

Supplemental Information for “Hardware-Efficient Quantum Principal Component Analysis for Medical Image Recognition”

Zidong Lin,^{1,*} Hongfeng Liu,^{1,*} Kai Tang,^{1,*} Yidai Liu,^{2,*} Liangyu Che,¹ Xinyue Long,¹ Xiangyu Wang,¹ Yu-ang Fan,¹ Keyi Huang,¹ Xiaodong Yang,¹ Tao Xin,¹ Xinfang Nie,^{1,3,†} and Dawei Lu^{1,3,4,‡}

¹*Shenzhen Institute for Quantum Science and Engineering and Department of Physics, Southern University of Science and Technology, Shenzhen 518055, China*

²*Department of Physics, Hong Kong University of Science and Technology, ClearWaterBay, Kowloon, HongKong, China*

³*Guangdong Provincial Key Laboratory of Quantum Science and Engineering, Southern University of Science and Technology, Shenzhen 518055, China*

⁴*International Quantum Academy, Shenzhen, 518055, China*

Appendix A: Dataset of the lung CT images

The CT images used in this study were sourced from the iCTFT database, a reliable resource that offers an extensive collection of thoracic CT images and clinical data from individuals diagnosed with COVID-19 pneumonia. To enable the application of the qPCA algorithm, it was necessary to preprocess both the training and testing sets of CT images. The preprocessing steps were carefully designed to extract the essential information required for incorporating the images into the quantum computer. These steps encompassed image cropping, computation of the covariance matrix, and normalization processes. By performing these procedures, the CT images were appropriately prepared for subsequent analysis using the qPCA algorithm.

Image cropping is utilized to effectively eliminate irrelevant information from the CT images, thereby reducing the dimension of the covariance matrix and simplifying the computation complexity. In our approach, we firstly divided the original lung CT image into two parts and retained only the left portion. Then, we cropped out the white regions, which corresponds to the skeleton or the heart which do not contribute to the diagnosis of COVID-19 in any way. This selective cropping process ensures that the remaining image region retains the crucial features necessary for accurate analysis. A comparison of the original CT and the processed CT is shown in Figure S1.

The test dataset are displayed in Figure S2 (negative) and Figure S3 (positive).

* These authors contributed equally to this work.

† niexf@sustech.edu.cn

‡ ludw@sustech.edu.cn




Figure S1. Preprocessing of CT images prior to qPCA involves removing unnecessary redundant information through cropping.




Figure S2. Negative CT images of the test dataset.




Figure S3. Positive CT images of the test dataset.

To compute the covariance matrix, the CT image underwent a series of preprocessing steps. Firstly, the image was discretized into a grayscale matrix consisting of 260×190 pixels, with each pixel representing a gray value. Consequently, the matrix was flattened into a $49,400 \times 1$ column vector. In the training data set, we selected two CT images: one positive and one negative image. By concatenating the column vectors of these two images, we obtained a $49,400 \times 2$ matrix. After applying a centralization step, as described in the Methods section, we computed the covariance matrix $C = \mathbf{X}\mathbf{X}^T$ and the 2×2 commuted covariance matrix $D = \mathbf{X}^T\mathbf{X}$. These matrices are Hermitian semidefinite-positive but not necessarily trace one. Notably, the dimension of C is $49,400 \times 49,400$, making it impractical to solve its eigenvectors and eigenvalues using current quantum techniques. Therefore, in our experiment, we focused on processing the lower-dimensional matrix D instead. Let us assume the eigenvectors and eigenvalues of D are denoted as $|e_i\rangle$ and λ_i , respectively, meaning that $\mathbf{X}^T\mathbf{X}|e_i\rangle = \lambda_i|e_i\rangle$. It follows that $C\mathbf{X}|e_i\rangle = \lambda_i\mathbf{X}|e_i\rangle$, indicating that λ_i is also an eigenvalue of matrix C . Moreover, the corresponding eigenvector can be obtained by applying the matrix \mathbf{X} to $|e_i\rangle$, i.e., $\mathbf{X}|e_i\rangle$. The most computationally intensive part of qPCA lies in the eigendecomposition of the matrix D . However, since the matrix D is not a quantum state and cannot be directly processed by a quantum computer, we need to encode it as a quantum state by normalizing the commuted covariance matrix as $\rho_D = D/\text{Tr}(D)$. This normalization operation only scales down the eigenvalues of matrix D by a factor of $\text{Tr}(D)$ but does not affect the eigenvectors.

Appendix B: Pseudocode for the image classification process

The fundamental concept of PCA revolves around computing the principal components (eigenvectors) and their corresponding weights (eigenvalues) by performing an eigendecomposition of the covariance matrix. Notably, in 1991, Sirovich and Kirby pioneered the application of PCA in image processing, specifically focusing on human faces. Eigen-decomposition, as the primary time-consuming part in PCA, is usually solved by the QR algorithm or power iteration algorithm numerically. Here, we explore the application in CT images classification by replacing the it with our iterative qPCA algorithm. The pseudocode for the image classification process is presented in Table S1.

