MITSUBISHI ELECTRIC
Changes for the Better

ICAD 2025

Quantum Diffusion Models for Few-Shot Learning

Ruhan Wang

Ye Wang

Jing Liu

Toshiaki Koike-Akino

MITSUBISHI ELECTRIC RESEARCH LABORATORIES (MERL)
Cambridge, Massachusetts, USA
http://www.merl.com

Outline

- Background of Generative AI
- Overview of Quantum AI
- Introduction of Quantum Diffusion Models
- Proposed Methods for Quantum Few-Shot Learning
- Experiments
  - 5 to 30% accuracy improvement
- Summary

Generative AI: Surpassing Human-Level Creativity

- VAE
- GAN
- Diffusion

[Karras et al, 2018]

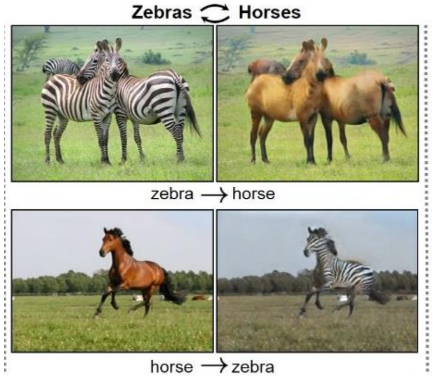

StyleGAN [Karras et al, 2019]

Diffusion [Dhariwal et al., 2021]

Dall-E 2 [Ramesh et al. 2022]

Imagen [Saharia et al. 2022]

StableDiffusion [Rombatch et al. 2022]

Generative AI: Surpassed Human-Level Creativity

星新一賞
THE HOSHI AWARD

- AI won some art prizes and inventions
  - Kamome Ashizawa’s **AI-generated novel**, “Are you there?”, took the Hoshi-Shinichi Literary Award in Nikkei press, Feb. 2022
  - Jason Allen’s **AI-generated painting**, “Théâtre D’opéra Spatial,” took first place in the digital category at the Colorado State Fair. Sep. 2022
  - DABUS **AI-generated patents** granted in South Africa, Jul. 2022
  - Boris Eldagsen’s **AI-generated photo**, “The Electrician” came top in open competition at the World Photography Organization’s Sony World Photography Awards, Apr. 2023.

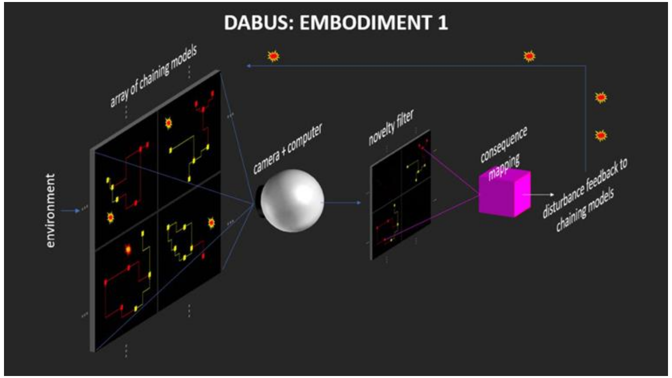

Green vs. Red AI

- Roy Schwartz, Jesse Dodge, Noah A. Smith, Oren Etzioni, “Green AI”, Communications of the ACM, December 2020, Vol. 63 No. 12, Pages 54-63 10.1145/3381831: https://cacm.acm.org/magazines/2020/12/248800-green-ai/

