



OPEN

Quantum annealing feature selection on light-weight medical image datasets

Merlin A. Nau¹✉, Luca A. Nutricati^{2,3}, Bruno Camino⁴, Paul A. Warburton^{2,5} & Andreas K. Maier¹

We investigate the use of quantum computing algorithms on real quantum hardware to tackle the computationally intensive task of feature selection for light-weight medical image datasets. Feature selection is often formulated as a k of n selection problem, where the complexity grows binomially with increasing k and n . Quantum computers, particularly quantum annealers, are well-suited for such problems, which may offer advantages under certain problem formulations. We present a method to solve larger feature selection instances than previously demonstrated on commercial quantum annealers. Our approach combines a linear Ising penalty mechanism with subsampling and thresholding techniques to enhance scalability. The method is tested in a toy problem where feature selection identifies pixel masks used to reconstruct small-scale medical images. We compare our approach against a range of feature selection strategies, including randomized baselines, classical supervised and unsupervised methods, combinatorial optimization via classical and quantum solvers, and learning-based feature representations. The results indicate that quantum annealing-based feature selection is effective for this simplified use case, demonstrating its potential in high-dimensional optimization tasks. However, its applicability to broader, real-world problems remains uncertain, given the current limitations of quantum computing hardware. While learned feature representations such as autoencoders achieve superior reconstruction performance, they do not offer the same level of interpretability or direct control over input feature selection as our approach.

Keywords Medical imaging, Machine learning, Quantum computing, Quantum annealing, Image reconstruction

Medical imaging plays a crucial role in modern clinical practice, providing essential insights for diagnosis, treatment planning, and monitoring. However, the increasing complexity and volume of imaging data present significant challenges for manual analysis. Machine learning (ML), particularly deep learning (DL), has emerged as a transformative tool to automate tasks such as disease classification, segmentation, and outcome prediction^{1,2}. By learning complex representations from large datasets, DL methods have substantially improved diagnostic performance. Nevertheless, most ML models, especially those involving high-dimensional inputs like images, scale computationally with both the number of samples and the number of features, leading to high resource demands.

Feature selection (FS) aims to mitigate these challenges by identifying a subset of relevant features that contribute most significantly to a target variable. Effective FS reduces input dimensionality, improves model interpretability, lowers computational complexity, and often enhances generalization performance³. In medical imaging, FS can also have direct clinical applications, such as minimizing radiation exposure by identifying the most informative measurements for acquisition^{4,5}.

Traditional FS methods are categorized into *filter methods*, which evaluate features based on statistical metrics like mutual information (MI), and *embedded methods*, which incorporate FS into model training. In particular, MI quantifies the dependence between variables without assuming linearity, making it a widely used and theoretically justified criterion for supervised FS^{6,7}.

¹Pattern Recognition Lab, Friedrich-Alexander-Universität Erlangen-Nürnberg, 91052 Erlangen, Germany.

²London Centre of Nanotechnology, University College London, London WC1H 0AH, UK. ³Rudolf Peierls Centre for Theoretical Physics, University of Oxford, Oxford OX1 3PU, UK. ⁴Department of Chemistry, University College London, London WC1H 0AJ, UK. ⁵Department of Electronic and Electrical Engineering, University College London, London WC1E 7JE, UK. ✉email: merlin.nau@fau.de

Identifying the optimal subset of features from a set of n candidates is combinatorially complex, requiring evaluation of $\binom{n}{k}$ possible subsets for k selected features. To manage this intractability, FS can be reformulated as a quadratic unconstrained binary optimization (QUBO) problem, enabling the use of powerful combinatorial optimization techniques.

In our study, we systematically compare a range of feature selection strategies. As a baseline, we include random feature selection, which provides a lower-bound for performance. To leverage the spatial structure inherent in imaging data, we implement a sampling-based method selecting pixels equally spaced throughout the image. For information-theoretic selection, we formulate a MI-based QUBO and solve it using the classical heuristic solver `qbsolv`. Additionally, we consider *unsupervised* dimensionality reduction via Sparse Principal Component Analysis (SPCA)⁸, which aims to find sparse latent representations of the input data, and *supervised* feature selection using the Lasso⁹, which identifies features predictive of the target through ℓ_1 -regularization. Finally, we evaluate an autoencoder (AE) approach, where a bottleneck neural network is trained to learn a compact representation of the input without supervision. Together, these methods span randomized, statistical, sparsity-based, and learned paradigms, providing a comprehensive benchmark for evaluating quantum annealing-based FS approaches in medical imaging.

Notably, QUBO formulations are well-suited for emerging quantum computing methods, including adiabatic quantum computing (AQC), which leverage quantum phenomena to potentially explore large solution spaces more efficiently than classical approaches^{10,11}. Quantum annealing (QA), a practical realization of AQC, solves QUBO problems by evolving a quantum system from an easily prepared ground state towards the ground state of a complex problem Hamiltonian^{12–14}. Commercial quantum annealers, such as those developed by D-Wave Systems¹⁵, offer access to hundreds to thousands of qubits connected in sparse topologies suitable for optimization tasks. Although current quantum devices fall under the category of noisy intermediate-scale quantum (NISQ) hardware¹⁶ and remain limited in qubit count and precision, they provide an opportunity to explore QUBO-based FS at non-trivial problem scales. While QUBOs can, in principle, also be solved on gate-based quantum computers using the Quantum Approximate Optimization Algorithm (QAOA)¹⁷, such approaches remain infeasible at present due to limited qubit availability and hardware noise and are not investigated in this manuscript.

Recent work has investigated MI-based QUBO models for FS, executed on simulated and real quantum hardware^{18,19}. D-Wave's tutorial¹⁸ introduced a basic framework for MI-based FS on toy datasets. Building on this, Muecke et al.¹⁹ demonstrated MI-QUBO FS for MNIST and synthetic datasets, including image compression tasks. However, due to hardware limitations, these studies were restricted to very small problem sizes ($n \leq 34$) on quantum devices. A recent study by Hellstern et al.²⁰ also investigates quantum and classical methods for feature selection formulated as a QUBO problem, with a focus on determining the optimal number k of selected features. In contrast, our work assumes a fixed k and investigates how well various methods, including QA, can select informative subsets under this constraint for use in compression and reconstruction tasks. Further efforts have extended QA-based FS to areas such as hyperspectral image classification²¹, RNA sequencing data analysis²², recommendation systems²³, and radiomics feature selection²⁴. Beyond FS, QA has also been explored in medical imaging applications including tomographic reconstruction^{25–27}, segmentation²⁸, and super-resolution tasks²⁹.