Appendix C: NMR implementation

Sample. – Our experiments were conducted using a nuclear magnetic resonance (NMR) quantum processor, the ^{13}C -labeled trans-crotonic acid dissolved in d6-acetone is utilized as a 4-qubit quantum processor. The qubits in this system correspond to the coupled spin-1/2 $^{13}\text{C}_1$ to $^{13}\text{C}_4$ in a molecule and the methyl

Algorithm qPCA for CT image classification

Input: 2 training CT images, 1 test CT image

Output: the CT image in training set that is the closest to the test CT image

1. Load, reshape and centralize the data

$$\text{Train} \rightarrow \mathbf{X} = [\mathbf{x}_1^{\text{train}}, \mathbf{x}_2^{\text{train}}], \text{Test} \rightarrow \mathbf{x}^{\text{test}}$$

 2. Construct the commuted covariance matrix: $\mathbf{D} \leftarrow \mathbf{X}^T \mathbf{X}$

 3. **Call** iterative-qPCA: input (\mathbf{D})

$$\text{output } |e_i\rangle, \lambda_i \text{ when } \mathbf{D}|e_i\rangle = \lambda_i|e_i\rangle$$

 4. Compute the principal components $\mathbf{e}_i = \mathbf{X}|e_i\rangle$

 5. Compute Ω^{train} by $\Omega_i^{\text{train}} \leftarrow (\mathbf{e}_1^T \mathbf{x}_i^{\text{train}}, \mathbf{e}_2^T \mathbf{x}_i^{\text{train}})$ for $i = 1, 2$

 6. Compute Ω^{test} by $\Omega^{\text{test}} \leftarrow (\mathbf{e}_1^T \mathbf{x}^{\text{test}}, \mathbf{e}_2^T \mathbf{x}^{\text{test}})$

 7. Compute the 2-norm $r_i = \|\Omega^{\text{test}} - \Omega_i^{\text{train}}\|$ for $i = 1, 2$

 Find the minimal r_i , so that $\mathbf{X}_i^{\text{train}}$ is the closest to \mathbf{X}^{test}

 End While **Return** the found i^{th} images in the training dataset.

Table S1. Pseudocode of the CT image classification problem using qPCA.

group M , H_1 and H_2 were decoupled throughout all experiments. The internal Hamiltonian of this system is

$$\mathcal{H}_{\text{int}} = \sum_{j=1}^4 \pi \nu_j \sigma_z^j + \sum_{j < k, =1}^4 \frac{\pi}{2} J_{jk} \sigma_z^j \sigma_z^k, \quad (\text{S1})$$

where ν_j represents the chemical shift of the j th spin and J_{jk} denotes the J-coupling strength between spins j and k . The $^{13}\text{C}_1$ and $^{13}\text{C}_2$ are designated as the probe qubit and trial qubits, respectively, while $^{13}\text{C}_3$ and $^{13}\text{C}_4$ serve as copy qubits. For further reference, Figure S4 provides the molecular structure and NMR parameters of the sample.

Pseudo-pure state preparation. – All experiments were performed on a Bruker Ascend 600 MHz spectrometer (14.1 T) at room temperature. The thermal equilibrium state of the NMR system is a highly mixed state given by:

$$\rho_{\text{eq}} = \frac{I}{2^4} + \varepsilon \sum_{i=1}^4 \sigma_z^i \quad (\text{S2})$$

Here I represents the 16×16 identity matrix, and the polarization ε is in the order of 10^{-5} . Various methods exist for initialization, including the spatial average method, selective transition method, time spatial method, and cat-state method. In our experiment, we utilized the spatial average method to initialize the NMR system, employing the pulse sequence depicted in Figure S5. The rectangles in the circuit symbolize single-qubit rotations achieved through the application of radio-frequency (rf) pulses. The spin Hamiltonian

