $u_1(x)$ and $u_2(x)$ as a function of variable x , which in this case corresponds to time. The associated rate equations read

$$F_1[d_x \mathbf{u}, \mathbf{u}, x] = \frac{du_1}{dx} - \lambda_1 u_2 - \lambda_2 u_1 = 0, \quad u_1(0) = u_{1,0}, \quad (33)$$

$$F_2[d_x \mathbf{u}, \mathbf{u}, x] = \frac{du_2}{dx} + \lambda_2 u_2 + \lambda_1 u_1 = 0, \quad u_2(0) = u_{2,0}, \quad (34)$$

where $\lambda_{1,2}$ are coupling parameters, $u_{1,0}, u_{2,0}$ are initial conditions. The larger $|\lambda_1|$ is in comparison to $|\lambda_2|$ the more strongly coupled the two equations will be. This can be intuitively seen considering $|\lambda_1| \geq |\lambda_2|$, leading to larger contribution of u_2 into the equation for du_1/dx derivative than u_1 , and vice versa.




FIG. 8. DQC-based solution for the strongly coupled equations example. (a) Results from the circuit trained to solve the system of differential equations Eqs. (33) and (34) for $u_{1,0} = 0.5$, $u_{2,0} = 0$, $\lambda_1 = 5$, and $\lambda_2 = 3$. We present DQC solutions $f_1(x)$ and $f_2(x)$ obtained using three different boundary evaluation techniques: pinned boundary (labeled as Pin), floating boundary (Float), and optimized boundary (Optim). Corresponding DQC solutions are shown by dashed curves. Analytical solutions $u_1(x)$ and $u_2(x)$ are shown by thick solid curves. (b) Full loss L_F (dashed curves 1, 3, 5) and quality of solution L_Q (solid curves 2, 4, 6) are shown as functions of iteration number n_j for optimization results displayed in (a).

To move from considering one differential equation to a system of equations, we need to encode multiple functions using quantum registers as described in the Sec. III Methods. For this specific example we choose a simple parallel encoding, where separate cost functions and ansatze are considered. In this case, each function has a separate set of parameters to be optimized. The loss is changed accordingly to include information on separate contributions from two coupled differential equation that are optimized simultaneously. We encode each function using the differentiable feature map combined with individual variational ansatz parametrized by the set of angles θ_1 and θ_2 and before deciding on the boundary evaluation type we have

$$f_1(x) = \langle \emptyset | \hat{U}_\phi^\dagger(x) \hat{U}_{\theta_1}^\dagger \hat{C}^{(1)} \hat{U}_{\theta_1} \hat{U}_\phi(x) | \emptyset \rangle, \quad (35)$$

$$f_2(x) = \langle \emptyset | \hat{U}_\phi^\dagger(x) \hat{U}_{\theta_2}^\dagger \hat{C}^{(2)} \hat{U}_{\theta_2} \hat{U}_\phi(x) | \emptyset \rangle, \quad (36)$$

where $\hat{C}^{(1,2)}$ are in principle different cost functions for each equation. For the loss function, we consider the sum of the MSE losses for the first and second differential equation. This loss is written as $\mathcal{L}_\theta[d_x f, f, x] = \sum_{i=1}^M L(F_1[d_x f, f, x_i], 0) + \sum_{i=1}^M L(F_2[d_x f, f, x_i], 0)$ with F_1 and F_2 as written in Eqs. (33), (34), and L depends on the loss choice as detailed in the Methods section. There can be additional boundary loss terms depending on boundary evaluation method chosen. If present, they contribute to the loss function as a sum of the boundary terms for u_1 and u_2 . Note that we consider the loss as a sum of individual contributions coupled together we are trying to minimise both simultaneously with equal weight. Due to the coupling between the two equations this could lead to competition between the two loss terms (a parameter update which causes a loss decrease for one DE's loss may lead to an increase in the other). This may result in increasing the chance of converging to a local minima rather than the global minimum; however, this effect can be mitigated if loss contributions are weighted in some way. Another solution is to use quantum kernel methods, where loss corresponds to the overlap between quantum feature states. Choosing a suitable loss function in this case is an important point to consider in the future.

We define the problem setting parameters to $\lambda_1 = 5$, $\lambda_2 = 3$ and initial conditions to $u_{1,0} = 0.5$, $u_{2,0} = 0$. We set up the training scheme as in Sec. IV A using a six-qubit register, cost choice of total magnetization in the Z direction for both u_1 and u_2 , hardware efficient variational ansatz with depth $d = 5$, Adam optimizer with learning rate 0.02, and feature map choice of the Chebyshev tower feature map. We test the performance for the three boundary evaluation types: *pinned boundary*, *floating boundary*, and *optimized boundary*. The results are shown in Fig. 8. The pinned and optimized boundary handlers perform similarly, slowly converging to the analytical solution $u_{1,2}(x)$ within 250 iterations [Fig. 8(a)]. The two approaches have similar convergence in terms of the full loss [Fig. 8(b), dashed curves labeled by 1 and 5], but differ in terms of quality of solution [Fig. 8(b), dashed curves labeled by 2 and 6]. When using floating boundary type a function close to the true solution is obtained [with L_Q value of approximately 10^{-5} , see curve 4 in Fig. 8(b)]. This difference in convergence rate is a result of boundary information having an impact on the loss, demanding the matching for pinned and optimized boundary handlers, whereas the floating boundary automatically matches the initial condition and no loss boundary term is needed. The consequence of competing terms in the loss function can be seen in the early oscillations of the full loss in Fig. 8(b).

D. Fluid dynamics applications

An area where solvers for complex differential equations are much required is fluid dynamics [2]. In this case, several outstanding models are hard to tackle due to their nonlinear nature. Examples include Burger's equation and Navier-Stokes equations. We concentrate on the latter and show how one can approach them with the DQC solver. Navier-Stokes equations describe a flow of incompressible fluids. This highly nonlinear set of partial differential equations is used to model fluids, magnetoplasma, turbulence etc. It is heavily used in the aerospace industry and weather forecasting. It can be derived from general principles. Namely, we consider fluid motion that obeys Newton's law and we simply track the fluid mass passing through a (infinitesimal) volume.




FIG. 9. Quasi-1d fluid dynamics. (a) We consider an example of fluid dynamics problem corresponding to a convergent-divergent nozzle. The air flows from converging part of the nozzle ($x < 0.5$), passes through the throat placed symmetrically in the middle, and exits to diverging part ($x > 0.5$). (b) System variables [density $\rho(t)$, temperature $T(t)$, and velocity $V(t)$] are shown as functions of time, near the center of the nozzle at $x_0 = 0.4$. For clarity, we label the corresponding curves by 1, 2, and 3. (c) Steady state solutions are plotted as functions of the spatial dimension x , with the same labeling.

