RESEARCH ARTICLE

Variational quantum classifiers through the lens of the Hessian

Pinaki Sen¹, Amandeep Singh Bhatia^{2*}, Kamalpreet Singh Bhangu², Ahmed Elbeltagi³

1 Department of Electrical Engineering, National Institute of Technology, Agartala, Tripura, India, **2** Chitkara University Institute of Engineering & Technology, Chitkara University, Rajpura, Punjab, India, **3** Agricultural Engineering Dept., Faculty of Agriculture, Mansoura University, Mansoura, Egypt

* amandeepbhatia.singh@gmail.com

Abstract

In quantum computing, the variational quantum algorithms (VQAs) are well suited for finding optimal combinations of things in specific applications ranging from chemistry all the way to finance. The training of VQAs with gradient descent optimization algorithm has shown a good convergence. At an early stage, the simulation of variational quantum circuits on noisy intermediate-scale quantum (NISQ) devices suffers from noisy outputs. Just like classical deep learning, it also suffers from vanishing gradient problems. It is a realistic goal to study the topology of loss landscape, to visualize the curvature information and trainability of these circuits in the existence of vanishing gradients. In this paper, we calculate the Hessian and visualize the loss landscape of variational quantum classifiers at different points in parameter space. The curvature information of variational quantum classifiers (VQC) is interpreted and the loss function's convergence is shown. It helps us better understand the behavior of variational quantum circuits to tackle optimization problems efficiently. We investigated the variational quantum classifiers via Hessian on quantum computers, starting with a simple 4-bit parity problem to gain insight into the practical behavior of Hessian, then thoroughly analyzed the behavior of Hessian's eigenvalues on training the variational quantum classifier for the Diabetes dataset. Finally, we show how the adaptive Hessian learning rate can influence the convergence while training the variational circuits.

OPEN ACCESS

Citation: Sen P, Bhatia AS, Bhangu KS, Elbeltagi A (2022) Variational quantum classifiers through the lens of the Hessian. PLoS ONE 17(1): e0262346. <https://doi.org/10.1371/journal.pone.0262346>

Editor: Thippa Reddy Gadekallu, Vellore Institute of Technology: VIT University, INDIA

Received: October 13, 2021

Accepted: December 22, 2021

Published: January 20, 2022

Copyright: © 2022 Sen et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: Data will be available on publication of the paper.

Funding: The authors received no specific funding for this work.

Competing interests: The authors have declared that no competing interests exist.

Introduction & motivation

In recent years, the enhancement of machine learning algorithms by noisy intermediate-scale quantum (NISQ) technology and mainly the variational quantum circuits have garnered significant attention among academic and research communities [1]. Researchers have applied the variational quantum algorithms in various applications in the NISQ era, mainly the ones related to quantum artificial intelligence. The variational quantum algorithms (VQAs) have shown great learning capability to counterbalance the errors in the device framework. They are considered to be the greatest hope for the journey toward quantum advantage. The first variational quantum eigensolver (VQE) was proposed as a state ansatz to determine the

ground state energy of physical systems [2]. Since the first VQE introduced, several VQE variants have been proposed with a plethora of alterations for computation of excited states such as orthogonality constrained VQE [3], subspace approach VQE [4, 5], adiabatically assisted VQE [6] and multistate contracted VQE [7]. These variational quantum circuits are constructed with an ansatz whose parameters are trained with several optimization methods to minimize the cost. For combinatorial optimization tasks, the quantum approximate optimization algorithm (QAOA) was proposed originally to attain approximate solutions [8]. These architectures of variational quantum circuits have shown to be computationally universal [9, 10].

Recently, the use of VQAs has got an incredible response and is widely applied in quantum machine learning applications. VQAs consist of a small number of qubits and quantum circuits, which make them resistant to noise. The VQAs are effective for a classification task in machine learning; the objective is to train a classifier and predict the label of each input accurately [11–13]. Suppose, a training data is given $\{x_i, y_i\}$, where x_i s and y_i s are inputs and labels respectively. The variational quantum circuit is used as a black box to predict the right label (y_i) for each input after embedding the classical data into the quantum states. VQAs are the quantum variants of neural networks, the most commonly used and highly successful machine learning model. Till now, several architectures for quantum neural networks (QNNs) have been proposed and applied in different areas [14–17]. Recently, Cong et al. [18] introduced the quantum convolutional neural networks (QCNNs) and used them to discriminate quantum states of distinct topological phases. Bhatia et al. [19] performed the simulation of several entangled states on a quantum computer. Romero et al. [20] proposed a quantum variational autoencoder for compressing the quantum data efficiently. Pepper et al. [21] demonstrated the experimental realization of a proposed quantum autoencoder, which will likely be an essential primordial in quantum machine learning. In recent years, generative adversarial networks (GANs) have been an exciting topic of research in classical machine learning. Romero and Aspuru-Guzik [22] introduced the quantum variant of GAN for learning continuous distributions and to speedup classical GAN using quantum systems. Kappen [23] proposed a method to represent the classical data distribution in a quantum system. Moreover, the potential of VQAs can be evaluated in several industry based applications such as supply chain management, intelligent healthcare, smart agriculture, manufacturing production, and cloud manufacturing [24].

The major downside of variational quantum circuits is the occurrence of a barren plateau, which vanishes the gradients of cost function exponentially with the increase in the number of qubits [25, 26]. It can be represented by flat plateaus of the loss function [27]. Firstly, in 2018, McClean et al. [28] studied the barren plateau phenomenon numerically for a class of random quantum circuits. It has been shown that variance of the gradients of modest size quantum circuits vanishes exponentially. Cerezo et al. [29] described a barren plateau phenomenon as a point in the parameter space and investigated it for shallow quantum neural networks (QNNs). It has been proven that global cost gradients vanish exponentially at all depths, whereas local costs exhibit non-vanishing gradients at shorter depth quantum circuits. Grant et al. [30] investigated the problem of a barren plateau of parametrized quantum circuits in the energy landscape. The approach is based on the initialization of parameters to avoid the barren plateaus. Initially, some parameter values are selected randomly, and the remaining values are chosen so that it becomes a chain of shallow circuits. The depth of circuits is used to determine the first updated parameter not to be lodged in a barren plateaus problem during training. Recently, it has investigated that the barren plateau is missing in QNNs and QCNNs with tree tensor network (TTN) architecture [31, 32].

