Quantum Machine Learning: Perspectives in Cybersecurity

Quantum Machine Learning: perspectives in Cybersecurity

Davide Pastorello[1,2][0000−0001−5915−6796]

[1] University of Bologna, Department of Mathematics, piazza di porta San Donato 5, 40126 Bologna (Italy)
[2] TIFPA-INFN, via Sommarive 14, 38123 Povo TN (Italy)
davide.pastorello3@unibo.it

Abstract. In this work, we give an overview on some recent results related to quantum machine learning (QML) regarding the training of quantum generative adversarial neural networks by means of classical shadows, and a parametric model implemented on a quantum annealer. Then, we argue that QML models can be robust against targeted data corruption and gradient-based attacks, motivating the exploration of the relationship between QML and cybersecurity.

Keywords: Quantum machine learning · cybersecurity · data corruption · gradient masking

1 Introduction

Machine learning models, in particular artificial neural networks (ANNs), are utilized in various cybersecurity contexts to detect and prevent threats. Some examples of applications are: anomaly detection, intrusion detection, anti-phishing filters, malware detection, biometric authentication systems, log analysis, optimization of cybersecurity defenses [1]. Since ANNs represent a powerful technology for enhancing cybersecurity, protecting neural networks themselves from cyberattacks is crucial to ensuring the security of systems based on them [2]. There are several cyberattacks which can be moved to ANNs and other machine learning models: data poisoning attacks to modify the training data used to train a model; input perturbation attacks to generate malicious inputs and adversarial examples; direct manipulations of model parameters; model explorations to extract sensitive information such as identifying specific samples in the training set or retrieving sensitive model parameters; reverse engineering attacks aimed to reconstruct the model from excessive output data or other artifacts generated by the model. Some common defense strategies against these attacks are: encryption of training data during storage and transfer; authentication and authorization mechanisms to control access to ANNs and related data; input traffic monitoring to detect any attempts of attack designed to manipulate the model behavior; gradient masking for protecting the model from adversarial attacks based on the access to the gradients during the training. The application of quantum computing to machine learning tasks offers some interesting solutions characterized by

a quantum advantage with respect to the classical counterparts in terms of time and space complexity, expressive power, generalization capability, at least on a theoretical level [4]. Nevertheless, from an empirical viewpoint, quantum machine learning (QML) leverages on the availability of working prototypes of quantum processing units (QPUs). QML algorithms can be developed to address specific tasks in cybersecurity, however we adopt a different viewpoint focusing on the protection of QML models against attacks. Encoding data into quantum states is a strategy of data encryption based on the no-cloning theorem which claims that an unknown quantum state cannot be copied. As a consequence, an attacker can access data only performing measurement processes over the quantum states producing observable alterations of data themselves (as a fundamental quantum property) revealing an attack occurred. In this way, quantum encryption of data is characterized by a paradigm of unconditional security that is the main aspect of quantum cryptography [3]. In quantum information processing schemes, like QML models, quantum encoding offers an effective data encryption since data can be accessed only by measurements. However, storing quantum states is highly demanding because of decoherence due to the high sensitivity of quantum systems to the interactions with the environment. Roughly speaking, the quantum encoding of data and the non-classical quantum operations make a QML model robust w.r.t. targeted attacks which require high knowledge on how the model processes data; on the other hand, the high sensitivity to noise and the lack of standardized error-correction procedures make QML models fragile under untargeted attacks based on random data corruption. The paper is structured as follows: in section 2, we explain the adopted notations, in particular for the reader that is not familiar with the quantum formalism. In section 3, we summarize how to train a quantum generative adversarial network exploiting quantum encryption without storing quantum data but only classical estimates. Then, in section 4, we describe a parametric model implemented on a quantum annealer which can be protected by gradient-based attacks moved over the training. In section 5, we discuss how the considered cases provide a suggestion about QML as a promising resource in cybersecurity.

2 Notation and formalism

Let us clarify the adopted notation based on the standard quantum formalism in finite dimension. Any physical quantity, or observable, which can be measured over a quantum system is mathematically described by a linear self-adjoint operator acting on a (finite-dimensional) Hilbert space H where the inner product is denoted by $\langle \ | \ \rangle$. Any physical state of a quantum system is mathematically described by a linear operator ρ on H such that $\rho \geq 0$ and $\mathrm{Tr}(\rho) = 1$. Let $|\psi\rangle$ be any unit vector in H, the state given by the projector $\rho = |\psi\rangle\langle\psi|$ is called *pure state*, it can be identified directly to $|\psi\rangle$ up to phase factors. Given an observable A and a quantum state ρ, the real number $\langle A\rangle_\rho = \mathrm{tr}(A\rho)$ is the expectation value of A related to repeated measurement processes of A when the quantum system is in the state ρ. In the case of a pure state $|\psi\rangle$ we have $\langle A\rangle_\psi = \langle\psi|A\psi\rangle$.

