

LSTM-QGAN: Scalable NISQ Generative Adversarial Network

Cheng Chu Aishwarya Hastak Fan Chen

Luddy School of Informatics, Computing, and Engineering, Indiana University Bloomington, USA

Abstract—Current quantum generative adversarial networks (QGANs) still struggle with practical-sized data. First, many QGANs use principal component analysis (PCA) for dimension reduction, which, as our studies reveal, can diminish the QGAN’s effectiveness. Second, methods that segment inputs into smaller patches processed by multiple generators face scalability issues. In this work, we propose LSTM-QGAN, a QGAN architecture that eliminates PCA preprocessing and integrates quantum long short-term memory (QLSTM) to ensure scalable performance. Our experiments show that LSTM-QGAN significantly enhances both performance and scalability over state-of-the-art QGAN models, with visual data improvements, reduced Fréchet Inception Distance scores, and reductions of $5\times$ in qubit counts, $5\times$ in single-qubit gates, and $12\times$ in two-qubit gates.

Index Terms—NISQ, Quantum Generative Adversarial Network, Long Short-Term Memory

I. INTRODUCTION

Current QGANs. Recent advancements in Noisy Intermediate-Scale Quantum (NISQ) platforms [1]–[3] have catalyzed intense research on Quantum Generative Adversarial Networks (QGANs) [4]–[13], which are well-suited to the constraints of NISQ systems, such as limited qubit counts and shallow circuit depths [14]. Building on the foundational work [4] that established the theoretical superiority of QGANs over classical counterparts, early QGAN implementations [5]–[7] only focused on low-dimensional inputs like single-bit data. Subsequent research introduced innovations such as Wasserstein loss [8] and novel architectures [9] to improve training stability. More recent work [10], [11] expanded QGANs to high-dimensional data, like the 28×28 MNIST dataset, by employing dimensionality reduction techniques like Principal Component Analysis (PCA). The state-of-the-art (SOTA) PatchGAN [12] further segments inputs into smaller patches, enabling efficient processing on practical NISQ devices.

Limitations. Despite recent developments, QGANs continue to face challenges in managing practical-sized data. First, while pre- and post-processing with PCA and inverse PCA [10], [11] enable QGANs to handle large-dimensional data, PCA often dominates the process, diminishing the contributions of the QGANs themselves. Second, although PatchGAN [12] facilitates the direct processing of practical-sized inputs through multiple small patches, its architectural limitations demand an increasing number of quantum resources as input size grows, leading to serious scalability challenges. For instance, generating a single MNIST image requires a prohibitively high 56 sub-quantum generators and 280 qubits. Third, and more concerning, our preliminary study shows a

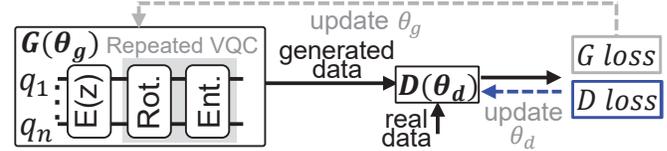


Fig. 1. A standard QGAN architecture.

significant decline in output quality as PatchGAN scales from its original 5-qubit design [12] to 8 qubits, severely limiting its effectiveness at larger scales.

Contributions. We introduce LSTM-QGAN, a novel architecture that eliminates the need for PCA when processing large-dimensional data. The design allows for the use of a constant amount of NISQ computing resources as input size increases. However, as additional hardware resources become available, the architecture scales efficiently, ensuring consistent and reliable performance. Our contributions include:

- **Preliminary Analysis.** We conduct experiments on the SOTA QGANs [10]–[12], revealing previously undisclosed limitations in PCA pre-processing and model scalability.
- **Scalable Architecture.** We present LSTM-QGAN, a scalable QGAN architecture inspired by recent advances in quantum long-short memory (QLSTM) [15]–[17]. LSTM-QGAN eliminates the need for PCA, maintains constant NISQ resources as input size grows, and efficiently scales with increasing quantum computing resources.
- **Enhanced Performance.** We conduct evaluations on NISQ computers. Experimental results show that LSTM-QGAN significantly enhances generative performance and improves scalability compared to SOTA QGANs.

