

A Review of Quantum-Based Diffusion Models in Generative AI

Glauco Lima¹, Ernestas Filatovas², Marco Marcozzi² and Remigijus Paulavičius²

¹ Laboratory of Bioinformatics and Computational Biology, Federal University of ABC, Av. dos Estados, 5001 - Bangú, Santo André, Brazil

² Institute of Data Science and Digital Technologies, Vilnius University, Akademijos str. 4, LT-08412 Vilnius, Lithuania

glauco.endrigo@hotmail.com,

{ernestas.filatovas, marco.marcozzi, remigijus.paulavicius}@amif.vu.lt

Abstract. In recent years, the application of generative AI in several areas has been increasing. Concurrently, quantum computing has been advancing at an accelerated pace, unlocking new possibilities across various fields. This article provides an overview of the integration of quantum computing with generative AI, focusing on diffusion model techniques. We explore use cases documented in recent literature, illustrating how quantum computing techniques, when combined with diffusion models, are being leveraged to drive innovation.

Keywords: Quantum machine learning; Quantum diffusion model; Quantum Computing; Generative AI.

1 Introduction

Diffusion models have become key drivers of progress in generative Artificial Intelligence (AI) [1]. Their capabilities to produce diverse and high-quality synthetic data have led to widespread adoption. Platforms such as AlphaFold3 [2] and RFDiffusion [3] have achieved significant success using this technology in the field of generative biology. In generative imaging, DALL-E 2/3 from OpenAI [4], [5] and Stable Diffusion from Stability AI are good examples [6].

However, traditional diffusion models demand significant computational resources and encounter challenges in scalability and speed. These considerable resource demands present a major bottleneck in the broad deployment of diffusion models, as highlighted in recent research that outlines the high computational costs and the difficulties associated with efficiently training these models [7].

Recent computational challenges have spurred a growing interest in integrating quantum computing with diffusion models. This surge is paralleled by the rapid development of quantum technologies, particularly the advent of Noisy Intermediate-Scale Quantum (NISQ) devices [8]. NISQ processors are built on various quantum physical systems, each leveraging unique strategies. A widely adopted approach in quantum computing relies on superconducting circuits, utilizing qubits designed to operate at cryogenic temperatures [9]. These systems harness Josephson junctions to enable quantum coherence and gate operations. Additionally, superconducting circuits can be employed within quantum annealers, where optimization problems are solved through energy minimization processes [10]. An alternative paradigm is based on photonic quantum computing, where quantum information is encoded within the properties of light, such as polarization and phase. They can operate at room temperature and enable high-speed data transmission [11]. Another class of quantum computing platforms operates by directly manipulating individual atomic systems. Trapped ion-based architectures utilize electromagnetic fields to confine charged atomic species, leveraging their internal states for qubit encoding and employing laser pulses for quantum gate implementation [12]. Similarly, neutral atom-based computing relies on optical tweezers to trap Rydberg atoms [13]. Another type of quantum computing uses topologically protected Majorana-based qubits, and leverages particles called Majorana fermions to store and process quantum information [14].

The primary motivation for utilizing quantum computing lies in its ability to harness the unique properties of quantum mechanical systems, such as superposition, interference and entanglement. These features have the potential to enable the execution of computational tasks that would either be impossible or significantly more challenging on a classical supercomputer [15]. To the best of our knowledge, there is currently no survey on Quantum Diffusion Models (QDMs). This paper aims to fill that gap by reviewing and analyzing recent advances in the field, with a focus on comparing various QDM variants and their underlying architectures.

The remainder of this paper is organized as follows. Section 2 outlines the fundamental principles underlying diffusion models, key quantum encoding techniques and parameterized quantum circuits. Section 3 presents a review of QDMs, highlighting methodological innovations. Finally, Section 4 concludes the paper by summarizing the main findings.

2 Background

2.1 Diffusion Models

This section introduces the core architecture of Diffusion Models (DMs) and two important techniques used in Quantum Machine Learning (QML) that are key to these models.