A 2-dimensional Hilbert space, where a qubit is described, identifies to $\mathbb{C}^2$ then, according to the quantum postulates of quantum mechanics, the Hilbert space of a system composed by n qubits is given by the tensor product space $\mathsf{H} = (\mathbb{C}^2)^{\otimes n}$ whose dimension scales exponentially in the number of qubits. The time evolution of a quantum state is described by the action of a unitary operator U acting on H in these terms $\rho \mapsto U\rho U^\dagger$. If $\mathsf{H} = (\mathbb{C}^2)^{\otimes n}$, any unitary operator on H is called *n-qubit gate*. We define a *quantum circuit* as the composition of the following operations: initialization of N qubits to a known state, application of a finite sequence of n-qubit gates with $n \leq N$, and a measurement process over $k \leq N$ qubits. Quantum circuits provide a universal model of quantum computation [5].

Let $\{|0\rangle, |1\rangle\}$ be the standard basis of $\mathbb{C}^2$, the orthonormal basis $\{|b\rangle : b \in \{0,1\}^n\}$ of $(\mathbb{C}^2)^{\otimes n}$ is called *computational basis*. The measurement process over n qubits, in the overall state ρ, w.r.t. the computational basis provides the outcome $b \in \{0,1\}^n$ with probability $\mathbb{P}_\rho(b) = \langle b|\rho\, b\rangle$. In other words, the observable given by the Pauli matrix σ_z is measured on each qubit.

A *Pauli operator* (or Pauli observable) acting on $(\mathbb{C}^2)^{\otimes n}$ is an operator of the form $P = P_1 \otimes \cdots \otimes P_n$ with $P_i \in \{\mathbb{I}, \sigma_x, \sigma_y, \sigma_z\}$, for every $i = 1, ..., n$, where $\mathbb{I}$ is the 2×2 identity matrix and $\sigma_x, \sigma_y, \sigma_z$ are the Pauli matrices. A *k-local Pauli operator* is a Pauli operator which acts as the identity on $n - k$ qubits at least. An example of k-local Pauli operator is $P = \mathbb{I}^{n-k} \otimes (\sigma_x)^k$.

3 Shadow protocol for training a quantum GAN

Quantum neural networks (QNNs) are defined as parametric quantum circuits which can be trained by backpropagation in analogy to classical feedforward ANNs. A general scheme can be summarized as follows: there is a parametric part, where parameters are functions of classical input, which encodes data into quantum states and a variational part where the parameters are optimized in the training framework of a parametric model (figure 1).

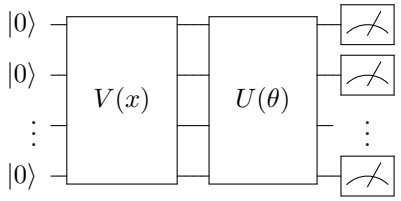

Figure 1. A simpliefied quantum neural network represented as a parametric circuit: V encodes a classical input x into a n-qubit state and U comprises the parameters θ of the model. Any qubit is initialized in $|0\rangle$ and every qubit is measured at the end.

For a QNN the corresponding model function is defined in terms of the ex-

pectation value of an n-qubit observable A:

$$f(x,\theta) = \langle 0|V^\dagger(x)U^\dagger(\theta)\, A\, U(\theta)V(x)|0\rangle \qquad x \in X\,,\ \theta \in \Theta, \tag{1}$$

where X is the input domain and Θ the parameter space.

