

Noise-Dependent Eavesdropping Strategies in BB84: A Quantum Circuit Learning Approach

Demerson N. Gonçalves, Lester de Abreu Faria, João T. Dias and Fernando M. Araújo-Moreira

Abstract—Quantum circuit learning (QCL) has been proposed to optimize eavesdropping attacks on BB84. We investigate the behavior of learned attacks under depolarizing, bit-flip, phase-flip, amplitude-damping, phase-damping, and combined noise. After validating the variational model against the noiseless phase-covariant cloning machine (PCCM) frontier, we compare the optimized QCL attacks with analytical PCCM reference points evaluated under the same noise conditions. Results show that deviations from PCCM behavior depend primarily on noise structure: basis-dependent channels induce asymmetric strategies, whereas depolarizing noise largely preserves symmetry. Consequently, aggregate figures of merit may hide basis-resolved asymmetries within the restricted attack model considered here.

Keywords—Quantum key distribution, BB84, quantum circuit learning, noisy quantum channels, phase-covariant cloning

I. INTRODUCTION

Quantum key distribution (QKD) is one of the most prominent applications of quantum information science because it enables secure key establishment based on physical principles rather than on computational assumptions. Since the introduction of the BB84 protocol [1] and the entanglement-based proposal E91 [2], the field has developed into a mature research area supported by rigorous security analyses [3], [4], [5], practical classical post-processing methods [6], [7], and major advances in quantum cryptographic theory and implementation [8]. Within this framework, the characterization of eavesdropping strategies remains a central problem, since security is operationally determined by the tradeoff between Eve’s information gain and the disturbance observed by the legitimate users.

For prepare-and-measure QKD, BB84 is the canonical setting in which this tradeoff can be analyzed. In the ideal noiseless regime, the study of individual attacks is closely connected to optimal cloning, and the phase-covariant cloning machine (PCCM) provides the standard symmetric benchmark for attacks on the BB84 signal states [9], [10], [11]. In realistic scenarios, however, the communication channel and the associated devices are affected by noise. Since different

noise processes may degrade the computational and diagonal bases unequally, the optimal information-disturbance balance available to an eavesdropper may change as the channel departs from the ideal symmetric setting.

From a telecommunications perspective, such noise processes are not abstract modeling artifacts but are directly associated with physical transmission media and hardware limitations, including fiber attenuation, polarization effects, atmospheric turbulence in free-space optical links, and detector imperfections. These impairments are typically non-uniform across different quantum states, potentially introducing asymmetries that are not captured by idealized symmetric models. Understanding how these physical-layer effects influence the optimal information-disturbance tradeoff is therefore essential for assessing the practical security of QKD systems deployed in real communication networks.

Quantum machine learning (QML) and variational quantum algorithms (VQAs) provide a natural framework for studying this problem because they allow one to optimize parameterized quantum processes directly in terms of operational figures of merit [13], [14], [15], [16]. In this spirit, Decker *et al.* [18] formulated eavesdropping on BB84 as a quantum circuit learning problem and showed that, in the noiseless regime, the learned attack reproduces the optimal PCCM tradeoff, while in the noisy regime it can depart from the symmetric PCCM structure and become basis dependent.

An important question that remains open is how general this departure from PCCM-like behavior is across different classes of quantum noise. In particular, it is not yet clear whether basis-dependent learned attacks arise only for specific asymmetric channels or whether they persist more broadly across qualitatively distinct error mechanisms, such as depolarizing noise, basis-biased flips, amplitude damping, phase damping, and combined error processes. This issue is relevant both conceptually and practically, since realistic QKD implementations are typically influenced by multiple imperfections rather than by a single idealized noise source. Addressing this point requires a systematic comparison across different noise models.

In this work, we investigate QCL-based attacks on BB84 under multiple noise models. Our study proceeds in three stages. First, we validate the noiseless regime against the analytical PCCM frontier, ensuring that the variational framework reproduces the expected optimal benchmark before any channel imperfection is introduced. Second, we numerically reproduce the bit-flip scenario emphasized in [18], namely a noisy channel acting before Eve’s attack, and compare the learned solutions with the corresponding noisy PCCM curves.