Diffusion models are a class of generative models gaining traction for their robust performance in data synthesis, particularly in image generation. They operate through a two-phase process rooted in statistical physics. As shown in Figure 1, a data sample undergoes a forward diffusion process, where noise is incrementally added over multiple steps, transforming it into a near-random state. This systematic degradation is intuitive: the data “spreads out” until its original structure is lost. A neural network is then trained to reverse this process, learning to denoise the sample step by step, reconstructing it from noise into a coherent output matching the target distribution [16].




Fig. 1 Illustration of the forward diffusion (top arrow) and backward denoising (bottom arrow) processes in a diffusion model. Starting from the clean image x_0 (left), noise is progressively added to obtain increasingly corrupted versions $\{x_1, x_2, \dots, x_T\}$ (right). During sampling (reverse diffusion), the model iteratively denoises the noisy image x_T back into a clean reconstruction x_0 .

2.2 Types of Encoding

In QMLs, classical data must be encoded into quantum states to enable quantum processing. Various encoding methods have been developed for this purpose [17]. Here we review the ones that have been used in QDMs.

Amplitude encoding, frequently used in QDMs, takes advantage of quantum superposition to efficiently represent high-dimensional data. The corresponding quantum state for a classical data point $x = (x_1, x_2, \dots, x_N)$ is given by:

$$|x\rangle = \frac{1}{\sqrt{\sum_{k=1}^N x_k^2}} \sum_{k=1}^N x_k |i_k\rangle$$

where $|i_k\rangle$ denotes the computational basis state. A key benefit here is that an N -dimensional vector can be encoded using only $\log_2 N$ qubits, which is exponentially more efficient than the classical representation [18].

Another approach, known as *quantum embedding*, has been explored for QDM research [19]. Quantum embeddings leverage quantum computers to map classical data into a high-dimensional Hilbert space. A quantum feature map, implemented as a quantum circuit $\Phi(x, \theta)$, transforms an input x into a quantum state $\Phi(x, \theta)|0\dots 0\rangle$ where θ are trainable parameters. These parameters are optimized—often using classical techniques like gradient descent—to maximize the separation between quantum states corresponding to different classes, measured via metrics such as the Hilbert-Schmidt distance [19].

Additionally, *angle encoding* embeds a classical N -dimensional vector $x = (x_1, x_2, \dots, x_N)$ into an N -qubit product state via the encoded state:

$$|x\rangle = \bigotimes_{k=1}^N R(x_k)|0\rangle$$

where R is typically a single qubit rotation (e.g., R_y). In practice, one first normalizes each feature x_k to $[0, 1]$, then sets $\theta_k = 2 \arcsin(\sqrt{x_k})$ and applies $R_y(\theta_k)$ to qubit k , so that measuring each qubit recovers statistics tied to the original data [18].

2.3 Parameterized Quantum Circuits

Parameterized Quantum Circuits (PQCs) are essential to hybrid quantum-classical machine learning approaches [20]. They involve the following steps: (1) state preparation, typically initializing qubits in $|0\rangle^{\otimes n}$; (2) application of parameterized unitary operations (e.g., variational circuits); (3) measurement of the quantum state; and (4) classical optimization, where measurement outcomes are used to update the circuit parameters iteratively. This framework allows exploring high-dimensional Hilbert spaces to learn complex data patterns. However, it faces challenges such as barren plateaus, NISQ noise, and trade-offs between expressibility and trainability [20].

3 Quantum Generative Diffusion Models

This section examines various QDMs, with Table 1 summarizing the typical variants and their characteristics.