II. BACKGROUND

QGAN Basics. Figure 1 illustrates a standard QGAN with two parameterized models: the Generator, $G(\theta_g)$, which generates synthetic data, and the Discriminator, $D(\theta_d)$, which evaluates the generated data against real data. G is implemented using a quantum neural network (QNN), typically composed of a data encoder $E(\cdot)$ and repeated layers of a variational quantum circuit (VQC) with one-qubit rotations (i.e., *Rot.*) and two-qubit entanglement (i.e., *Ent.*). D in SOTA QGANs [10]–[12] can be implemented with either classical or quantum models. The objective is to optimize the predefined minmax loss \mathcal{L} , as outlined in Equation 1, where z represents the latent variable. The specific loss function can be implemented using various specified functions [6]–[8], [12], [13]. The overall goal

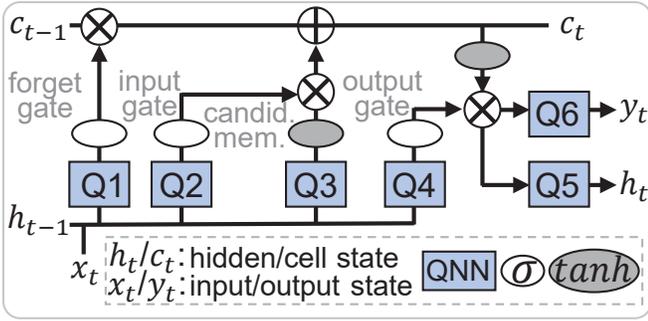


Fig. 2. A standard QLSTM architecture.

is to enable G to generate data indistinguishable from real data, while D improves its ability to differentiate between them.

$$\min_{\theta_g} \max_{\theta_d} \mathcal{L}\{D_{\theta_d}(G_{\theta_g}(z)), D_{\theta_d}(x)\} \quad (1)$$

SOTA QGANs. To manage larger-dimensional data with limited qubits on NISQ computers, SOTA QGANs [10]–[12]) primarily utilize the following two techniques:

- **Pre- and Post-Processing.** Several recent QGANs [10], [11] utilize principal component analysis to reduce input dimensions (e.g., from 784 to 4 in [10]) to fit within the limitations of NISQ computers with constrained qubits. The key steps in PCA involve: (1) standardizing the data to have zero mean and unit variance, and (2) calculating the covariance matrix \mathbf{C} and the matrix \mathbf{V}_k , which contains the top k eigenvectors (principal components). For any data matrix \mathbf{X} with mean μ , the data can be reduced to the top k principal components by $\mathbf{Z}=\mathbf{X}\mathbf{V}_k$. The reduced-dimensional data \mathbf{Z}^* can then be reconstructed to approximate the original data through inverse PCA: $\mathbf{X}^*=\mathbf{Z}^*\mathbf{V}_k^\top + \mu$.
- **Patched Input.** PatchGAN [12] segments the input into small regional patches and trains a dedicated sub-generator for each, capable of generating synthesized data that follows the pattern of the corresponding patch. This approach makes it a resource-efficient QGAN framework. The number of sub-generators scales with the input size; for instance, the 5-qubit design in the original work [12] requires 56 sub-generators to process the 784-pixel MNIST dataset, and doubling the input size would proportionally increase the number of sub-generators needed.