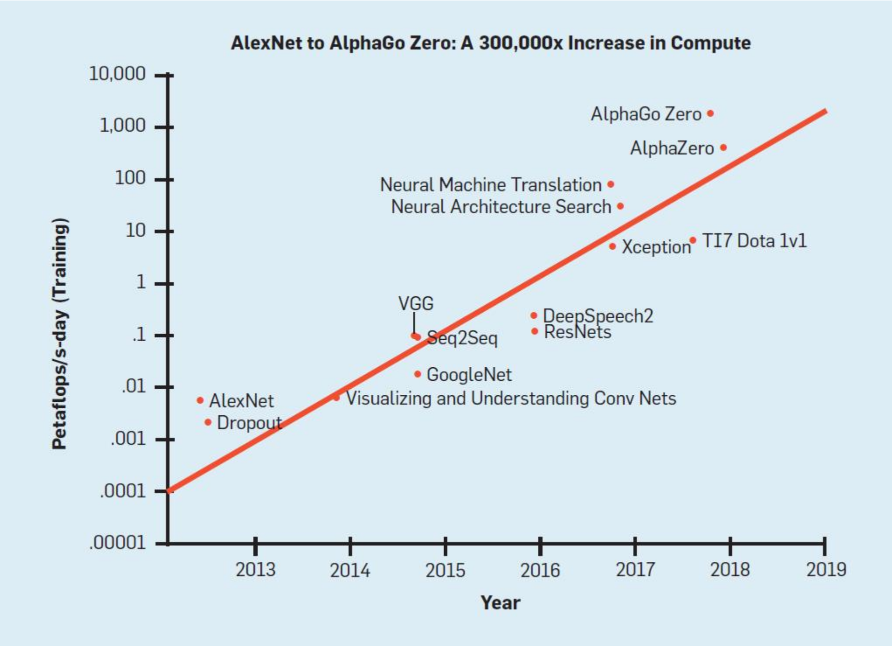

The computation used to train deep learning models has increased 300,000x in six years: nearly 10x annually

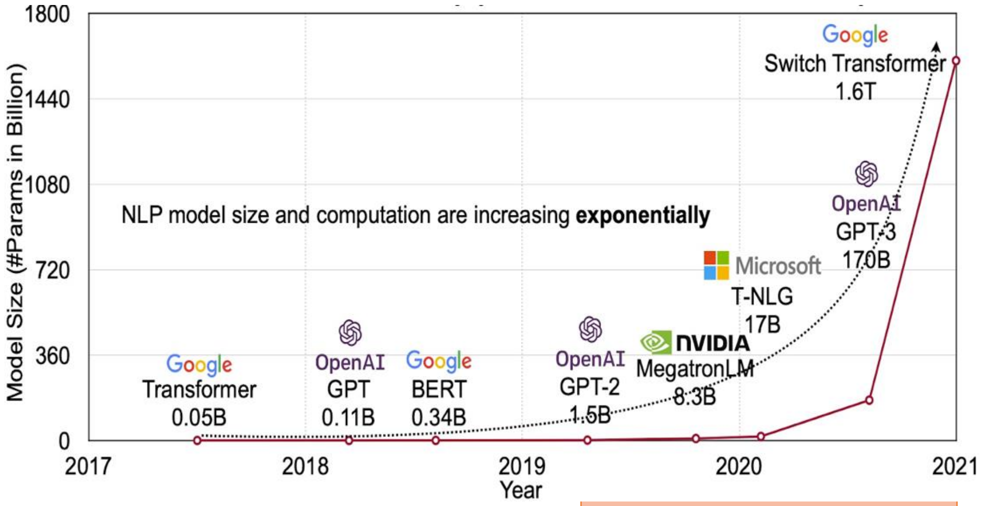

GPT-3: 175B params
GPT-4: ~1.7T params

Deep Learning Crisis for Sustainable Growth

- Escalating power consumption of DNN training
  - [Strubell et al. Energy and policy considerations for deep learning in NLP. 2019]
  - DNN training with network architecture search (NAS) on GPUs requires **5-fold higher** carbon emission of single car lifetime!
- Therefore, we should consider **Green AI**
  - Efficient, fast, low-power, lightweight AI
  - New computing modality alternative to CPU/GPU/TPU: Natural computing (**Quantum**)

Consumption	CO_2e (lbs)
Air travel, 1 passenger, NY↔SF	1984
Human life, avg, 1 year	11,023
American life, avg, 1 year	36,156
Car, avg incl. fuel, 1 lifetime	126,000
Training one model (GPU)	
NLP pipeline (parsing, SRL)	39
w/ tuning & experimentation	78,468
Transformer (big)	192
w/ neural architecture search	626,155