In this work, we explore how QA can be applied to FS in medical imaging using light-weight datasets. Our focus is on demonstrating FS at larger scales than previously achieved on real quantum hardware, by carefully adapting the QUBO construction to current architectural limitations. Specifically, our experiments are conducted on the MedMNIST benchmark³⁰, which comprises standardized 28×28 pixel images across multiple imaging modalities. We treat each pixel as an individual feature, recognizing that pixels are not optimal descriptors but offering a manageable and standardized starting point for proof-of-concept experiments.

Our contributions are threefold: First, we implement MI-based QUBO FS on six light-weight medical imaging datasets and demonstrate feasibility on a classical solver. Second, we introduce hardware-aware adaptations including subsampling, thresholding, and a sparsity-preserving linear Ising penalty to enable embedding of larger QUBO instances on a QA system. Finally, we demonstrate the utility of the selected features by training a convolutional encoder for lossy image compression, illustrating potential applications of FS beyond classification. Although constrained by current hardware limitations, this study provides a proof of concept for applying quantum optimization methods to FS tasks in medical imaging. While this study focuses on a simplified toy problem, the approach could inform potential clinical applications such as radiation dose reduction during image acquisition^{4,5} and improved diagnostic interpretability by identifying key image regions.

Quantum annealing

We focus on the particular implementation of AQC known as QA, in which the problem to be solved is mapped to the minimization of an Ising Hamiltonian, in order to be embeddable in a quantum annealer. FS is also suitable for this technique since it can be formulated as a QUBO problem, which in turn can easily be mapped to an Ising system. QA operates on the fundamental concept of quantum tunneling and the quantum adiabatic theorem to explore the vast solution space of a problem in order to find the optimal solution of the Ising Hamiltonian, which corresponds to a global minimum of the optimization function. It is important to note that while QA has shown promise in a variety of fields^{31–38}, it is not a universal quantum computing approach like gate-based quantum computers. Indeed, future gate-based quantum computers are versatile and capable of performing a wide range of computations, including optimization of QUBOs via algorithms such as the QAOA¹⁷. In contrast, quantum annealers are specialized devices tailored for optimization using adiabatic quantum evolution.

This is not a limitation for our analysis, as our problem can be cast into a form embeddable in a quantum annealer. The QA device we use in this study is developed by D-Wave Systems, featuring specialized hardware that generates and sustains the necessary quantum states. These devices use superconducting qubits and magnetic

fields to create controlled quantum environments where the annealing process takes place. In particular, we shall perform our analysis on the Advantage_system4.1 architecture¹⁵: this annealer contains 5627 qubits, connected in a Pegasus structure, but only has a total of 40279 couplings between them. This topology is depicted in Fig. 2.

The full Hamiltonian in QA comprises a time-dependent mixture of a known driver Hamiltonian, responsible for inducing quantum fluctuations, and a problem Hamiltonian, which encodes the classical Ising objective to be minimized. The original idea behind QA (and AQC more generally) is to begin in the ground state of the trivial system and adiabatically replace the trivial Hamiltonian with the problem Hamiltonian, while remaining in the ground state throughout. In particular, the mixed Hamiltonian takes the following form:

$$\mathcal{H} = A(s) \sum_i \sigma_i^x + B(s) \left(\sum_{ij} J_{ij} \sigma_i^z \sigma_j^z + \sum_i h_i \sigma_i^z \right), \quad (1)$$

where ij label the qubits, σ_i^z are the z -spin Pauli matrices, and σ_i^x are the transverse field components, while the couplings h_i and J_{ij} between the qubits are set and kept constant. These couplings define the classical cost function and the part of the Hamiltonian multiplied by $B(s)$. The parameter $s(t)$ (with t being time) is a user-defined control-parameter that can be adjusted, while $A(s)$ and $B(s)$ describe the consequent change in the quantum characteristics of the annealer. The network of qubits starts in a global superposition over all possible classical states and, as s approaches a value of 1, the system localises into a single classical state once a measurement of σ_i^z on all sites has been performed. The anneal schedule increases linearly with time, with $s(0) = 0$ and $s(t_f) = 1$, where t_f is the total annealing time. At early anneal times ($s \approx 0$), the system is governed by the transverse field $A(s)$, promoting quantum superposition and tunneling. As s approaches 1, the system transitions to the classical problem Hamiltonian. To solve the optimization problem, the problem is encoded into the classical component of the annealing Hamiltonian, specifically, the term multiplied by $B(s)$ in Eq. (1). This term, referred to as the problem Hamiltonian H_P , defines a classical energy landscape in the form of a generalized Ising model:

$$H_P = \sum_{ij} J_{ij} \mu_i \mu_j + \sum_i h_i \mu_i, \quad (2)$$

where $\mu_i \in \{-1/2, 1/2\}$ are the eigenvalues of the spin operator σ_i^z . The coefficients J_{ij} and h_i encode the binary quadratic objective function as Ising couplings and local fields, respectively. If the process remains adiabatic, the final state of the system corresponds to a lowest-energy configuration solving the original problem. Since real-world QA may not always reach the global minimum due to noise and non-adiabatic effects, the process is typically repeated many times to gather a distribution of solutions, from which the best one is selected.

Methods

Mutual information feature selection QUBO

The goal for our FS method is to define an optimization objective tailored to a QUBO and complying with the constraints introduced by the current hardware. Consider an image dataset consisting of square, two-dimensional images with width W , paired with their labels. The training dataset can be represented as $D = \{(\mathbf{X}^i, \mathbf{y}^i)\}_{i \in [N]}$, where each $\mathbf{X}^i \in \mathbb{R}^{W \times W}$ is an image, and $y^i \in \mathbb{N}$ is its corresponding class label. By flattening the images, the dataset transforms into $D_f = \{(\mathbf{x}^i, \mathbf{y}^i)\}_{i \in [N]}$, where $\mathbf{x}^i \in \mathbb{R}^n$ represents the n -dimensional vectorized data, and y^i remains the class label.