QLSTM. Long short-term memory [18] effectively captures spatiotemporal information, enabling task-specific regulation of data flow. Recent work [15] has introduced quantum LSTM, extended to various sequential learning tasks [16], [17]. As shown in Figure 2, QLSTM retains the classical LSTM gating mechanism, with the key distinction being the integration of QNNs. Due to page limit, we refer readers to [15], [18] for detailed insights into QLSTM. LSTM has already been applied in classical GANs [19], [20], demonstrating enhanced generative power and reduced computational cost. Building on this, we aim to leverage LSTM’s ability to selectively retain relevant patterns within a QGAN by training a QLSTM-based

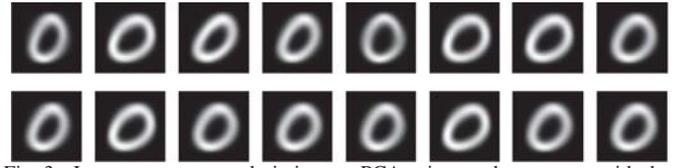


Fig. 3. Images reconstructed via inverse PCA using random vectors with the covariance matrix from the MNIST dataset.

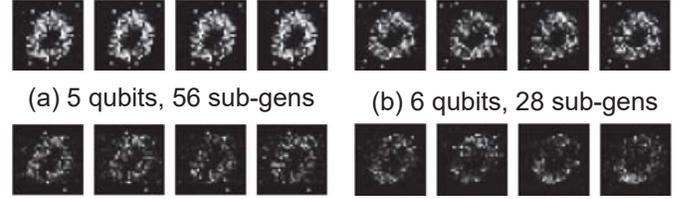


Fig. 4. Image quality degradation in scaling PatchGAN [12].

generator using different patched inputs, rather than separate sub-generators for different patches as in [12].

III. PRELIMINARY STUDY AND MOTIVATION

A. Preliminary Study

PCA Overshadows QGANs. QGANs [10], [11] on MNIST use PCA and inverse PCA for dimensionality reduction and reconstruction. To assess PCA’s impact, we reduced 28×28 MNIST images to 1×2 vectors using `scikit-learn`, generating the corresponding \mathbf{C} , \mathbf{V}_2 , and μ . We then randomly generated 1×2 vectors, applied inverse PCA, and present the reconstructed images in Figure 3. The reconstructed images closely resemble the original MNIST data and are comparable to those generated by QGANs [10], [11], suggesting that PCA pre- and post-processing may dominate, potentially overshadowing QGAN effectiveness. *These results highlight concerns about the independent validity of QGANs when PCA is used, emphasizing the need for evaluation with unprocessed data.*

Scalability for PatchGAN. PatchGAN [12] claims effectiveness with patch-based processing of high-dimensional inputs but originally reports results with only 5 qubits. To evaluate scalability, we increased the qubit count from 5 to 8. Since PatchGAN employs amplitude encoding and processes one patch at a time, we adjusted the number of sub-generators to cover all 784 pixels in an MNIST image. As show in Figure 4, the generated images degrade rapidly with increasing qubits, with sub-figure titles indicating qubit counts and required sub-generators (sub-gens). *These findings underscore PatchGAN’s poor scalability, suggesting limited potential for handling larger-scale inputs even with additional qubits.*

B. Motivation

Our preliminary results highlight the critical need for a QGAN model capable of directly processing real-world data without PCA preprocessing, as well as a more scalable architecture to overcome the limitations of existing QGANs. Motivated by these findings, we are exploring the integration of patched inputs inspired by PatchGAN [12] to enable direct input processing without PCA. Specifically, we are investigating a scalable QGAN framework that leverages QLSTM as the generator’s backbone, utilizing QLSTM’s ability to

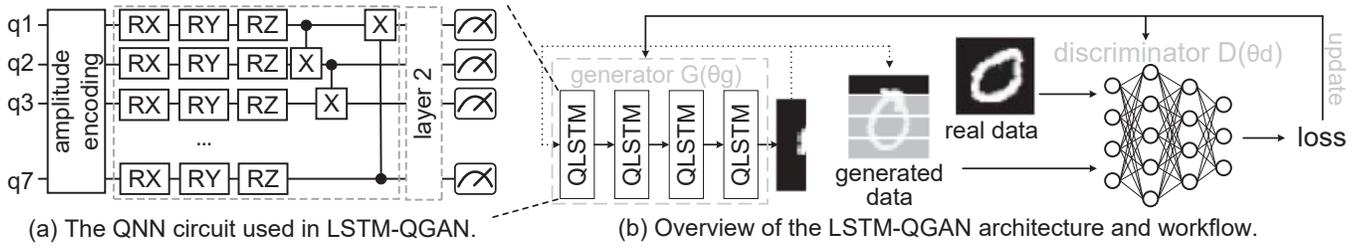


Fig. 5. The proposed LSTM-QGAN framework.