In classical machine learning, the loss landscapes of neural networks and their characterization via the lens of Hessian have been well investigated. Several works have explored the flatness of local minima using the eigenvalues of Hessian [32–36]. During training, the generalization ability of the neural networks depends on batch size. It has explored that mini-batch size tends to favor flat minima basins of the loss function with several eigenvalues of the Hessian $\lambda_j = 0$ and hardly any $\lambda_j < 0$ [37]. Perez-Salinas et al. attained fast convergence for some quantum classification problems using Hessian based optimization method called batched optimization. Rebentrost et al. [38] proposed the quantum variants of two popular iterative optimization algorithms, Newton's and gradient descent. At time t , numerous copies of quantum state $|\psi\rangle_t$ are used to generate other numerous copies of quantum state $|\psi\rangle_{t'}$ at $t' = t+1$ by using the Hessian of an objective function and in consideration of a gradient vector. Recently, Huembeli and Dauphin [39] calculated the Hessian of a loss function and characterized the loss landscape of variational quantum circuits. It showed that the Hessian helps escape the flat surface of loss function for certain data-driven variational quantum circuits.

The VQCs have been extensively employed in a wide array of new applications. Fig 1 shows the schematic representation of a variational quantum circuit. It involves evaluating a cost function or its gradient on a quantum system [40, 41]. A classical optimization loop trains the parameters (θ) of a variational quantum circuit $V(\theta)$ to reduce the cost. It is well-known that the feature map encodings in variational parameterized circuit architecture designs produce loss functions that train easier, and well-selected training parameters “optimizer” that generalize well. Nevertheless, the effects of parameters on the entire loss landscape are not well studied. It has not received the attention it deserves. The occurrence of barren plateaus issue could abolish the quantum advantage with a parameterized quantum circuit [9]. The visualization and understanding of the loss landscape of classical machine learning algorithms remains a vital and highly active area of research. Due to issues of computational complexity, the structure of loss landscape of variational quantum classifiers is not well visualized and recognized as compared to classical neural networks. A better visualization can really work for the advancement of optimization algorithms and can highlight the shortcomings of quantum circuit designs. Hence, it is a natural goal to study the loss landscape of variational quantum classifiers

Fig 1. Schematic representation of a variational quantum circuit (VQC). VQC is created by combining the feature map, variational circuit and measurement component. It is followed by classical optimization model. It includes the iterative implementation of quantum and classical components of the circuit.

<https://doi.org/10.1371/journal.pone.0262346.g001>

with the eigenvalues of Hessian to recognize when the quantum speedup is achievable. The following contributions are claimed:

- Visualized the loss landscape of variational quantum classifiers at different points in parameter space using Hessian matrices.
- Analyzed the behavior of Hessian's eigenvalues on training the variational quantum classifier for different datasets.
- Investigated that how the adaptive Hessian learning rate can influence the convergence while training the variational circuits.
- Observed that adaptive Hessian learning rate can help to overshoot the cost if it gets stuck into local minima and converge quickly.

The eigenvectors and eigenvalues of Hessian present a clear interpretation of the loss landscape of a VQC. We started with a 4-bit parity problem to provide perception about the behavior of Hessian and then studied the VQC trained on diabetes classical data acting as a classifier. The organization of the rest of this paper is as follows: Sect. 2 is devoted to preliminaries and Hessian's computation of variational quantum classifier. In Sect. 3, we computed the Hessian on a quantum simulator and visualized the curvature information of a parity function. In Sect. 3 (B), we characterized the loss landscape (i.e., curvature information) of data-driven variational quantum classifiers via Hessian of the loss function, and the experimental results are plotted for the diabetes dataset. In Sect. 4, we show how the adaptive Hessian learning rate can help to overshoot the cost that helps it avoid getting stuck in local minima during training of variational circuits. Finally, Sect. 5 is the conclusion.

Preliminaries

In this Section, some basic concepts of loss function visualization with the Hessian matrix are given. Here, we give the background required to understand our results. Consider a real-valued function $f(\theta) = f(\theta_1, \theta_2, \dots, \theta_n)$ with $\theta = (\theta_1, \theta_2, \dots, \theta_n)$. The Hessian matrix ($H^2 f(\theta)$) of $f(\theta)$ is represented as the square matrix of the second derivatives of a real-valued function of n variables, which help us to characterize the loss landscape (i.e. curvature information) [42, 43]. The gradient ($\nabla f(\theta)$) gives the partial derivatives of the function. The Hessian operator $H^2 f(\theta)$ gives the partial derivatives of the gradient [44].

$$\nabla f(\theta) = \begin{bmatrix} \frac{\partial f}{\partial \theta_1} \\ \vdots \\ \frac{\partial f}{\partial \theta_n} \end{bmatrix}, H^2 f(\theta) = \begin{bmatrix} \frac{\partial^2 f}{\partial \theta_1^2} & \cdots & \frac{\partial^2 f}{\partial \theta_1 \theta_n} \\ \vdots & \ddots & \vdots \\ \frac{\partial^2 f}{\partial \theta_n \theta_1} & \cdots & \frac{\partial^2 f}{\partial \theta_n^2} \end{bmatrix} \quad (1)$$

If the second derivatives are continuous, then the Hessian is symmetric. The Hessian matrix consists of information about the geometric information of the function. Its eigenvalues ($\lambda_1, \lambda_2, \dots, \lambda_n$)'s are real and used to get this curvature information. So the only thing to examine is whether the eigenvalues λ_i 's are negative or positive. Suppose, x_i is an eigenvector associated with λ_i , then it is represented as i^{th} eigenpair (λ_i, x_i) of H . Using the Hessian matrix, we can determine whether the certain point θ on a surface is locally positive or negative. The eigenvalues of the Hessian matrix give the directions of the derivatives. Suppose the Hessian evaluated at a given point θ consists of all positive eigenvalues $\lambda_i > 0$ (i.e., positive-definite matrix). In

that case, it shows a local positive curvature, and θ is a local minimum of f . Similarly, if all the eigenvalues are negative, it shows a locally negative curvature, and θ is a local maximum of f . The zero eigenvalues indicate the zero curvature of the function or flat directions. If the eigenvalues are mixed (some positive, some negative), then the surface has a saddle point θ of f . Thus, the Hessian can be used to determine the convexity and concavity of a function of one or two variables [45].