While the class label y^i is not used in the subsequent image reconstruction process, it plays a crucial role in feature selection. Specifically, we propose a model-agnostic FS task, where we want to maximize the mutual information (MI) of the features with the class label to determine feature relevance. The selected features are then used the image reconstruction task, independent of y^i . In our simplified example, we will treat the image pixels as features and formulate the pixel selection problem as a QUBO model:

$$\min(f_Q(\hat{\mathbf{x}})) = \hat{\mathbf{x}}^T \mathbf{Q} \hat{\mathbf{x}} = \sum_{i=1}^n Q_{i,i} \hat{x}_i + \sum_{i=1}^n \sum_{j>i}^n Q_{i,j} \hat{x}_i \hat{x}_j. \quad (3)$$

Here, \hat{x}_i indicates whether a feature is selected (1) or not (0), where $\hat{\mathbf{x}} \in \{0, 1\}^n$. The linear terms originate from the binary nature of $\hat{x}_i^2 = \hat{x}_i$. In turn, our matrix $\mathbf{Q} \in \mathbb{R}^{n \times n}$ describes our optimization problem. In the subsequent steps, we show how we construct the QUBO matrix from the information contained in the image datasets. Intuitively, feature x_i is more likely to be selected if $Q_{i,i}$ is low. Similarly, we can increase the chances of choosing feature x_i and feature x_j together if the $Q_{i,j}$ term is small. The diagonal elements $Q_{i,i}$ of the matrix \mathbf{Q} encode the importance of each feature and are defined as the negative MI between feature x_i and the class label y . The off-diagonal elements $Q_{i,j}$ (for $i \neq j$) represent the redundancy between features, calculated as the MI between features x_i and x_j . Together, these define a binary quadratic objective function that balances feature relevance and redundancy. This model closely resembles the form of an Ising problem, and in fact one can translate a QUBO into an Ising problem by mapping $x_i \rightarrow \mu_i + \frac{1}{2}$, where μ_i is defined as in Eq. (2).

In FS, MI is widely used as it captures nonlinear dependencies and does not assume specific distributions^{6,7}. We closely follow the MI FS QUBO described in^{18,19}. To calculate MI, we estimate the joint and marginal probabilities of the features and labels. For efficient computation with continuous, real-valued features, this

requires discretization. We divide each feature dimension into $n_{bins} = 20$ bins using quantiles, assigning each feature value to its corresponding bin. Labels, being discrete, do not require binning. This discretization and probability estimation allows for efficient computation of MI between features and labels. Please refer to¹⁹ for a detailed explanation of this method. From the discretized dataset $\hat{D} = \{(\mathbf{b}^i, \mathbf{y}^i)\}_{i \in [N]}$, the empirical joint and marginal probabilities are defined as follows:

$$\hat{p}_{X,Y}(b, y) = \frac{1}{N} \sum_{j=1}^N \mathbb{1}(b_j = b \wedge y_j = y) \quad (4)$$

$$\hat{p}_{X_i, X_j}(b_i, b_j) = \sum_y \sum_{b_k \neq i, j} \hat{p}_{X,Y}(b, y) \quad \text{and} \quad \hat{p}_{X_i, Y}(b_i, y) = \sum_{b_k \neq i} \hat{p}_{X,Y}(b, y) \quad (5)$$

$$\hat{p}_Y(y) = \sum_{b_k} \hat{p}(b, y) \quad \text{and} \quad \hat{p}_{X_i}(b_i) = \sum_y \sum_{b_k \neq i} \hat{p}_{X,Y}(b, y) \quad (6)$$

We initialize our linear QUBO terms with the negative MI of a feature with its class label, which we label as importance:

$$Q_{i,i} = -I_{i,i} = -I(x_i, y) \approx \sum_b \sum_y \hat{p}_{X_i, Y}(b, y) \log \frac{\hat{p}_{X_i, Y}(b, y)}{\hat{p}_{X_i}(b) \hat{p}_Y(y)}. \quad (7)$$

We want to avoid choosing features that share a lot of information. This is achieved by populating the off diagonal term $Q_{i,j}$ with the MI between feature x_i and x_j . A high MI value for $Q_{i,j}$ will increase the energy of the solution that selects both the x_i and x_j features decreasing the probability of it being returned by the annealer.

$$Q_{i,j} = R_{i,j} = R(x_i, x_j) \approx \sum_b \sum_{b'} \hat{p}_{X_i, X_j}(b, b') \log \frac{\hat{p}_{X_i, X_j}(b, b')}{\hat{p}_{X_i}(b) \hat{p}_{X_j}(b')}. \quad (8)$$

Finally, we have to enforce our constraint of choosing k of n features. To enforce such constraint, an option is to use a quadratic penalty $\alpha(\sum_{i=1}^n \hat{x}_i - k)^2$, where α must be tuned to weight the constraint appropriately. The quadratic constraint in the form of a QUBO \mathbf{C} is constructed by $C_{i,i} = 1 - 2k$ and $C_{i,j} = -2n + 2$, resulting in a fully connected problem graph. We construct our FS QUBO problem by additively combining the individual matrices: $\mathbf{Q} = -\mathbf{I} + \mathbf{R} + \alpha\mathbf{C}$, where \mathbf{I} is the diagonal importance matrix, \mathbf{R} is the redundancy matrix filling the quadratic couplings, and \mathbf{C} enforces the feature count constraint. The full procedure of creating the QUBO is depicted in Fig. 3 and an illustration of the quadratic constraint is shown in Fig. 4.