Demerson N. Gonçalves, Department of Mathematics, Federal Center for Technological Education Celso Suckow da Fonseca (CEFET-RJ), Petrópolis-RJ, Brazil, and Military Institute of Engineering (IME), Rio de Janeiro-RJ, Brazil, e-mail: demerson.goncalves@cefet-rj.br; Lester de Abreu Faria, Aeronautics Institute of Technology (ITA), São José dos Campos-SP, Brazil, e-mail: lester@ita.br; João T. Dias, Department of Telecommunications Engineering, Federal Center for Technological Education Celso Suckow da Fonseca (CEFET-RJ), Rio de Janeiro-RJ, Brazil, and Military Institute of Engineering (IME), Rio de Janeiro-RJ, Brazil, e-mail: joao.dias@cefet-rj.br; Fernando M. Araújo-Moreira, Military Institute of Engineering (IME), Rio de Janeiro-RJ, Brazil, e-mail: fernando.manuel@ime.br.

Third, we extend the analysis to several single-qubit noise families, including depolarizing, bit-flip, phase-flip, amplitude damping, phase damping, and combined bit-and-phase flip errors. Our goal is to determine how the learned attack generalizes across noise channels, which regimes remain close to PCCM-like behavior, and which channels induce stronger basis asymmetry.

II. QCL ATTACK MODEL AND NOISELESS PCCM VALIDATION

We study individual attacks on BB84 in a variational framework. The attack is individual in the sense that the same interaction is applied independently to each transmitted signal and no joint quantum operation is performed across different rounds. Eve may retain each individual probe until basis disclosure, after which she applies the corresponding readout unitary and measures it. We formalize this scenario as follows.

Alice prepares one of the four BB84 states,

$$|\psi_b^Z\rangle = |b\rangle, \quad |\psi_0^X\rangle = |+\rangle, \quad |\psi_1^X\rangle = |-\rangle, \quad (1)$$

with $b \in \{0, 1\}$. Eve couples the signal qubit to a single ancilla initialized in $|0\rangle_E$ through a parameterized interaction unitary $U(\theta)$, where θ denotes the 12 trainable parameters of the circuit depicted in Fig. 2. After basis disclosure, Eve applies a basis-conditioned single-qubit readout unitary on her probe:

$$C_{\text{readout}}^{(c)}(\Lambda_c) = R_x(\lambda_{c,1})R_y(\lambda_{c,2})R_z(\lambda_{c,3}), \quad c \in \{Z, X\}, \quad (2)$$

where $\Lambda_c = (\lambda_{c,1}, \lambda_{c,2}, \lambda_{c,3})$ is an independent parameter vector for each announced basis. The rotation gates are defined by

$$R_\mu(\vartheta) = \exp\left(-i\frac{\vartheta}{2}\sigma_\mu\right), \quad \mu \in \{x, y, z\}. \quad (3)$$

Let $\rho_E^{(b,c)}(\theta)$ denote Eve's reduced state conditioned on Alice's bit b and basis c , after the interaction $U(\theta)$. Eve's conditional readout probability is

$$P_E(\hat{b}|b, c) = \left[\Pi_{\hat{b}} C_{\text{readout}}^{(c)}(\Lambda_c) \rho_E^{(b,c)}(\theta) C_{\text{readout}}^{(c)\dagger}(\Lambda_c) \right], \quad (4)$$

where $\Pi_{\hat{b}} = |\hat{b}\rangle\langle\hat{b}|$.

Figure 1 summarizes the individual-attack model. Alice prepares a BB84 signal state according to her bit x and basis b_A , while Eve couples the signal qubit to an ancilla initialized in $|0\rangle$ through $U(\theta)$. After basis reconciliation, the ancilla is processed by the basis-conditioned readout head C_{readout} and then measured immediately. No cross-round quantum memory is used.




Fig. 1. General scheme of the individual attack model.

The parameterized circuit architecture for $U(\theta)$ is shown in Fig. 2. It consists of two layers, each composed of three

single-qubit rotations followed by a CNOT gate, totaling 12 trainable parameters (θ_1 to θ_{12}). The readout head C_{readout} is applied exclusively to Eve's ancilla and uses 3 independent parameters for each announced basis (Z and X), totaling 18 trainable parameters overall. This structure is sufficient to optimize individual attacks with per-signal delayed readout in both noiseless and noisy regimes.




Fig. 2. Parameterized circuit architecture (ansatz) used by the QCL algorithm to optimize Eve's attack.

The operational figures of merit are the average fidelities between Alice and Bob and between Alice and Eve,

$$F_{AB} = \frac{1}{4} \sum_{c,b} \Pr(\hat{b}_B = b | b, c), \quad (5)$$

$$F_{AE} = \frac{1}{4} \sum_{c,b} \Pr(\hat{b}_E = b | b, c),$$

where $b \in \{0, 1\}$ denotes Alice's bit and $c \in \{Z, X\}$ denotes the preparation and measurement basis.