The general form Navier-Stokes equation can be presented in the form

$$\frac{\partial(\rho v_x)}{\partial t} + \nabla \cdot (\rho v_x \mathbf{V}) = -\frac{\partial p}{\partial x} + \frac{\partial \tau_{xx}}{\partial x} + \frac{\partial \tau_{yx}}{\partial y} + \frac{\partial \tau_{zx}}{\partial z} + \rho f_x, \tag{37}$$

$$\frac{\partial(\rho v_y)}{\partial t} + \nabla \cdot (\rho v_y \mathbf{V}) = -\frac{\partial p}{\partial y} + \frac{\partial \tau_{xy}}{\partial x} + \frac{\partial \tau_{yy}}{\partial y} + \frac{\partial \tau_{zy}}{\partial z} + \rho f_y, \tag{38}$$

$$\frac{\partial(\rho v_z)}{\partial t} + \nabla \cdot (\rho v_z \mathbf{V}) = -\frac{\partial p}{\partial z} + \frac{\partial \tau_{xz}}{\partial x} + \frac{\partial \tau_{yz}}{\partial y} + \frac{\partial \tau_{zz}}{\partial z} + \rho f_z, \tag{39}$$

where (v_x, v_y, v_z) are instantaneous velocities in the (x, y, z) directions, τ is the stress tensor, f the body force per unit mass acting on the fluid element, ρ the density, and p the pressure. Here \mathbf{V} is a velocity field.

The building blocks for the Navier-Stokes equations described above are the continuity equation for the density ρ subjected to energy conservation and the momentum conservation rules. Finally, as we use energy conservation, we can rewrite equations in terms of one of the thermodynamic state functions, being temperature, or pressure, or enthalpy etc.

While general by itself, Navier-Stokes equations are usually solved in relevant limiting cases. Specifically, this can correspond to space reduction (2D, quasi-1D, 1D), isotropic/anisotropic media properties, and fluid properties (viscous or inviscid flow). The specific example we choose to start with is the flow through a convergent-divergent nozzle, being a paradigmatic task in the aerospace industry [Fig. 9(a)] [2,142]. The Navier-Stokes equations can be rewritten for the inviscid fluid in quasi-1D approximation. They read [2,142]

$$\frac{\partial \rho}{\partial t} = -\rho \frac{\partial V}{\partial x} - \rho V \frac{\partial(\ln A)}{\partial x} - V \frac{\partial \rho}{\partial x}, \tag{40}$$

$$\frac{\partial T}{\partial t} = -V \frac{\partial T}{\partial x} - (\gamma - 1)T \left(\frac{\partial V}{\partial x} + V \frac{\partial(\ln A)}{\partial x} \right), \tag{41}$$

$$\frac{\partial V}{\partial t} = -V \frac{\partial V}{\partial x} - \frac{1}{\gamma} \left(\frac{\partial T}{\partial x} + \frac{T}{\rho} \frac{\partial \rho}{\partial x} \right), \tag{42}$$

where Eq. (40) corresponds to the continuity equation, Eq. (41) describes the energy conservation, and Eq. (42) stems from momentum conservation. $A(x)$ corresponds to the spatial shape of the nozzle, and is a function of the lateral coordinate x . γ describes the ratio of specific heat capacities, and is equal to 1.4 for the relevant case of air flow. Here we used nondimensional variables [2].

We set the problem with nozzle shape

$$A(x) = 1 + 4.95(2x - 1)^2, \quad 0 \leq x \leq 1, \tag{43}$$

and for simplicity specify boundary conditions as

$$\rho(x = 0) = 1, \quad T(x = 0) = 1, \quad V(x = 0) = 0.1, \tag{44}$$

when solving the initial value problem for the steady-state flow. We also need to specify the initial conditions if solving the dynamical problem. These are chosen as [2]

$$\rho(x, t = 0) = 1 - 0.944x, \tag{45}$$

$$T(x, t = 0) = 1 - 0.694x, \tag{46}$$

$$V(x, t = 0) = (0.1 + 3.27x)T(x, t = 0)^{1/2}. \tag{47}$$

Let us consider the stationary state problem represented by Eqs. (40)–(42) with conditions above and equate the time derivatives to zero. Interestingly, when trying to solve the system for the steady state solution using various classical methods as, for example, implemented in MATHEMATICA’s NDSOLVE, the calculations do not converge. This proves to be challenging as the system is stiff. Solutions may become unstable depending on initial values, and specifically the input velocity. To understand the problem, it is instructive to rewrite the system of stationary Navier-Stokes equations in the form

$$\frac{d\rho}{dx} = \frac{\rho V^2 d_x(\ln A)}{T - V^2}, \tag{48}$$

$$\frac{dT}{dx} = \frac{TV^2(\gamma - 1)d_x(\ln A)}{T - V^2}, \tag{49}$$

$$\frac{dV}{dx} = -\frac{TV d_x(\ln A)}{T - V^2}. \tag{50}$$




FIG. 10. DQC-based solution for the Navier-Stokes convergent-divergent nozzle problem. (a) Intermediate solution in training where dashed curves correspond to solutions drawn from trained DQC and solid curves to the true solution. Circuits are trained over twenty points in range $(0, 0.4)$. (b) Final solution. Solid curves with labels 1, 2, and 3 show the true solutions for the density, temperature, and velocity, respectively. Dashed curves with labels 4, 5, and 6 show the DQC solutions for the density, temperature, and velocity, respectively. The DQC solution matches the known solution and also what would be physically expected. As the air goes through the nozzle it accelerates and cools down. (c) Full loss (L_F , curve 7), quality of solution (L_Q , curve 8), and differential loss (L_D , curve 9) are shown as functions of the iteration number n_j for the training resulting in solution (b).

Immediately we observe that each function at the RHS diverges at the point x s.t. $T(x) = V(x)^2$. This leads to singular behavior and breaks classical solvers, including the ones with stiff handling. At the same time, full dynamical solution and its $t \rightarrow \infty$ extrapolation are possible, shown in Figs. 9(b) and 9(c) once the *well-suited* initial conditions (45)–(47) are chosen allowing to avoid instability, as usually done in computational fluid dynamics [2,142].

E. DQC solution

We proceed to solve the stationary system of Navier-Stokes equations for the convergent-divergent nozzle by constructing optimized DQC. We consider the case of subsonic-supersonic transition, where flow velocity increases after the center of the nozzle and the qualitative behavior of other state variables (temperature and density drop) is known. However, getting the quantitative results is difficult. We show that derivative quantum circuits can find solutions despite the challenge for the classical solution for the continuous grid.

To construct the solution we employ a two-stage optimization approach, where the solution is first obtained at smaller x , and later generalized until the end of the nozzle. For the circuit we again consider a six-qubit quantum register with a Chebyshev quantum feature map, keeping in mind that $x \in [0, 1)$. We choose the cost functions as a total magnetization $\hat{C} = \sum_{j=1}^N \hat{Z}_j$. As we consider three functions $\{\rho(x), T(x), V(x)\}$, three equations contribute to the loss function in the combined manner. We use *floating boundary* handling, where each curve is adjusted according to starting values, chosen as $\rho(0) = 1$, $T(0) = 1$, and $V(0) = 0.1$. The variational ansatz is taken in the standard hardware efficient form with $d = 6$ depth. ADAM is used as a classical optimizer, and the learning rate is set to 0.01.