Mapping the problem to quantum hardware

In the previous section, we described how to construct a MI-based QUBO model for feature selection. While this formulation can be solved using classical algorithms such as simulated annealing or tabu search, executing it on a quantum annealer presents additional challenges due to the limited connectivity of current hardware. Specifically, although modern quantum annealers such as D-Wave's Advantage system feature thousands of physical qubits, their underlying Pegasus topology imposes constraints on qubit-to-qubit connectivity, necessitating the use of embedding strategies. These embeddings often require multiple physical qubits to represent a single logical variable, creating an overhead that scales with the problem's connectivity. To adapt our problem to the hardware, we applied a series of structured reductions aimed at decreasing both the dimensionality and the connectivity of the QUBO graph:

First, we partition the image into non-overlapping 2×2 neighborhoods and select the pixel with the highest MI with the class label from each block. This reduces the number of features from $28 \times 28 = 784$ to $14 \times 14 = 196$ while preserving local spatial diversity and maximizing information retention. The block size of 2×2 was chosen as the smallest non-overlapping subsampling unit compatible with the original 28×28 image size. It offers a trade-off between granularity and QUBO tractability: larger blocks would further reduce dimensionality but risk discarding important local information.

Despite the dimensionality reduction, the resulting QUBO remains densely connected due to the presence of pairwise redundancy terms between features. To address this, we apply a thresholding strategy that discards weak quadratic couplings, defined as those with low MI between feature pairs. Specifically, we retain the top 2000 largest $Q_{i,j}$ terms in absolute value. This threshold level was empirically determined to balance two competing objectives: maintaining enough pairwise structure to preserve problem fidelity, while reducing graph density to ensure embeddability within the limited connectivity of the hardware. The resulting sparsified QUBO retains essential interactions while substantially reducing embedding overhead.

Rather than using a fully connected quadratic penalty term to enforce the cardinality constraint $\sum_i \hat{x}_i = k$, which would reintroduce dense connectivity, we use a sparsity-preserving linear penalty approach. This introduces an adjustable weight α_l on the diagonal of the QUBO, as described in the following section, and allows for consistent enforcement of the feature count constraint with reduced connectivity overhead. Together, these steps result in a QUBO instance that can be embedded in the D-Wave Advantage system4.1 architecture. The final annealing parameters, including annealing time, number of reads, number of physical qubits and average chain length are described in the experimental section.

Linear Ising penalties

To limit the connectivity of the QUBO model while still enforcing the constraint, we propose using a linear penalty term, $|\sum_{i=1}^n \hat{x}_i - k|$, similar to the approach presented in^{39,40}. This linear penalty, denoted as \mathbf{L} , introduces an offset α_l on the diagonal of the QUBO matrix. The parameter α_l is tuned to achieve the desired Hamming weight of the solution vector $\hat{\mathbf{x}}$, where the Hamming weight corresponds to the number of selected features in the solution.

For each dataset, we tune the linear penalty coefficient α_l through a deterministic search procedure to ensure that the QUBO solution satisfies the hard constraint $\sum_i x_i = k$. Starting from a small initial value, α_l is incrementally increased until the returned solution has exactly k selected features. The resulting values are dataset-specific and remain fixed for all subsequent experiments.

When constructing the QUBO we substitute \mathbf{C} for \mathbf{L} : $\mathbf{Q} = -\mathbf{I} + \mathbf{R} + \alpha_l \mathbf{L}$. An overview and visual comparison of the constraint creation process for the QUBO is provided in Fig. 4. While the linear penalty constraint may be less effective for certain problem instances, our experiments show it consistently enforces the desired behavior.

Reconstruction decoder

Our experiment focuses on reconstructing images from the selected subset of pixels. Once the features are extracted for each dataset, we train a convolutional decoder to reconstruct the original image from the selected pixels. An illustration of the procedure is displayed in Fig. 1. The reconstruction network consists of a linear layer followed by two two-dimensional transposed convolutional layers, each followed by a ReLU activation function and then a sigmoid activation function.

Datasets

For our experiments, we used the MedMNIST dataset³⁰, which contains 18 standardized datasets used for biomedical image classification. The collections comprise 12 two-dimensional and 6 three-dimensional datasets. The collection contains data scales from a few hundred to 100,000 and binary and multi-class classification tasks. In particular, the dataset collection should facilitate light-weight machine learning research in medical imaging without directly facing clinical challenges. Due to the size restrictions of our QA device, we only analyze the two-dimensional grayscale image datasets. We note that ChestMNIST, OCTMNIST, PneumoniaMNIST and BreastMNIST consist of images that are semi-registered. This plays an important role when discussing pixels as feature descriptors. The datasets analyzed are summarized in Table 1.

Results

Reconstruction experiment

We compare different FS methods for selecting subsets of features in our reconstruction experiments. The FS methods evaluated include: (1) a random sampling approach that selects k features at random, (2) a subsampled approach that evenly spreads the pixels across the image in a grid-like fashion, (3) supervised feature selection using Lasso regression, where features are selected based on their regression coefficients with respect to the target labels, (4) the full QUBO formulation solved using the classical tabu-search solver `qbsolv`⁴¹, (5) our proposed QA method, which solves a sparsified QUBO adapted for hardware constraints, (6) unsupervised feature compression using SPCA, which generates a sparse representation of the data without supervision, and (7) a learned feature extractor, specifically an AE with a latent dimension of $k = 25$, whose encoder resembles the presented decoder architecture.

