

Quantum Diffusion Models for Few-Shot Learning

Wang, Ruhan; Wang, Ye; Liu, Jing; Koike-Akino, Toshiaki

TR2026-042 April 15, 2026

Abstract

Modern quantum machine learning (QML) methods involve the variational optimization of parameterized quantum circuits on training datasets, followed by predictions on testing datasets. Most state-of-the-art QML algorithms currently lack practical advantages especially in few-shot learning tasks. In this work, we propose three new frameworks employing quantum diffusion model (QDM) as a solution for the few-shot learning: label-guided generation inference (LGGI); label-guided denoising inference (LGDI); and label-guided noise addition inference (LGNAI). Experimental results demonstrate that our proposed algorithms significantly outperform existing methods.

Springer Nature 2026

Quantum Diffusion Models for Few-Shot Learning

Ruhan Wang^{1*}, Ye Wang²[0000–0001–5220–1830], Jing Liu²[0000–0002–1712–2966],
and Toshiaki Koike-Akino²[0000–0002–2578–5372]

¹ Indiana University,
107 S Indiana Ave, Bloomington, IN 47405, USA
ruhwan@iu.edu

² Mitsubishi Electric Research Laboratories (MERL),
201 Broadway, Cambridge, MA 02139, USA
{yewang, jiliu, koike}@merl.com

Abstract. Modern quantum machine learning (QML) methods involve the variational optimization of parameterized quantum circuits on training datasets, followed by predictions on testing datasets. Most state-of-the-art QML algorithms currently lack practical advantages especially in few-shot learning tasks. In this work, we propose three new frameworks employing quantum diffusion model (QDM) as a solution for the few-shot learning: label-guided generation inference (LGGI); label-guided denoising inference (LGDI); and label-guided noise addition inference (LGNAI). Experimental results demonstrate that our proposed algorithms significantly outperform existing methods.

Keywords: Quantum Machine Learning (QML) · Diffusion Models · Few-Shot Learning (FSL) · Quantum Neural Networks (QNN).

1 Introduction

Quantum machine learning (QML) has emerged as a transformative tool for automated decision-making, influencing diverse fields such as finance, healthcare, and drug discovery [9–12]. However, in the specific context of few-shot learning, where only limited data is available for training, QML has yet to reach its full potential, often demonstrating suboptimal performance. In classical machine learning, diffusion models have exhibited significant promise in tackling few-shot learning challenges [4, 5, 26, 27]. Nevertheless, the use of quantum diffusion models (QDMs) for few-shot learning remains largely unexplored [1]. This gap is mainly attributed to the limitations of current quantum computing resources and the intrinsic noise that hinders the performance of QDMs, despite their demonstrated success in generative tasks [13].

In this work, we propose a suite of novel algorithms based on the QDM framework to address the quantum few-shot learning (QFSL) problem. Our contributions are as follows:

* This work was conducted while R.W. was an intern at MERL.




Fig. 1. Variational quantum circuits (VQC) used for quantum neural networks.

- **Introduction of QDM-Based Few-shot Learning Algorithms:** Leveraging the QDM’s unique features, we introduce three new algorithms:
 - Leveraging the QDM’s generative strengths, we present the **QDM-based Label-Guided Generation Inference (QDM-LGGI)** algorithm, by harnessing label-guided generative capabilities.
 - To fully exploit the diffusion stage, we introduce **QDM-based Label-Guided Noise Addition Inference (QDM-LGNAI)**.
 - We propose yet another algorithm: **QDM-based Label-Guided Denoising Inference (QDM-LGDI)**, tailored to perform test inference during the denoising stage.
- **Extensive Comparative Analysis:** We evaluate our proposed algorithms across multiple benchmark datasets and compare them against established baselines, demonstrating the superior performance and robustness of our approaches. Additionally, we analyze the effect of quantum noise on the model’s effectiveness for a real quantum processor.
- **Comprehensive Ablation Study:** We conduct a thorough ablation study to assess the influence of various components and hyperparameters on algorithmic performance. This analysis also extends to examining the potential of our proposed approaches in zero-shot learning, providing insights into their adaptability beyond few-shot contexts.

2 Background

Quantum Neural Network (QNN): Quantum Neural Networks (QNNs) represent a class of models that extend traditional neural network concepts into the quantum domain, offering powerful computational tools for diverse machine learning tasks. A typical QNN consists of three core components: a data encoder $E(x)$ that maps classical input data x into a quantum state $|x\rangle$, a variational quantum circuit (VQC) Q that manipulates the quantum state to generate a processed output, and a measurement layer M that projects the quantum output back to a classical vector. Various VQC ansatzs commonly used in QNNs are depicted in Fig. 1 [14, 17, 16, 15]. The model is trained by minimizing a loss function that quantifies the difference between the QNN’s output and the target label y . This training process involves hybrid quantum-classical optimization, where parameters in the VQC are iteratively updated to minimize the loss.