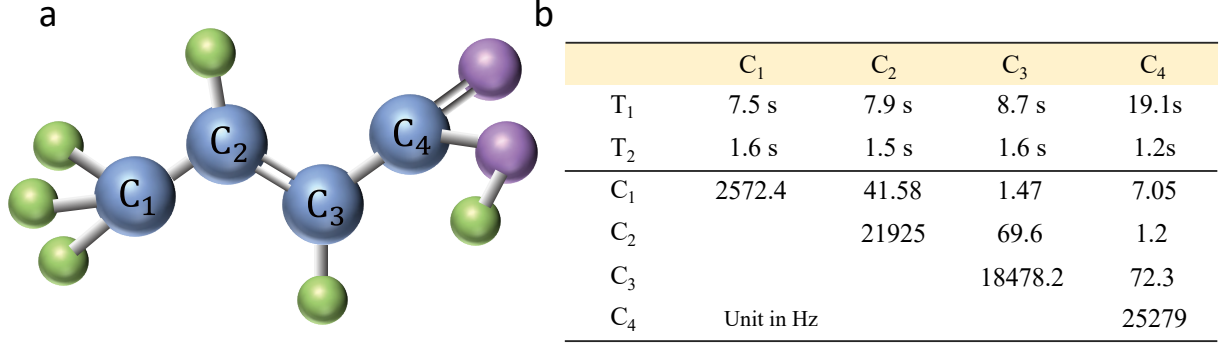


Figure S4. **a**, Molecular structure of ¹³C-labeled trans-crotonic acid. C₁, C₂, C₃ and C₄ are used as four qubits in the experiment. **b**, Molecular properties and the Hamiltonian parameters of the sample. Top part: the encoding and relaxation time T_2 of C₁ – C₄. Bottom part: the diagonal and off-diagonal elements give the chemical shifts ν_j and the scalar coupling strengths J_{jk} (in units of Hz), respectively.

of the rf pulse is given by:

$$H_{\text{rf}} = \pi\omega_1 \sum_{i=1}^4 (\sigma_x^i \cos(\omega t + \phi) + \sigma_y^i \sin(\omega t + \phi)). \quad (\text{S3})$$

Here ω_1 , ω_{rf} and ϕ denote the amplitude, frequency and phase of the rf pulse, respectively, while γ represents the gyromagnetic ratio of the ¹³C nuclei. The two-qubit gates are realized by using the scalar coupling between different spins. By employing this pulse sequence, the equilibrium state described in Eq. S2 is converted into a pseudo-pure state (PPS) given by

$$\rho_{\text{PPS}} = \frac{1 - \epsilon}{16} I + \epsilon |0000\rangle\langle 0000|. \quad (\text{S4})$$

Here, the polarization ϵ is also in the order of 10^{-5} . The large identity matrix does not evolve under any unitary propagator and cannot be observed in NMR. Therefore, the quantum system can be treated as the deviation part $|0000\rangle\langle 0000|$ up to a scale factor. In our experimental setup, we combined each segment of the quantum circuit, separated by three gradient pulses, into one unitary operation. We then utilized the optimal-control algorithm to search for the corresponding rf pulse. The shaped pulses used in the experiments had lengths of 2 ms, 20 ms, 15 ms and 8 ms, respectively. All the pulses exhibited fidelities over 99.5%.

Initialization. – To initialize the spin system for the qPCA algorithm, we begin with the PPS state and prepare the system in the initial state $\rho_0 = |0\rangle\langle 0| \otimes |+\rangle\langle +| \otimes \rho_D \otimes \rho_D$. The trial qubit C₂ can easily be prepared to the superposition state $|+\rangle$ using the Hadamard gate. However, loading the information of the training dataset into the quantum circuit, i.e., prepare C₃ and C₄ to ρ_D , requires an efficient implementation of quantum RAM, which is a highly challenging task beyond the scope of our experiment. For our purpose,

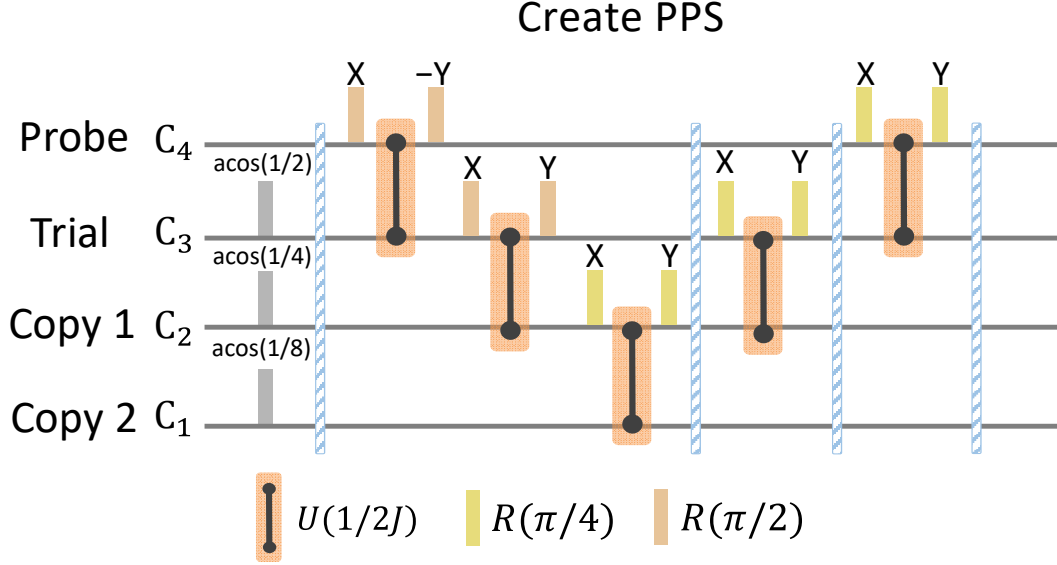


Figure S5. NMR pulse sequence to prepare the 4-qubit system to the PPS. Experimentally, we packed the unitary evolutions separated by the gradient pulses together, so that the PPS is prepared by 4 pulses searched by the optimal-control algorithm, which are 2 ms, 20 ms, 15 ms and 8 ms, respectively.