Hessian computation of VQC

In classical machine learning, the neural networks are trained over a dataset consisting of feature vectors $\{x_i\}$ and labels $\{y_i\}$ by reducing the cost function as

$$C(\theta) = \frac{1}{n} \sum_{j=1}^n c(l(\theta, x_j), y_j), \quad (2)$$

where the $l(\theta, x_j)$ is the prediction parameterized by weights (θ) of the neural network, $c(\cdot)$ denotes the loss function that measures how well the neural network predicts the label by calculating its difference with the neural network prediction, and n denotes the size of data samples. The loss functions exist in a high-dimensional space due to the presence of several parameters in neural networks. Therefore, its visualization is not possible in higher-dimensional space. Analogously, the loss landscape of VQC has not been extensively examined as compared to classical neural networks. In a quantum layer, the classical data ($\vec{x}$) is encoded into the quantum state $|f(\theta, \vec{x})\rangle$ using a feature map consisting of quantum gates with parameters. In classical layer, the parameterized function $f(\theta, \vec{x})$ is evaluated depending upon the learning parameters (θ) in variational quantum circuit on performing measurement. The tuning of parameters (θ) is executed by minimizing a loss function on a classical computer.

The gradient of a quantum circuit can be evaluated by estimating the expectation value of an observable with reference to θ . It consists a series of unitary transformations. The gradient of an expectation value of an observable $\langle f(\theta, \vec{x}) \rangle = \langle \vec{x} | V^\dagger(\theta) f V(\theta) | \vec{x} \rangle$ can be given as

$$\frac{\partial \langle f(\theta, \vec{x}) \rangle}{\partial \theta_j} = \frac{1}{2} \left[\langle f \left(\left(\theta_j, \theta_j + \frac{\pi}{2} \right), \vec{x} \right) \rangle - \langle f \left(\left(\theta_j, \theta_j - \frac{\pi}{2} \right), \vec{x} \right) \rangle \right] \quad (3)$$

where θ_j denotes the all parameters except θ_j , V is a product of unitary matrices, θ_i is an angle which parametrized the $V(\theta) = e^{-i\theta_j P_j/2}$, where P_j is a Hermitian operator with eigenvalues ± 1 . Let us now define the Hessian matrix elements. The Hessian (H) of a quantum circuit can be computed by performing the parameter shift rule two times [39, 46–48].

$$H_{ij} = \frac{\partial^2 \langle f(\theta, \vec{x}) \rangle}{\partial \theta_i \partial \theta_j} = \partial_i \partial_j \langle f(\theta, \vec{x}) \rangle \quad (4)$$

$$\begin{aligned} H_{ij} = & \frac{1}{4} \left[\langle f \left(\left(\theta_{j,i}, \theta_j + \frac{\pi}{2}, \theta_i + \frac{\pi}{2} \right), \vec{x} \right) \rangle \right. \\ & - \langle f \left(\left(\theta_{j,i}, \theta_j + \frac{\pi}{2}, \theta_i - \frac{\pi}{2} \right), \vec{x} \right) \rangle \\ & - \langle f \left(\left(\theta_{j,i}, \theta_j - \frac{\pi}{2}, \theta_i + \frac{\pi}{2} \right), \vec{x} \right) \rangle \\ & \left. + \langle f \left(\left(\theta_{j,i}, \theta_j - \frac{\pi}{2}, \theta_i - \frac{\pi}{2} \right), \vec{x} \right) \rangle \right] \end{aligned} \quad (5)$$

Thus, the loss function $c(f(\theta, \vec{x}))$ curvature can be studied via second-order derivative of the loss. The chain rule is performed twice to obtain the Hessian matrix elements of c .

$$\begin{aligned} \partial\theta_i\partial\theta_j c(f(\theta, \vec{x}), y) &= \partial\theta_i\partial\theta_j f(\theta, \vec{x}) c'(f(\theta, \vec{x})) + \partial\theta_j \\ &\quad f(\theta, \vec{x}) \partial\theta_i f(\theta, \vec{x}) c''(f(\theta, \vec{x})) \end{aligned} \quad (6)$$

where c' and c'' denote the first order and second order derivatives of the loss function, respectively with respect to the parameter (θ). In this paper, we considered the cost (or loss) function $c(f(\theta, \vec{x}), y) = (y - f(\theta, \vec{x}))^2$, which will be minimized. On calculating the derivatives, the Eq (6) becomes:

$$\begin{aligned} \partial\theta_i\partial\theta_j c(f(\theta, \vec{x}), y) &= \partial\theta_i\partial\theta_j f(\theta, \vec{x}) (-2(y - f(\theta, \vec{x}))) \\ &\quad + \partial\theta_j f(\theta, \vec{x}) \partial\theta_i f(\theta, \vec{x}) \cdot 2 \end{aligned} \quad (7)$$

If the cost displays barren plateau, then its variance is vanishing exponentially as $\text{Var}[\partial f(\theta, \vec{x})] \leq F(n)$, where $F(n) \in \mathcal{O}(1/a^n)$, for $a > 1$ [28, 48, 49]. Then, the Hessian matrix elements are vanishing exponentially.

Experiment settings

We have used the PyTorch library [50], and pennylane package [51] for developing and training the variational quantum classifiers. The code is written in pennylane with a large number of parameters to expedite the experiments. The implementation is performed using PyTorch to accelerate the simulation using algebraic manipulation. All the quantum simulations are performed using python framework on the PennyLane platform for quantum differentiable programming [51]. The 2D and 3D graphs are plotted using a Plotly i.e. graphing library available in python.