At the first stage we train DQC in the region $(x_{\min}, x_{\max}) = (0, 0.4)$, such that the circuit represents the region close to initial point $x = 0$. As expected, in the subsonic region flow velocity grows towards the middle of the nozzle, while temperature and density drop slowly [see Fig. 10(a), dashed curves 4, 5, 6]. The training is set for 20 equally distributed

points of x , with no prior regularization and using floating boundary. We optimize DQC for $n_{\text{iter}} = 200$ iterations for Adam with the learning rate of $\alpha = 0.01$. We find a high-quality solution based on the gradient information in the imposed solution region $x < 0.4$ shown in Fig. 10(a).

We proceed to search for the full solution that includes the divergent nozzle part for $x > 0.5$ at the second training stage. At this session we choose the grid of 40 points, where two regions of $(0, 0.4)$ and $(0.6, 0.9)$ with 20 points each are used. As we discussed before, the key problem of the convergent-divergent nozzle in the subsonic-supersonic transition case is the divergence around the middle of the nozzle. This causes a major problem to classical solvers, that are unable to find a steady state solution directly. Divergent contributions from this region also impact the loss function, and makes training complicated. However, by excluding the region around the nozzle throat, $(0.4, 0.6)$, in the training we alleviate this problem. Proceeding with the use of the same ansatz and boundary handling, we feed the variational angles from stage 1 as initial parameters for stage 2 training. We employ weak regularization to ensure that the required subsonic-supersonic solution is made. Namely, we use 20 points in the $(0, 0.4)$ region benefiting from the previously found solution (in general we have access to as many points as we need), and also add 5 points in the $x \in (0.6, 0.9)$ region representing weak bias towards supersonic solution type. The training is performed for $n_{\text{iter}} = 600$, learning rate of $\alpha = 0.005$, and regularization switch-off function $\zeta(n_j)$ set to be removed smoothly around $n_j = 150$. The full solution is shown in Fig. 10(b). It converges to the expected long-time behavior for the system, where function derivatives from DQC match the nonlinear contributions. The increase of speed for the air flow in the convergent part and decrease of temperature and density is reproduced quantitatively. The details of the optimization run can be inferred from the loss plotted in Fig. 10(c). This shows a distinct region in the presence of regularization ($n_j < 150$), where quality of solution L_Q remains low [Fig. 10(c), curve 8], while the full loss L_F improves [Fig. 10(c), curve 7] thanks to regularization contribution and weakly improved derivative contribution L_D [Fig. 10(c), curve 9]. For the region with

switched off regularization we observe steady improvement of all metrics, showing that DQCs are efficiently trained to approach true solution, also evidenced by L_Q decrease. Notably, as compared to many methods relying on sparse discretization of x , we have found the solution along the full nozzle length.

V. SUMMARY AND OUTLOOK

We presented a general framework for solving general (systems of) nonlinear differential equations using differentiable quantum circuits on gate-based quantum hardware. The method makes use of quantum feature map circuits to encode function values into a latent space. This allows us to consider spectral decompositions of the trial solutions to differential equations. We showed how our method can accurately represent nonlinear solutions, using the high-dimensional Hilbert space of a qubit register. For this we exploit as an example a large spectral basis set of Chebyshev or Fourier polynomials. We also showed how analytical circuit differentiation can be used to represent function derivatives appearing in differential equations, and constructed loss functions which aim to improve the prepared trial solution. This method opens up a new way of solving general, complex, (non)linear systems; as an example we presented solutions to the Navier-Stokes equations, but the same method can be applied across all disciplines where differential equations arise. Although we showed how the method works explicitly for a spectral (global) treatment, the scheme is general and one may instead consider a *digital quantum feature map* which allows a grid-based (local) treatment. We did not yet consider the impact of noise on the algorithm performance; however, it is well-known that the analytical gradients, which are used here, are much less susceptible to noise than numerical gradients evaluated on variational quantum circuits. We expect recent developments in noisy black-box optimization algorithms for NISQ application would perform well in this method, as well as error mitigation techniques. The presented quantum feature maps are a good start, but could be improved upon further, where the goal is to find efficient representations of functions. This choice will likely be problem-specific and in some cases problem-motivated, allowing to make optimal use of (quantum) resources. A good choice of variational ansatz is crucial to loss function convergence success and speed. We presented some examples, but an active area of research is to improve upon these using for example adaptively growing circuits or stronger interplay with intelligent classical optimization protocols, such as Bayesian optimization.

An important question to discuss is the potential of achieving a quantum advantage, specifically for solving differential equations with a hybrid workflow. Similarly to classical ML,

variational quantum machine learning uses quantum circuits as universal function approximators [173]. While VQA inherently remains a heuristic, similarly to machine learning-type approaches, they are highly suitable to problems where traditional methods fail or are inefficient. For instance, classical neural networks based on Fourier representation were recently shown to outperform traditional PDE solvers by three orders of magnitude [147]. This can be pushed further by quantum variational protocols. For expressive feature maps, the representation power of quantum circuits grows superpolynomially with the number of qubits. In classification tasks this leads to increased accuracy even for small networks [174]. For differential equations, rich basis sets allow for accurate solution representation in a memory-frugal form—a saved list of optimal circuit parameters for a given DQC structure. Having the expressibility advantage, the success of operation then depends on (1) trainability of quantum circuits and (2) absolute operation rates. The former is well-known to suffer from a phenomenon commonly referred to as barren plateaus. That can be mitigated by a carefully chosen circuit structure, as recently shown for convolutional quantum circuits [115,175]. The latter requires quantum hardware specifications with high sampling rates and a low-latency hybrid operational mode. In the pre-FTQC era with large-scale quantum devices, yet limited by noise, parallel operation may offer this possibility.

Finally, differential quantum circuits can enjoy another type of computational advantage related to training. In this paper we considered capturing the function similarity through a loss function based on a cost Hamiltonian suitable for near-term operation. Anticipating the appearance of quantum devices with deep quantum circuits and problems with multidimensional functions, we consider possible strategies for parallel DQC training at different nodes of the training grid. One possibility is training of differential equation solutions as a superposition of quantum circuits (feature maps applied for different values of variables). The DE is then solved by matching differentiated version at all points (using linear combinations of unitaries), and loss can be measured from quantum kernel estimation. This was used recently to speed up generative modeling with quantum batch training [124]. We will further investigate that approach in future work. To conclude, we suggest that the DQC method could be a potential contender for exploring the boundaries of quantum advantage in industrially relevant problems, and expect many improvements can be made building on the presented idea.

ACKNOWLEDGMENT

We thank Benno Broer for useful discussions on the subject and reading the manuscript.

-
- [1] G. F. Simmons, *Differential Equations with Applications and Historical Notes* (McGraw-Hill, Singapore, 1972).
 [2] J. D. Anderson, Jr., *Computational Fluid Dynamics: The Basics with Applications* (McGraw-Hill, Singapore, 1995).