we assume that we have access to a sufficient number of pre-existing copies of ρ_D . In our experiment, we begin by applying a single-qubit rotation $R_y(\theta)$ to adjust the populations of the ground state and excited state to match the eigenstates of ρ_D , where $\theta = \arccos(\lambda_1 - \lambda_2)$. Following that, we utilize a gradient pulse to remove quantum coherence. Finally, a unitary evolution $U_1 = \{|e_1\rangle, |e_2\rangle\}$ is applied to transition from the computational basis to the eigenbasis. At this stage, the density matrix ρ_D is loaded onto the quantum circuit.

Preparation of ρ_k with PQC. – Our input state at the k -th step of the iterative qPCA is $\rho_{k-1} \otimes \rho_D \otimes \rho_D$, where ρ_{k-1} represents the final states of the two-qubit system from the previous $(k-1)$ -iteration. We have developed a parameterized quantum circuit (PQC) shown in Figure S6, which allows us to prepare ρ_{k-1} starting from the initial state $|00\rangle$. Unlike previous PQCs that primarily focused on unitary operations, this PQC incorporates non-unitary evolution by introducing two gradient pulses to eliminate quantum coherence in the system. The PQC comprises seven parameters, θ_j ($j = 1, 2, \dots, 7$). The goal now is to determine the optimal values for these parameters, θ_j .

To determine the optimal values, we input $\rho_0 = |00\rangle\langle 00|$ into the PQC, resulting in an output state denoted as:

$$\rho_{\text{out}} = u_3((u_2((u_1\rho_0u_1^\dagger) \odot G_z) * u_2^\dagger) \odot G_z) * u_3^\dagger, \quad (\text{S5})$$

where $G_z = I + |01\rangle\langle 10| + |10\rangle\langle 01|$ represents the gradient pulse. Here, $u_1 = u_1(\theta_1, \theta_2, \theta_3)$, $u_2 = R_y(\pi)$

and $u_3 = u_3(\theta_4, \theta_5, \theta_6, \theta_7)$ represent the propagators for the three segments of the quantum circuit separated by the two gradient pulses. The symbol \odot represents the element-wise multiplication. We experimentally reconstruct the final quantum state using tomography technology, and the objective function measures the fidelity between the output state ρ_{out} and the target state ρ_f :

$$f(\theta) = \frac{\text{Tr}(\rho_f \rho_{\text{out}}^\dagger)}{\sqrt{\text{Tr}(\rho_f^2) \text{Tr}(\rho_{\text{out}}^2)}}. \quad (\text{S6})$$

By maximizing this objective function across the parameter space using suitable optimization algorithms, we can obtain a good approximation of the parameter vector θ . When the target state ρ_{k-1} is highly entangled, modifications may be required in the circuit segment for realizing u_3 to enhance its entangling capability. Figure S6b demonstrates the convergence of the objective function in preparing an arbitrary mixed state in the form of

$$\rho_f = \begin{pmatrix} 0.1684 & 0.0122 & 0.0420 & 0.0816 \\ 0.0122 & 0.3316 & -0.0816 & -0.0176 \\ 0.0420 & -0.0816 & 0.1684 & -0.0122 \\ 0.0816 & -0.0176 & -0.0122 & 0.3316 \end{pmatrix}. \quad (\text{S7})$$

We observe that the maximum quantum state fidelity is achieved within 10 optimization iterations.

Measurement. – In a 4-qubit NMR quantum processor, the signal from each spin is typically split into 8 peaks, which arise from the couplings between different nuclei. In NMR spin dynamics, these peaks consist of both real and imaginary parts. The real part encodes the expectation values of the observable spin's Pauli matrix σ_x , while the imaginary part encodes the expectation values of σ_y . The remaining spins are measured in the computational basis. For example, considering the probe qubit in our experiment, the peaks correspond to the other three qubits being in the states $\{000, 001, 010, 011, 100, 101, 110, 111\}$. We can express $|0\rangle\langle 0|$ as $(I + \sigma_z)/2$ and $|1\rangle\langle 1|$ as $(I - \sigma_z)/2$. This allows us to measure the expectation values of single-quantum coherence operators involving σ_x or σ_y in the target qubit, and σ_z or I in the remaining qubits. Examples of such operators include $\sigma_x III$ and $\sigma_x II\sigma_z$. By applying readout pulses, we can measure other operators by converting them into single-quantum coherences. For instance, a readout pulse of $\exp(-i\sigma_x^1\pi/4)$ can be applied to measure Pauli observables like $\sigma_z III, \sigma_z II\sigma_z, \sigma_z I\sigma_z I, \sigma_z \sigma_z II$, and so on. This measurement mechanism plays a crucial role in measuring observables and performing quantum state tomography in NMR experiments.