Warm-up example

In this section, we characterized the loss landscape curvature of the four-bit parity problem with the Hessian. We begin exploring the concept of Hessian of loss functions of VQCs with a warm-up activity of solving a four-bit parity problem. Before illustrating the variational quantum classifier, the partial function is defined as a Boolean function whose output is 1 if and only if the input vector has an uneven number of ones. It is also called the XOR function of two inputs. The n -bit parity function is given as [52]

$$f : v \in \{0, 1\}^{\otimes n} \rightarrow w \begin{cases} 1 & \text{if odd number of 1's in } v \\ 0 & \text{else} \end{cases} \quad (8)$$

The first step is to encode the input vectors into a quantum state. In a warm-up example, the inputs are 4-bit strings that are encoded into the state of qubits. The feature mapping and variational circuit of the four-bit parity problem are shown in Fig 2. A single qubit rotation is defined as:

$$R(\phi, \theta, \omega) = RZ(\omega)RY(\theta)RZ(\phi) = \begin{bmatrix} e^{-i(\phi+\omega)/2} \cos(\theta/2) & -e^{i(\phi-\omega)/2} \sin(\theta/2) \\ e^{-i(\phi-\omega)/2} \sin(\theta/2) & e^{i(\phi+\omega)/2} \cos(\theta/2) \end{bmatrix}. \quad (9)$$

The initial layer of R_x gates prepares the initial state, which is also known as feature map circuit. Later, there are three layers of variational gates, where each layer consists of rotational

Fig 2. Feature map and variational quantum circuits for the parity problem. A four-qubit feature map consisting rotation around x-axis to prepare the initial state. It is followed by three layers of variational gates, where each layer consists of rotational gates with three trainable parameters ($R(\omega_i) = R(\phi, \theta, \omega)$) in each of the four qubits, followed by a set of CNOT gates. For two classes, the measurement is performed on one qubit, which is enough to have orthogonal measurements for the classes.

<https://doi.org/10.1371/journal.pone.0262346.g002>

gates with three trainable parameters in each of the four qubits, followed by a set of CNOT gates. It is to be noted that the combination of CNOT gates in each layer has a different structure. The variation circuit consists of 36 parameterized gates and 12 non-parameterized gates.

The input data $\vec{v}$ is encoded into quantum state using feature map function as $\psi : \vec{v} \rightarrow |\psi(\vec{v})\rangle\langle\psi(\vec{v})|$. Initially, one qubit rotations are performed around x-axis. It is followed by U_1 single-qubit gate to apply a quantum phase to the qubit. Furthermore, a controlled-NOT multi-qubit gate is applied to flip the target qubit when the control qubit is in $|1\rangle$. The purpose of the feature map circuit is to map the classical input data into the quantum state.

The final n -qubit feature quantum state becomes

$$|\psi(\vec{v})\rangle = U_{\psi(\vec{v})}|0\rangle^{\otimes n} \quad (10)$$

Second, a short depth quantum circuit $V(\theta)$ is applied to the feature state. It depends upon the selection of parameterization for the gates and the number of d layers. The classical optimizer handles the parameters during training to reduce the value of a loss function. Before returning a final classifier outcome, classical postprocessing is performed to the expectation value of the circuit. The aim is to determine the optimal classifying circuit $V(\theta)$ that separates the dataset with distinct labels. In variational quantum circuit, we used the $R_y(\theta)$ and $R_z(\theta)$ parameterized gates that are applied to rotate the qubits by angle (θ) around y-axis and z-axis, respectively. CNOT gates follow it. The objective is to find a sequence of gates that forms a final state $|\psi\rangle_0$. A cost function (C_f) is defined as the square of trace distance (D_t) between final $|\psi\rangle_0$ and initial state $|\psi\rangle = V(\theta)^\dagger|0\rangle$, which is determined as

$$C_f = \text{Tr}[O_f V(\theta)|\psi\rangle_0\langle\psi|_0 V(\theta)^\dagger] \quad (11)$$

where $O_f = 1 - |0\rangle\langle 0|$. It is equivalent to $C_f = D_t(|\psi\rangle_0\langle\psi|)^2$. As a warm-up to our study, the performance of a variational quantum classifier is tested for a parity problem. It is learnable in a quantum setting, a binary classification task $w = \{0, 1\}$ and its outcomes are measured on the computational basis. We selected the third qubit in the Pauli-Z direction to perform the measurement and thresholding (Δ) the expectation value $\langle Z \rangle \leq \Delta$ ($\langle Z \rangle > \Delta$) to classify the input vectors in to one of the labels $w = 0$ ($w = 1$), respectively. We utilized the gradient descent optimizer (GD) to iterate a parameter update based on the gradient of the loss function. It is used to minimize an objective function to its local minimum by adjusting the parameters repeatedly. The partial derivatives of the cost function are determined with respect to each parameter and store the outcome in a gradient. A step of the GD optimizer determines the new values via the rule $\theta^{(t+1)} = \theta^{(t)} - \eta \nabla f(\theta^{(t)})$, where η is a user-defined hyperparameter

Fig 3. Loss landscape of the parity problem for (θ_3 and θ_7). In fig (a, c), the loss landscape is visualized with a local loss function for two parameters θ_3 and θ_7 , where rest of the parameters are set to the optimal values. Moreover, it cannot be visualize for more than 2 parameters for the 3D Loss, because it cannot obtain the full range of loss between (0 and 1). In fig (b, d), the points in contour plots show the direction of improvement in the optimization process i.e. how the parameters are optimized during iterations. Fig (b) is just the zoom-in version of Fig (d) for better clarification. The green color point depicted the optimal value of θ_3 and θ_7 .

<https://doi.org/10.1371/journal.pone.0262346.g003>

relating to step size. We now consider the concept of Hessian that how it helps to realize the loss landscape (or curvature information) for the parity problem. Firstly, the parameters are initialized randomly and determine the set of parameters that can produce the target state. The optimization problem is translated into the loss function minimization.

It is presented as a function of θ_3 and θ_7 , and rest of the parameters are set to the optimal values. Fig 3 shows the loss landscape of a parity problem with a local loss function of parameters, where ($\theta_1 = \theta_3$ and $\theta_2 = \theta_7$). We started with some random initialization. The contour plots Fig 3(b) and 3(d) show the direction of improvement in the optimization process that how the parameters (θ_3 and θ_7) are optimized during 100 iterations. The green color point denotes the optimal value. The zoom-out and zoom-in versions of contour plots are shown in Fig 3(b) and 3(d) for a better view of optimal values.