In the study [21], the authors propose a quantum adaptation of generative diffusion models, replacing traditional neural networks with PQCs. Two key variations were explored: a latent model and a conditioned model. The latent model employs a classical autoencoder to encode data into a lower-dimensional space before quantum processing, enabling the use of smaller PQCs and improving sample quality. Furthermore, the QDM was adapted into a conditioned version by increasing the Hilbert space dimension with additional qubits to encode labels, allowing for the generation of samples based on specific input conditions. According to the authors, increasing the number of measurements initially improves the quality of generated samples by introducing nonlinearity. One drawback they mention is that beyond a certain threshold, increasing measurements worsens performance by reducing sample variability. Excessive measurements reset qubits to a fixed state, erasing the initial noise information.

The work [22] collects measurement samples from multi-qubit states and then trains a variational diffusion model that progressively corrupts those samples with noise and employs a denoising network of residual and attention blocks to reverse the process. A key highlight of the approach is that it achieves high fidelity while using significantly less memory than RNNs and transformers—for learning the distribution of a 2-qubit quantum state. One limitation, however, is that all experiments are restricted to W and GHZ-type distributions.

The study [23] introduces a novel quantum generative model. This model transforms a target quantum state into a completely mixed state—a state where all possible quantum states are equally probable—through a non-unitary forward process, which increases entropy by not preserving state purity. This transformation is driven by a depolarization channel, a noise model that probabilistically replaces the state with the completely mixed state to simulate realistic quantum noise. The trainable backward process efficiently reconstructs the original state using parameter sharing, which reduces the number of parameters by reusing them across the model, and partial trace operations, which trace out auxiliary subsystems to maintain

non-unitarity while simplifying computations. Timestep embedding enhances this process by integrating information about the diffusion stage into the quantum state, guiding denoising. A pro of the work is its timestep embedding technique, which boosts performance by adaptively distributing quantum states on the Bloch sphere, enhancing temporal information learning over traditional qubit encoding. On the other hand, a con is the high gate complexity, $O(N^2)$, of the denoising circuit.

One notable work is the proposal [24], which introduces three distinct quantum approaches that leverage noise as a beneficial resource in generative modeling. In the first approach, the Classical-Quantum Generative Diffusion Model (CQGDM), the forward diffusion is executed using classical methods. However, the denoising phase is achieved by a Quantum Neural Network (QNN) that can be realized as a fully PQC or as a hybrid quantum-classical network. In the second approach, known as the Quantum-Classical Generative Diffusion Model (QCGDM), the diffusion process itself is performed in a quantum framework by subjecting quantum data to noise via quantum channels—such as depolarizing channels or dynamics modeled by Stochastic Schrödinger Equations—which gradually transform an initial quantum state into a maximally mixed state. The backward denoising is then handled by classical neural networks. The third approach, the Quantum-Quantum Generative Diffusion Model (QQGDM), fully embeds both the diffusion and denoising phases within a quantum domain. The forward process employs quantum noise channels to degrade an initial quantum state into a completely mixed state. The reverse process is carried out using PQCs that incorporate interactions with ancillary qubit systems that act as an environment and are then traced out. This fully quantum implementation enables the exploration and manipulation of complex quantum probability distributions. A highlight is that QQGDM generates non-classical probability distributions, achieving a high average quantum fidelity of 0.997 ± 0.013 in simulations for reconstructing one-qubit states. However, a drawback of their methodology is that its reliance on timestep-specific PQCs without time embedding complicates training, and the use of depolarizing channels limits noise diversity, risking barren plateaus and reduced scalability for broader quantum data distributions.

A thesis [25] explores the foundational principles of QDM using PQCs, highlighting key advancements in the field. This research introduces four

model variants: a base model defining the core architecture, a temporal model incorporating timestep embedding to encode temporal information, a conditional model enabling targeted image generation via label embedding and a hybrid model combining both embedding techniques. Additionally, this thesis presents a model capable of generating full-color images—an achievement that, to the best of our knowledge, is the first of its kind. A constraint is the reliance on single-qubit embedding for timestep and label information, which restricts the model's ability to scale to tasks requiring diverse or numerous labels, as the limited angular range reduces label distinguishability.