The attack parameters are optimized with the loss

$$\mathcal{L}(\Theta, \Lambda) = \alpha(F_{AB} - f)^2 - F_{AE}, \quad (6)$$

where f is a target fidelity for Bob.

The optimization was performed using density-matrix simulation in PennyLane to compute analytical gradients. The Adam optimizer was configured with a learning rate of $\eta = 0.05$ and 300 steps for the noiseless simulation, and $\eta = 0.1$ and 100 steps for the noisy simulations. Model parameters were initialized from a zero-mean normal distribution with standard deviation 0.05.

To ensure reproducibility and mitigate local minima, we performed multiple independent training runs: 6 restarts for the noiseless validation and 10 restarts for each noisy scenario, with a base random seed of 7 and distinct offsets per configuration. For the systematic multi-channel study, the target fidelity was fixed at $F_{AB} = 0.85$; for the noiseless validation and bit-flip reproduction, target fidelities were sampled uniformly in $[0.5, 1.0]$. Noise strengths were swept over $\{0.00, 0.05, \dots, 0.25\}$. Reported QCL points correspond to the lowest-loss run; error bars in Figs. 5 and 6 indicate the standard deviation across runs.

In the noiseless regime, the optimal individual attack is given by the phase-covariant cloning machine (PCCM). Its analytical tradeoff is parameterized by an angle θ , yielding the fidelities:

$$F_{AB} = \frac{1 + \cos(\theta/2)}{2}, \quad F_{AE} = \frac{1 + \sin(\theta/2)}{2}, \quad (7)$$




Fig. 3. Noiseless validation of the learned BB84 attack against the analytical PCCM reference. The solid curve is the PCCM frontier, crosses mark the reference PCCM points, and colored markers denote the best QCL solutions.

which implies the explicit analytical boundary:

$$F_{AE}^{\text{PCCM}}(F_{AB}) = \frac{1 + \sqrt{1 - (2F_{AB} - 1)^2}}{2}. \quad (8)$$

This curve is used as the main validation benchmark for the variational model in the ideal regime.

Our validation is performed by optimizing the QCL attack for several target fidelities f and comparing the resulting pairs (F_{AB}, F_{AE}) with the analytical PCCM frontier. In addition, we verify that the learned attack remains effectively symmetric across the two BB84 bases by checking whether

$$F_{AB}^Z \approx F_{AB}^X, \quad F_{AE}^Z \approx F_{AE}^X. \quad (9)$$

Figure 3 shows the noiseless validation of the learned attack against the explicit PCCM reference. The solid curve represents the analytical PCCM frontier, the crosses mark the reference PCCM operating points associated with the target Bob fidelities, and the colored markers denote the best QCL solutions. The final QCL solutions closely match the explicit PCCM reference points along the analytical frontier, indicating that the variational model reproduces the expected noiseless tradeoff with good accuracy. Although the optimization trajectories remain irregular, the agreement between the final learned solutions and the PCCM benchmark is sufficient to support the use of the framework in the subsequent noisy-channel analysis.

Only after this noiseless benchmark is reproduced do we proceed to the multi-noise analysis.

III. LEARNED BB84 ATTACKS UNDER MULTIPLE NOISE MODELS

After validating the noiseless regime against the analytical PCCM frontier, we turn to the noisy setting. The channel is divided into an Alice–Eve segment and an Eve–Bob segment. In the experiments reported here, the noise acts exclusively before Eve’s interaction,

$$\mathcal{N}_{AE} = \mathcal{N}_p, \quad \mathcal{N}_{EB} = \mathcal{I}. \quad (10)$$

Accordingly, Eve is modeled as a constrained adversary located after a non-adversarial noisy segment. This setting is intended to characterize noise-dependent explicit attacks and does not constitute a worst-case security model for BB84.

A. Noise Channels and Kraus Operators

To model the physical imperfections acting on the signal qubit prior to Eve’s attack, we employ standard single-qubit quantum channels described in the Kraus operator formalism, $\mathcal{E}(\rho) = \sum_k E_k \rho E_k^\dagger$. The evaluated noise models and their respective Kraus operators E_k , parameterized by a noise strength or probability $p \in [0, 1]$, are defined as follows:

- **Bit-Flip Channel:** Models the accidental bit inversion of the qubit with probability p . Its Kraus operators are given by:

$$E_0 = \sqrt{1-p} I, \quad E_1 = \sqrt{p} X \quad (11)$$

where I is the identity matrix and X is the Pauli- X operator.