The minimum value of a loss can be recovered with the gradient descent optimizer due to the point of local convexity in the landscape. Fig 4 depicts the behavior of eigenvalues of the Hessian for parity problem during training. During the optimization process, the Hessian matrix is calculated on all the trainable parameters and eigenvalues are recorded after each iteration. In Fig 4, the eigenvalues are plotted for some specific iterations in ascending order to

Fig 4. The evolution of behavior of the eigenvalues of the Hessian during training for the parity problem. Initially, it shows the mixture of non-negative and negative values. Finally, a well converged loss is observed where the most of the Hessian's eigenvalues are non-negative. During the training, we separate the curves for better clear view as the difference between smallest and biggest eigenvalues is changing for different epochs.

<https://doi.org/10.1371/journal.pone.0262346.g004>

observe how the behaviour of the eigenvalues is changing with the trainable parameters being updated during the progression of an optimization process. Fig 4(a)–4(d) shows the variations between the minimum and maximum eigenvalues at each iteration. The distribution of eigenvalues for the randomly initialized quantum circuit shows the mixture of negative and non-negative values close to zero at iterations (0–7). Fig 4(b) and 4(c) shows some of the eigenvalues are positive, some are negative, and the bulk of them is zero. Fig 4(d) shows a loss for the well converged variational circuit where we left with a single negative eigenvalue of the Hessian and rest all are non-negative (at 100th iteration). In fact, the zero gradient ensures that it is a global minimum. Moreover, the zero eigenvalues correlate to directions where variations in parameters do not alter the curvature information. Thus, Hessian's behavior of the eigenvalues helps to obtain the minimum and maximum stability in the directions of loss landscape.

Classification of diabetes

Let us consider the case of the classical diabetes dataset for the classification of diabetes in supervised learning. In this section, we analyzed the loss landscape through the Hessian for diabetes dataset and investigated how VQC will perform to predict diabetes. The diabetes dataset can be download from UCI machine learning repository [53]. It consists of 8 input features

Fig 5. Variational quantum circuit. A general 8-qubit variational quantum circuit is constructed for diabetes dataset with a local cost function. In feature map, Hadamard gates, rotation around y-axis and control-Z entangling gates are applied. The variational part of circuit consisting parameterized rotations $R(\omega_j) = R(\varphi_1, \varphi_2, \varphi_3)$ applied on each qubit.

<https://doi.org/10.1371/journal.pone.0262346.g005>

(age, glucose, insulin, pregnancies, body mass index (BMI), skin thickness, diabetes pedigree function, and blood pressure) and one binary output feature.

Firstly, the 8 input features are encoded into the state of qubits by a quantum feature map. Consider a classical dataset $D = \{(x^n, y^n)\}_{n=1}^S$ for binary classification, where $y^n \in \{0, 1\}$ i.e. 0

Fig 6. Loss landscape of the diabetes dataset for $(\theta_0$ and $\theta_{24})$. (a) The loss landscape is demonstrated with a local loss function for two qubits θ_0 and θ_{24} . Fig (b-c) The direction of improvement in the optimization process is shown using contour plots i.e. how the parameters are optimized during iterations.

<https://doi.org/10.1371/journal.pone.0262346.g006>

for no diabetes and 1 is for diabetes. Each segment of classical data is encoded into an amplitude of a qubit using single-qubit rotations. Afterward, a variational quantum circuit $V(\theta)$ is applied to the feature quantum state for training and classification of diabetes. The 8-qubit variational quantum circuit with a local cost function is constructed, as shown in Fig 5. The feature map consists of Hadamard gates, rotation around y-axis and control-Z entangling gates. It is followed by a variational part of the circuit containing single-qubit rotations $R(\omega_j) = R(\varphi_1, \varphi_2, \varphi_3)$ on each qubit.

The feature map and variational quantum circuits are repeated two times. It is constructed simple enough so that it can be executed on real quantum systems and complex to separate the input data after mapping. For our experiments, the feature map and variational quantum circuits were constructed to a fixed number of depth and qubits. To evaluate the most likely state of an outcome qubit, the quantum circuit can be determined for several iterations by considering a similar input to calculate the probability distribution among the basis states. We used the gradient descent optimization method to find the parameters which can determine the probabilities closest to reality. The measurement is performed on a qubit by employing the Pauli operators in a specific direction. The outcome qubit discovers the predicted value of an input, i.e., the class label values allocated. The loss landscape of VQC for a diabetes dataset with a loss function of parameters (θ_0) and (θ_{24}) is visualized in Fig 6, where θ_1 and θ_2 are set to θ_0 and θ_{24} , respectively. Here, the other parameters are set to the optimal value after each iteration. In each iteration, the workflow of the optimization process consists of three steps feature map, variational circuit and observation. These are performed on quantum circuits for each of the samples on the dataset and the loss is calculated by taking the distance of prediction from the label for each sample and then averaging them. It is further used in the classical optimization process of the trainable parameters of the variational circuit and finally the trainable parameters are updated as the final step of each iteration. During the observation purposes, the hessian matrix is calculated on the trainable parameters in each iteration and then the eigenvalues of the hessian matrix are recorded only for the data visualization and need not to be incorporated while being implemented on practical purpose. The input vectors are normalized to lie in $[-\pi, \pi]$ and plotted the loss landscape for two qubits in a range. If we try to plot the loss landscape for more than two parameters, then it cannot obtain the full range between (0 and 1). The contour plots (b) and (c) show the evolution of optimized parameters during 30 iterations in Fig 7. Furthermore, the loss landscape with a loss function of parameters (θ_1) and (θ_{25}) has been analyzed in Fig 7. It has been observed from the landscape that optimum cost is single with good local minima. It shows the prediction map of VQC with Z-measurement for the diabetes dataset.

Fig 7. Loss landscape of the diabetes dataset for (θ_1) and (θ_{25}) . (a) The loss landscape of the diabetes dataset is visualized with a local loss function for two parameters θ_1 and θ_{25} . Fig (b-c) give a clear view of the optimal values during iterations.