Quantum Computing

- Morgan Stanley: Quantum tech. can drive **4th industrial revolution**
- Escalating government funds: National Quantum Initiative
- Quantum processor providers: **IBM**, **Google**, **Microsoft**, **Honeywell**, Intel, IONQ, rigetti, …
- Quantum cloud services: IBMQ, Amazon Bracket, …
- Free libraries to evaluate quantum computing on realistic simulators or real devices

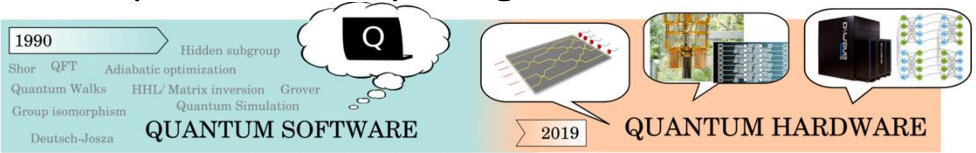

QISKit
Quantum Information Software Kit
PYQUIL
Cirq
QuTiP
Quantum Toolbox in Python
PENNYLANE
PROJECTQ

intel
Microsoft
rigetti
Honeywell
IBM
Google

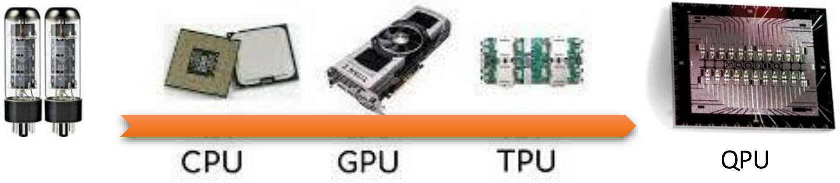

Evolution of Quantum Processing Unit (QPU)

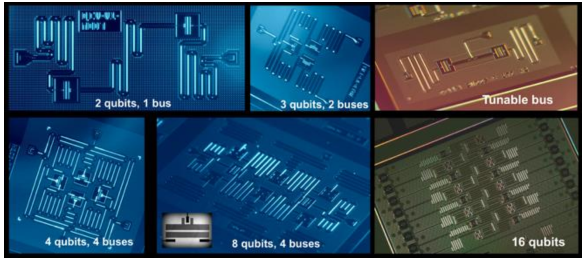

- Rapid QPU development to allow many qubits
  - IBM released **127-qubit** QPUs in Nov. 2021
  - IBM released **433-qubit** QPUs in Nov. 2022
  - IBM released **1121-qubit** QPUs in Dec. 2023 (**4158-qubits** by 2025)

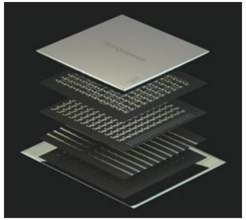

IBM 127-qubit QPU (Nov. 2021)

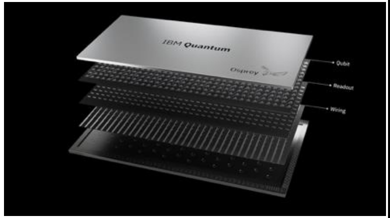

IBM 433-qubit QPU (Nov. 2022)

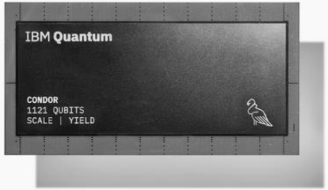

IBM 1121-qubit QPU (Dec. 2023)