Few-Shot Learning (FSL): FSL addresses the challenge of performing supervised learning with limited labeled data [23, 24]. It is typically structured

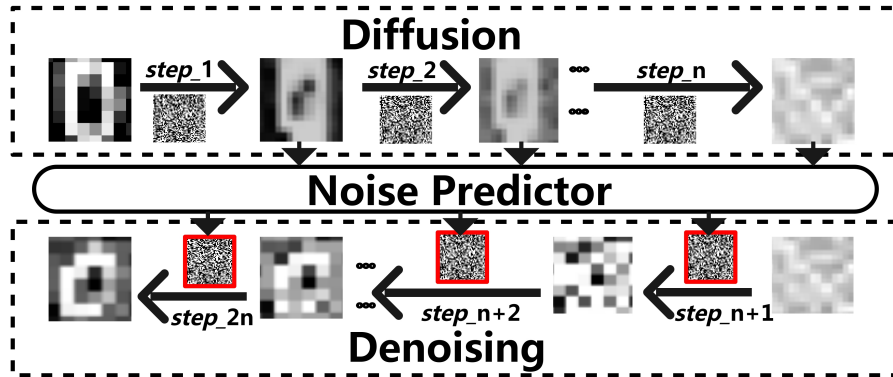


Fig. 2. Framework of Quantum Denoising Diffusion Model (QDDM). The **Noise Predictor** estimates noise present in the noisy image data.

around two main sets: a support set, containing a small number of labeled examples across n classes with k examples per class (known as n -way k -shot learning), and a query set, which contains unlabeled examples that the model must classify into one of the n classes. Solutions to FSL generally fall into data-based, model-based, and algorithm-based approaches [6].

Quantum Few-Shot Learning (QFSL): QFSL leverages QNNs as classifiers to tackle FSL tasks within the quantum realm [8, 7]. However, conventional QFSL algorithms often struggle with performance limitations due to the constraints of quantum computing resources and the presence of noise in quantum hardware. Overcoming these challenges is key to advancing the applicability of QFSL in practical QML scenarios.

Diffusion Model (DM): Diffusion models [18, 19] are generative models designed to produce realistic high-dimensional data, such as images. A diffusion model operates through two primary stages: the diffusion process, which gradually adds noise to the data over multiple steps, transforming it into a simpler distribution; and the denoising process, which reverses the diffusion noise to reconstruct the original data.

Quantum Diffusion Model (QDM): DMs were extended to quantum domain as QDMs for generative tasks, such as quantum state generation and quantum circuit design. The Quantum Denoising Diffusion Model (QDDM) [1] is a leading approach in this field, surpassing classical models of similar size by harnessing the computational advantages of quantum mechanism. Fig. 2 presents the QDDM framework, while Fig. 3 illustrates its image generation process under label guidance.

Our work augments the QDM with a label-guided mechanism, enhancing its capability to tackle QFSL tasks. This extension incorporates an additional qubit and applies a Pauli-X rotation by an angle of $\frac{2\pi y}{n}$, where y denotes the target label and n represents the total number of classes. Leveraging QDDM’s state-of-the-art performance, we adopt it as the foundational QDM framework



Fig. 3. Generated images using QDDM under different label guidance. The model takes random noise as input, and the noisy image undergoes a gradual denoising process guided by the specified label.

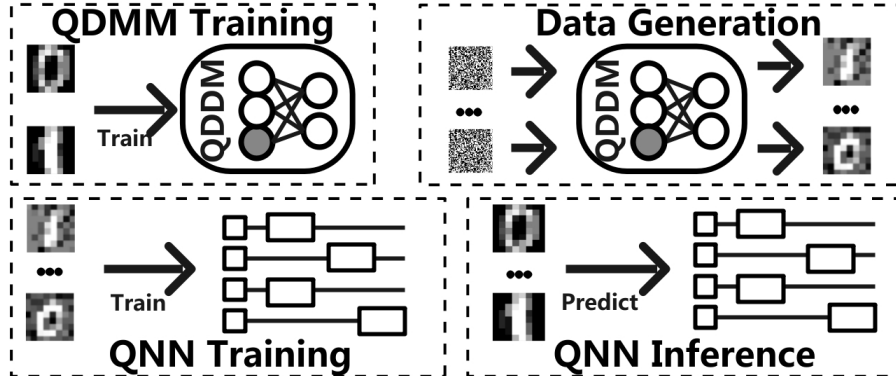


Fig. 4. Framework of QDDM-based Label-Guided Generation Inference (QDM-LGGI). A small set of labeled data is used to train the label-guided QDDM, which is then utilized to generate additional synthetic data. The synthetic data are subsequently used to train the QNN for the inference task. The gray-filled circle represents the embedded label.

in this study. Our implementation includes multiple strongly entangling layers for denoising diffusion.