- **Phase-Flip Channel:** Represents phase errors occurring with probability p without altering the computational basis populations:

$$E_0 = \sqrt{1-p} I, \quad E_1 = \sqrt{p} Z \quad (12)$$

where Z denotes the Pauli- Z operator.

- **Depolarizing Channel:** Encapsulates symmetric errors where the state is completely depolarized with probability p , uniformly contracting the Bloch sphere:

$$E_0 = \sqrt{1 - \frac{3p}{4}} I, \quad E_1 = \sqrt{\frac{p}{4}} X, \quad (13)$$

$$E_2 = \sqrt{\frac{p}{4}} Y, \quad E_3 = \sqrt{\frac{p}{4}} Z \quad (14)$$

where Y is the Pauli- Y operator.

- **Amplitude Damping Channel:** Simulates energy relaxation (dissipation) effects with damping parameter p :

$$E_0 = \begin{pmatrix} 1 & 0 \\ 0 & \sqrt{1-p} \end{pmatrix}, \quad E_1 = \begin{pmatrix} 0 & \sqrt{p} \\ 0 & 0 \end{pmatrix} \quad (15)$$

- **Phase Damping (Dephasing) Channel:** Describes loss of quantum phase coherence without energy exchange, governed by parameter p :

$$E_0 = \begin{pmatrix} 1 & 0 \\ 0 & \sqrt{1-p} \end{pmatrix}, \quad E_1 = \begin{pmatrix} 0 & 0 \\ 0 & \sqrt{p} \end{pmatrix} \quad (16)$$

Additionally, the combined bit-and-phase flip model is implemented via the sequential composition of the individual bit-flip and phase-flip channels.

We first reproduce the bit-flip scenario emphasized in [18], considering a bit-flip channel with probability 0.25 acting before Eve’s interaction and comparing the learned QCL solutions with the corresponding noisy PCCM curves in the Z basis, the X basis, and their average. We then extend the analysis to several single-qubit noise models acting on the signal qubit, namely depolarizing, bit-flip, phase-flip, amplitude damping, phase damping, and a combined bit-and-phase flip model. The noisy channels were

implemented using PennyLane’s built-in single-qubit noise operations, namely `DepolarizingChannel`, `BitFlip`, `PhaseFlip`, `AmplitudeDamping`, and `PhaseDamping`, while the combined bit-and-phase flip model was implemented as a sequential composition of `BitFlip` and `PhaseFlip` [19].

For each target fidelity f , the corresponding PCCM reference fidelities, F_{AB}^{PCCM} and F_{AE}^{PCCM} , are determined analytically from the PCCM tradeoff in Eqs. (4) and (5). The selected noise channel $\mathcal{N}_\lambda$ is applied to the signal qubit before Eve’s interaction, consistently with the QCL simulations. The PCCM reference is not re-optimized during training: its operating point is fixed for each value of f , while only the parameters of the QCL circuit are iteratively optimized. Therefore, any deviation between the learned QCL solutions and the PCCM reference reflects exclusively the optimization of the variational QCL model under the prescribed target fidelity and noise condition, since no optimization is performed on the PCCM benchmark itself.

Our analysis focuses on two complementary aspects. First, we evaluate the deviation of the learned attack from the noisy PCCM reference through the fidelities F_{AB} and F_{AE} , together with their basis-resolved counterparts $F_{AB}^Z, F_{AB}^X, F_{AE}^Z, F_{AE}^X$. Second, we monitor the basis asymmetry

$$|F_{AB}^Z - F_{AB}^X|, \quad (17)$$

which indicates whether the channel affects the two BB84 bases differently from Bob’s perspective.

Figure 4 reproduces the noisy bit-flip scenario considered in prior work and compares the final QCL solutions with the corresponding noisy PCCM curves in the Z basis, the X basis, and their average. The figure shows that the learned attack no longer follows the same symmetric tradeoff in both bases, sacrificing performance in the affected basis while improving it in the less affected one. To compare different channels on equal footing, Fig. 5 reports the noise-induced change in the QCL–PCCM Alice-Eve fidelity gap, measured relative to the zero-noise case for each channel. Figure 6 shows the corresponding basis asymmetry in Bob’s fidelity.