capture spatiotemporal information across patches with a single generator, rather than separate sub-generators as in [12]. Additionally, we are reengineering the quantum circuit ansatz within the QLSTM structure to improve hardware efficiency, fully addressing the NISQ constraints overlooked in previous QLSTM studies like [15].

IV. LSTM-QGAN

A. Overall Architecture

As illustrated in Figure 5(b), LSTM-QGAN utilizes QLSTM at the core of the generator to enhance scalability and resource efficiency. Like PatchGAN [12], the discriminator in LSTM-QGAN can be implemented using either a classical or quantum neural network, depending on the available quantum computing resources. The following outlines the key components and configurations in LSTM-QGAN.

- **Patch Inputs without PCA.** In line with [12], [13], LSTM-QGAN processes patched inputs to generate corresponding output patches, which are then recombined into a complete output. Unlike [10], [11], LSTM-QGAN eliminates the need for PCA and inverse PCA, processing the original data directly. This introduces a trade-off between resources (i.e., qubit number N) and processing latency (i.e., steps T). With an N -qubit implementation, LSTM-QGAN generates 2^N measured probabilities at each step as the output vector for each synthetic patched output. These vectors are then compared to the real patched input data in the discriminator. The total number of steps, T , is determined by $D/2^N$, where D represents the size of the real data.
- **Scalable QGAN with LSTM.** The generator in LSTM-QGAN consists of four QLSTM cells, as shown in Figure 5(b). Normally distributed noise z is input to the generator, producing the initial sub-image $G_{\theta_g}(z)$. The discriminator evaluates both synthetic and real input patches, computing the loss \mathcal{L} . Unlike PatchGAN [12], which requires separate generators for each patch—drastically increasing NISQ resource overhead with input size—LSTM-QGAN utilizes QLSTM’s ability to retain relevant patterns while discarding irrelevant information, independent of patch indices. To achieve this, gradients from all patches within a single input are averaged and applied to update the model parameters collectively, resulting in an image-adaptive generator that scales with increasing data dimensions while maintaining fixed resource usage.
- **Training Optimization.** Convergence in QGAN training is a critical challenge, significantly influenced by the choice of quantum loss function. Within the LSTM-QGAN framework, we evaluated both the conventional

TABLE I
DESIGN COST COMPARISON.

	PatchGAN [12]	LSTM-QGAN (Δ)
Qubits per QNN	5	7
Number of QNNs	56	8
Total Number of Qubits	280	56 (5 \times ↓)
Total Number of 1QG	1680	336 (5 \times ↓)
Total Number of 2QG	1344	112 (12 \times ↓)

binary cross entropy loss [6], [7], [12] and the Wasserstein loss [8], [13]. The specific Wasserstein loss used for LSTM-QGAN is detailed in Equation 2, where $\mathbb{L}_{\hat{x}} = \mathbb{E}_{\hat{x} \sim P_{\hat{x}}} [(\|\nabla_{\hat{x}} D_{\theta_d}(\hat{x})\|_2 - 1)^2]$, P_r and P_g represent the real data (i.e., x) and generated data (i.e., $\tilde{x} \in D_{\theta_d}(G_{\theta_g}(z))$) distributions, respectively. The distribution $P_{\hat{x}}$ is uniformly sampled between P_r and P_g , and λ is a constant. Experimental results on the impact of QGAN loss functions are discussed in Section V.