3 Method

To address QFSL challenges, we propose methods from both data and algorithmic perspectives. From the data perspective, we leverage QDDM to augment the training samples, using the generated data to enhance QNN training and thereby improve its prediction accuracy on real data. From an algorithmic standpoint, we employ two strategies to guide QDDM in completing the inference process across two distinct stages: diffusion and denoising.

3.1 QDM-Based Label-Guided Generation Inference (QDM-LGGI)

In QFSL, the limited availability of labeled training data often hampers the performance of QNNs. To address this limitation, expanding the training dataset through data augmentation techniques is essential for enhancing model accuracy




Fig. 5. Framework of QDDM-based Label-Guided Denoising Inference (QDDM-LGDI). The term \tilde{c}_m^n represents the pre-LGDI. Solid circles in different colors dictated noise at step m associated with label L_i . The term L_0/L_1 -loss denotes the difference between the true noise and the predicted noise under the guidance of different labels L_i .




Fig. 6. Framework of QDDM-based Guided Noise Addition Inference (QDDM-LGNAI). Solid circles in different colors represent distinct embedded labels. The images, each framed by a square of varying colors, indicate the generated noise under the guidance of different labels L_i .

and robustness. Leveraging the QDDM, known for its strong generative capabilities, provides a powerful solution to augment the dataset effectively in QFSL tasks.

Our QDDM-LGDI algorithm is designed to maximize QFSL outcomes by generating high-quality, label-guided synthetic data. Initially, we use a separate pre-trained dataset to conduct the QDDM’s pretraining, followed by fine-tuning on a small subset of labeled data. Once fine-tuned, the QDDM acts as a generative model, producing additional samples that closely approximate the distribution of the original data. These synthetic samples are incorporated into the training set, expanding the dataset used to train the QNN. This augmentation process leads to a more comprehensive dataset, which enhances the QNN’s inference accuracy on real-world data by improving its generalization to unseen examples.

To further refine the data generated by the QDDM, we employ a label-guided generation technique. During QDDM training, amplitude encoding is applied to represent classical data in the quantum state space, while angle encoding is used to embed label information. In the generation phase, random noise combined with a specified label is provided as input to the QDDM, which generates data aligned with that label. This method ensures that the synthetic data is label-specific, which increases the QNN’s ability to distinguish between classes accurately.

Fig. 3 illustrates the data generation process for different label-guidance scenarios, showing how label-specific data distributions are created. Fig. 4 presents the framework of the QDDM-LGDI algorithm, highlighting the integration of label guidance into the QDDM’s generative process to produce targeted data for QFSL.

3.2 QDDM-Based Label-Guided Noise Addition Inference (QDDM-LGNAI)

The learning objective of the QDDM relies on the accuracy of a noise predictor to estimate and reduce noise within noisy data, aiming to minimize the discrepancy

between predicted noise and actual noise. In QDDM, this noise prediction process is guided by label information, where each label is associated with a unique noise prediction that reflects the characteristics of that label. By applying the correct label guidance, the error between the predicted and actual noise can be minimized, thus improving inference accuracy. Based on this principle, we introduce the QDM-Based Label-Guided Noise Addition Inference (QDM-LGNAI) method, illustrated in Fig. 5.

The QDM-LGNAI process begins with finetuning the QDDM on a small subset of labeled data. Once trained, the noise predictor \mathcal{P} within the QDDM is used for inference. Given an input x_0 , with potential labels $\{L_1, L_2, \dots, L_m\}$, noise is progressively added to x_0 across \mathcal{T} iterations to simulate various levels of data degradation. At each time step t , the noisy data x_t is calculated as $x_{t-1} + \epsilon_t$, where $\epsilon_t \sim \mathcal{N}(x_{t-1}, \mathcal{W}[t])$, and \mathcal{W} is a weight matrix that controls the magnitude of noise added at each step.

The noise predictor \mathcal{P} is then used to estimate the noise in x_t for each label, producing a set of label-specific noise predictions $\{\mathcal{P}(x_t|L_1), \dots, \mathcal{P}(x_t|L_m)\}$. This label-guided noise prediction allows QDM-LGNAI to evaluate how well each label aligns with the observed noisy data.