Parametric circuits can be applied to construct generators and discriminators within the quantum version of generative adversarial networks (GAN). It is well-known that classical GANs are relevant tools in several tasks of cybersecurity [6] and it is crucial to protect them from data poisoning in this context. In quantum generative adversarial networks (QGANs), the generator is a quantum circuit, which can be implemented using a series of quantum gates that manipulate the quantum state of a set of qubits, this circuit is designed to generate data resembling those from the training dataset. The discriminator is also implemented as a quantum circuit, this circuit evaluates the likelihood of the data generated by the generator, comparing it with the real data from the training set [7]. The loss function used to train a QGAN is often defined using quantum concepts, such as quantum state overlap or quantum divergence, rather than traditional loss metrics like cross-entropy. During the training, the parameters of the generator and discriminator quantum circuits are optimized using variational algorithms within an adversarial framework. In the quantum architecture, the training set is made by quantum states (which may encode classical data) assumed to be stored in a quantum memory that is generally susceptible to noise and non-correctable errors with severe drawbacks concerning security. Below, we argue how to represent quantum states in terms of classical estimates with the advantage of storing a training set for a QGAN in a classical memory.

In [8], we considered the so-called shadow protocol that is a procedure to construct classical estimates of quantum states, called *classical shadows*, by means of measurements and quantum/classical processing. Let us summarize the shadow protocol. Let $\mathcal{U}$ be an ensemble of unitary operators over a n-qubit Hilbert space, where any $U \in \mathcal{U}$ has a statistical weight attached, that is tomographically complete[3] and an unknown quantum state ρ acting on $(\mathbb{C}^2)^{\otimes n}$, consider the following procedure [9]:

1. Sample $U \in \mathcal{U}$ and evolve the state $\rho \mapsto U\rho U^\dagger$;
2. Measure the state $U\rho U^\dagger$ in the computational basis. Let $\hat{b} \in \{0,1\}^n$ be the outcome of the measurement.
3. Apply the inverse evolution to the state $|\hat{b}\rangle\langle\hat{b}|$ obtaining $U^\dagger|\hat{b}\rangle\langle\hat{b}|U$ which can be saved as classical information.

The procedure above produces an ensemble of states $\{U^\dagger|\hat{b}\rangle\langle\hat{b}|U\}_{U,\hat{b}}$, by the expectation of this ensemble we can define the map:

$$\mathcal{M}(\rho) := \mathbb{E}_{U,\hat{b}}[U^\dagger|\hat{b}\rangle\langle\hat{b}|U], \tag{2}$$

[3] $\mathcal{U}$ is said to be *tomographically complete* if for each $\rho \neq \sigma$ there are $U \in \mathcal{U}$ and an element $|b\rangle$ of the computational basis $\{|b\rangle\ :\ b \in \{0,1\}^n\}$ such that: $\langle b|U\rho U^\dagger|b\rangle \neq \langle b|U\sigma U^\dagger|b\rangle$.

which can be extended to the space of self-adjoint operators on $(\mathbb{C}^2)^{\otimes n}$ by linearity. As a linear map, $\mathcal{M}$ is invertible and the **classical shadow** of the quantum state ρ is defined by:

$$\hat{\rho} := \mathcal{M}^{-1}\left(U^\dagger|\hat{b}\rangle\langle\hat{b}|U\right). \tag{3}$$

The classical shadow is computed classically and stored as classical information, then a closed analytic form for $\mathcal{M}^{-1}$ is required. If the ensemble of unitaries $\mathcal{U}$ is defined by random Pauli matrices measured on each qubit, the classical shadow returned from $U|\hat{b}\rangle\langle\hat{b}|U^\dagger$, where $U = P_1 \otimes \cdots \otimes P_n$ is [9]:

$$\hat{\rho} = \bigotimes_{i=1}^{n}(3P_i^\dagger|\hat{b}_i\rangle\langle\hat{b}_i|P_i - \mathbb{I}) \quad \text{where} \quad \hat{b}_i \in \{0,1\} \quad \forall i = 1,...,n. \tag{4}$$

The classical shadows of a quantum state can be used to efficiently estimates expectation values of observables [9]. Moreover, for any n-qubit quantum state ρ, the computation of a number of classical shadows that is logarithmic in n provides an accurate estimate of ρ w.r.t. the *local quantum Wasserstein distance of order 1* that is a notion from the quantum optimal mass transport [8]. Let us briefly introduce this concept from the definition of the local quantum norm of an n-qubit observable.

Definition 1. *Let $1 = c_1 \leq \ldots \leq c_n$. For any self-adjoint operator H on $(\mathbb{C}^2)^{\otimes n}$ (equivalently, for any n-qubit observable H), we define the* **local quantum norm** *of H as:*

$$\|H\|_{\text{loc}} := 2\min\left\{\max_{x\in[n]}\sum_{\Lambda\ni x} c_{|\Lambda|}\,\|H_\Lambda\|_\infty : H = \sum_{\Lambda\subseteq[n]} H_\Lambda\,,\ H_\Lambda \in \mathcal{O}_\Lambda\right\}, \tag{5}$$

where $\Lambda \subseteq [n] := \{1,...,n\}$ and $\mathcal{O}_\Lambda$ is the set of all the self-adjoint operators acting on $\bigotimes_{x\in\Lambda}\mathbb{C}^2$.