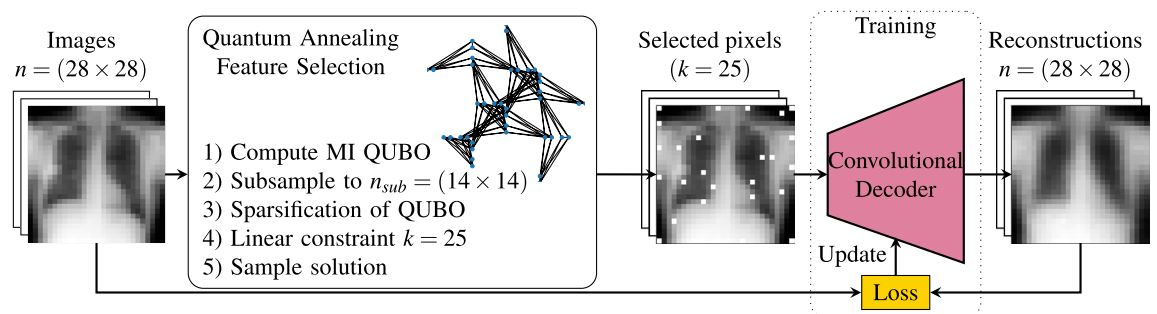


Fig. 1. Illustration of the feature selection process using quantum annealing to extract pixels and train a convolutional decoder for reconstruction. The image dataset is used to compute mutual information (MI) between pixels and class labels (for importance) and between pixels (for redundancy). These statistics define the linear and quadratic terms of the quadratic unconstrained binary optimization (QUBO) formulation. To make the problem compatible with quantum hardware, the QUBO is downsampled spatially and sparsified to reduce connectivity. A soft linear constraint enforces the selection of a fixed number of features, and the QUBO is then submitted to the quantum annealer. The selected pixels are treated as a compressed representation, which is used to train a decoder for image reconstruction.




Fig. 2. (Left) A part of the Pegasus topology implemented in the quantum annealing D-Wave Advantage_system4.1 architecture, where qubits (blue circles) are connected with couplings (black lines) to a maximum of 15 other qubits. (Right) General example of anneal schedule parameters, where $A(s)$ and $B(s)$ scale the transverse field and Ising contributions, respectively. These coefficients are functions of the parameter $s \in [0, 1]$, which depends on the physical time t . In specific hardware implementations, like the D-Wave Advantage_system4.1, these parameters may take different forms.

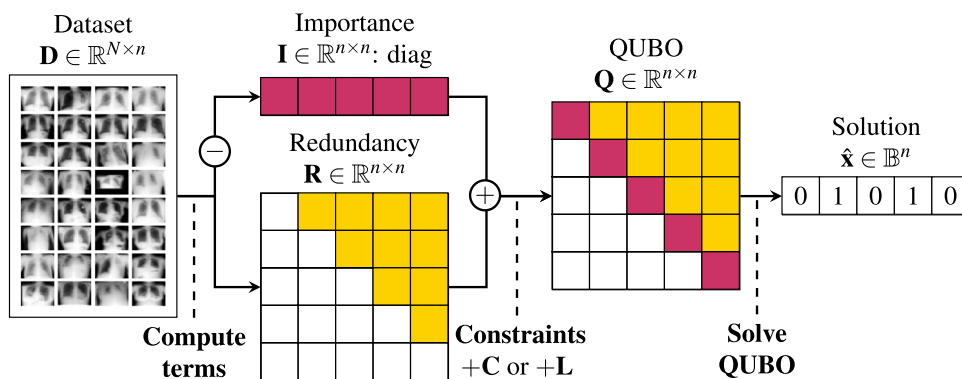


Fig. 3. Feature selection pipeline: Images are flattened to compute the importance and redundancy terms, which are combined into the QUBO. The k of n constraint is enforced via a linear penalty or a quadratic constraint (Fig. 4). Then, the QUBO is solved using classical or quantum solvers.

Beyond these approaches, numerous FS techniques exist, including random forest-based selection⁴², simulated annealing-based selection⁴³ or exact solvers like Gurobi⁴⁴. While these were not explicitly compared in our experiments, they represent alternative strategies for feature selection that may be explored in future work.

We trained a convolutional encoder, using feature sets of size $k = 25$, that takes the selected pixels as input to reconstruct the original image. The encoder was trained using the Adam optimizer with a learning rate of 0.001 for 20 epochs, minimizing mean squared error (MSE) loss. Reconstruction performance was validated first on the MNIST dataset, showing results consistent with those reported by Muecke et al.¹⁹. The experiments were then extended to the MedMNIST dataset. Each selection and training process was repeated five times to calculate the mean and standard deviation. The reconstruction results, expressed in terms of MSE on the test set, are summarized in Table 2. Visual examples of reconstructed images for each dataset using QA-selected pixels are shown in Fig. 5.

Simulation and hardware experiments

In the classical simulation setup, we performed FS on the full QUBO of size 784×784 , incorporating a quadratic constraint. The QUBO was solved using the `qbsolv` algorithm, which partitions the problem into smaller subproblems. From the generated sampleset, the solution with the lowest energy was selected. The runtime of `qbsolv` on the full size QUBO was 5.6 ± 0.56 s.

To validate the solutions on real quantum hardware, we executed the experiments on the D-Wave Advantage_system4.1 through LeapTM using the associated Ocean Python API⁴⁵. As discussed above, we used a series of steps to reduce the size and connectivity of the QUBO model. The original QUBO was defined on a 28×28 pixel grid, resulting in a size of 784×784 . To downscale this, we used a subsampling strategy,

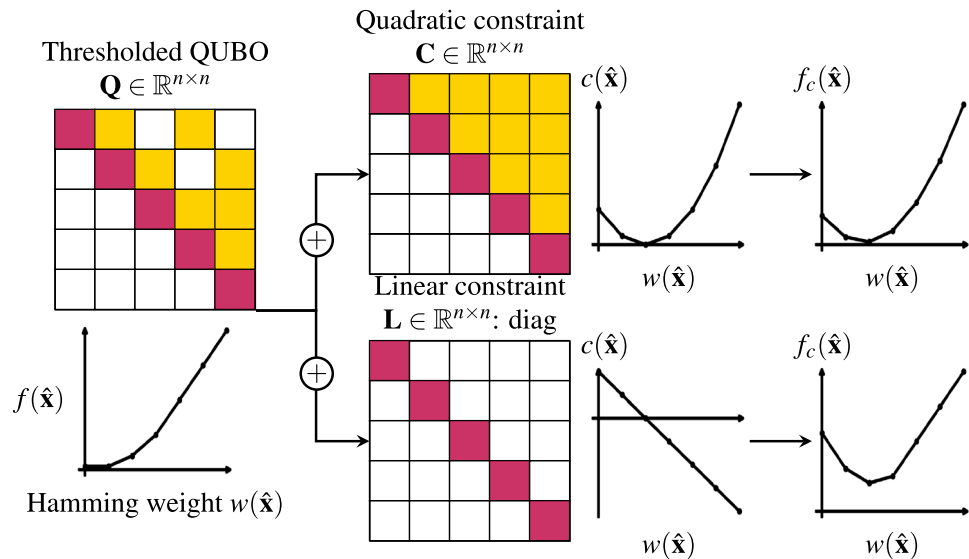


Fig. 4. Illustration of the conventional quadratic constraint to enforce selecting k of n features, which is infeasible due to limited connectivity on the annealer. When dealing with a sparsified QUBO, we propose a linear Ising penalty to enforce the constraint. The QUBO and its constraint are shown, along with a plot of the optimization energy (y-axis) against the Hamming weight (x-axis), indicating how many features are selected $\sum_i \hat{x}_i$.