- [3] J. C. Butcher, *Numerical Methods for Ordinary Differential Equations* (Wiley, Chichester, UK, 2016).
 [4] J. P. Boyd, *Chebyshev and Fourier Spectral Methods*, 2nd ed. (Dover, Mineola, NY, 2001).

- [5] E. Hairer and G. Wanner, *Solving Ordinary Differential Equations II: Stiff and differential-Algebraic Problems*, 2nd ed. (Springer-Verlag, Berlin, 1996).
- [6] C. W. Gear, Numerical solution of ordinary differential equations: Is there anything left to do?, *SIAM Rev.* **23**, 10 (1981).
- [7] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, New York, 2010).
- [8] A. Montanaro, Quantum algorithms: An overview, *npj Quant. Inf.* **2**, 15023 (2016).
- [9] A. W. Harrow, A. Hassidim, and S. Lloyd, Quantum Algorithm for Linear Systems of Equations, *Phys. Rev. Lett.* **103**, 150502 (2009).
- [10] A. M. Childs, R. Kothari, and R. D. Somma, Quantum algorithm for systems of linear equations with exponentially improved dependence on precision, *SIAM J. Comput.* **46**, 1920 (2017).
- [11] L. Wossnig, Z. Zhao, and A. Prakash, Quantum Linear System Algorithm for Dense Matrices, *Phys. Rev. Lett.* **120**, 050502 (2018).
- [12] Y. Subasi, R. D. Somma, and D. Orsucci, Quantum Algorithms for Systems of Linear Equations Inspired by Adiabatic Quantum Computing, *Phys. Rev. Lett.* **122**, 060504 (2019).
- [13] L. Lin and Y. Tong, Optimal quantum eigenstate filtering with application to solving quantum linear systems, *Quantum* **4**, 361 (2020).
- [14] A. C. Vazquez, R. Hiptmair, and S. Woerner, Enhancing the quantum linear systems algorithm using richardson extrapolation, [arXiv:2009.04484](https://arxiv.org/abs/2009.04484).
- [15] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, Quantum machine learning, *Nature (London)* **549**, 195 (2017).
- [16] M. Schuld, I. Sinayskiy, and F. Petruccione, An introduction to quantum machine learning, *Contemp. Phys.* **56**, 172 (2015).
- [17] S. Lloyd, M. Mohseni, and P. Rebentrost, Quantum principal component analysis, *Nat. Phys.* **10**, 631 (2014).
- [18] P. Rebentrost, M. Mohseni, and S. Lloyd, Quantum Support Vector Machine for Big Data Classification, *Phys. Rev. Lett.* **113**, 130503 (2014).
- [19] S. Lloyd, S. Garnerone, and P. Zanardi, Quantum algorithms for topological and geometric analysis of data, *Nat. Commun.* **7**, 10138 (2016).
- [20] V. Dunjko, J. M. Taylor, and H. J. Briegel, Quantum-Enhanced Machine Learning, *Phys. Rev. Lett.* **117**, 130501 (2016).
- [21] M. Schuld, I. Sinayskiy, and F. Petruccione, Prediction by linear regression on a quantum computer, *Phys. Rev. A* **94**, 022342 (2016).
- [22] N. Liu and P. Rebentrost, Quantum machine learning for quantum anomaly detection, *Phys. Rev. A* **97**, 042315 (2018).
- [23] P. Rebentrost, M. Schuld, L. Wossnig, F. Petruccione, and S. Lloyd, Quantum gradient descent and Newton's method for constrained polynomial optimization, *New J. Phys.* **21**, 073023 (2019).
- [24] C. Blank, D. K. Park, J.-K. K. Rhee, and F. Petruccione, Quantum classifier with tailored quantum kernel, *npj Quantum Inf.* **6**, 41 (2020).
- [25] S. K. Leyton and T. J. Osborne, A quantum algorithm to solve nonlinear differential equations, [arXiv:0812.4423](https://arxiv.org/abs/0812.4423).
- [26] Yudong Cao, A. Daskin, S. Frankel, and S. Kais, Quantum circuit design for solving linear systems of equations, *Mol. Phys.* **110**, 1675 (2012).
- [27] D. W. Berry, High-order quantum algorithm for solving linear differential equations, *J. Phys. A: Math. Theor.* **47**, 105301 (2014).
- [28] A. Montanaro and S. Pallister, Quantum algorithms and the finite element method, *Phys. Rev. A* **93**, 032324 (2016).
- [29] D. W. Berry, A. M. Childs, A. Ostrander, and G. Wang, Quantum algorithm for linear differential equations with exponentially improved dependence on precision, *Commun. Math. Phys.* **356**, 1057 (2017).
- [30] P. C. S. Costa, S. Jordan, and A. Ostrander, Quantum algorithm for simulating the wave equation, *Phys. Rev. A* **99**, 012323 (2019).
- [31] A. M. Childs, J.-P. Liu, and A. Ostrander, High-precision quantum algorithms for partial differential equations, [arXiv:2002.07868](https://arxiv.org/abs/2002.07868).
- [32] S. Wang, Z. Wang, W. Li, L. Fan, G. Cui, Z. Wei, and Y. Gu, A quantum Poisson solver implementable on NISQ devices, [arXiv:2005.00256](https://arxiv.org/abs/2005.00256).
- [33] N. Linden, A. Montanaro, and C. Shao, Quantum vs. classical algorithms for solving the heat equation, [arXiv:2004.06516](https://arxiv.org/abs/2004.06516).
- [34] A. M. Childs and J.-P. Liu, Quantum spectral methods for differential equations, *Commun. Math. Phys.* **375**, 1427 (2020).
- [35] V. Giovannetti, S. Lloyd, and L. Maccone, Quantum Random Access Memory, *Phys. Rev. Lett.* **100**, 160501 (2008).
- [36] A. Scherer, B. Valiron, S.-C. Mau, S. Alexander, E. van den Berg, and T. E. Chapuran, Concrete resource analysis of the quantum linear system algorithm used to compute the electromagnetic scattering cross section of a 2D target, *Quantum Inf. Process.* **16**, 60 (2017).
- [37] J. Preskill, Quantum computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).
- [38] F. Arute *et al.*, Quantum supremacy using a programmable superconducting processor, *Nature (London)* **574**, 505 (2019).
- [39] S. McArdle, S. Endo, A. Aspuru-Guzik, S. C. Benjamin, and X. Yuan, Quantum computational chemistry, *Rev. Mod. Phys.* **92**, 015003 (2020).
- [40] Y. Cao, J. Romero, J. P. Olson, M. Degroote, P. D. Johnson, M. Kieferová, I. D. Kivlichan, T. Menke, B. Peropadre, N. P. D. Sawaya, S. Sim, L. Veis, and A. Aspuru-Guzik, Quantum chemistry in the age of quantum computing, *Chem. Rev.* **119**, 10856 (2019).
- [41] V. E. Elfving, B. W. Broer, M. Webber, J. Gavartin, M. D. Halls, K. P. Lorton, and A. Bochevarov, How will quantum computers provide an industrially relevant computational advantage in quantum chemistry?, [arXiv:2009.12472](https://arxiv.org/abs/2009.12472).
- [42] 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, *Nat. Commun.* **5**, 4213 (2014).
- [43] J. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, The theory of variational hybrid quantum-classical algorithms, *New J. Phys.* **18**, 023023 (2016).