$$\min_{\theta_g} \max_{\theta_d} \mathbb{E}_{x \sim P_r} [D_{\theta_d}(x)] - \mathbb{E}_{\tilde{x} \sim P_g} [D_{\theta_d}(\tilde{x})] - \lambda \mathbb{L}_{\hat{x}} \quad (2)$$

B. NISQ Implementation

LSTM-QGAN offers flexibility in implementing G and D . For fair comparison, D is implemented as a classical neural network, as in PatchGAN [12]. In the QLSTM cells for G , we employed a hardware-efficient ansatz inspired by recent QNNs [21]–[23], instead of the generic circuit from [15]. Figure 5(a) shows the QNN circuit, which utilizes seven qubits. Each VAC block includes R_x , R_y , and R_z layers, followed by 2-qubit CX entanglement layer, with the VQC layers repeated twice. The measurement layer converts the quantum state into classical vectors. Although the gate count matches that in [15], our circuit uses native gates, while the $R(\alpha, \beta, \gamma)$ gate in [15] requires synthesis into multiple native gates.

Design Overhead. Table I compares the hardware resources required by PatchGAN and LSTM-QGAN for the MNIST dataset. Due to architectural differences, a QNN in PatchGAN refers to the quantum generator used for each input patch, while in LSTM-QGAN, it refers to the quantum module within the QLSTM. The last three rows of Table I highlight that LSTM-QGAN achieves a significant reduction: a 5 \times decrease in qubit counts, a 5 \times decrease in one-qubit gates (1QG), and a 12 \times decrease in two-qubit gates (2QG).

V. EXPERIMENTS AND RESULTS

A. Experimental Setup

Schemes and Benchmarks. We compare LSTM-QGAN with PatchGAN [12] using the MNIST dataset, which

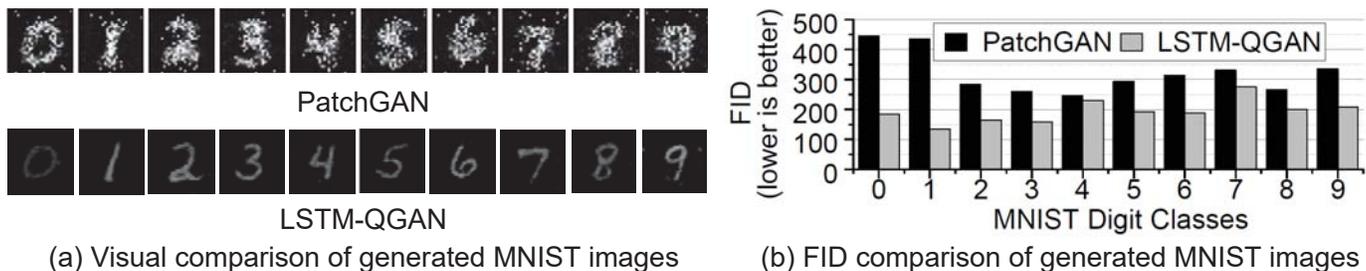


Fig. 6. Comparisons between images generated by PatchGAN and LSTM-QGAN.

consists of 28×28 grayscale images of handwritten digits 0~9. PatchGAN is implemented according to its original design [12], utilizing 5 qubits and 56 sub-generators. Each sub-generator produces a 14-pixel patch, and together, the 56 sub-generators generate the entire 784-pixel MNIST image. For LSTM-QGAN, we implement the generator with two QLSTM layers, each containing 4 QNNs with 7 qubits. At each time step, the LSTM-QGAN generates a 196-pixel patch, requiring 4 time steps to produce a complete MNIST image.

Simulation. All QGANs are implemented with the PennyLane and Torchquantum libraries. PatchGAN and LSTM-QGAN are trained using the ADAM optimizer with a $2e-4$ learning rate, a 128 batch size, and 1000 epochs. Quantum circuits are run on the NISQ `IBM_kyoto` computer [24].