As illustrated in the main text, the eigenvalues of the commuted covariance matrix D can be obtained by performing a Fourier transform on the x -magnetization of the first spin. The expression for the x -magnetization is given by: $M_x(t) = |a_1|^2 \cos(\lambda_1 t) + |a_2|^2 \cos(\lambda_2 t)$. To measure $M_x(t)$, we can read out

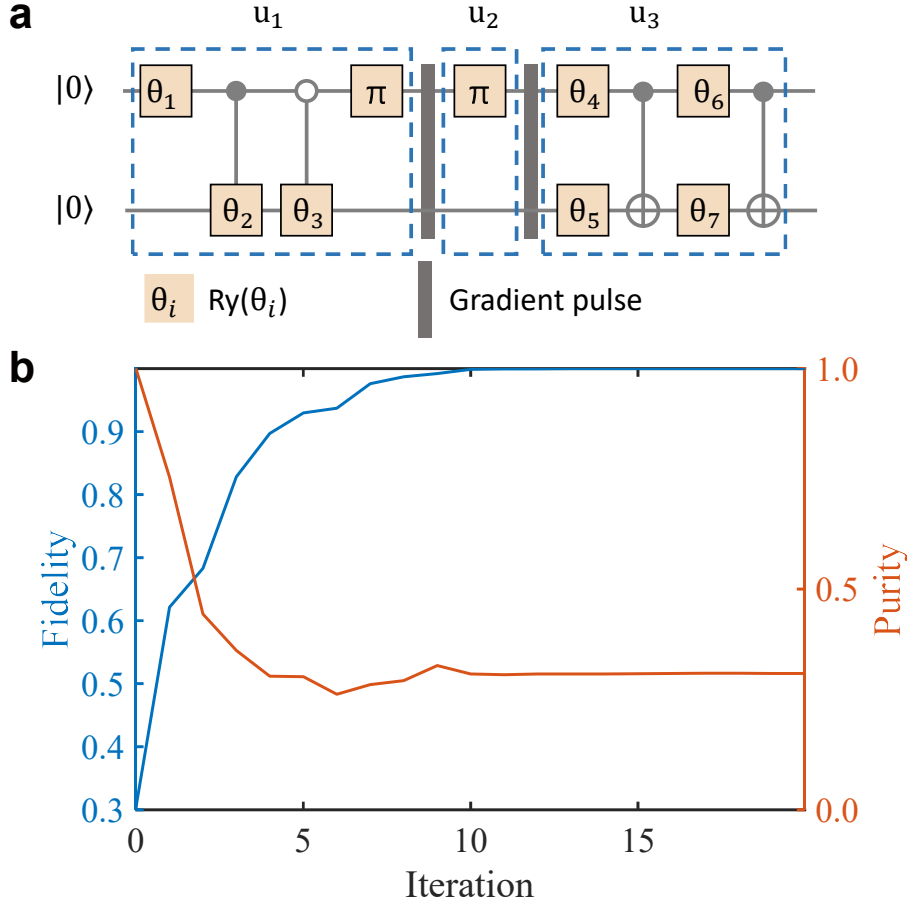


Figure S6. **a** The PQC used to initialize the probe-trial two-qubit subsystem to the mixed two-spin state ρ_{k-1} in the k -th step of the iterative qPCA. **b** The achieved maximum state fidelity and the purity with different iterations in the preparation of a randomized mixed state ρ_f given in Eq. (S7).

the observable $\sigma_x III$. It is important to note that $III = |000\rangle\langle 000| + |001\rangle\langle 001| + \dots + |111\rangle\langle 111|$. Therefore, $\sigma_x III$ represents the sum of the strengths of all the peaks of the probe qubit.

Quantum state tomography. – The experimental results of the eigenvectors are solely associated with the quantum state of the subspace spanned by the probe and trial qubits, as described in the main text. These results can be obtained by performing a full quantum state tomography on the 4-qubit system and subsequently post-processing the data by tracing out the two copy qubits. To expedite the experimental process, we directly measured the two-qubit subsystem. A quantum state can be decomposed in the complete Pauli basis $\sum_{i=1}^{2^n} \otimes \sigma_{0,x,y,z}^i$, where $\sigma_0 = I$, and σ_x , σ_y , and σ_z are the Pauli matrices. In comparison, conducting full quantum state tomography on the 4-qubit system requires 17 measurements, whereas 2-qubit tomography only necessitates 4 experiments, significantly reducing the experimental time. The quantum operations and the corresponding measured Pauli matrices are outlined in Table S2. The operations X and Y represent $X = \exp(-i\sigma_x\pi/4)$ and $Y = \exp(-i\sigma_y\pi/4)$, respectively.