<https://doi.org/10.1371/journal.pone.0262346.g007>

Fig 8. The evolution of behaviour of the eigenvalues of the Hessian during training for the diabetes dataset. It shows the eigenvalues evolution during the training for different epochs. Initially, the Hessian's eigenvalues consisting of negative and positive values equally, and many of them are zero. After the convergence, it shows maximum positive eigenvalues.

<https://doi.org/10.1371/journal.pone.0262346.g008>

The distribution of Hessian's eigenvalues is determined over the training process of a variational circuit to locate one of the minima. It is used to investigate whether a particular stationary point is a saddle point or not. At the beginning of the training, the gradient descent method is struggling to break the symmetry. Due to the small gradients, it faces a problem in training small quantum circuits. Fig 8(a) shows a distribution of the Hessian's eigenvalues consisting of equally possible negative and positive values, and most of them are zero for iterations (0, 2, and 4). We observed that the negative eigenvalues gradually started to disappear with the increase in number of iterations, as shown in Fig 8(c). After the convergence at 30th iteration, a single negative eigenvalue is left and rest all became non-negative, in Fig 8(d). The bulk of zero eigenvalues shows a flat direction of the $f(\theta, \vec{x})$ in the loss landscape, where any alterations in circuit parameters do not disturb the loss landscape. The positive semi-definite behavior of the Hessian's eigenvalues signifies a very steady result. Although, it is not practical to visualize the loss landscape of VQC in 3-dimensional due to the problem of fixing the other parameters. Nevertheless, it is feasible to visualize the loss landscape of variational quantum classifiers through the lens of the Hessian.

Convergence via adaptive Hessian learning rate

In this section, we show how the Hessian can adjust adaptively to the learning rate (LR) for each parameter. Adaptive learning rate has been a popular practice in classical machine learning. In any gradient based optimization method, a very large learning rate can cause an overshoot in cost while coming close to smaller gradients, whereas very small LR will cause an extremely slow approach towards the lowest point in loss landscape. Therefore, a smart trade off between both leads to the requirement of tuning the Learning rate properly. It is not as

simple as it looks. We used an adaptive learning rate, which initially starts with the higher values of learning rate and gradually reduces it with reducing gradient values.

Adaptive Hessian learning rate (A-HLR) is something similar to the previous concept with an add on, specific to the quantum machine learning model's loss landscape analysis. Here, instead of using continuous variation of learning rates, we consider a set of discrete values of learning rates. We begin with the largest LR and gradient descent optimization process. After each step, we compare the updated cost with the previous cost. After few consecutive steps, if the difference between updated and previous cost is below a threshold level, then LR is set to the next smaller discrete value in the set of LR's. The complete process is repeated until we reach the lowest LR in the set of discrete learning rates. Since the loss landscape of quantum machine learning models is quite different from the classical machine learning models. There can be observed a specific repetitive manner of landscape unlike classical model, which leads to the higher probability of getting stuck in local minima.

In case, if the execution reached lowest LR among the set of discrete learning rates, there can be two possible cases (i) either the optimization stuck in a local minima, we need to get out of this, (ii) or the optimization problem reached the global minima. To tackle these scenarios, we considered the concept of Hessian matrix and used it with a set of decreasing learning rates. The implementation is started with the largest learning rate, for each of values of the learning rate in the set and the process of optimization continues until the difference of loss values in two consecutive iterations goes below tolerance. Once it occurs, then the next learning rate is considered from the predefined set and the process is repeated until it reaches the least learning rate value. If the count of negative eigenvalues is lesser than a threshold value, then the solution obtained is a global minima. In case, if the count is greater, then it got stuck in local minima. It is to be noted that the threshold value of negative eigenvalue count is a hyperparameter, which has to be tuned properly depending on the type of data in a given dataset. (Practically it should be zero, but in reality its not always possible to reach that point). If the model is stuck at local minima, i.e. the learning rate has already reached the lowest of its all possible given values. Then, the optimization will be again started with the higher learning rate with an objective that the cost will overshoot and come out of the local minima. The process is repeated until the LR comes to its lowest value. Finally, it is evaluated whether the optimization has reached global minima or not using the Hessian matrix.

Fig 9(a) shows the comparison that how the cost is evolving throughout the optimization process with an adaptive Hessian learning rate (A-HLR) method and gradient descent method with constant learning rates. It has been determined that A-HLR converged very well within 25 iterations. The gradient descent method with LR = 0.5 learning rate also converges, but not as quick as with A-HLR. Fig 9(b) shows how the value of adaptive learning rate evolves on using A-HLR during first 25 iterations and depicts how it overshoots the local minima of the cost function using lowest to highest LRs. Fig 9(c) shows the comparison between A-HLR and gradient descent methods with different learning rates for 100 iterations. Fig 9(d) depicts how the value of adaptive learning rate evolves on using A-HLR during the first 100 iterations. A-HLR fits the local shape of gradient to the loss landscape very well, provides a faster convergence than gradient descent method with constant learning rates. It permits one to select a descent direction for faster convergence during the training of variational circuits. Therefore, the local traps in the loss landscape can be avoided by using the adaptive hessian learning rate approach.

Conclusion

In this paper, the curvature information of the loss landscape of variational quantum classifiers has been visualized via the lens of Hessian eigenvalues. We developed a simple theoretical

Fig 9. Training cost of four bit parity dataset with adaptive Hessian learning rate (A-HLR). In fig (a, c), the comparison is shown between the adaptive Hessian learning rate and gradient descent method with constant learning rates. A-HLR shown faster convergence than gradient descent with fixed learning rates. It has been shown the stable and efficient convergence of the cost function for parity dataset during training. Fig (b, d) shows how A-HLR evolves during training of 25 and 100 iterations. It depicts how to overshoot the local minima of the cost function using adaptive Hessian lower to higher learning rates.

<https://doi.org/10.1371/journal.pone.0262346.g009>

quantum model of Hessians and gradients, as justified by datasets for numerical justifications on VQCs. The parity function problem is considered as a warm study to show the behavior of Hessian's eigenvalues. It has observed that VQC has an exceptional ability to generalize small datasets. Furthermore, we visualized the cost function landscape of VQC designed for the diabetes dataset. It converges efficiently for data-driven problems. We identified some differences in the convergence with adaptive Hessian learning rate and gradient descent method using fixed learning rate. It has been observed that adaptive Hessian learning rate helps to overshoot the cost if it gets fall into local minima and converge quickly. It seems beneficial to study the local curvature information of VQC through the Hessian. The integration of gradient-based methods and Noisy Intermediate-Scale Quantum (NISQ) devices is still a young area and potentially has a lot more to offer. In the future, this work will open up new avenues of research in solving classical and quantum optimization problems and framework design. It will help the research communities to accelerate the analysis of variational quantum algorithms based on Hessian.