- [44] 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).
- [45] 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 (London)* **549**, 242 (2017).
- [46] J. I. Colless, V. V. Ramasesh, D. Dahlen, M. S. Blok, M. E. Kimchi-Schwartz, J. R. McClean, J. Carter, W. A. de Jong, and I. Siddiqi, Computation of Molecular Spectra on a Quantum Processor with an Error-Resilient Algorithm, *Phys. Rev. X* **8**, 011021 (2018).
- [47] 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 the lattice Schwinger model, *Nature (London)* **569**, 355 (2019).
- [48] M. Ganzhorn, D. J. Egger, P. Barkoutsos, P. Ollitrault, G. Salis, N. Moll, M. Roth, A. Fuhrer, P. Mueller, S. Woerner, I. Tavernelli, and S. Filipp, Gate-Efficient Simulation of Molecular Eigenstates on a Quantum Computer, *Phys. Rev. Appl.* **11**, 044092 (2019).
- [49] A. J. McCaskey, Z. P. Parks, J. Jakowski, S. V. Moore, T. Morris, T. S. Humble, and R. C. Pooser, Quantum chemistry as a benchmark for near-term quantum computers, *npj Quantum Inf.* **5**, 99 (2019).
- [50] Y. Nam, J.-S. Chen, N. C. Panti, K. Wright, C. Delaney, D. Maslov, K. R. Brown, S. Allen, J. M. Amini, J. Apisdorf, K. M. Beck, A. Blinov, V. Chaplin, M. Chmielewski, C. Collins, S. Debnath, A. M. Ducore, K. M. Hudek, M. Keesan, S. M. Kreikemeier *et al.*, Ground-state energy estimation of the water molecule on a trapped ion quantum computer, *npj Quantum Inf.* **6**, 33 (2020).
- [51] Q. Gao, H. Nakamura, T. P. Gujarati, G. O. Jones, J. E. Rice, S. P. Wood, M. Pistoia, J. M. Garcia, and N. Yamamoto, Computational investigations of the lithium superoxide dimer rearrangement on noisy quantum devices, *J. Phys. Chem. A* **125**, 1827 (2021).
- [52] F. Arute *et al.* (Google AI Quantum and Collaborators), Hartree-Fock on a superconducting qubit quantum computer, *Science* **369**, 1084 (2020).
- [53] B. T. Gard, L. Zhu, G. S. Barron, N. J. Mayhall, S. E. Economou, and E. Barnes, Efficient symmetry-preserving state preparation circuits for the variational quantum eigensolver algorithm, *npj Quantum Inf.* **6**, 10 (2020).
- [54] H. R. Grimsley, S. E. Economou, E. Barnes, and N. J. Mayhall, An adaptive variational algorithm for exact molecular simulations on a quantum computer, *Nat. Commun.* **10**, 3007 (2019).
- [55] H. L. Tang, V. O. Shkolnikov, G. S. Barron, H. R. Grimsley, N. J. Mayhall, E. Barnes, and S. E. Economou, qubit-ADAPT-VQE: An adaptive algorithm for constructing hardware-efficient ansatzes on a quantum processor, [arXiv:1911.10205](https://arxiv.org/abs/1911.10205).
- [56] D. Claudino, J. Wright, A. J. McCaskey, and T. S. Humble, Benchmarking adaptive variational quantum eigensolvers, *Front. Chem.* **8**, 606863 (2020).
- [57] W. J. Huggins, J. McClean, N. Rubin, Z. Jiang, N. Wiebe, K. B. Whaley, and R. Babbush, Efficient and noise resilient measurements for quantum chemistry on near-term quantum computers, *npj Quantum Inf.* **7**, 23 (2021).
- [58] T.-C. Yen, V. Verteletskyi, and A. F. Izmaylov, Measuring all compatible operators in one series of single-qubit measurements using unitary transformations, *J. Chem. Theory Comput.* **16**, 2400 (2020).
- [59] J. M. Kübler, A. Arrasmith, L. Cincio, and P. J. Coles, An adaptive optimizer for measurement-frugal variational algorithms, *Quantum* **4**, 263 (2020).
- [60] K. M. Nakanishi, K. Mitarai, and K. Fujii, Subspace-search variational quantum eigensolver for excited states, *Phys. Rev. Res.* **1**, 033062 (2019).
- [61] M. Metcalf, N. P. Bauman, K. Kowalski, and W. A. de Jong, Resource efficient chemistry on quantum computers with the variational quantum eigensolver and the double unitary coupled-cluster approach, *J. Chem. Theory Comput.* **16**, 6165 (2020).
- [62] N. P. Bauman, J. Chládek, L. Veis, J. Pittner, and K. Kowalski, Variational quantum eigensolver for approximate diagonalization of downfolded Hamiltonians using generalized unitary coupled cluster ansatz, *quantum*, [arXiv:2011.01985](https://arxiv.org/abs/2011.01985).
- [63] T. E. O'Brien, B. Senjean, R. Sagastizabal, X. Bonet-Monroig, A. Dutkiewicz, F. Buda, L. DiCarlo, and L. Visscher, Calculating energy derivatives for quantum chemistry on a quantum computer, *npj Quantum Inf.* **5**, 113 (2019).
- [64] R. M. Parrish, E. G. Hohenstein, P. L. McMahon, and T. J. Martinez, Quantum Computation of Electronic Transitions Using a Variational Quantum Eigensolver, *Phys. Rev. Lett.* **122**, 230401 (2019).
- [65] S. Yalouz, B. Senjean, J. Günther, F. Buda, T. E. O'Brien, and L. Visscher, A state-averaged orbital-optimized hybrid quantum-classical algorithm for a democratic description of ground and excited states, *Quantum Sci. Technol.* **6**, 024004 (2021).
- [66] V. E. Elfving, J. A. Gámez, and C. Gogolin, Simulating quantum chemistry in the restricted Hartree-Fock space on a qubit-based quantum computing device, *Phys. Rev. A* **103**, 032605 (2021).
- [67] M. Benedetti, E. Lloyd, S. Sack, and M. Fiorentini, Parameterized quantum circuits as machine learning models, *Quantum Sci. Technol.* **4**, 043001 (2019).
- [68] A. Perdomo-Ortiz, M. Benedetti, J. Realpe-Gomez, and R. Biswas, Opportunities and challenges for quantum-assisted machine learning in near-term quantum computers, *Quantum Sci. Technol.* **3**, 030502 (2018).
- [69] W. Huggins, P. Patel, K. B. Whaley, and E. M. Stoudenmire, Towards quantum machine learning with tensor networks, *Quantum Sci. Technol.* **4**, 024001 (2019).
- [70] M. Cerezo, A. Sone, T. Volkoff, L. Cincio, and P. J. Coles, Cost-function-dependent barren plateaus in shallow quantum neural networks, *Nat. Commun.* **12**, 1791 (2021).
- [71] L. Lamata, Quantum machine learning and quantum biomimetics: A perspective, *Mach. Learn.: Sci. Technol.* **1**, 033002 (2020).
- [72] E. Farhi and H. Neven, Classification with quantum neural networks on near term processors, [arXiv:1802.06002](https://arxiv.org/abs/1802.06002).
- [73] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii, Quantum circuit learning, *Phys. Rev. A* **98**, 032309 (2018).
- [74] V. Havlíček, A. D. Corcoles, K. Temme, A. W. Harrow, A. Kandala, J. M. Chow, and J. M. Gambetta, Supervised learning with quantum-enhanced feature spaces, *Nature (London)* **567**, 209 (2019).