To identify the correct label, we calculate the mean squared error (MSE) between each predicted noise value $\mathcal{P}(x_t|L_i)$ and the actual noise ϵ_t . By minimizing this discrepancy, QDM-LGNAI aims to select the label that most accurately explains the noisy data across all \mathcal{T} iterations. The label with the minimum cumulative error over \mathcal{T} time steps is chosen as the predicted label, as follows: $\arg \min_{L_i \in \mathcal{L}} \sum_{t=1}^{\mathcal{T}} \text{MSE}(\mathcal{P}(x_t|L_i), \epsilon_t)$.

This approach leverages the QDDM’s ability to capture label-specific noise patterns, allowing QDM-LGNAI to enhance classification accuracy even in the presence of noisy inputs. By aligning predicted noise with actual noise under different label conditions, QDM-LGNAI improves the robustness of QFSL applications, especially where noisy data is prevalent.

3.3 QDM-Based Label-Guided Denoising Inference (QDM-LGDI)

Our QDM-Based Label-Guided Denoising Inference (QDM-LGDI) method leverages the denoising capabilities of the QDDM to iteratively restore noisy data, steering it toward its original structure. In QDM-LGDI, the QDDM’s noise predictor estimates the noise in each sample and iteratively subtracts it over a series of steps. By introducing label-specific guidance at each denoising step, QDM-LGDI ensures that the data generated under the correct label aligns most closely with the original data, enhancing the accuracy of inference.

Starting with an initial input x_0 , noise is incrementally added across \mathcal{T} iterations, producing a sequence of progressively noisier data points $\{x_1, x_2, \dots, x_{\mathcal{T}}\}$. For each noisy data point, the QDDM’s noise predictor \mathcal{P} estimates the noise under a specified label L_i , with $\mathcal{P}(x_{\mathcal{T}}|L_i)$ representing the label-guided noise estimate. This predicted noise is then subtracted iteratively, resulting in progressively denoised data points $\{x_{\mathcal{T}+1}, x_{\mathcal{T}+2}, \dots, x_{2\mathcal{T}}\}$, where each denoising step is defined by: $x_{\mathcal{T}+t+1}|L_i = x_{\mathcal{T}+t} - \mathcal{P}(x_{\mathcal{T}+t}|L_i)$.

This iterative denoising sequence ultimately reconstructs a fully restored version of the data, with the goal of identifying the label L_i that provides the most accurate denoising guidance. This label selection ensures that the denoised output aligns as closely as possible with the original data state prior to noise addition.

To determine the correct label, QDM-LGDI calculates the MSE between each denoised data point $x_{2\mathcal{T}-t}|L_i$ and the corresponding noisy data point x_t , aggregated across all \mathcal{T} iterations. This cumulative MSE provides a measure of discrepancy between the generated data and the original data for each label. The label that minimizes this cumulative error over all denoising steps is chosen as the most accurate.

By aligning noise prediction with label-specific guidance, QDM-LGDI empowers the QDDM to generate data that is both accurate and closely aligned with the correct label. This approach is particularly advantageous in scenarios requiring fine distinctions between closely related classes, as it leverages subtle label-specific variations in noise characteristics. This label-guided denoising technique refines the QFSL algorithm’s classification performance, enhancing its capacity to discern nuanced patterns within noisy data environments.

4 Experiments

In this section, we outline the foundational settings of our experimental setup, detailing the datasets, baseline algorithms, and parameter configurations. We further conduct a series of experiments to address the following research questions, with each question explored in its respective subsection:

- What are the performance advantages of our proposed QDM-based algorithms compared to existing baseline methods?
- Which factors most significantly influence the performance of our algorithms?
- How effectively do our algorithms address the zero-shot learning problem?

4.1 Basic Experimental Settings

We begin by providing a detailed description of the datasets utilized, the baseline algorithms chosen, and the specific configurations applied in our experiments.

Dataset: We use the Digits MNIST [25], MNIST [21], and Fashion MNIST [22] datasets, which allow for an assessment of our algorithms across varied data characteristics. For the 2-way k -shot tasks, we select classes 0 and 1 from both Digits MNIST and MNIST, and the T-shirt and Trouser classes from Fashion MNIST. In the 3-way k -shot tasks, we include classes 0, 1, and 2 from Digits MNIST and MNIST, and the T-shirt, Trouser, and Pullover classes from Fashion MNIST. The remaining categories in each dataset are reserved for QDDM pretraining; for example, classes 4–9 in MNIST are used to pretrain QDDM for evaluating QDM-based QFSL algorithms on MNIST. To simulate few-shot learning, we use one image per class for one-shot tasks and ten images per class

Table 1. QNN architectures used in the few-shot learning task. The variable n denotes the number of layers in the QNN, with a default value of $n = 1$. **1QG** refers to single-qubit gates, while **2QG** represents two-qubit gates.