In the definition of this norm, we consider all the decompositions of an observable H as a sum of local operators acting on the regions $\Lambda \subseteq [n]$. We define the dependence of any such decomposition on a qubit x as the sum of the operator norm of each k-local operator that acts on x weighted by a coefficient c_k. We then define the local norm of such decomposition as the maximum dependence on a qubit, and the local norm of H as the minimum local norm of all its possible decompositions.

Definition 2. *The* **local quantum W_1 norm** *of a traceless operator Δ in $\mathcal{O}_\Lambda$ is the dual to the local quantum norm:*

$$\|\Delta\|_{W_1\text{loc}} := \max\left\{\text{Tr}\left[\Delta\, H\right] : H \in \mathcal{O}_{[n]}\,, \|H\|_{\text{loc}} \leq 1\right\}. \tag{6}$$

The distance $d(\rho,\sigma) := \|\rho - \sigma\|_{W_1\text{loc}}$, induced by the local quantum W_1 norm, is a measure of distinguishability between quantum states of a n-qubit system and it can be used to evaluate the convergence of the shadow protocol.

As proved in [8], for any n-qubit state the empirical mean of $O(\log n)$ classical shadows provides an estimate of the state which is accurate with respect to the local quantum W_1 distance. The accuracy in estimating a quantum state with classical shadows in this metric has a remarkable consequence in the training of a QGAN [8], as summarized below.

In [10], a QGAN capable of learning complex quantum states is proposed. The discriminator computes the following distance which captures the distinguishability of quantum states with respect to k-local Pauli operators:

$$D^{(k)}(\rho,\sigma) := \max\left\{\mathrm{Tr}\left[(\rho-\sigma)\,H\right] : H \in \mathcal{O}^{(k)}_{[n]}\,,\ \|H\|_{\tilde{L}} \leq 1\right\}\,, \tag{7}$$

for any two n-qubit quantum states ρ, σ where $\mathcal{O}^{(k)}_{[n]}$ is the set of the linear combinations of the k-local Pauli operators and the Lipschitz constant of $H = \sum_P w_P\,P$ is:

$$\|H\|_{\tilde{L}} := 2\max_{x\in[n]} \sum_{P\in\mathcal{P}_{n,x}} |w_P|\,, \tag{8}$$

where $\mathcal{P}_{n,x}$ is the set of the n-qubit Pauli operators acting nontrivially on the qubit x. The following proposition is proved in [8].

Proposition 1. *For any two n-qubit quantum states ρ, σ, we have*

$$D^{(k)}(\rho,\sigma) \leq c_k\,\|\rho-\sigma\|_{W_1\mathrm{loc}}\,. \tag{9}$$

Let us consider a QGAN where the discriminator generates a classical estimate of the true state constructed as the empirical mean of $O(\log n)$ classical shadows, as proved in [8]. Thanks to Proposition 1, such estimate is accurate also in the distance $D^{(k)}$. Therefore, no more copies of the true state will be needed and the information contained in its classical shadow will be sufficient. The generator and the discriminator are trained against each other in the adversarial scenario, and the expectation value of the discriminator observable on the true state is estimated via its classical estimate without needing further copies of the true state. After enough iterations, the generated state will be close to the classical shadow of the true state in the k-local quantum W_1 distance. Thanks to the properties of the classical shadow, the generated state will be close also to the true state in the $D^{(k)}$ distance. In other words, we conclude that a QGAN can be equivalently trained over classical shadows in place of true quantum states, if no prior information about the state is available. Therefore, assuming the training quantum states encode classical data, there is a scenario where data required to train the model are locked by a quantum encryption but classically represented allowing a robust storage against noise. In this sense, we have a QML model trained over quantum data which is capable to process quantum data avoiding the inconvenience of a fragile quantum memory that is a shortcoming for security.