Dataset (28 × 28)	Modality	Tasks (# Classes)	# Samples	Train/Test
ChestMNIST	Chest X-ray	Binary (2)	112,120	78,468/22,433
OCTMNIST	Retinal OCT	Multi (4)	109,309	97,477/1,000
PneumoniaMNIST	Chest X-ray	Binary (2)	5,856	4,708/624
BreastMNIST	Breast ultrasound	Binary (2)	780	546/156
TissueMNIST	Microscope	Multi (8)	236,386	165,466/47,280
OrganAMNIST	Abdominal CT	Multi (11)	58,830	34,561/17,778

Table 1. MedMNIST v2 2D datasets used for the feature selection experiments in this manuscript. Table adapted from³⁰.

Dataset	Random × 10 ^{−3}	Sampled × 10 ^{−3}	Lasso × 10 ^{−3}	qbsolv × 10 ^{−3}	QA × 10 ^{−3}	AE × 10 ^{−3}	SPCA × 10 ^{−3}
ChestMNIST	6.7 ± 0.6	6.0 ± 0.0	9.3 ± 0.1	5.3 ± 0.2	5.5 ± 0.2	2.3 ± 0.0	3.0 ± 0.0
OCTMNIST	6.7 ± 1.8	5.9 ± 0.1	17.3 ± 0.1	4.7 ± 0.1	5.3 ± 0.5	1.4 ± 0.1	2.6 ± 0.0
PneumoniaMNIST	6.9 ± 0.5	5.9 ± 0.4	7.3 ± 0.2	5.5 ± 0.1	5.3 ± 0.2	3.1 ± 0.1	2.9 ± 0.0
BreastMNIST	22.3 ± 3.5	23.5 ± 1.3	22.1 ± 2.6	21.8 ± 3.0	21.2 ± 2.6	14.8 ± 1.7	4.5 ± 0.0
TissueMNIST	3.6 ± 0.2	3.1 ± 0.0	4.4 ± 0.0	3.3 ± 0.1	3.0 ± 0.1	1.6 ± 0.0	1.9 ± 0.0
OrganAMNIST	39.4 ± 0.9	37.2 ± 0.3	40.1 ± 0.1	40.3 ± 0.4	37.4 ± 0.5	24.7 ± 0.2	25.2 ± 0.0

Table 2. Evaluation of the image compression task using selected pixels and learned features, measured via test set mean square error. Methods compared include random pixel selection, sampled pixel selection (uniform spacing over the image grid), supervised Lasso selection, solving the full QUBO with qbsolv, solving a sparsified QUBO with quantum annealing (QA), and compression using an autoencoder (AE) and sparse principal component analysis (SPCA). Results are reported as mean ± standard deviation over five independent runs. For QA experiments, sparsified QUBOs were used to match current hardware embedding constraints, while all classical methods, including qbsolv, operated on the full feature space.

selecting the pixel with the highest MI from each 2 × 2 neighborhood in the image. This reduced the QUBO size to 196 × 196. Despite this reduction, the connectivity of the QUBO still exceeded the hardware’s constraints. To address this, we applied a thresholding technique that removed weaker quadratic couplings, retaining 2000 couplings in the final problem representation. This sparsification step aligned with an observed average chain




Fig. 5. Visual comparison of images from the test set, overlaid with the selected pixels (top row) and the decoder reconstructed images from the selected pixels.

length of 4.51 ± 0.36 physical qubits per logical variable and used around 883.17 ± 82.80 qubits on the QA hardware. The average chain length and qubit count was calculated as the mean over the six datasets with five independent runs per sampling.

Minor embeddings were generated dynamically using the `EmbeddingComposite` from D-Wave's Ocean SDK, employing heuristic algorithms to map logical variables to the Pegasus P16 topology. Chain strengths were automatically selected by the embedding composite, proportional to the problem QUBO coefficients, to maintain intra-chain consistency without manually tuning. Chain break resolution was handled via majority vote, where the most common qubit value within a chain determines the logical variable's assignment.

The k of n constraint was enforced using a sparsity-preserving linear penalty, ensuring compatibility with the hardware while preserving the structure of the problem. The subsampled, thresholded, and linear Ising-constraint-enforced QUBO was mapped to the annealing hardware, with the annealing time set to $t_f = 20 \mu s$. We performed 1000 reads to form a sample set and selected the solution with the lowest energy observed. The runtime on the quantum annealer was 260.84 ± 12.64 ms. Overhead times, such as queueing the problem to the quantum annealer and preparing the QUBO, were not included in the reported runtime, as these are shared with the `qbsolv` workflow.

Discussion

In this work, we presented a method to encode a FS problem that can be implemented on commercially available quantum computing hardware. Our approach focuses on selecting the k most important features, as measured by MI, from six light-weight medical image datasets. We evaluated the selected features by training a convolutional reconstruction decoder and measuring the MSE of reconstructed samples compared to ground truth test set images.

To address the limitations of quantum hardware connectivity, we enforced a linear penalty to reduce the connectivity of the problem graph. This allowed us to generate a subsampled, thresholded QUBO formulation, that reduces the connectivity and computational complexity of the problem. Despite these simplifications, the solutions obtained on the quantum hardware were quantitatively (see Table 2) and qualitatively (see Fig. 5) comparable to those derived from a simulated solver operating on the complete problem description. This demonstrates that our method effectively balances the trade-off between hardware limitations and solution quality. The performance of our quantum-based FS approach was comparable to that achieved using a classical solver. However, we note that the approximated QUBO may yield suboptimal global solutions due to discarded weak interactions, and future hardware developments or hybrid embedding strategies may alleviate this.