QNN	# Qubits	1QG	2QG	# Param.
QMLP	6	ROT	CRX	$24 \times n$
C14	6	RY	CRX	$12 \times n$
OPTIC	6	ROT	CNOT	$18 \times n$
Quantumnat	6	U3	CU3	$36 \times n$

Table 2. Parameter configurations used in QDDM training. **LR** denotes the learning rate, while **Diff-Step** specifies the number of diffusion steps.

Dataset	Model Shape	LR	Diff-Step
Digit MNIST	47 layers	0.00097	10
MNIST	60 layers	0.00211	10
Fashion MNIST	121 layers	0.00014	10

for ten-shot tasks. During inference, each class is represented by 200 images in the evaluation dataset. For computational efficiency, all images are resized to 8×8 .

Baselines and Parameter Settings. To evaluate the performance of our QDM-based algorithms against established benchmarks, we select four prominent QNN architectures that are widely employed in QML for QFSL tasks [14, 17, 16, 15]. These architectures are depicted in Fig. 1, whose detailed specifications, including their structural configurations and parameter settings, are summarized in Table 1. During training, classical data is mapped to quantum states using amplitude encoding, ensuring efficient representation within the quantum framework. Optimization is conducted via the Adam optimizer, with a learning rate fixed at 0.001, to minimize the cross-entropy loss function over 40 iterations, providing a robust foundation for comparative analysis.

QDDM Training: Before applying the QDM-based algorithms to the QFSL tasks, we pre-train the QDDM using a label-guided, quantum-dense architecture. Labels are encoded through RX rotations, and strongly entangling layers [20] are used to effectively process the data. The training process employs the Adam optimizer over 10,000 iterations. To further enhance the QDDM’s performance, the architecture is fine-tuned to extract critical features from the input images.

The depth of the QNN within QDDM significantly influences its feature extraction capabilities. While deeper architectures often provide superior feature representations, excessive depth may lead to barren plateaus that impede effective training. To address this challenge, we carefully design the model architecture and optimize the learning rate for each dataset. Table 2 summarizes the specific training configurations applied across datasets, forming the basis of our evaluation of the QDM-based algorithms in FSL tasks.

Table 3. Performance comparison of QDM-based algorithms across various tasks, with $\mathcal{T} = 5$. Each algorithm is evaluated using 5 random seeds to report mean performance and standard error. The best-performing algorithm for each task is highlighted in blue.

Dataset	Tasks	QDM-LGDI	QDM-LGNAI	QDM-LGGI	QMLP	C14	OPTIC	Quantumnat
Digits	2w-01s	0.975 \pm 0.059	0.978 \pm 0.003	0.992 \pm 0.009	0.764 \pm 0.108	0.505 \pm 0.175	0.525 \pm 0.133	0.751 \pm 0.147
	2w-10s	0.983 \pm 0.006	0.997 \pm 0.002	0.984 \pm 0.012	0.892 \pm 0.086	0.627 \pm 0.086	0.886 \pm 0.193	0.722 \pm 0.186
	3w-01s	0.525 \pm 0.001	0.635 \pm 0.007	0.573 \pm 0.069	0.338 \pm 0.087	0.447 \pm 0.193	0.475 \pm 0.021	0.555 \pm 0.013
	3w-10s	0.857 \pm 0.015	0.801 \pm 0.008	0.632 \pm 0.035	0.355 \pm 0.059	0.481 \pm 0.183	0.698 \pm 0.121	0.687 \pm 0.156
MNIST	2w-01s	0.943 \pm 0.002	0.965 \pm 0.003	0.805 \pm 0.093	0.675 \pm 0.067	0.567 \pm 0.064	0.845 \pm 0.149	0.701 \pm 0.162
	2w-10s	0.953 \pm 0.011	0.978 \pm 0.005	0.915 \pm 0.079	0.817 \pm 0.048	0.810 \pm 0.152	0.807 \pm 0.173	0.727 \pm 0.151
	3w-01s	0.475 \pm 0.003	0.505 \pm 0.007	0.428 \pm 0.035	0.325 \pm 0.027	0.503 \pm 0.122	0.477 \pm 0.159	0.501 \pm 0.012
	3w-10s	0.720 \pm 0.016	0.825 \pm 0.008	0.405 \pm 0.022	0.547 \pm 0.085	0.607 \pm 0.142	0.770 \pm 0.191	0.527 \pm 0.078
Fashion	2w-01s	0.738 \pm 0.007	0.768 \pm 0.007	0.898 \pm 0.036	0.688 \pm 0.064	0.581 \pm 0.187	0.765 \pm 0.149	0.583 \pm 0.181
	2w-10s	0.755 \pm 0.020	0.805 \pm 0.002	0.895 \pm 0.066	0.731 \pm 0.035	0.773 \pm 0.099	0.793 \pm 0.157	0.887 \pm 0.129
	3w-01s	0.453 \pm 0.008	0.433 \pm 0.001	0.483 \pm 0.012	0.331 \pm 0.098	0.332 \pm 0.172	0.473 \pm 0.128	0.622 \pm 0.063
	3w-10s	0.655 \pm 0.018	0.735 \pm 0.004	0.585 \pm 0.025	0.647 \pm 0.015	0.527 \pm 0.173	0.593 \pm 0.139	0.653 \pm 0.032
Average		0.754 \pm 0.015	0.795 \pm 0.004	0.719 \pm 0.045	0.574 \pm 0.060	0.546 \pm 0.140	0.678 \pm 0.150	0.666 \pm 0.120