The authors of the work [26] proposed two quantum hybrid diffusion models for image synthesis by integrating variational quantum circuits (VQCs) into classical U-Net architectures. The first model, Quantum Vertex U-Net (QVU-Net), replaces ResNet convolutional layers at the U-Net's vertex with quantum layers. The second model, Quantvolutional U-Net (QuanvU-Net), incorporates quantum layers in the encoder for feature extraction. The PQCs are strategically integrated into the U-Net architecture at points where image dimensions are reduced, such as the vertex or second encoder level. This approach minimizes qubit requirements. However, as the authors themselves acknowledge, increasing the number of quantum circuits significantly raises computational demands, which slows training and inference times and creates scalability challenges for larger datasets or more complex models.

The work [27] introduces three novel quantum diffusion based algorithms, Label Guided Generation Inference (LGGI), Label Guided Denoising Inference (LGDI) and Label Guided Noise Addition Inference (LGNAI) to address few-shot learning challenges. The authors leverage QDMs to enhance data generation and inference under limited training samples. LGGI generates synthetic data to augment training for quantum neural networks (QNNs), while LGNAI and LGDI guide noise addition and removal during diffusion/denoising stages using label information. They strategically perform amplitude encoding on classical features and angle encoding on labels during training. They point out that if too many diffusion steps are applied, the original information may degrade excessively into noise, causing the denoising process to overemphasize the label and reconstruct a generic class prototype.

Table 1. Quantum Diffusion Model Studies and Their Contributions

Study	Contribution	Platform (qubits used)	Dataset
Cacioppo et al.(2023)	Proposes QDM with PQCs, introducing latent and conditioned variants for improved sample quality and conditional generation.	PennyLane, 27q - IBM Hanoi (8q full model, 3q latent, 7q conditioned)	Quantum Simulator: MNIST digits {0,1} (16×16), Latent MNIST digits 0–9 (28×28). Quantum Hardware: Reduced latent MNIST digits {0,1} (dimension 4).
Wang et al., (2023)	Uses forward noise-corruption and a parallel reverse denoising network of residual and attention blocks.	Classical system	Measurement outcomes from multi-qubit quantum states.
Chen and Zhao (2024)	Non-unitary forward process (depolarization channel) with parameter-shared, timestep-embedded backward reconstruction.	Tensorcircuit framework (1–8q)	Random quantum states
Parigi et al. (2024)	Classical-Quantum, Quantum-Classical, and Quantum-Quantum Generative DMs using quantum noise, PQCs, and classical NNs.	PennyLane; CQGDM: 4q, QCGDM: 1q, QQGDM: 2q	CQGDM: 1,000 points uniformly in [-1,1]. QCGDM/QQGDM: Random 1-qubit pure states.
Kivijervi (2024)	Explores foundational QDM variants (base, temporal, conditional, hybrid) achieving full-color image generation.	PennyLane (6q base, 7q temporal/ conditional, 8q hybrid)	MNIST: 1,024 images resized to 16×16
De Falco et al. (2024)	Improved image quality, faster convergence, fewer parameters via transfer learning.	PennyLane/Flax (12q for VQC, 4q final channel)	MNIST and Fashion MNIST (28×28, 60,000 images)
Wang et al (2024)	Few-shot quantum diffusion algorithms: LGGI, LGNAI, LGDI.	IBM Almaden	MNIST (28×28), Digits MNIST (8×8), Fashion MNIST (28×28)
Shah and Vatsa (2025)	Pairwise Bell-state entangling to reduce qubits required.	PennyLane (6q/8q/10q for 8×8/16×16/32×32 images)	MNIST (28×28, 70,000 images); CIFAR-10 (32×32, 60,000 images). Tested at 8×8, 16×16, 32×32.
Wang et al. (2025)	Adapts DPM ODEs with Carleman linearization for quantum implementation.	Classical system	ImageNet-100 (128×128, 256×256, 512×512) in latent spaces (16×16×4, 32×32×4, 64×64×4)
Han and Patel (2025)	Quantum noise in generative diffusion models enhancing randomness.	127q - IBM Naxca	MNIST: 5,000 images