Author Contributions

Conceptualization: Amandeep Singh Bhatia, Kamalpreet Singh Bhangu.

Data curation: Pinaki Sen, Kamalpreet Singh Bhangu.

Formal analysis: Pinaki Sen, Amandeep Singh Bhatia, Kamalpreet Singh Bhangu.

Investigation: Pinaki Sen, Amandeep Singh Bhatia, Kamalpreet Singh Bhangu.

Methodology: Pinaki Sen, Amandeep Singh Bhatia.

Project administration: Amandeep Singh Bhatia.

Resources: Amandeep Singh Bhatia, Ahmed Elbeltagi.

Software: Amandeep Singh Bhatia, Kamalpreet Singh Bhangu, Ahmed Elbeltagi.

Supervision: Amandeep Singh Bhatia.

Visualization: Amandeep Singh Bhatia, Ahmed Elbeltagi.

Writing – original draft: Amandeep Singh Bhatia, Kamalpreet Singh Bhangu.

Writing – review & editing: Amandeep Singh Bhatia, Ahmed Elbeltagi.

References

1. Preskill J., Quantum computing in the nisq era and beyond, *Quantum* 2 (2018) 79.
2. Peruzzo A., McClean J., Shadbolt P., Yung M.-H., Zhou X.-Q., Love P. J., et al. A variational eigenvalue solver on a photonic quantum processor, *Nature communications* 5 (2014) 4213. <https://doi.org/10.1038/ncomms5213> PMID: 25055053
3. Higgott O., Wang D., Brierley S., Variational quantum computation of excited states, *Quantum* 3 (2019) 156.
4. Nakanishi K. M., Mitarai K., Fujii K., Subspace-search variational quantum eigensolver for excited states, *Physical Review Research* 1 (3) (2019) 033062. <https://doi.org/10.1103/PhysRevResearch.1.033062>
5. J. R. McClean, M. P. Harrigan, M. Mohseni, N. C. Rubin, Z. Jiang, S. Boixo, et al. Low depth mechanisms for quantum optimization, *arXiv preprint:2008.08615* (2020).
6. A. Garcia-Saez, J. Latorre, Addressing hard classical problems with adiabatically assisted variational quantum eigensolvers, *arXiv preprint:1806.02287* (2018).
7. Parrish R. M., Hohenstein E. G., McMahon P. L., Martínez T. J., Quantum computation of electronic transitions using a variational quantum eigensolver, *Physical review letters* 122 (23) (2019) 230401. <https://doi.org/10.1103/PhysRevLett.122.230401> PMID: 31298869
8. E. Farhi, J. Goldstone, S. Gutmann, A quantum approximate optimization algorithm, *arXiv preprint:1411.4028* (2014).
9. Morales M. E., Biamonte J., Zimborás Z., On the universality of the quantum approximate optimization algorithm, *Quantum Information Processing* 19 (9) (2020) 1–26. <https://doi.org/10.1007/s11128-020-02748-9>
10. Wong R., Bhatia A. S., Quantum Algorithms: Application Perspective, Limitations and Future Applications of Quantum Cryptography, 82–101 (2021).
11. Stokes J., Izaac J., Killoran N., Carleo G., Quantum natural gradient, *Quantum* 4 (2020) 269. <https://doi.org/10.22331/q-2020-05-25-269>
12. Bhatia A. S., Saggi M. K., Kumar A., and Jain S., “Matrix product state–based quantum classifier,” *Neural computation*, vol. 31, no. 7, pp. 1499–1517, 2019. https://doi.org/10.1162/neco_a_01202
13. Bhatia A. S., Wong R., Recent Progress in Quantum Machine Learning, Limitations and Future Applications of Quantum Cryptography, 232–256 (2021).
14. E. Farhi, H. Neven, Classification with quantum neural networks on near term processors, *arXiv preprint:1802.06002* (2018).
15. M. Altaisky, Quantum neural network, *arXiv preprint quant-ph/0107012* (2001).
16. Schuld M., Sinayskiy I., Petruccione F., The quest for a quantum neural network, *Quantum Information Processing* 13 (11) (2014) 2567–2586. <https://doi.org/10.1007/s11128-014-0809-8>
17. A. Abbas, D. Sutter, C. Zoufal, A. Lucchi, A. Figalli, S. Woerner, The power of quantum neural networks, *arXiv preprint:2011.00027* (2020).