Operations	Pauli Operators
II	XI, YI, XZ, YZ, IX, IY, ZX, ZY
XI	YX, YY, ZI, ZZ, XZ, XI
IX	XY, YY, IZ, ZZ, ZX, IX
YI	XX, XY, ZI, ZZ, YZ, YI

Table S2. Quantum operations to perform tomography on the probe and trial qubits. The operators X_i and Y_i represent the $\pi/2$ local operations on the i -th qubit. The right column shows the Pauli operators that can be obtained by current readout pulses.

Appendix D: Experimental results using different training datasets

To validate the effectiveness and stability of the protocol, we conducted experiments using four pairs of different positive and negative CT images as training sets to determine the eigenvalues and eigenvectors of the covariance matrix. While the main text only presented two sets of experimental results due to space limitations, we now present all the results from the four training sets in Figure S7a-d. In Figure S7, each part of the figure represents the following information: The first line displays the cropped CT images used for each training set in each group, consisting of one negative and one positive CT image. The second row shows the corresponding density matrix ρ_D , which is normalized from the commuted covariance matrix D . The third row presents the experimental transverse magnetization $M_x(t) = \sum_{k=1} |a_k|^2 \cos(\lambda_k t)$ obtained by reading out the off-diagonal elements of the probe qubit. The last row depicts the Fourier transformation results of $M_x(t)$ at the end of 200, 400, 600, 800, and 1000 iterations for each group. In the Fourier transformation results, the peak positions corresponding to the eigenvalues remain stable with different iteration numbers, while the half-width of the peaks noticeably decreases. This indicates that the eigenvalue estimation becomes more precise as the number of iterations increases. Thus, the results demonstrate that the protocol achieves higher accuracy with an increased number of iterations.

After obtaining the experimental eigenvectors $|e_i^{\text{exp}}\rangle$ of ρ_D through iterative qPCA, we can compute the weight vectors of the training CT images and testing CT images using two steps, both of which can be efficiently performed on a classical computer.

The first step is to compute the eigenvectors of the covariance matrix C . From the relationship $\rho_D |e_i\rangle = \lambda_i |e_i\rangle$ and $\rho_D = D/\text{Tr}(D)$, we have $D |e_i\rangle = \text{Tr}(D) \lambda_i |e_i\rangle$. This implies that $|e_i\rangle$ is also an eigenvector of the matrix D with the corresponding eigenvalue $\text{Tr}(D) \lambda_i$. By applying a matrix \mathbf{X} on both sides of the aforementioned eigenfunction and considering $\text{Tr}(D) = \text{Tr}(C)$, we obtain $C \mathbf{X} |e_i\rangle = \text{Tr}(C) \lambda_i \mathbf{X} |e_i\rangle$. Consequently, we can compute the two principal components of the covariance matrix as $\mathbf{e}_1 = \mathbf{X} |e_2\rangle$ and $\mathbf{e}_2 = \mathbf{X} |e_1\rangle$, which are two $49,400 \times 1$ vectors.