- [75] M. Schuld, A. Bocharov, K. M. Svore, and N. Wiebe, Circuit-centric quantum classifiers, *Phys. Rev. A* **101**, 032308 (2020).
- [76] Y. Suzuki, H. Yano, Q. Gao, S. Uno, T. Tanaka, M. Akiyama, and N. Yamamoto, Analysis and synthesis of feature map for kernel-based quantum classifier, *Quantum Mach. Intell.* **2**, 1 (2020).
- [77] H. Yano, Y. Suzuki, R. Raymond, and N. Yamamoto, Efficient discrete feature encoding for variational quantum classifier, [arXiv:2005.14382](https://arxiv.org/abs/2005.14382).
- [78] Y. Du, M.-H. Hsieh, T. Liu, D. Tao, and N. Liu, Quantum noise protects quantum classifiers against adversaries, [arXiv:2003.09416](https://arxiv.org/abs/2003.09416).
- [79] J.-G. Liu and L. Wang, Differentiable learning of quantum circuit born machine, *Phys. Rev. A* **98**, 062324 (2018).
- [80] S. Lloyd and C. Weedbrook, Quantum Generative Adversarial Learning, *Phys. Rev. Lett.* **121**, 040502 (2018).
- [81] J. Zeng, Y. Wu, J.-G. Liu, L. Wang, and J. Hu, Learning and inference on generative adversarial quantum circuits, *Phys. Rev. A* **99**, 052306 (2019).
- [82] G. Verdon, M. Broughton, and J. Biamonte, A quantum algorithm to train neural networks using low-depth circuits, [arXiv:1712.05304](https://arxiv.org/abs/1712.05304).
- [83] M. Benedetti, E. Grant, L. Wossnig, and S. Severini, Adversarial quantum circuit learning for pure state approximation, *New J. Phys.* **21**, 043023 (2019).
- [84] M. Benedetti, D. Garcia-Pintos, O. Perdomo, V. Leyton-Ortega, Y. Nam, and A. Perdomo-Ortiz, A generative modeling approach for benchmarking and training shallow quantum circuits, *npj Quantum Inf.* **5**, 45 (2019).
- [85] C. Zoufal, A. Lucchi, and S. Woerner, Quantum generative adversarial networks for learning and loading random distributions, *npj Quantum Inf.* **5**, 103 (2019).
- [86] J. Romero and A. Aspuru-Guzik, Variational quantum generators: Generative adversarial quantum machine learning for continuous distributions, *Adv. Quantum Technol.* **4**, 2000003 (2021).
- [87] B. Coyle, D. Mills, V. Danos, and E. Kashefi, The Born supremacy: quantum advantage and training of an Ising Born machine, *npj Quantum Inf.* **6**, 60 (2020).
- [88] J. S. Otterbach *et al.*, Unsupervised machine learning on a hybrid quantum computer, [arXiv:1712.05771](https://arxiv.org/abs/1712.05771).
- [89] S. S. Mendelson, R. W. Strand, G. B. Oldaker IV, and J. M. Farinholt, Quantum-assisted clustering algorithms for NISQ-era devices, [arXiv:1904.08992](https://arxiv.org/abs/1904.08992).
- [90] V. Dunjko and H. J. Briegel, Machine learning & artificial intelligence in the quantum domain: A review of recent progress, *Rep. Prog. Phys.* **81**, 074001 (2018).
- [91] 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, *Proc. Natl. Acad. Sci. USA* **115**, 1221 (2018).
- [92] F. Albarran-Arriagada, J. C. Retamal, E. Solano, and L. Lamata, Measurement-based adaptation protocol with quantum reinforcement learning, *Phys. Rev. A* **98**, 042315 (2018).
- [93] A. A. Melnikov, P. Sekatski, and N. Sangouard, Setting Up Experimental Bell Tests with Reinforcement Learning, *Phys. Rev. Lett.* **125**, 160401 (2020).
- [94] J. Walln fer, A. A. Melnikov, W. D r, and H. J. Briegel, Machine learning for long-distance quantum communication, *PRX Quantum* **1**, 010301 (2020).
- [95] J. Olivares-Sanchez, J. Casanova, E. Solano, and L. Lamata, Measurement-based adaptation protocol with quantum reinforcement learning in a rigetti quantum computer, *Quantum Rep.* **2**, 293 (2020).
- [96] J. M. Kreula, S. R. Clark, and D. Jaksch, Non-linear quantum-classical scheme to simulate non-equilibrium strongly correlated fermionic many-body dynamics, *Sci. Rep.* **6**, 32940 (2016).
- [97] C. Cirstoiu, Z. Holmes, J. Iosue, L. Cincio, P. J. Coles, and A. Sornborger, Variational fast forwarding for quantum simulation beyond the coherence time, *npj Quantum Inf.* **6**, 82 (2020).
- [98] S. Endo, J. Sun, Y. Li, S. C. Benjamin, and X. Yuan, Variational Quantum Simulation of General Processes, *Phys. Rev. Lett.* **125**, 010501 (2020).
- [99] D. Paulson, L. Dellantonio, J. F. Haase, A. Celi, A. Kan, A. Jena, C. Kokail, R. van Bijnen, K. Jansen, P. Zoller, and C. A. Muschik, Towards simulating 2D effects in lattice gauge theories on a quantum computer, [arXiv:2008.09252](https://arxiv.org/abs/2008.09252).
- [100] J. J. Meyer, J. Borregaard, and J. Eisert, A variational toolbox for quantum multi-parameter estimation, [arXiv:2006.06303](https://arxiv.org/abs/2006.06303).
- [101] Y. Ouyang, N. Shettell, and D. Markham, Robust quantum metrology with explicit symmetric states, [arXiv:1908.02378](https://arxiv.org/abs/1908.02378).
- [102] B. Koczor, S. Endo, T. Jones, Y. Matsuzaki, and S. C. Benjamin, Variational-state quantum metrology, *New J. Phys.* **22**, 083038 (2020).
- [103] M. Schuld, V. Bergholm, C. Gogolin, J. Izaac, and N. Killoran, Evaluating analytic gradients on quantum hardware, *Phys. Rev. A* **99**, 032331 (2019).
- [104] K. Mitarai, Y. O. Nakagawa, and W. Mizukami, Theory of analytical energy derivatives for the variational quantum eigensolver, *Phys. Rev. Res.* **2**, 013129 (2020).
- [105] R. M. Parrish, E. G. Hohenstein, P. L. McMahon, and T. J. Martinez, Hybrid quantum/classical derivative theory: analytical gradients and excited-state dynamics for the multistate contracted variational quantum eigensolver, [arXiv:1906.08728](https://arxiv.org/abs/1906.08728).
- [106] S. McArdle, T. Jones, S. Endo, Y. Li, S. Benjamin, and X. Yuan, Variational ansatz-based quantum simulation of imaginary time evolution, *npj Quantum Inf.* **5**, 75 (2019).
- [107] J. Stokes, J. Izaac, N. Killoran, and G. Carleo, Quantum natural gradient, *Quantum* **4**, 269 (2020).
- [108] N. Yamamoto, On the natural gradient for variational quantum eigensolver, [arXiv:1909.05074](https://arxiv.org/abs/1909.05074).
- [109] B. Koczor and S. C. Benjamin, Quantum natural gradient generalised to non-unitary circuits, [arXiv:1912.08660](https://arxiv.org/abs/1912.08660).
- [110] R. Sweke, F. Wilde, J. Meyer, M. Schuld, P. K. F hrmann, B. Meynard-Piganeau, and J. Eisert, Stochastic gradient descent for hybrid quantum-classical optimization, *Quantum* **4**, 314 (2020).
- [111] B. Koczor and S. C. Benjamin, Quantum analytic descent, [arXiv:2008.13774](https://arxiv.org/abs/2008.13774).
- [112] G. Verdon, J. Pye, and M. Broughton, A universal training algorithm for quantum deep learning, [arXiv:1806.09729](https://arxiv.org/abs/1806.09729).
- [113] K. Beer, D. Bondarenko, T. Farrelly, T. J. Osborne, R. Salzmann, D. Scheiermann, and R. Wolf, Training deep quantum neural networks, *Nat. Commun.* **11**, 808 (2020).
- [114] S. Garg and G. Ramakrishnan, Advances in quantum deep learning: An overview, [arXiv:2005.04316](https://arxiv.org/abs/2005.04316).
- [115] Iris Cong, S. Choi, and M. D. Lukin, Quantum convolutional neural networks, *Nat. Phys.* **15**, 1273 (2019).