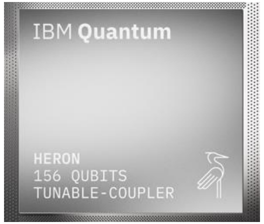

IBM 156-qubit QPU (Nov. 2024): 5000 gates

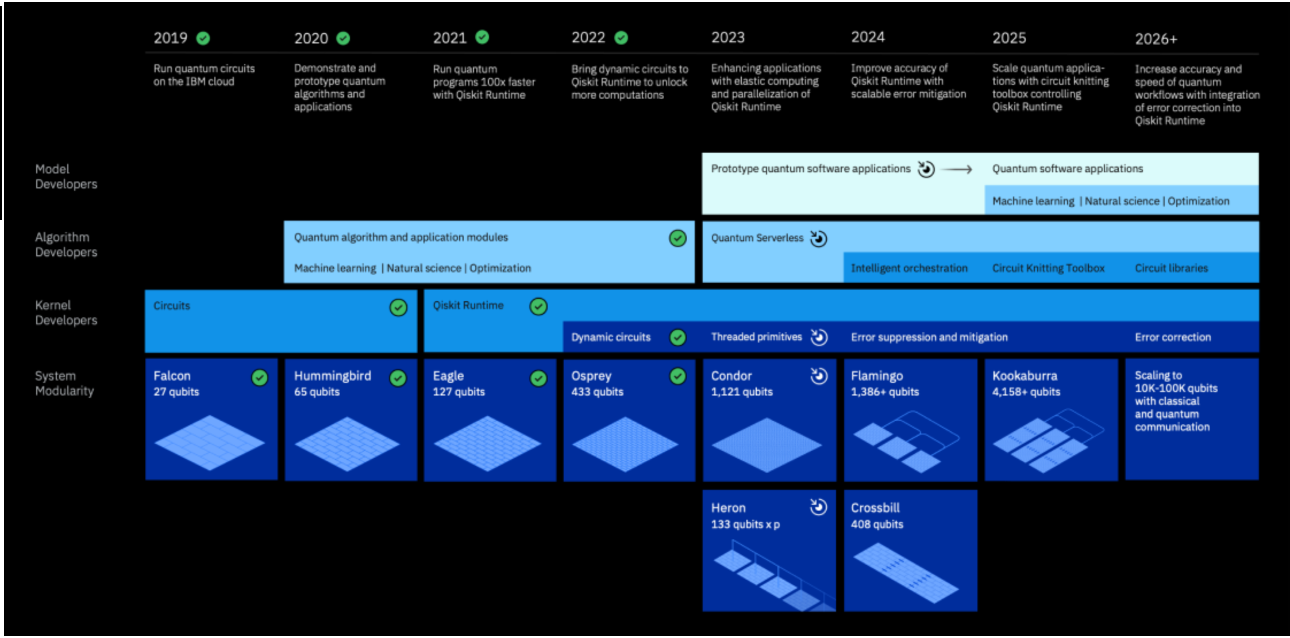

IBM QPU development roadmap (as of 2022)

MERL's QML Research

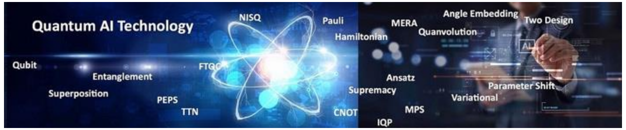

- MERL QML research highlight:
  – https://www.merl.com/research/highlights/quantum-ai
  – QHack: 2022 AWS award; IBM award; 2023 Nvidia award; 3rd prize

- Various Industrial QML Applications
  – Matsumine, T., Koike-Akino, T., Wang, Y., **"Channel Decoding with Quantum Approximate Optimization Algorithm"**, *ISIT*, July 2019.
  – Koike-Akino, T., Matsumine, T., Wang, Y., Millar, D.S., Kojima, K., Parsons, K., **"Variational Quantum Demodulation for Coherent Optical Multi-Dimensional QAM"**, *OFC/NFOEC*, March 2020, pp. T3D.6.
  – Koike-Akino, T., Wang, P., Wang, Y., "**Quantum Transfer Learning for Wi-Fi Sensing**", ICC, June 2022.
  – Liu, B., Koike-Akino, T., Wang, Y., Parsons, K., "**Variational Quantum Compressed Sensing for Joint User and Channel State Acquisition in Grant-Free Device Access Systems** ", ICC, June 2022.
  – Koike-Akino, T., Wang, P., Wang, Y., "**AutoQML: Automated Quantum Machine Learning for Wi-Fi Integrated Sensing and Communications**", SAM, Aug. 2022.
  – Koike-Akino, T., et al., "**Quantum Feature Extraction for THz Multi-Layer Imaging**", IRMMW-THz, Aug. 2022.
  – Koike-Akino, T., Wang, Y., "**quEEGNet: Quantum AI for Biosignal Processing**", BHI, Sep. 2022.
  – Koike-Akino, T., Wang. P., "**Post-Deep Learning Era: Emerging Quantum Machine Learning for Sensing and Communications**", GLOBECOM, Dec. 2022.
  – Liu, B., Koike-Akino, T., Wang, Y., Parsons, K., "**Learning to Learn Quantum Turbo MIMO Detection**", arXiv, 2022.
  – Koike-Akino, T., "COVID-19 Quantum Forecasting", QHack, Mar. 2022.
  – Koike-Akino, T., "Quantum mixed reality (XR)", QHack, Mar. 2023.
  – Ahmed, M.R., Koike-Akino, T., Parsons, K., Wang, Y., "**AutoHLS: Learning to Accelerate Design Space Exploration for HLS Designs**", MWSCAS, Aug. 2023.
  – Fujihashi, T., Koike-Akino, T., "**Quantum implicit neural compression**", AAAI Workshop, Mar. 2025