The study [28] presents an advancement over earlier methods, enabling the processing of higher-dimensional images with intricate pixel structures, even on platforms with limited qubits. They first flatten the 2D image data into a 1D vector, normalize it, and then apply amplitude encoding. Following this, the quantum circuit implements pairwise Bell-state preparation: Hadamard gates are applied to the first half of the qubits (excluding any ancilla qubits) to create superposition, and CNOT gates are used to entangle each qubit from the first half (control qubits from 0 to $n/2 - 1$) with a corresponding qubit in the second half (indices $n/2$ to $n - 1$). This entanglement establishes strong correlations, and this allows the subsequent PQC to operate on a reduced set of qubits while still accessing information from the entire input state. A highlight is that unlike hybrid models, which rely on classical autoencoders, their work reduces parameter count directly within the quantum circuit. A limitation, however, is their work's high computational time compared to classical models; the paper notes an exponential increase during classical simulations of the quantum circuits.

Recent research [29] has transformed the mathematical equations underlying Denoising Diffusion Probabilistic Models (DPMs), particularly the ordinary differential equations (ODEs), into a form suitable for processing by quantum computers through a technique known as “Carleman linearization”. This work focuses on two key approaches: DPM-solver-k, which utilizes precise mathematical derivatives to approximate the model's behavior, and UniPC, which employs measurements at various points to estimate the model's evolution. The authors provide theoretical proof that their quantum algorithms would perform efficiently on future fault-tolerant quantum computers. Therefore, the authors take explicit ‘out-of-the-box steps’ to establish a new avenue for demonstrating quantum computing's utility in machine learning tasks. A critical downside, however, is that the truncation used in their method introduces approximation errors.

A recent work [30] introduces an approach involving three main steps. First, they encode each image into a quantum circuit by applying PCA followed by an angle embedding, transforming the image data into a format suitable for quantum processing. Next, they introduce noise using quantum gates: each qubit receives a rotation via an R_x gate to add noise, and then additional paired rotations that cancel out net movement but compound the noise effect. Finally, after running the circuit on quantum hardware, they measure the qubits to obtain marginal probability distributions,

reverse the angle embedding, and apply an inverse PCA to reconstruct the noisy image for training the diffusion model. Ingeniously, they repurposed quantum decoherence as a source of true physics-driven noise, bypassing classical pseudo-randomness. A downside of their method is that PCA and angle embedding discard spatial details during encoding, which degrades reconstruction fidelity.

Overall, the works on QDMs reviewed here operate within a modest qubit range of 1 to 16. A central trend across these studies is the optimization of qubit usage to overcome the limitations of current quantum hardware. This is achieved through techniques such as encoding data into lower-dimensional latent spaces prior to quantum processing [21], employing pairwise Bell-state preparation to exploit entanglement for efficient information access with fewer qubits [28], and resizing images to conserve computational resources [25].

Another prominent trend is the strategic use of quantum noise, stemming from inherent quantum fluctuations. Studies like [24] and [30] highlight how noise, introduced via quantum channels (e.g., depolarization channels) or gate-based rotations, facilitates the creation of complex, entanglement-driven probability distributions that are computationally infeasible to replicate classically. Predominantly, PQC are employed to reverse the diffusion process [21], [24], [25], [28]. Additionally, there is a clear trajectory toward integrating temporal and conditional information into QDMs, with timestep embedding [23], [25] and label-based conditioning [21], [25] enabling more precise control over the generative process. These trends suggest a field moving toward resource-efficient, noise-augmented, and hybrid models capable of tackling increasingly sophisticated generative tasks.