- [116] M. Cerezo and P. J. Coles, Impact of barren plateaus on the hessian and higher order derivatives, [arXiv:2008.07454](#).
- [117] Y. Herasymenko and T. E. O'Brien, A diagrammatic approach to variational quantum ansatz construction, [arXiv:1907.08157](#).
- [118] A. Skolik, J. R. McClean, M. Mohseni, P. van der Smagt, and M. Leib, Layerwise learning for quantum neural networks, *Quantum Mach. Intell.* **3**, 5 (2021).
- [119] S. Sim, J. Romero, J. F. Gonthier, and A. A. Kunitsa, Adaptive pruning-based optimization of parameterized quantum circuits, *Quantum Sci. Technol.* **6**, 025019 (2021).
- [120] Y. Du, T. Huang, S. You, M.-H. Hsieh, and D. Tao, Quantum circuit architecture search: error mitigation and trainability enhancement for variational quantum solvers, [arXiv:2010.10217](#).
- [121] Shi-Xin Zhang, Chang-Yu Hsieh, S. Zhang, and H. Yao, Differentiable quantum architecture search, [arXiv:2010.08561](#).
- [122] Kaining Zhang, Min-Hsiu Hsieh, L. Liu, and D. Tao, Toward trainability of quantum neural networks, [arXiv:2007.11921](#).
- [123] 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, *Sci. Adv.* **5**, eaaw9918 (2019).
- [124] He-Liang Huang, Y. Du, M. Gong, Y. Zhao, Y. Wu, C. Wang, S. Li, F. Liang, J. Lin, Y. Xu, R. Yang, T. Liu, Min-Hsiu Hsieh, H. Deng, H. Rong, Cheng-Zhi Peng, Chao-Yang Lu, Yu-Ao Chen, D. Tao, X. Zhu, and Jian-Wei Pan, Experimental quantum generative adversarial networks for image generation, [arXiv:2010.06201](#).
- [125] Ying Li and S. C. Benjamin, Efficient Variational Quantum Simulator Incorporating Active Error Minimization, *Phys. Rev. X* **7**, 021050 (2017).
- [126] K. Temme, S. Bravyi, and J. M. Gambetta, Error Mitigation for Short-Depth Quantum Circuits, *Phys. Rev. Lett.* **119**, 180509 (2017).
- [127] S. Endo, S. C. Benjamin, and Y. Li, Practical Quantum Error Mitigation for Near-Future Applications, *Phys. Rev. X* **8**, 031027 (2018).
- [128] X. Bonet-Monroig, R. Sagastizabal, M. Singh, and T. O'Brien, Low-cost error mitigation by symmetry verification, *Phys. Rev. A* **98**, 062339 (2018).
- [129] S. McArdle, X. Yuan, and S. Benjamin, Error-Mitigated Digital Quantum Simulation, *Phys. Rev. Lett.* **122**, 180501 (2019).
- [130] M. Otten and S. K. Gray, Accounting for errors in quantum algorithms via individual error reduction, *npj Quantum Inf.* **5**, 11 (2019).
- [131] D. Bondarenko and P. Feldmann, Quantum Autoencoders to Denoise Quantum Data, *Phys. Rev. Lett.* **124**, 130502 (2020).
- [132] T. Giurgica-Tiron, Y. Hindy, R. LaRose, A. Mari, and W. J. Zeng, Digital zero noise extrapolation for quantum error mitigation, [arXiv:2005.10921](#).
- [133] Hsin-Yuan Huang, K. Bharti, and P. Rebentrost, Near-term quantum algorithms for linear systems of equations, [arXiv:1909.07344](#).
- [134] C. Bravo-Prieto, R. LaRose, M. Cerezo, Y. Subasi, L. Cincio, and P. J. Coles, Variational quantum linear solver, [arXiv:1909.05820](#).
- [135] X. Xu, J. Sun, S. Endo, Y. Li, S. C. Benjamin, and X. Yuan, Variational algorithms for linear algebra, [arXiv:1909.03898](#).
- [136] Chih-Chieh Chen, Shiue-Yuan Shiau, Ming-Feng Wu, and Yuh-Renn Wu, Hybrid classical-quantum linear solver using noisy intermediate-scale quantum machines, *Sci. Rep.* **9**, 16251 (2019).
- [137] M. R. Perelshtein, A. I. Pakhomchik, A. A. Melnikov, A. A. Novikov, A. Glatz, G. S. Paraoanu, V. M. Vinokur, G. B. Lesovik, Advanced quantum supremacy using a hybrid algorithm for linear systems of equations, [arXiv:2003.12770](#).
- [138] A. Borle, V. E. Elfving, and S. J. Lomonaco, Quantum approximate optimization for hard problems in linear algebra, [arXiv:2006.15438](#).
- [139] J. Pan, Y. Cao, X. Yao, Z. Li, C. Ju, H. Chen, X. Peng, S. Kais, and J. Du, Experimental realization of quantum algorithm for solving linear systems of equations, *Phys. Rev. A* **89**, 022313 (2014).
- [140] T. Xin, S. Wei, J. Cui, J. Xiao, I. Arrazola, L. Lamata, X. Kong, D. Lu, E. Solano, and G. Long, Quantum algorithm for solving linear differential equations: Theory and experiment, *Phys. Rev. A* **101**, 032307 (2020).
- [141] M. Lubasch, J. Joo, P. Moinier, M. Kiffner, and D. Jaksch, Variational quantum algorithms for nonlinear problems, *Phys. Rev. A* **101**, 010301(R) (2020).
- [142] F. Gaitan, Finding flows of a Navier-Stokes fluid through quantum computing, *npj Quantum Inf.* **6**, 61 (2020).
- [143] J. Sirignano and K. Spiliopoulos, DGM: A deep learning algorithm for solving partial differential equations, *J. Comp. Phys.* **375**, 1339 (2018).
- [144] I. E. Lagaris, A. Likas, and D. I. Fotiadis, Artificial neural networks for solving ordinary and partial differential equations, *IEEE Trans. Neural. Netw. Learn. Syst.* **9**, 987 (1998).
- [145] M. Raissi, P. Perdikaris, and G. E. Karniadakis, Physics informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations, *J. Comput. Phys.* **378**, 686 (2019).
- [146] M. Raissi, A. Yazdani, and G. E. Karniadakis, Hidden fluid mechanics: Learning velocity and pressure fields from flow visualizations, *Science* **367**, 1026 (2020).
- [147] Zongyi Li, N. Kovachki, K. Azizzadenesheli, B. Liu, K. Bhattacharya, A. Stuart, and A. Anandkumar, Fourier neural operator for parametric partial differential equations, [arXiv:2010.08895](#).
- [148] 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:2001.04385v3](#).
- [149] C. Rackauckas, M. Innes, Y. Ma, J. Bettencourt, L. White, and V. Dixit, DiffEqFlux.jl - A julia library for neural differential equations, [arXiv:1902.02376v1](#).
- [150] M. Bachmayr, R. Schneider, and A. Uschmajew, Tensor networks and hierarchical tensors for the solution of high-dimensional partial differential equations, *Found. Comput. Math.* **16**, 1423 (2016).
- [151] J. J. Garcia-Ripoll, Quantum-inspired algorithms for multivariate analysis: From interpolation to partial differential equations, *Quantum* **5**, 431 (2021).
- [152] A patent application for the method described in this manuscript has been submitted by Qu & Co BV with OK, AP and VE as inventors.
- [153] M. Schuld and N. Killoran, Quantum Machine Learning in Feature Hilbert Spaces, *Phys. Rev. Lett.* **122**, 040504 (2019).