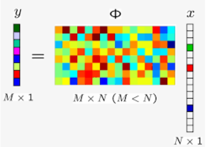

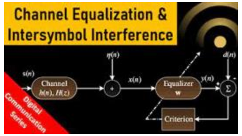

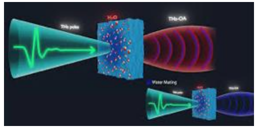

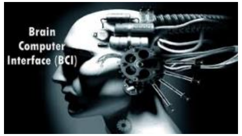

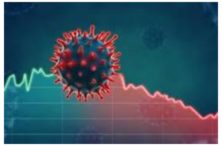

Quantum Diffusion Models

- Quantum denoising diffusion models (QDDM)

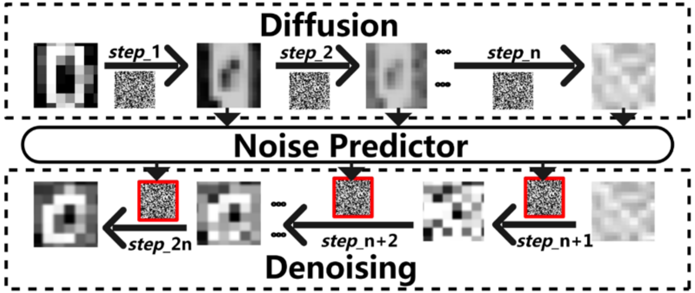

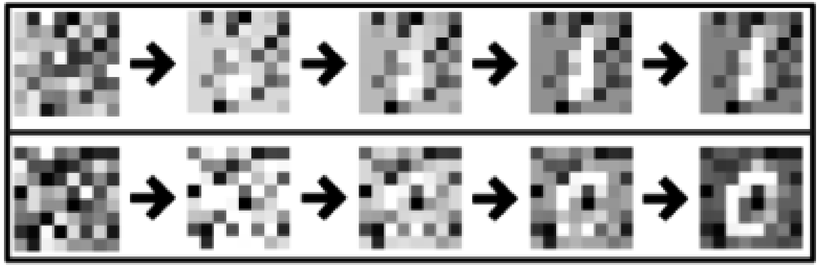

- Quantum few-shot learning: 3 different approaches

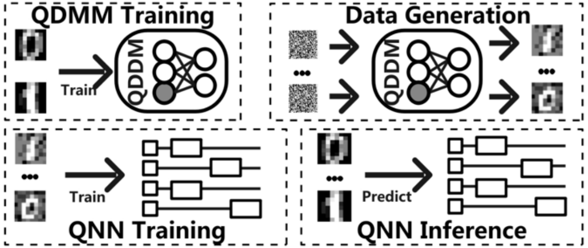

Generation inference

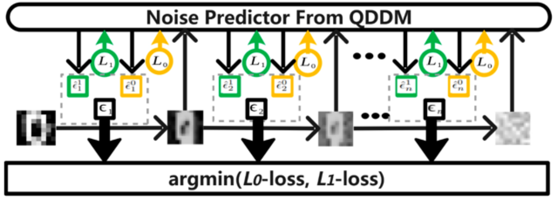

Diffusion inference

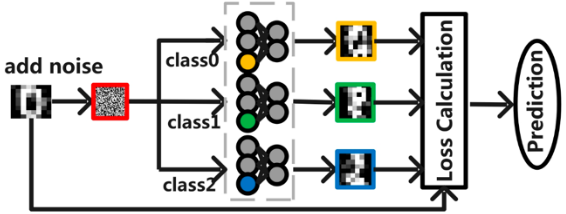

Denoising inference

- 5-30% accuracy improvement

Diffusion Models

- Denoising Diffusion Probabilistic Model (DDPM) [Ho 2020] is a pioneering generative AI model, outperforming VAE and GAN
  - 2 processes: Diffusion steps; Denoising steps