4.2 Performance Analysis of QDM-based QFSL Algorithms

In n -way, k -shot configuration, we systematically sample k images from each of n distinct categories, amassing a comprehensive dataset of $n \times k$ training images. These images are subsequently utilized to complete the fine-tuning of the QDDM. The fine-tuned model is then integrated into the QDM-based algorithms, enhancing their performance and adaptability in diverse quantum few-shot learning scenarios. During the fine-tuning phase, we use a small dataset to refine the model over 100 iterations, maintaining the same learning rate as used in the initial training phase.

Table 3 presents a detailed performance comparison of the QFSL algorithms against established baseline methods within varied scenarios: 2-way 1-shot, 2-way 10-shot, 3-way 1-shot, and 3-way 10-shot. The outcomes indicate that our QDM-based algorithms consistently deliver state-of-the-art performance, thereby confirming their superior efficacy in QFSL applications. Furthermore, we assess the resilience of QDM-based algorithms in a practical quantum circuit having quantum noise. The experiments are conducted on the IBM_Almaden quantum processor, with each circuit executed 1,000 times. As shown in Fig. 8, despite the inherent noise challenges of quantum hardware, the performance degradation observed is marginal. This finding underscores the robustness of our algorithms, maintaining high accuracy in adverse quantum environments and demonstrating their strong potential for real-world deployment on quantum computing platforms.

4.3 Factors Impacting the Effectiveness of QDM-based QFSL Algorithms

This section investigates key factors that affect the performance of QDM-based algorithms, including the influence of diffusion and denoising step counts, the quantity and diversity of training data, and the selection of QNNs used in QDM-LGGI.




Fig. 7. Performance of QDM-based algorithms on the 3-way, 1-shot task under varying diffusion and denoising step configurations.

Impact of Diffusion and Denoising Steps: The calibration of diffusion and denoising steps is crucial in shaping the quality of the generated images, which in turn significantly impacts the inference accuracy of QDM-based algorithms. As shown in Fig. 7, modifying the step count leads to significant performance variations on the Digits MNIST and MNIST datasets. Our experiments highlight that QDM-LGGI is notably responsive to changes in these steps. Increasing the number of steps generally improves the fidelity of the images to the target data distribution, thereby enhancing inference performance. However, an excessive number of steps can result in substantial degradation of the initial image into noise during the diffusion process, complicating the subsequent reconstruction in the denoising phase. This can result in the reconstruction algorithm overly emphasizing the guiding label, creating discrepancies with the original image—a challenge particularly pronounced in QDM-LGNAI and QDM-LGDI. This observation underscores the critical need for a judicious balance in step count to optimize inference accuracy effectively.

Effect of Training Data Quantity: The volume of training data critically affects the performance of QDM-based QFSL algorithms, directly influencing the quality and efficacy of the QDDM model. This study compares one-shot and ten-shot learning scenarios across various datasets, with detailed performance metrics consolidated in Table 3. The findings affirm that an increased volume of training data substantially benefits the training of the QDDM model, fostering more nuanced representations and enabling the generation of more precise and varied samples. Consequently, QDM-based algorithms exhibit enhanced inference accuracy with a well-trained QDDM model, accentuating the significance of ample and representative training data in QFSL tasks.

Choice of QNN Architecture: The QNN architecture employed within QDM-LGGI profoundly influences inference performance by determining the model’s representational and generalization capabilities. QDM-LGGI utilizes images generated by QDDM to train the QNN, which is then employed for inference tasks. As shown in Fig. 9, the inference accuracy notably varies across different




Fig. 8. Performance of QDM-based algorithms for the 3-way, 1-shot task on IBM_Almaden.