Our experiments showed that the QUBO-based FS method identified plausible features for training the reconstruction encoder. However, the effectiveness of FS was dataset-dependent. For medical imaging datasets with relatively aligned images, such as ChestMNIST, PneumoniaMNIST, BreastMNIST, and OCTMNIST, the method performed well, resulting in meaningful feature subsets that supported accurate reconstructions. In contrast, for datasets where the image content was misaligned or highly heterogeneous, such as cell images, or multi-organ images, the selected features did not outperform simple subsampling. This is consistent with existing knowledge that localized pixel features are suboptimal for these tasks, as they fail to capture global spatial or contextual information.

Beyond evaluating the QUBO-based FS method, we systematically compared it against a range of classical and learning-based FS strategies to establish a broader performance baseline. Random pixel selection served as a lower-bound baseline, providing a reference for performance without structured feature selection. Subsampling, based on uniformly spaced pixel selection, acted as a natural baseline for image tasks, maintaining spatial coverage while ignoring label information. Lasso regression, despite its strong performance in tabular settings, underperformed in our experiments, often selecting spatially clustered pixels without considering redundancy, leading to poor reconstruction quality. In contrast, `qbsolv`-based solutions consistently outperformed random and subsampling approaches across several datasets, highlighting the value of structured MI-based optimization

even with heuristic solvers. Although global solvers such as Gurobi⁴⁴ are not incorporated in the current study due to the computational complexity of solving the complete QUBOs with 784 binary variables, they remain an important benchmark for future work.

SPCA and AE based feature compression methods achieved the lowest reconstruction errors across all datasets. However, as noted by Hellstern et al.²⁰, such unsupervised methods generate new transformed features rather than selecting existing ones, necessitating access to the full set of original features at deployment. This limitation highlights why neural networks, which excel at learning hierarchical and spatially invariant representations, are particularly effective for such tasks. Our QUBO-based method, by contrast, identifies explicit subsets of raw pixels, offering direct applicability to tasks such as optimizing imaging acquisition protocols, reducing radiation dose, and enhancing diagnostic interpretability. It is important to note that the objective of this study was not to propose a new feature extractor but rather to select a subset of features from a large set that adequately describes the data distribution. Learned feature descriptors, such as those produced by neural networks, offer unparalleled performance due to their ability to generalize across large datasets, but they often lack interpretability. In contrast, our QUBO-based approach provides interpretable FS grounded in statistical measures such as MI and redundancy, offering insights into the most relevant features for specific tasks.

Integrating learned feature representations with interpretable FS methods could present a powerful hybrid approach. For instance, features extracted from foundation models or other pre-trained neural networks could be used as inputs to our QUBO-based framework, combining the strengths of deep learning with the interpretability and sparsity benefits of quantum-inspired FS. This would enable both effective and interpretable solutions, bridging the gap between data-driven learning and human-comprehensible FS. Additionally, future work could investigate the incorporation of more advanced quantum algorithms or optimizing hardware configurations to further improve scalability and performance.

Conclusion

In this work, we provided an introduction to using quantum annealers for feature selection in an image compression task. We presented a method to perform feature selection on currently available quantum annealing hardware and applied it on light-weight medical imaging datasets. The method outperforms other simple feature selection techniques, but cannot compete against trainable deep learning based feature extractors. By leveraging an adapted QUBO formulation with thresholding, subsampling and hardware optimized constraints, we demonstrated how quantum hardware can be used effectively despite its current limitations. This work highlights potentials of quantum feature selection as a foundation for future explorations in interpretable and scalable feature selection methodologies, possibly combining deep learning based feature extractors and quantum based feature selection.

Data availability

The datasets analyzed during this work are publicly available and can be found at: <https://medmnist.com/>.

Received: 25 February 2025; Accepted: 1 August 2025

Published online: 07 August 2025

References

1. Litjens, G. et al. A survey on deep learning in medical image analysis. *Med. Image Anal.* **42**, 60–88 (2017).
2. Shen, D., Wu, G. & Suk, H.-I. Deep learning in medical image analysis. *Annu. Rev. Biomed. Eng.* **19**, 221–248 (2017).
3. Chandrashekar, G. & Sahin, F. A survey on feature selection methods. *Comput. Electr. Eng.* **40**, 16–28 (2014).
4. Fuchs, T. O. et al. Optimization of computed tomography data acquisition by means of quantum computing. In *Proceedings of the European Conference on Non-Destructive Testing (ECNDT)* (ECNDT, 2023).
5. Prjamkov, D. et al. Comparison of different quantum computing devices for optimization of computed tomography data acquisition. *e-J. Nondestruct. Test.* (2024).
6. Peng, H., Long, F. & Ding, C. Feature selection based on mutual information criteria of max-dependency, max-relevance, and min-redundancy. *IEEE Trans. Pattern Anal. Mach. Intell.* **27**, 1226–1238 (2005).
7. Battiti, R. Using mutual information for selecting features in supervised neural net learning. *IEEE Trans. Neural Netw.* **5**, 537–550 (1994).
8. Zou, H., Hastie, T. & Tibshirani, R. Sparse principal component analysis. *J. Comput. Graph. Stat.* **15**, 265–286 (2006).
9. Tibshirani, R. Regression shrinkage and selection via the lasso. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **58**, 267–288 (1996).
10. Schuld, M., Sinayskiy, I. & Petruccione, F. An introduction to quantum machine learning. *Contemp. Phys.* **56**, 172–185 (2015).
11. Biamonte, J. et al. Quantum machine learning. *Nature* **549**, 195–202 (2017).
12. Born, M. & Fock, V. Beweis des adiabatsatzes. *Z. Phys.* **51**, 165–180 (1928).
13. Farhi, E. et al. A quantum adiabatic evolution algorithm applied to random instances of an np-complete problem. *Science* **292**, 472–475 (2001).
14. McGeoch, C. C. Adiabatic quantum computation and quantum annealing: Theory and practice. *Synth. Lect. Quantum Comput.* **5**, 1–93 (2014).
15. Lanting, T. The D-Wave 2000Q Processor. In Presented at AQC 2017 (2017).
16. Preskill, J. Quantum computing in the Nisq era and beyond. *Quantum* **2**, 79 (2018).
17. Farhi, E., Goldstone, J. & Gutmann, S. A quantum approximate optimization algorithm. *arXiv preprint arXiv:1411.4028* (2014).
18. Preskill, J. Quantum computing in the Nisq era and beyond. *Quantum* **2**, 79 (2018).
19. Mücke, S., Heese, R., Müller, S., Wolter, M. & Piatkowski, N. Feature selection on quantum computers. *Quantum Mach. Intell.* **5**, 11 (2023).
20. Hellstern, G., Dehn, V. & Zaefferer, M. Quantum computer based feature selection in machine learning. *IET Quantum Commun.* **5**, 232–252 (2024).
21. Otgonbaatar, S. & Datcu, M. A quantum annealer for subset feature selection and the classification of hyperspectral images. *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.* **14**, 7057–7065 (2021).
22. Romero, S., Gupta, S., Gatlin, V., Chapkin, R. S. & Cai, J. J. Quantum annealing for enhanced feature selection in single-cell RNA sequencing data analysis. *arXiv preprint arXiv:2408.08867* (2024).