4 Conclusions

Quantum generative diffusion models draw inspiration from classical diffusion models, which have recently demonstrated state-of-the-art performance in a variety of generative tasks. In this work, we have reviewed recent advances in quantum diffusion models, focusing on aspects such as model architectures. As highlighted in the introduction, diffusion models are pivotal across diverse domains, making their enhancement through quantum computing crucial to overcoming classical limitations in scalability

and efficiency. This review, the first survey in the field of QDMs up to this date, consolidates advancements and provides foundational insights to propel the discipline forward.

Future research should prioritize rigorous evaluations of QDM architectures and methodologies. One suggested avenue for exploration involves leveraging quantum diffusion models in the context of graph-structured data. Such data, with their inherent relational complexity, could rigorously evaluate QDMs' ability to model interconnected systems.

References

- [1] Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Wentao Zhang, Bin Cui, and Ming-Hsuan Yang. Diffusion models: A comprehensive survey of methods and applications. *ACM Computing Surveys*, 56(4):1–39, 2023.
- [2] Josh Abramson, Jonas Adler, Jack Dunger, Richard Evans, Tim Green, Alexander Pritzel, Olaf Ronneberger, Lindsay Willmore, Andrew J Ballard, Joshua Bambrick, et al. Accurate structure prediction of biomolecular interactions with alphafold 3. *Nature*, 630(8016):493–500, 2024.
- [3] Joseph L Watson, David Juergens, Nathaniel R Bennett, Brian L Trippe, Jason Yim, Helen E Eisenach, Woody Ahern, Andrew J Borst, Robert J Ragotte, Lukas F Milles, et al. De novo design of protein structure and function with rdiffusion. *Nature*, 620(7976):1089–1100, 2023.
- [4] Gary Marcus, Ernest Davis, and Scott Aaronson. A very preliminary analysis of dall-e 2. *arXiv preprint arXiv:2204.13807*, 2022.
- [5] Mohamad-Hani Temsah, Abdullah N Alhuzaimi, Mohammed Almansour, Fadi Al-jamaan, Khalid Alhasan, Munirah A Batarfi, Ibraheem Altamimi, Amani Alharbi, Adel Abdulaziz Alsuhaibani, Leena Alwakeel, et al. Art or artifact: evaluating the accuracy, appeal, and educational value of ai-generated imagery in dall-e 3 for illustrating congenital heart diseases. *Journal of Medical Systems*, 48(1):54, 2024.
- [6] Johanna Karras, Aleksander Holynski, Ting-Chun Wang, and Ira Kemelmacher-Shlizerman. Dreampose: Fashion video synthesis with stable diffusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22680–22690, 2023.
- [7] Tianshuo Xu, Peng Mi, Ruilin Wang, and Yingcong Chen. Towards faster training of diffusion models: An inspiration of a consistency phenomenon. *arXiv preprint arXiv:2404.07946*, 2024.
- [8] Kishor Bharti, Alba Cervera-Lierta, Thi Ha Kyaw, Tobias Haug, Sumner Alperin-Lea, Abhinav Anand, Matthias Degroote, Hermanni Heimonen, Jakob S Kottmann, Tim Menke, et al. Noisy intermediate-scale quantum algorithms. *Reviews of Modern Physics*, 94(1):015004, 2022.
- [9] Wei-Yang Liu, Dong-Ning Zheng, and Shi-Ping Zhao. Superconducting quantum bits. *Chinese Physics B*, 27(2):027401, 2018.
- [10] Sheir Yarkoni, Elena Raponi, Thomas Bäck, and Sebastian Schmitt. Quantum annealing for industry applications: Introduction and review. *Reports on Progress in Physics*, 85(10):104001, 2022.