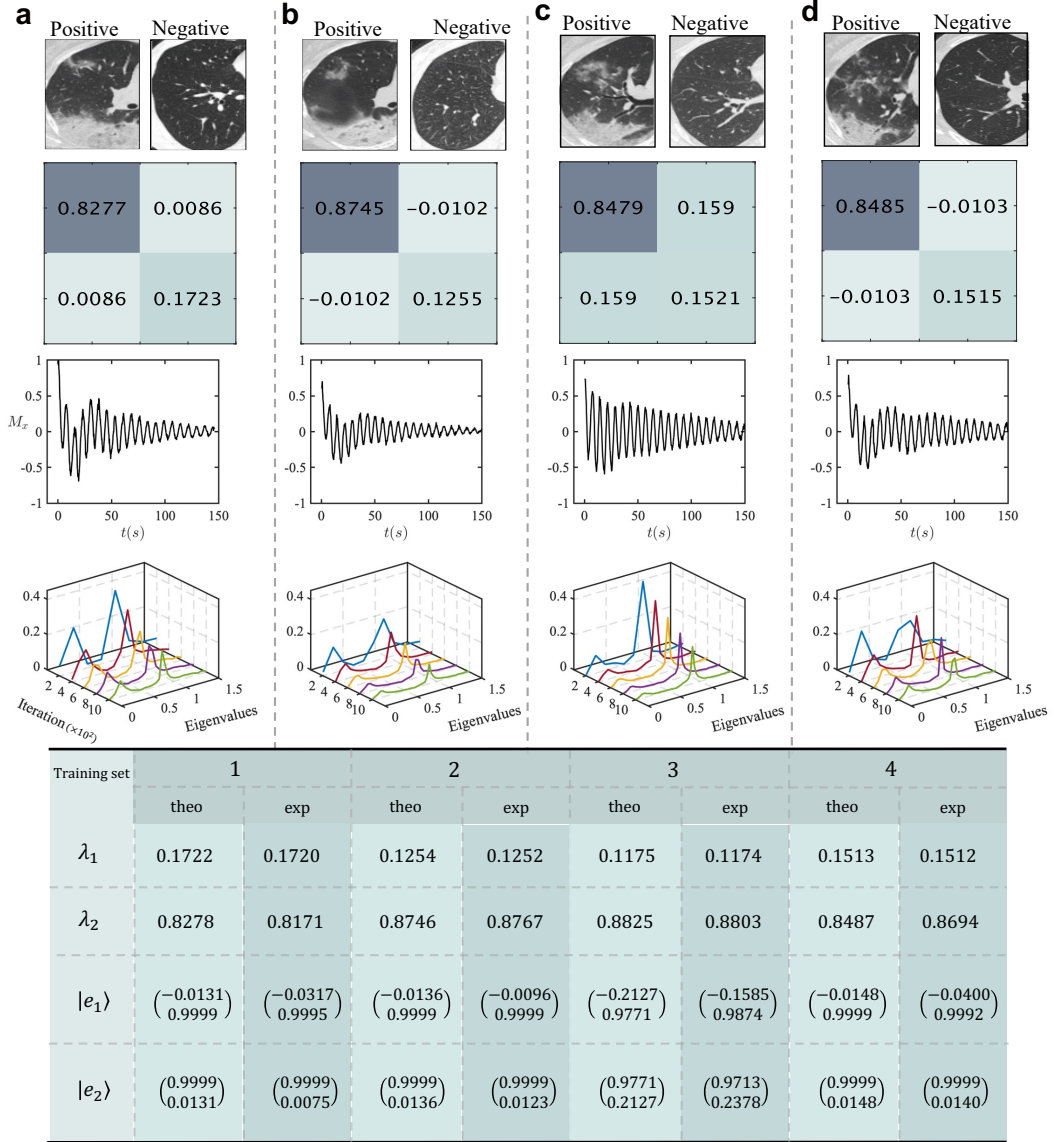


Figure S7. Experimental results of the eigenvalues and eigenstates of the 4 groups of training datasets (a-d). From the top to bottom are the CT images of the 4 training datasets, the corresponding commuted covariance matrix D , the experimental transverse magnetization as a function of the sampling time $M_x(t)$ and the Fourier transformation results of $M_x(t)$ at the end of different iterations. The table gives the experimental and theoretical results of the eigenvalues and eigenvectors of the 4 experiments.

The second step involves computing the weight vectors of the training and testing images by projecting the vectors $\mathbf{x}_i^{\text{train}}$ ($i = 1, 2$) or $\mathbf{x}_j^{\text{test}}$ ($j = 1, \dots, 78$) onto the two principal components \mathbf{e}_1 and \mathbf{e}_2 . This results in two-dimensional weight vectors for the testing and training CT images, given by $\omega_i^{\text{train}} = (\mathbf{e}_1^T \mathbf{x}_i^{\text{train}}, \mathbf{e}_2^T \mathbf{x}_i^{\text{train}})$ and $\omega_j^{\text{test}} = (\mathbf{e}_1^T \mathbf{x}_j^{\text{test}}, \mathbf{e}_2^T \mathbf{x}_j^{\text{test}})$, respectively. We visualize all the two weight vectors of the training and testing images on a plane, as shown in Figure S8. The results for Group 1 to Group 4 are depicted in **a**, **b**, **c**, and **d**, respectively. By employing the iterative qPCA process, most of