Fig. 9. Performance of QDM-LGGI on the 3-way, 1-shot task across different QNNs.




Fig. 10. Performance of QDM-based algorithms on the zero-shot, two-class classification task.




Fig. 11. Performance of QDM-based algorithms on the zero-shot, three-class classification task.

QNN architectures, potentially attributable to variations in circuit expressibility and entangling capabilities, as noted in [17]. Certain QNN architectures may offer more robust representations of the input data, capturing complex relationships with greater fidelity and thus yielding superior performance in QFSL tasks. This insight highlights the importance of selecting an appropriate QNN architecture that aligns with the complexity of the task and the characteristics of the data to maximize the effectiveness of QDM-LGGI.

4.4 Zero-Shot Learning with QDM-based QFSL Algorithms

We explore the efficacy of our QDM-based algorithms in addressing “zero-shot” learning tasks, where no prior examples from the evaluation classes are available during training. To assess this capability, we perform evaluations between two datasets with closely aligned data distributions: Digit MNIST and MNIST. Given the challenges posed by quantum diffusion models and limited hardware

resources, these datasets provide an ideal testbed to verify cross-domain generalization in zero-shot settings. We conduct the following verifications:

- **Verification of QDM-LGDI Performance in Zero-Shot Learning:** This setup aims to assess the generalization capacity of QDM-LGDI when a QNN is trained on data generated by QDDM from one domain and subsequently applied to another.
 - **Digit MNIST:** The QDDM model is trained on the MNIST dataset, where an augmented dataset is generated for QNN training. After training the QNN, it is used to perform inference on the Digit MNIST dataset to evaluate cross-domain generalization.
 - **MNIST:** The QDDM model is trained on the Digit MNIST dataset, with the augmented dataset used for QNN training. The trained QNN is then applied to the MNIST dataset for inference, allowing us to evaluate generalization in the reverse direction.
- **Verification of QDM-LGNAI and QDM-LGGI Performance in Zero-Shot Learning:** This evaluation tests the direct applicability of QDM-LGNAI and QDM-LGGI algorithms in zero-shot settings without retraining the QNN on new data.
 - **Digit MNIST:** The QDDM model is trained on the MNIST dataset and directly applied for inference on the Digit MNIST dataset, examining the robustness of the QDM-LGNAI and QDM-LGGI algorithms in adapting to unseen yet similar data distributions.
 - **MNIST:** Similarly, the QDDM model is trained on the Digit MNIST dataset and subsequently applied directly for inference on the MNIST dataset, assessing the generalization potential of these algorithms in the reverse direction.

Figs. 10 and 11 show the performance results for the zero-shot tasks across these configurations. The results demonstrate that QDM-based algorithms achieve strong zero-shot learning performance, particularly when the source and target datasets exhibit similar data structures and distributions. These findings underscore the robustness and adaptability of QDM-based QFSL algorithms, showcasing their potential to generalize effectively across related domains even in the absence of target domain data during training.

5 Conclusions and Future Work

In this work, we presented the QDM as a pioneering approach to overcome the unique challenges of QFSL. We introduced three novel algorithms—QDM-LGDI, QDM-LGNAI, and QDM-LGGI—developed from complementary data-driven and algorithmic perspectives to enhance the performance of QFSL tasks. Through extensive experimentation, we demonstrated that these algorithms achieve significant performance gains over traditional baselines, underscoring the potential of QDM to advance QFSL by effectively leveraging quantum noise modeling and label guidance.

While the results are promising, our current implementation of QDM faces limitations, particularly in its scalability to more complex datasets and real-world quantum applications. These constraints highlight the need for further

research to unlock the full potential of QDM. Future work could focus on extending the capabilities of QDM through improvements in model architecture and optimization techniques, enabling it to handle more intricate datasets with diverse and high-dimensional features.