Taken together, Fig. 5 and Fig. 6 show that the departure from PCCM-like behavior is strongly channel dependent. Bit-flip and phase-flip produce the largest growth of basis asymmetry, indicating a pronounced imbalance between the two BB84 bases, whereas amplitude and phase damping induce a weaker but still non-negligible asymmetry. Depolarizing noise remains nearly symmetric throughout the range considered. At the same time, the noise-induced change in the QCL–PCCM gap is not uniform across channels: phase-flip yields the clearest positive increase, phase damping produces a smaller but systematic enhancement, depolarizing and amplitude damping tend to reduce the relative QCL advantage, and the combined bit-and-phase flip channel exhibits a non-monotonic transition. These results indicate that basis asymmetry is an important but not sufficient indicator of the QCL advantage, and that different noise families define qualitatively distinct regimes for learned BB84 attacks.

From an engineering standpoint, these results suggest that different noise families define distinct operational regimes




Fig. 4. Bit-flip 25% before Eve’s attack: QCL versus PCCM. The curves show the noisy PCCM fidelities in the Z basis, the X basis, and their average, while the diamond markers denote the final QCL solutions.




Fig. 5. Noise-induced change in the QCL–PCCM Alice-Eve fidelity gap across noise channels, measured relative to the zero-noise value for each channel. Positive values indicate that the channel increases the QCL advantage over PCCM, whereas negative values indicate that it reduces it.

for QKD systems. In particular, channels that introduce strong basis-dependent effects, such as bit-flip and phase-flip processes, enable more effective adaptive eavesdropping strategies, whereas more symmetric noise processes, such as depolarizing channels, tend to preserve the structure of the optimal symmetric attack. This distinction indicates that security assessments based solely on aggregate error rates, such as QBER, may be insufficient to fully characterize system vulnerability in practical deployments. This suggests that more refined metrics, such as basis-resolved error analysis or channel-aware security models, may be required to accurately capture the vulnerability of practical QKD systems.

On the other side, from a physical standpoint, the pronounced asymmetry induced by bit-flip and phase-flip channels aligns with their inherent basis dependence. A pure bit-flip error heavily degrades the computational (Z) basis while leaving the diagonal (X) basis relatively undisturbed, prompting the QCL algorithm to adapt Eve’s strategy to exploit the less noisy basis. Conversely, depolarizing noise acts symmetrically across the Bloch sphere, which explains why the learned attack




Fig. 6. Basis asymmetry in Bob's fidelity across noise channels, measured by $|F_{AB}^Z - F_{AB}^X|$. The curves show how different noise families affect the two BB84 bases unequally as the noise strength increases.

retains an almost zero basis asymmetry ($|F_{AB}^Z - F_{AB}^X| \approx 0$). This adaptive behavior demonstrates that the variational model does not merely degrade with noise, but actively redistributes the eavesdropping effort to maximize information gain along the path of least disturbance.

IV. CONCLUSIONS

Beyond the numerical characterization, our results highlight that the structure of noise, rather than its magnitude alone, plays a central role in determining effective eavesdropping strategies within the class of individual attacks considered. Since the PCCM reference remained fixed throughout the experiments, the observed deviations are attributable to the adaptive optimization performed by the QCL model under different noise conditions rather than to modifications of the analytical benchmark. This observation has direct implications for the design and evaluation of quantum communication links, where different physical-layer impairments may lead to qualitatively distinct security regimes. In particular, the presence of basis-dependent noise can enable adaptive attacks that are not captured by symmetric security models. This suggests that practical security assessments should move beyond aggregate error rates such as QBER and explicitly account for the underlying noise mechanisms, for instance via basis-resolved error analysis. These findings contribute to bridging the gap between theoretical QKD security analyses and the constraints of real-world telecommunication environments.

Nevertheless, it is important to emphasize that the performance and convergence of the learned attacks should not be interpreted as strict proofs of global optimality, but rather as approximations bounded by the expressivity and ansatz architecture of the variational quantum circuit employed. Furthermore, the present study is strictly limited to single-qubit noise models. While these single-qubit channels capture essential physical impairments, practical quantum devices are also subject to multi-qubit and correlated noise effects, which represent significant challenges in real-world quantum hardware. Investigating how such correlated and multi-qubit error mechanisms influence adaptive quantum circuit learning strategies constitutes an important direction for future work.