4 Parametric model on a quantum annealer

Quantum annealing (QA) is a heuristic search used to solve optimization problems [11, 12]. In practice, a *quantum annealer* can be considered a specific-purpose quantum computer designed to return the absolute or approximate ground state of the *Ising model*. The latter is described by the energy function of a spin glass system under the action of an external field, namely,

$$\mathsf{E}(\mathbf{z}) = \sum_{i=1}^{N} \theta_i z_i + \sum_{(i,j)} \Gamma_{ij} z_i z_j, \quad \text{with } \mathbf{z} \in \{-1,1\}^N, \theta_i \in \mathbb{R}, \text{and } \Gamma_{ij} \in \mathbb{R}, \quad (10)$$

where the sum $\sum_{(i,j)}$ is taken over the pairs of connected spins, counting each pair only once. The ground state is the spin configuration $\mathbf{z}^* \in \{-1,1\}^N$ that minimizes the function (10). Therefore, in practice, a quantum annealer solves a combinatorial optimization problem that can be represented as a quadratic unconstrained binary optimization (QUBO) problem, which is an NP-hard problem, by means of the change of variables $x_i = \frac{z_i+1}{2} \in \{0,1\}$. From the physical viewpoint, the spin glass is realized by a network of qubits in which quantum effects allow an efficient search of the global optimum that does not stops in local minima [13].

In [14], a general parametric learning model based on QA is proposed: real input are encoded into the nodes of the qubit network and the weights are given by the strength of the couplings in direct analogy to ANNs. More precisely, the model function is defined as follows: given an input vector $\theta = (\theta_1, \ldots, \theta_n) \in X \subset \mathbb{R}^n$, and the real parameters $\{\Gamma_{ij}\}$ with $i, j = 1 \ldots n$ (the nonzero Γ_{ij} are specified by the topology graph of the machine), one can define a parametric model F based on the ground state energy of an Ising model as

$$F(\theta|\Gamma, \lambda, \epsilon) := \lambda \min_{\mathbf{z}\in\{-1,1\}^n} \mathsf{E}(\theta, \Gamma, \mathbf{z}) + \epsilon \equiv \lambda\, \mathsf{E}_0(\theta, \Gamma) + \epsilon\,, \quad (11)$$

where $\lambda \in \mathbb{R}$ and $\epsilon \in \mathbb{R}$ are additional tunable parameters that do not influence the Ising model energy. Given a training set $\mathcal{D} = \{(\theta^{(a)}, y^{(a)})\}_{a=1,\ldots,N}$, with $y^{(a)} = f(\theta^{(a)})$, where $f : X \to \mathbb{R}$, with $X \subset \mathbb{R}^n$, is an unknown function to approximate, the model (11) can be trained by minimizing the MSE loss function

$$\mathcal{L}(\Gamma, \lambda, \epsilon) = \frac{1}{N} \sum_{a=1}^{N} [F(\theta^{(a)}|\Gamma, \lambda, \epsilon) - y^{(a)}]^2. \quad (12)$$

The objective is to address this minimization task employing a gradient descent approach, iteratively updating the parameters Γ, λ, and ϵ by taking steps in the direction opposite to the gradient of the loss function $\mathcal{L}$. Let us remark that each parameter is assumed to take values into a compact interval in $\mathbb{R}$; consequently, the parameter space is a hyperrectangle. On one hand, the partial derivatives of $\mathcal{L}$ with respect to λ and ϵ are well-defined and trivial to calculate. On the other hand, the following theorem, which provides the update rules for the optimization of $\mathcal{L}$ by gradient descent, implies that the gradient $\nabla_\Gamma \mathcal{L}$ is defined almost everywhere in the parameter hyperrectangle.

Theorem 1 ([14]). *Let F be the model defined in (11), $\mathcal{D} = \{(\theta^{(a)}, y^{(a)})\}_{a=1,\dots,N}$ be a training set for F, $\mathcal{L}$ be the MSE loss function defined in (12), and $\eta > 0$ be the learning rate. Then, the partial derivatives of F with respect to the couplings Γ are defined almost everywhere in the parameter space, and the update rule for the gradient descent of $\mathcal{L}$ are:*

$$\Gamma_{ij}^{(k+1)} = \Gamma_{ij}^{(k)} - \eta \frac{2\lambda^{(k)}}{N} \sum_{a=1}^{N} [\lambda^{(k)} \, \mathsf{E}_0(\theta^{(a)}, \Gamma^{(k)}) + \epsilon^{(k)} - y^{(a)}] z_i^* z_j^*, \tag{13}$$

where $\Gamma^{(k)}$ are the values of the parameters within the k-th iteration of the gradient descent, and $\mathbf{z}^* = \operatorname{argmin}_{\mathbf{z}} \mathsf{E}(\theta^{(a)}, \Gamma^{(k)}, \mathbf{z})$.