23. Nembrini, R., Ferrari Dacrema, M. & Cremonesi, P. Feature selection for recommender systems with quantum computing. *Entropy* **23**, 970 (2021).
24. Felefly, T. et al. An explainable MRI-radiomic quantum neural network to differentiate between large brain metastases and high-grade glioma using quantum annealing for feature selection. *J. Digit. Imaging* **36**, 2335–2346 (2023).
25. Jun, K. A highly accurate quantum optimization algorithm for CT image reconstruction based on sinogram patterns. *Sci. Rep.* **13**, 14407 (2023).
26. Nau, M. A., Vija, A. H., Gohn, W., Reymann, M. P. & Maier, A. K. Exploring the limitations of hybrid adiabatic quantum computing for emission tomography reconstruction. *J. Imaging* **9**, 221 (2023).
27. Nau, M. A., Vija, A. H., Reymann, M. P., Gohn, W. & Maier, A. K. Improving hybrid quantum annealing tomographic image reconstruction with regularization strategies. In *BVM Workshop*. 3–8 (Springer, 2024).
28. Jun, K. Quantum optimization algorithms for CT image segmentation from X-ray data. *arXiv preprint arXiv:2306.05522* (2023).
29. Choong, H. Y., Kumar, S. & Van Gool, L. Quantum annealing for single image super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 1150–1159 (2023).
30. Yang, J. et al. Medmnist v2-a large-scale lightweight benchmark for 2d and 3d biomedical image classification. *Sci. Data* **10**, 41 (2023).
31. Kadowaki, T. & Nishimori, H. Quantum annealing in the transverse Ising model. *Phys. Rev. E* **58**, 5355 (1998).
32. Heim, B., Rønnow, T. F., Isakov, S. V. & Troyer, M. Quantum versus classical annealing of Ising spin glasses. *Science* **348**, 215–217 (2015).
33. Harris, R. et al. Phase transitions in a programmable quantum spin glass simulator. *Science* **361**, 162–165 (2018).
34. King, A. D. et al. Observation of topological phenomena in a programmable lattice of 1,800 qubits. *Nature* **560**, 456–460 (2018).
35. Albash, T. & Marshall, J. Comparing relaxation mechanisms in quantum and classical transverse-field annealing. *Phys. Rev. Appl.* **15**, 014029 (2021).
36. King, A. D. et al. Scaling advantage over path-integral Monte Carlo in quantum simulation of geometrically frustrated magnets. *Nat. Commun.* **12**, 1113 (2021).
37. Abel, S., Nutricati, L. A. & Spannowsky, M. A genetic quantum annealing algorithm. *arXiv preprint arXiv:2209.07455* (2022).
38. Abel, S., Nutricati, L. A. & Rizos, J. String model building on quantum annealers. *Fortsch. Phys.* **71**, 2300167 (2023).
39. Ohzeki, M. Breaking limitation of quantum annealer in solving optimization problems under constraints. *Sci. Rep.* **10**, 3126 (2020).
40. Mirkarimi, P., Shukla, I., Hoyle, D. C., Williams, R. & Chancellor, N. Quantum optimization with linear Ising penalty functions for customer data science. *arXiv preprint arXiv:2404.05467* (2024).
41. Booth, M., Reinhardt, S. P. & Roy, A. Partitioning optimization problems for hybrid classical/quantum execution (2017). https://github.com/dwavesystems/qbsolv/blob/master/qbsolv_techReport.pdf.
42. Menze, B. H. et al. A comparison of random forest and its Gini importance with standard chemometric methods for the feature selection and classification of spectral data. *BMC Bioinform.* **10**, 1–16 (2009).
43. Mafarja, M. M. & Mirjalili, S. Hybrid whale optimization algorithm with simulated annealing for feature selection. *Neurocomputing* **260**, 302–312 (2017).
44. Gurobi Optimization, LLC. *Gurobi Optimizer Reference Manual* (2024).
45. D-Wave. D-Wave Leap. <https://cloud.dwavesys.com/leap/> (2023). Accessed 11 Dec 2024.

Author contributions

M.A.N., L.A.N., and B.C. conceived the experiments, M.A.N. and B.C. conducted the experiments, M.A.N., L.A.N., B.C., P.A.W. and A.K.M. analysed the results. All authors reviewed the manuscript.

Funding

Open Access funding enabled and organized by Projekt DEAL. This study was partially funded by the exchange visit program of the EPSRC International Network on Quantum Annealing (EP/W027003/1). Open Access funding enabled and organized by Projekt DEAL.

Declarations

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to M.A.N.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2025