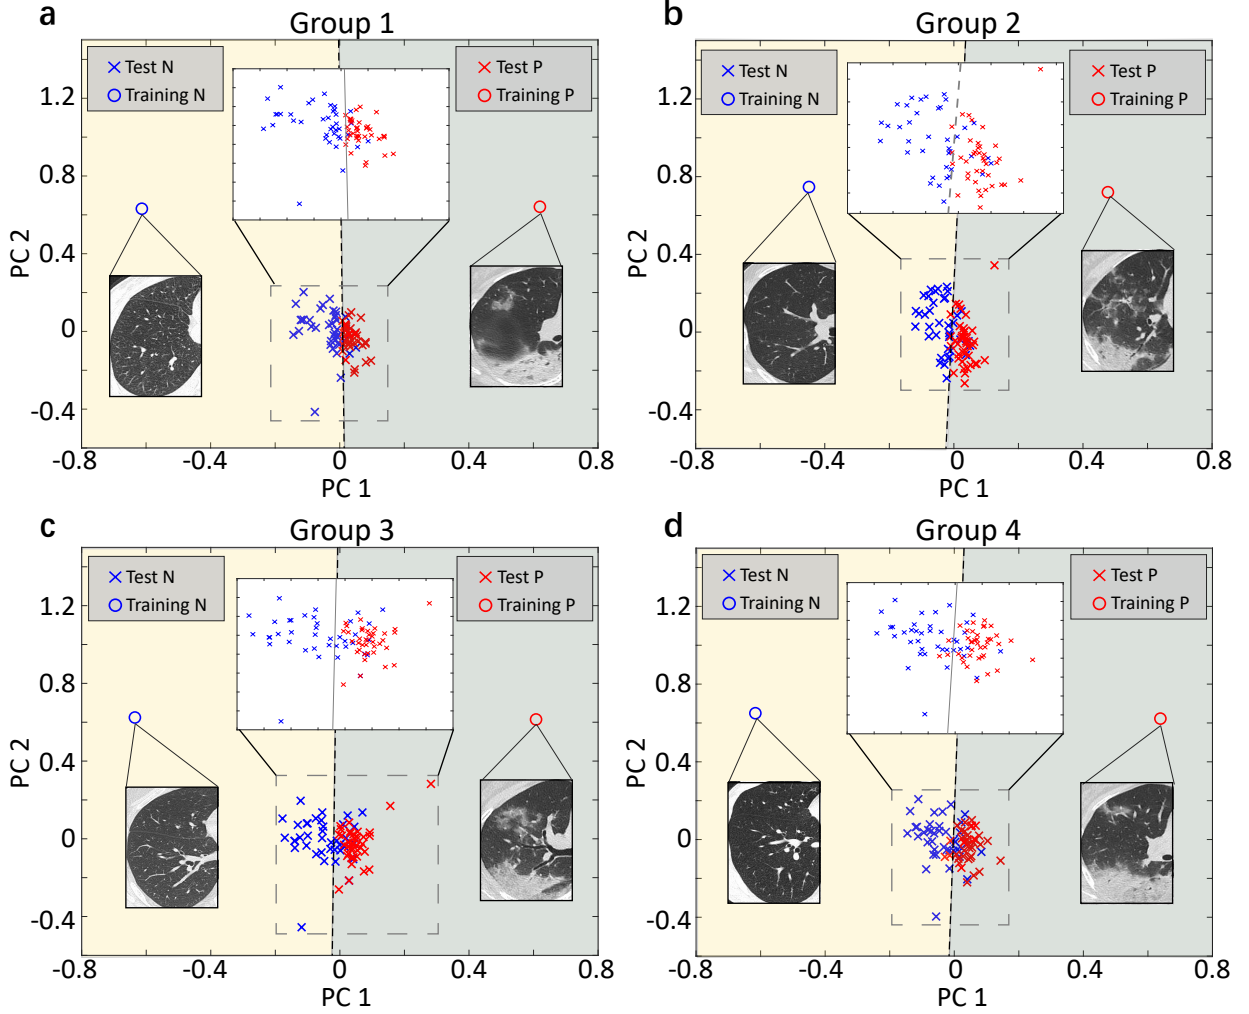


Figure S8. QPCA-aided medical classification results using four different training datasets are analyzed in Figure S7. After performing qPCA, each CT image in the dataset is projected onto a two-dimensional space spanned by the two experimental eigenvectors $|e_1^{\text{exp}}\rangle$ and $|e_2^{\text{exp}}\rangle$. The two training images are denoted by circles, while the test images are represented by crosses. The test dataset consists of 39 virus-negative images (blue) and 39 virus-positive images (red). The decision boundary is depicted by the dashed line. The success rates for each group are as follows: 92.31% for Group 1, 87.17% for Group 2, 80.83% for Group 3, and 82.36% for Group 4.

the images are successfully classified in all four groups, as demonstrated by the classification results of the weight vectors. In summary, after obtaining the experimental eigenvectors $|e_i^{\text{exp}}\rangle$ of ρ_D , we compute the weight vectors of the training and test CT images in two steps. This process can be efficiently executed on a classical computer, and the resulting weight vectors successfully classify the images in all four groups using iterative qPCA.

Appendix E: Complexity analysis of the algorithm

The multiplication of a $V \times L$ matrix and an $L \times W$ matrix has a complexity of $O(VLW)$. Consequently, the complexity for computing the covariance matrix is $O(m_1^2 d)$, where d is the dimensionality of the grayscale vectors, and m_1 and m_2 represents the number of training images and test images, respectively. Similarly, projecting the grayscale vectors onto the m_1 -dimensional principal components has a complexity of $O(m_1(m_1 + m_2)d)$. Since the dimensionality, d , is typically constant, the overall complexity can be approximated as $O(m_1(m_1 + m_2))$. This analysis demonstrates that the computation of weight vectors for the images can be efficiently performed with polynomial complexity. Thus, the process is computationally efficient.