Evaluation metrics. We evaluate the generated images using both qualitative (e.g., visual inspection) and quantitative methods. For quantitative assessment, we employ the Fréchet Inception Distance (FID), a widely recognized metric for measuring image similarity in GANs [25]. A lower FID score indicates a closer feature distance between real and generated images, signifying higher quality. In our experiments, we randomly select 500 real images and 500 generated images for comparison.

B. Results and Analysis

Comparison of Image Visual Quality. Figure 6(a) presents a visual comparison between the images generated by PatchGAN and LSTM-QGAN. PatchGAN demonstrates limited generation capabilities, as the outlines of the digits (0~9) are only vaguely identifiable, with noticeable white noise in the background. Additionally, the clarity of more complex digits, such as 4, 5, and 9, is particularly low, further highlighting its deficiencies. In contrast, LSTM-QGAN demonstrates superior image generation, producing sharper and more distinct digits with minimal noise, underscoring its enhanced capability in generating high-quality images.

Comparison of Image FID Scores. Figure 6(b) compares the FID scores of images generated by PatchGAN and LSTM-QGAN across different digit classes. The FID scores vary between the two models depending on the complexity and distinctiveness of each digit. Overall, LSTM-QGAN achieves lower FID scores than PatchGAN, indicating higher quality in the generated images. Specifically, PatchGAN shows significant variability, with its highest FID score at 445.22 (class

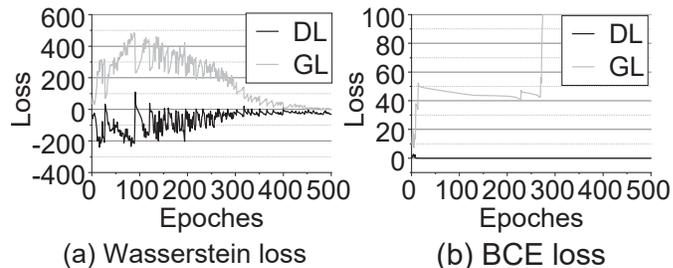


Fig. 7. Impact of loss functions on convergence.

0) and its lowest at 246.56 (class 4). In contrast, LSTM-QGAN consistently outperforms PatchGAN, with its highest FID score at 275.58 (class 7) and its lowest at 134.31 (class 1). On average, PatchGAN's FID score is 318.02, while LSTM-QGAN achieves a significantly lower average FID score of 193.28.

Impact of Loss Function. Figure 7 illustrates the impact of Wasserstein loss and binary cross-entropy (BCE) loss on the training convergence of LSTM-QGAN, comparing both generator loss (i.e., GL) and discriminator loss (i.e., DL). With Wasserstein loss, the DL initially decreases while GL increases, ultimately leading to convergence as training progresses. Conversely, with BCE loss, GL rapidly increases after several training cycles and stabilizes around 100, while DL drops sharply—indicating mode collapse, a known issue in GAN training. Although LSTM-QGAN with BCE loss can stabilize, achieving full convergence may require more sophisticated techniques. In contrast, Wasserstein loss offers greater training stability, resulting in smoother convergence.

VI. CONCLUSION

This work presents LSTM-QGAN, a quantum generative adversarial network (QGAN) architecture that overcomes key limitations in existing models. By eliminating reliance on principal component analysis (PCA) and integrating quantum long short-term memory (QLSTM), LSTM-QGAN achieves scalable performance with efficient resource use. As the first QGAN to incorporate QLSTM, this approach represents a significant advancement likely to inspire further research.

ACKNOWLEDGMENT

We thank the anonymous reviewers for their constructive and insightful comments. This work was supported in part by NSF OAC-2417589 and NSF CNS-2143120. We thank the IBM Quantum Researcher & Educators Program for their support of Quantum Computing Credits.