- [11] Shuntaro Takeda and Akira Furusawa. Toward large-scale fault-tolerant universal photonic quantum computing. *APL Photonics*, 4(6), 2019.
- [12] James D Siverns and Qudsia Quraishi. Ion trap architectures and new directions. *Quantum Information Processing*, 16:1–42, 2017.
- [13] Karen Wintersperger, Florian Dommert, Thomas Ehmer, Andrey Hoursanov, Johannes Klepsch, Wolfgang Mauere, Georg Reuber, Thomas Strohm, Ming Yin, and Sebastian Luber. Neutral atom quantum computing hardware: performance and end-user perspective. *EPJ Quantum Technology*, 10(1):32, 2023.
- [14] David Aasen, Morteza Aghaee, Zulfi Alam, Mariusz Andrzejczuk, Andrey Antipov, Mikhail Astafev, Lukas Avilovas, Amin Barzegar, Bela Bauer, Jonathan Becker, et al. Roadmap to fault tolerant quantum computation using topological qubit arrays. *arXiv preprint arXiv:2502.12252*, 2025.
- [15] Yazhen Wang. *Quantum computation and quantum information*. 2012.
- [16] Paul A Geroski. Models of technology diffusion. *Research policy*, 29(4-5):603–625, 2000.
- [17] Deepak Ranga, Aryan Rana, Sunil Prajapat, Pankaj Kumar, Kranti Kumar, and Athanasios V Vasilakos. Quantum machine learning: Exploring the role of data encoding techniques, challenges, and future directions. *Mathematics*, 12(21):3318, 2024.
- [18] Tuan A Ngo, Tuyen Nguyen, and Truong Cong Thang. A survey of recent advances in quantum generative adversarial networks. *Electronics*, 12(4):856, 2023.
- [19] Seth Lloyd, Maria Schuld, Aroosa Ijaz, Josh Izaac, and Nathan Killoran. Quantum embeddings for machine learning. *arXiv preprint arXiv:2001.03622*, 2020.
- [20] Daniel T Chang. Parameterized quantum circuits with quantum kernels for machine learning: A hybrid quantum-classical approach. *arXiv preprint arXiv:2209.14449*, 2022.
- [21] Andrea Cacioppo, Lorenzo Colantonio, Simone Bordoni, and Stefano Giagu. Quantum diffusion models. *arXiv preprint arXiv:2311.15444*, 2023.
- [22] Yong Wang, Shuming Cheng, Li Li, and Jie Chen. Learning quantum distributions with variational diffusion models. *IFAC-PapersOnLine*, 56(2):5888–5893, 2023.
- [23] Chuangtao Chen and Qinglin Zhao. Quantum generative diffusion model. *arXiv e-prints*, pages arXiv–2401, 2024.
- [24] Marco Parigi, Stefano Martina, and Filippo Caruso. Quantum-noise-driven generative diffusion models. *Advanced Quantum Technologies*, page 2300401, 2024.
- [25] Nikolai Theien Kivijervi. Quantum diffusion model. Master's thesis, 2024.
- [26] Francesca De Falco, Andrea Ceschini, Alessandro Sebastianelli, Bertrand Le Saux, and Massimo Panella. Quantum hybrid diffusion models for image synthesis. *KI-Künstliche Intelligenz*, pages 1–16, 2024.
- [27] Ruhan Wang, Ye Wang, Jing Liu, and Toshiaki Koike-Akino. Quantum diffusion models for few-shot learning. *arXiv preprint arXiv:2411.04217*, 2024.
- [28] Shivalee RK Shah and Mayank Vatsa. Enhancing quantum diffusion models with pairwise bell state entanglement. In *International Conference on Pattern Recognition*, pages 347–361. Springer, 2025.
- [29] Yunfei Wang, Ruoxi Jiang, Yingda Fan, Xiaowei Jia, Jens Eisert, Junyu Liu, and Jin-Peng Liu. Towards efficient quantum algorithms for diffusion probability models. *arXiv preprint arXiv:2502.14252*, 2025.
- [30] Jason Han and Tirthak Patel. Turning quantum noise on its head: Using the noise for diffusion models to generate images. *ACM SIGMETRICS Performance Evaluation Review*, 52(4):23–24, 2025.