The learning model is characterized by classical input and output and its main aspect is that the training can be performed by backpropagation without the explicit computation of the gradient but by means of an estimation provided by the quantum annealer. Experimental results on the learning performances related to the model on a real quantum annealer are reported in [14]. Remarkably, the quantum resources are used for both the execution and training of the model. This peculiarity makes the model secure against gradient-based attacks because gradients are never calculated explicitly but estimated by means quantum annealing. In other words, the quantum process implementing the gradient descent of the MSE loss realizes a gradient-masking. Moreover, the model does not require that the quantum annealer always returns the global optimum of the energy function, then also a noisy quantum machine is able to learn within this scheme. In this case, the noise introduces additional unpredictability to the quantum process increasing gradient obfuscation.

5 Conclusions

In the present work we summarized some recent techniques in QML concerning the training of quantum generative adversarial networks and the definition of a parametric model which can be trained and executed efficiently by a quantum annealer. In particular, we discussed how the possibility of training QGANs by means of classical estimates of quantum states allows to apply quantum encryption of data without the requirement of storing a dataset into a quantum memory which is fragile against data corruption. Then, we described a parametric model that can be trained by quantum annealing implementing the gradient descent of a MSE loss that is robust against gradient-based attacks.

In summary, machine learning models represent an important resource for cybersecurity and a crucial issue is then protecting these models in turn from cyberattacks. Here, focusing on specific examples, we argued that quantum encoding of data and quantum information processing allow the development of machine learning models that can be robust against attacks that are effective on classical systems. In conclusion, we suggest that QML schemes are not promising only for well-known reasons like: quantum speedups, space efficiency, capability of handling high-dimensional data, but also for robustness against cyberattacks.

Acknowledgements

This work was partially supported by project SERICS (PE00000014) under the MUR National Recovery and Resilience Plan funded by the European Union - NextGenerationEU. The author is a member of the "Gruppo Nazionale per la Fisica Matematica (GNFM)" of the "Istituto Nazionale di Alta Matematica "Francesco Severi" (INdAM)".

References

1. Macas, M., Wu, C., Fuertes, W.: A survey on deep learning for cybersecurity: Progress, challenges, and opportunities. Computer Networks, vol 212, 109032 (2022)
2. Pawlicki, M., Kozik, R., Choras, M.: A survey on neural networks for (cyber-) security and (cyber-) security of neural networks. Neurocomputing, vol. 500, 1075-1087 (2022)
3. Gisin, N., Ribordy, G., Tittel, W., Zbinden, H.: Quantum cryptography. Reviews of Modern Physics, 74(1), 145-195 (2002)
4. Pastorello, D.: Concise Guide to Quantum Machine Learning. Springer Singapore (2023)
5. Nielsen, M. A., Chuang, I. L.: Quantum Computation and Quantum Information. 10th Anniversary Edition, Cambridge University Press (2010)
6. Mirsky, A., Shabtai, A.: Generative Adversarial Networks in Security: A Survey. IEEE Communications Surveys & Tutorials, 2020
7. Lloyd, S., Weedbrook, C.: Quantum Generative Adversarial Learning. Physics. Rev. Lett. 121, 040502 (2018)
8. De Palma, G., Klein, T., Pastorello, D.: Classical shadows meet quantum optimal mass transport. Preprint arXiv:2309.08426
9. Huang, H. Y., Kueng, R.,Preskill, J.: Predicting many properties of a quantum system from very few measurements. Nature Physics 16, 1050-1057 (2020).
10. Kiani, B. T. et al.: Learning quantum data with the quantum earth mover's distance. Quantum Science and Technology, vol. 7, n. 4 (2022)
11. Kadowaki, T., Nishimori, H.: Quantum annealing in the transverse Ising model. Physical Review E 58(5), 5355 (1998)
12. Hauke, P., Katzgraber, H.G., Lechner, W., Nishimori, H. and Oliver, W. D.: Perspectives quantum annealing: Methods and implementations. Reports on Progress in Physics 83(5),608 054401 (2020)
13. McGeoch, C.C.: Adiabatic Quantum Computation and Quantum Annealing, Springer Cham (2014)
14. Schmid, L., Zardini, E., Pastorello, D.: A general learning scheme for classical and quantum Ising machines. Scipost Physics Core 7, 013 (2024)