18. Cong I., Choi S., Lukin M. D., Quantum convolutional neural networks, *Nature Physics* 15 (12) (2019) 1273–1278. <https://doi.org/10.1038/s41567-019-0648-8>
19. A. S. Bhatia and M. K. Saggi, "Implementing entangled states on a quantum computer," *arXiv preprint:1811.09833*, 2018.
20. Romero J., Olson J. P., Aspuru-Guzik A., Quantum autoencoders for efficient compression of quantum data, *Quantum Science and Technology* 2 (4) (2017) 045001. <https://doi.org/10.1088/2058-9565/aa8072>
21. Pepper A., Tischler N., Pryde G. J., Experimental realization of a quantum autoencoder: The compression of qutrits via machine learning, *Physical review letters* 122 (6) (2019) 060501. <https://doi.org/10.1103/PhysRevLett.122.060501> PMID: 30822053
22. J. Romero, A. Aspuru-Guzik, Variational quantum generators: Generative adversarial quantum machine learning for continuous distributions, *arXiv preprint:1901.00848* (2019).
23. Kappen H. J., Learning quantum models from quantum or classical data, *Journal of Physics A: Mathematical and Theoretical* 53 (21) (2020) 214001. <https://doi.org/10.1088/1751-8121/ab7df6>
24. P. K. Maddikunta, Q. V. Pham, B. Prabadevi, N. Deepa, K. Dev, T. R. Gadekallu, et al. Industry 5.0: a survey on enabling technologies and potential applications, *Journal of Industrial Information Integration* 100257 (2021).
25. Uvarov A. V., Biamonte J. D., On barren plateaus and cost function locality in variational quantum algorithms, *Journal of Physics A: Mathematical and Theoretical* 54 (24) (2021) 245301. <https://doi.org/10.1088/1751-8121/abfac7>
26. J. Kim, J. Kim, D. Rosa, Universal effectiveness of high-depth circuits in variational eigenproblems, *arXiv preprint:2010.00157* (2020).
27. E. Campos, A. Nasrallah, J. Biamonte, Abrupt transitions in variational quantum circuit training, *arXiv preprint:2010.09720* (2020).
28. McClean J. R., Boixo S., Smelyanskiy V. N., Babbush R., Neven H., Barren plateaus in quantum neural network training landscapes, *Nature communications* 9 (1) (2018) 1–6. <https://doi.org/10.1038/s41467-018-07090-4> PMID: 30446662
29. M. Cerezo, A. Sone, T. Volkoff, L. Cincio, P. J. Coles, Cost-function-dependent barren plateaus in shallow quantum neural networks, *arXiv preprint:2001.00550* (2020).
30. Grant E., Wossnig L., Ostaszewski M., Benedetti M., An initialization strategy for addressing barren plateaus in parametrized quantum circuits, *Quantum* 3 (2019) 214. <https://doi.org/10.22331/q-2019-12-09-214>
31. Li H., Xu Z., Taylor G., Studer C., Goldstein T., Visualizing the loss landscape of neural nets, in: *Advances in neural information processing systems*, 2018, pp. 6389–6399.
32. Y. Cooper, The loss landscape of overparameterized neural networks, *arXiv preprint:1804.10200* (2018).
33. Rotskoff G. M., Vanden-Eijnden E., Neural networks as interacting particle systems: Asymptotic convexity of the loss landscape and universal scaling of the approximation error, *stat* 1050 (2018) 22.
34. A. Pesah, M. Cerezo, S. Wang, T. Volkoff, A. T. Sornborger, P. J. Coles, Absence of barren plateaus in quantum convolutional neural networks, *arXiv preprint:2011.02966* (2020).
35. K. Zhang, M.-H. Hsieh, L. Liu, D. Tao, Toward trainability of quantum neural networks, *arXiv preprint:2011.06258* (2020).
36. L. Sagun, U. Evci, V. U. Guney, Y. Dauphin, L. Bottou, Empirical analysis of the hessian of over-parameterized neural networks, *arXiv preprint:1706.04454* (2017).
37. N. S. Keskar, D. Mudigere, J. Nocedal, M. Smelyanskiy, P. T. P. Tang, On large-batch training for deep learning: Generalization gap and sharp minima, *arXiv preprint:1609.04836* (2016).
38. Rebstroff P., Schuld M., Wossnig L., Petruccione F., Lloyd S., Quantum gradient descent and newton's method for constrained polynomial optimization, *New Journal of Physics* 21 (7) (2019) 073023. <https://doi.org/10.1088/1367-2630/ab2a9e>
39. Huembeli P. and Dauphin A., "Characterizing the loss landscape of variational quantum circuits," *Quantum Science and Technology*, vol. 6, no. 2, p. 025011, 2021. <https://doi.org/10.1088/2058-9565/abdbc9>
40. M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, et al., Variational quantum algorithms, *arXiv preprint:2012.09265* (2020).
41. Y. Du, T. Huang, S. You, M.-H. Hsieh, D. Tao, Quantum circuit architecture search: error mitigation and trainability enhancement for variational quantum solvers, *arXiv preprint:2010.10217* (2020).
42. Thacker W. C., The role of the hessian matrix in fitting models to measurements, *Journal of Geophysical Research: Oceans* 94 (C5) (1989) 6177–6196. <https://doi.org/10.1029/JC094iC05p06177>

43. C. Bishop, Exact calculation of the hessian matrix for the multilayer perceptron (1992).
44. Van Den Bos A., Complex gradient and hessian, *IEEE Proceedings-Vision, Image and Signal Processing* 141 (6) (1994) 380–382. <https://doi.org/10.1049/ip-vis:19941555>
45. Yuille A. L., Rangarajan A., The concave-convex procedure, *Neural computation* 15 (4) (2003) 915–936. <https://doi.org/10.1162/08997660360581958> PMID: 12689392
46. Mitarai K., Negoro M., Kitagawa M., Fujii K., Quantum circuit learning, *Physical Review A* 98 (3) (2018) 032309. <https://doi.org/10.1103/PhysRevA.98.032309>
47. Schuld M., Bergholm V., Gogolin C., Izaac J., Killoran N., Evaluating analytic gradients on quantum hardware, *Physical Review A* 99 (3) (2019) 032331. <https://doi.org/10.1103/PhysRevA.99.032331>
48. M. Cerezo, P. J. Coles, Impact of barren plateaus on the hessian and higher order derivatives, *arXiv preprint:2008.07454* (2020).
49. K. Sharma, M. Cerezo, L. Cincio, P. J. Coles, Trainability of dissipative perceptron-based quantum neural networks, *arXiv preprint:2005.12458* (2020).
50. Paszke A., Gross S., Massa F., Lerer A., Bradbury J., Chanan G., et al., Pytorch: An imperative style, high-performance deep learning library, in: *Advances in neural information processing systems*, 2019, pp. 8026–8037.
51. V. Bergholm, J. Izaac, M. Schuld, C. Gogolin, M. S. Alam, S. Ahmed, et al. PennyLane: Automatic differentiation of hybrid quantum-classical computations, *arXiv preprint:1811.04968* (2018).
52. Arslanov M., Ashigaliev D., Ismail E., N-bit parity ordered neural networks, *Neurocomputing* 48 (1-4) (2002) 1053–1056. [https://doi.org/10.1016/S0925-2312\(02\)00610-0](https://doi.org/10.1016/S0925-2312(02)00610-0)
53. Pima Indians Diabetes Database, <https://www.kaggle.com/uciml/pima-indians-diabetes-database>, [Online; accessed 14-Jan-2021].