ACKNOWLEDGEMENTS

This work was partially supported by Coordenação de Aperfeiçoamento de Pessoal de Nível Superior (CAPES), the Brazilian Ministry of Defense through the PRÓ-DEFESA V Program – IATQ Network (grant No. 88887.961681/2024-00), and Financiadora de Estudos e Projetos (FINEP) through the projects PDI-DQBRN/1545-22, CMLabIME/0160-22, and QUANTUM II/3310-24.

REFERENCES

- [1] C. H. Bennett and G. Brassard, "Quantum Cryptography: Public Key Distribution and Coin Tossing," in *Proceedings of the IEEE International Conference on Computers, Systems and Signal Processing*, Bangalore, India, pp. 175–179, 1984.
- [2] A. K. Ekert, "Quantum Cryptography Based on Bell's Theorem," *Physical Review Letters*, vol. 67, no. 6, pp. 661–663, 1991.
- [3] P. W. Shor and J. Preskill, "Simple Proof of Security of the BB84 Quantum Key Distribution Protocol," *Physical Review Letters*, vol. 85, no. 2, pp. 441–444, 2000.
- [4] R. Renner, "Security of Quantum Key Distribution," *International Journal of Quantum Information*, vol. 6, no. 1, pp. 1–127, 2008.
- [5] R. Renner, N. Gisin, and B. Kraus, "Information-Theoretic Security Proof for Quantum-Key-Distribution Protocols," *Physical Review A*, vol. 72, no. 1, p. 012332, 2005.
- [6] G. Brassard and L. Salvail, "Secret-Key Reconciliation by Public Discussion," in *Advances in Cryptology — EUROCRYPT '93*, T. Helleseth, Ed., Lecture Notes in Computer Science, vol. 765, Berlin, Heidelberg: Springer, pp. 410–423, 1994.
- [7] C. H. Bennett, G. Brassard, and J.-M. Robert, "Privacy Amplification by Public Discussion," *SIAM Journal on Computing*, vol. 17, no. 2, pp. 210–229, 1988.
- [8] S. Pirandola, U. L. Andersen, L. Banchi, M. Berta, D. Bunandar, R. Colbeck, D. Englund, T. Gehring, C. Lupo, C. Ottaviani, J. L. Pereira, M. Razavi, J. Shaari, M. Tomamichel, G. Vallone, P. Villoresi, and P. Wallden, "Advances in Quantum Cryptography," *Advances in Optics and Photonics*, vol. 12, no. 4, pp. 1012–1236, 2020.
- [9] W. K. Wootters and W. H. Zurek, "A Single Quantum Cannot Be Cloned," *Nature*, vol. 299, no. 5886, pp. 802–803, 1982.
- [10] D. Dieks, "Communication by EPR Devices," *Physics Letters A*, vol. 92, no. 6, pp. 271–272, 1982.
- [11] S. M. Barnett and S. Croke, "Quantum State Discrimination," *Advances in Optics and Photonics*, vol. 1, no. 2, pp. 238–278, 2009.
- [12] V. Scarani, H. Bechmann-Pasquinucci, N. J. Cerf, M. Dušek, N. Lütkenhaus, and M. Peev, "The Security of Practical Quantum Key Distribution," *Reviews of Modern Physics*, vol. 81, no. 3, pp. 1301–1350, 2009.
- [13] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, "Quantum Machine Learning," *Nature*, vol. 549, no. 7671, pp. 195–202, 2017.
- [14] M. Schuld and F. Petruccione, *Machine Learning with Quantum Computers*. Cham: Springer, 2021.
- [15] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii, "Quantum Circuit Learning," *Physical Review A*, vol. 98, no. 3, p. 032309, 2018.
- [16] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. J. Coles, "Variational Quantum Algorithms," *Nature Reviews Physics*, vol. 3, pp. 625–644, 2021.
- [17] J. Bowles, S. Ahmed, and M. Schuld, "Better than Classical? The Subtle Art of Benchmarking Quantum Machine Learning Models," *arXiv preprint arXiv:2403.07059*, 2024.
- [18] T. Decker, M. Gallezot, S. F. Kerstan, A. Paesano, A. Ginter, and W. Wormsbecher, "Quantum Key Distribution as a Quantum Machine Learning Task," *npj Quantum Information*, vol. 11, p. 140, 2025.
- [19] V. Bergholm, J. Izaac, M. Schuld, C. Gogolin, C. Blank, K. McKiernan, and N. Killoran, "PennyLane: Automatic differentiation of hybrid quantum-classical computations," *arXiv preprint arXiv:1811.04968*, 2018.