REFERENCES

- [1] Frank Arute et al., “Quantum supremacy using a programmable superconducting processor,” *Nature*, vol. 574, no. 7779, pp. 505–510, 2019.
- [2] Juan M Pino et al., “Demonstration of the trapped-ion quantum ccd computer architecture,” *Nature*, vol. 592, no. 7853, pp. 209–213, 2021.
- [3] Han-Sen Zhong et al., “Quantum computational advantage using photons,” *Science*, vol. 370, no. 6523, pp. 1460–1463, 2020.
- [4] Seth Lloyd and Christian Weedbrook, “Quantum Generative Adversarial Learning,” *Phys. Rev. Lett.*, vol. 121, pp. 040502, Jul 2018.
- [5] Pierre-Luc Dallaire-Demers and Nathan Killoran, “Quantum Generative Adversarial Networks,” *Phys. Rev. A*, vol. 98, pp. 012324, Jul 2018.
- [6] Christa Zoufal et al., “Quantum generative adversarial networks for learning and loading random distributions,” *npj Quantum Information*, vol. 5, no. 1, pp. 103, 2019.
- [7] Elton Yechao Zhu et al., “Generative quantum learning of joint probability distribution functions,” *Physical Review Research*, vol. 4, no. 4, pp. 043092, 2022.
- [8] Shouvanik Chakrabarti et al., “Quantum wasserstein generative adversarial networks,” *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [9] Murphy Y. Niu et al., “Entangling Quantum Generative Adversarial Networks,” *arXiv:2105.00080*, 2021.
- [10] Samuel A. Stein et al., “QuGAN: A Quantum State Fidelity based Generative Adversarial Network,” in *IEEE International Conference on Quantum Computing and Engineering*, 2021, pp. 71–81.
- [11] Cheng Chu et al., “Iqgan: Robust quantum generative adversarial network for image synthesis on nisq devices,” in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2023, pp. 1–5.
- [12] He-Liang Huang et al., “Experimental quantum generative adversarial networks for image generation,” *Physical Review Applied*, vol. 16, no. 2, pp. 024051, 2021.
- [13] Shu Lok Tsang et al., “Hybrid quantum-classical generative adversarial network for high resolution image generation,” *IEEE Transactions on Quantum Engineering*, 2023.
- [14] Kishor Bharti et al., “Noisy Intermediate-scale Quantum Algorithms,” in *Reviews of Modern Physics*, 2022, vol. 94, p. 015004.
- [15] Samuel Y. Chen et al., “Quantum long short-term memory,” in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2022, pp. 8622–8626.
- [16] Riccardo Di Sipio et al., “The dawn of quantum natural language processing,” in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2022, pp. 8612–8616.
- [17] Iordanis Kerenidis El Amine Cherrat et al., “Quantum vision transformers,” *arXiv:2209.08167*, 2022.
- [18] Alex Graves and Alex Graves, “Long short-term memory,” *Supervised sequence labelling with recurrent neural networks*, pp. 37–45, 2012.
- [19] Eoin Brophy et al., “Generative adversarial networks in time series: A systematic literature review,” *ACM Computing Surveys*, vol. 55, no. 10, pp. 1–31, 2023.
- [20] Yi Yu et al., “Conditional lstm-gan for melody generation from lyrics,” *ACM Transactions on Multimedia Computing, Communications, and Applications*, vol. 17, no. 1, pp. 1–20, 2021.
- [21] Sukin Sim et al., “Expressibility and entangling capability of parameterized quantum circuits for hybrid quantum-classical algorithms,” in *Advanced Quantum Technologies*, 2019, vol. 2, p. 1900070.
- [22] Cheng Chu et al., “QMLP: an error-tolerant nonlinear quantum MLP architecture using parameterized two-qubit gates,” in *ACM/IEEE International Symposium on Low Power Electronics and Design*, 2022, pp. 4:1–4:6.
- [23] Tirthak Patel et al., “OPTIC: A practical quantum binary classifier for near-term quantum computers,” in *Design, Automation & Test in Europe Conference & Exhibition*, 2022, pp. 334–339.
- [24] IBMQ, “Ibmq,” <https://quantum.ibm.com>.
- [25] Tuomas Kynkäänniemi et al., “The role of imagenet classes in fréchet inception distance,” *arXiv:2203.06026*, 2022.