- [154] S. Lloyd, M. Schuld, A. Ijaz, J. Izaac, and N. Killoran, Quantum embeddings for machine learning, [arXiv:2001.03622](https://arxiv.org/abs/2001.03622).
- [155] M. Schuld, R. Sweke, and J. J. Meyer, The effect of data encoding on the expressive power of variational quantum machine learning models, *Phys. Rev. A* **103**, 032430 (2021).
- [156] A. Mari, T. R. Bromley, and N. Killoran, Estimating the gradient and higher-order derivatives on quantum hardware, *Phys. Rev. A* **103**, 012405 (2021).
- [157] J. G. Vidal and D. O. Theis, Input redundancy for parameterized quantum circuits, *Front. Phys.* **8**, 297 (2020).
- [158] K. Mitarai and K. Fujii, Methodology for replacing indirect measurements with direct measurements, *Phys. Rev. Res.* **1**, 013006 (2019).
- [159] O. Kyriienko, Quantum inverse iteration algorithm for programmable quantum simulators, *npj Quantum Inf.* **6**, 7 (2020).
- [160] M. Abramowitz and I. A. Stegun, *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*, 10th ed. (National Bureau of Standards, Washington, DC, 1972).
- [161] E. Ya. Remez and V. T. Gavrilyuk, Computer development of certain approaches to the approximate construction of solutions of Chebyshev problems nonlinearly depending on parameters, *Ukr. Mat. Zh.* **12**, 324 (1960).
- [162] M. J. D. Powell, *Approximation Theory and Methods* (Cambridge University Press, New York, 1981).
- [163] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma, Simulating Hamiltonian Dynamics with a Truncated Taylor Series, *Phys. Rev. Lett.* **114**, 090502 (2015).
- [164] K. Nakaji and N. Yamamoto, Expressibility of the alternating layered ansatz for quantum computation, *Quantum* **5**, 434 (2021).
- [165] D. Chivilikhin, A. Samarin, V. Ulyantsev, I. Iorsh, A. R. Oganov, and O. Kyriienko, MoG-VQE: Multiobjective genetic variational quantum eigensolver, [arXiv:2007.04424](https://arxiv.org/abs/2007.04424).
- [166] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization, in *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*, (ICLR, San Diego, 2015).
- [167] A. Neumaier, Solving ill-conditioned and singular linear systems: A tutorial on regularization, *SIAM Rev.* **40**, 636 (1998).
- [168] F. Girosi, M. Jones, and T. Poggio, Regularization theory and neural networks architectures, *Neural Comput.* **7**, 219 (1995).
- [169] X.-Z. Luo, J.-G. Liu, P. Zhang, and L. Wang, Yao.jl: Extensible, efficient framework for quantum algorithm design, *Quantum* **4**, 341 (2020).
- [170] J.-G. Liu, Y.-H. Zhang, Y. Wan, and L. Wang, Variational quantum eigensolver with fewer qubits, *Phys. Rev. Res.* **1**, 023025 (2019).
- [171] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, Barren plateaus in quantum neural network training landscapes, *Nat. Commun.* **9**, 4812 (2018).
- [172] T. A. Bespalova and O. Kyriienko, Hamiltonian operator approximation for energy measurement and ground state preparation, [arXiv:2009.03351](https://arxiv.org/abs/2009.03351).
- [173] T. Goto, Q. H. Tran, and K. Nakajima, Universal approximation property of quantum feature map, [arXiv:2009.00298](https://arxiv.org/abs/2009.00298).
- [174] S. Y.-C. Chen, T.-C. Wei, C. Zhang, H. Yu, and S. Yoo, Quantum convolutional neural networks for high energy physics data analysis, [arXiv:2012.12177](https://arxiv.org/abs/2012.12177).
- [175] A. Pesah, M. Cerezo, S. Wang, T. Volkoff, Andrew T. Sornborger, Patrick J. Coles, Absence of barren plateaus in quantum convolutional neural networks, [arXiv:2011.02966](https://arxiv.org/abs/2011.02966).