Moreover, expanding QDM’s applicability across additional QML domains presents exciting opportunities. Potential directions include quantum object detection, quantum semantic segmentation, and other advanced generative tasks where QDM’s noise-guided diffusion and label-conditional modeling could provide significant advantages. Integrating QDM into these areas would not only broaden its utility but could also position it as a versatile and foundational framework for future quantum learning architectures. By advancing QDM’s robustness and adaptability, this line of research has the potential to transform QML applications and contribute meaningfully to the field of quantum artificial intelligence.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Kölle, Michael, et al., “Quantum denoising diffusion models,” arXiv preprint arXiv:2401.07049 (2024).
2. Parigi, Marco, Stefano Martina, and Filippo Caruso, “Quantum-noise-driven generative diffusion models,” *Advanced Quantum Technologies* (2024): 2300401.
3. Caro, Matthias C., et al., “Generalization in quantum machine learning from few training data,” *Nature communications* 13.1 (2022): 4919.
4. Li, Alexander C., et al., “Your diffusion model is secretly a zero-shot classifier,” *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2023.
5. Clark, Kevin, and Priyank Jaini, “Text-to-image diffusion models are zero shot classifiers,” *Advances in Neural Information Processing Systems* 36 (2024).
6. Wang, Yaqing, et al., “Generalizing from a few examples: A survey on few-shot learning,” *ACM computing surveys (csur)* 53.3 (2020): 1-34.
7. Wang, Ruhan, Philip Richerme, and Fan Chen, “A hybrid quantum-classical neural network for learning transferable visual representation,” *Quantum Science and Technology* 8.4 (2023): 045021.
8. Liu, Minzhao, et al., “Embedding learning in hybrid quantum-classical neural networks,” *2022 IEEE International Conference on Quantum Computing and Engineering (QCE)*. IEEE, 2022.
9. Wang, Ruhan, Fahiz Baba-Yara, and Fan Chen, “JustQ: Automated deployment of fair and accurate quantum neural networks,” *2024 29th Asia and South Pacific Design Automation Conference (ASP-DAC)*. IEEE, 2024.
10. Focardi, Sergio, Frank J. Fabozzi, and Davide Mazza, “Quantum option pricing and quantum finance,” *Journal of Derivatives* 28.1 (2020): 79-98.
11. Parsons, Donald Frederick, “Possible medical and biomedical uses of quantum computing,” *Neuroquantology* 9.3 (2011).
12. Cao, Yudong, Jhonathan Romero, and Alán Aspuru-Guzik, “Potential of quantum computing for drug discovery,” *IBM Journal of Research and Development* 62.6 (2018): 6-1.

13. Preskill, John, "Quantum computing in the NISQ era and beyond," *Quantum* 2 (2018): 79.
14. Chu, Cheng, et al., "QMLP: An error-tolerant nonlinear quantum MLP architecture using parameterized two-qubit gates," *Proceedings of the ACM/IEEE International Symposium on Low Power Electronics and Design*. 2022.
15. Wang, Hanrui, et al., "Quantumnat: quantum noise-aware training with noise injection, quantization and normalization," *Proceedings of the 59th ACM/IEEE design automation conference*. 2022.
16. Patel, Tirthak, Daniel Silver, and Devesh Tiwari, "OPTIC: A practical quantum binary classifier for near-term quantum computers," *2022 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, 2022.
17. Sim, Sukin, Peter D. Johnson, and Alán Aspuru-Guzik, "Expressibility and entangling capability of parameterized quantum circuits for hybrid quantum-classical algorithms," *Advanced Quantum Technologies* 2.12 (2019): 1900070.
18. Ho, Jonathan, Ajay Jain, and Pieter Abbeel, "Denoising diffusion probabilistic models," *Advances in neural information processing systems* 33 (2020): 6840-6851.
19. Song, Jiaming, Chenlin Meng, and Stefano Ermon, "Denoising diffusion implicit models," *arXiv preprint arXiv:2010.02502* (2020).
20. Bergholm, Ville, et al., "Pennylane: Automatic differentiation of hybrid quantum-classical computations," *arXiv preprint arXiv:1811.04968* (2018).
21. LeCun, Yann, et al., "Gradient-based learning applied to document recognition," *Proceedings of the IEEE* 86.11 (1998): 2278-2324.
22. Xiao, Han, Kashif Rasul, and Roland Vollgraf, "Fashion-MNIST: A novel image dataset for benchmarking machine learning algorithms," *arXiv preprint arXiv:1708.07747* (2017).
23. Sun, Qianru, et al., "Meta-transfer learning for few-shot learning," *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 2019.
24. Parnami, Archit, and Minwoo Lee, "Learning from few examples: A summary of approaches to few-shot learning," *arXiv preprint arXiv:2203.04291* (2022).
25. Alimoğlu, Fevzi, and Ethem Alpaydin, "Combining multiple representations for pen-based handwritten digit recognition," *Turkish Journal of Electrical Engineering and Computer Sciences* 9.1 (2001): 1-12.
26. Hu, T., Zhang, J., Yi, R., Du, Y., Chen, X., Liu, L., ... & Wang, C. (2024, March). "AnomalyDiffusion: Few-shot anomaly image generation with diffusion model," In *Proceedings of the AAAI Conference on Artificial Intelligence* (Vol. 38, No. 8, pp. 8526-8534).
27. Zhang, Baoquan, et al., "MetaDiff: Meta-learning with conditional diffusion for few-shot learning," *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 38. No. 15. 2024.