  - Variants: Implicit Diffusion [Song 2020]; Latent Diffusion [Rombach 2022]; Guided Diffusion [Ho 2022]

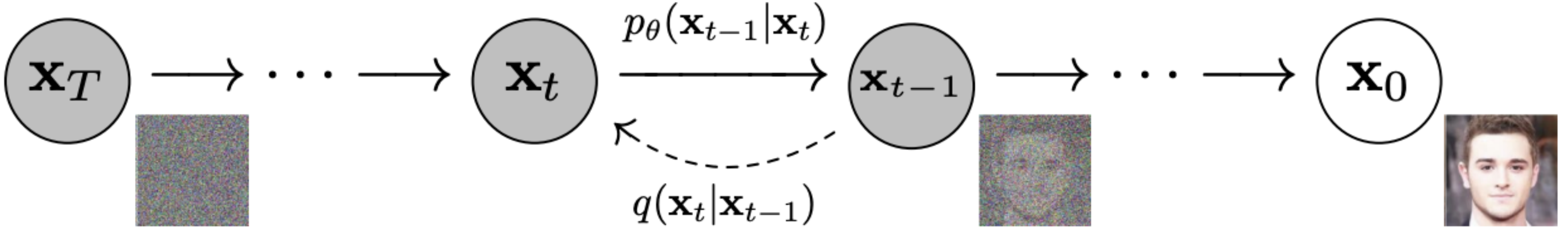

Quantum Denoising Diffusion Model (QDDM)

- The diffusion models were migrated to the quantum domain
  - Cacioppo 2023: Quantum diffusion models
  - Kölle 2024: Quantum denoising diffusion models
  - Zhang 2024: Generative quantum machine learning via denoising diffusion probabilistic models
  - Parigi 2024: Quantum-Noise-Driven Generative Diffusion Models
  - Kwun 2024: Mixed-State Quantum Denoising Diffusion Probabilistic Model
  - Chen 2024 Quantum generative diffusion model: a fully quantum-mechanical model for generating quantum state ensemble
- We use **QDDM** as a foundational tool for **few-shot learning (FSL)**

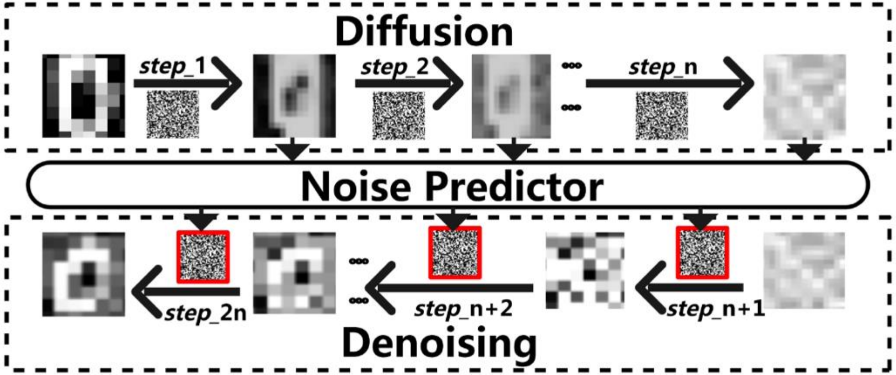

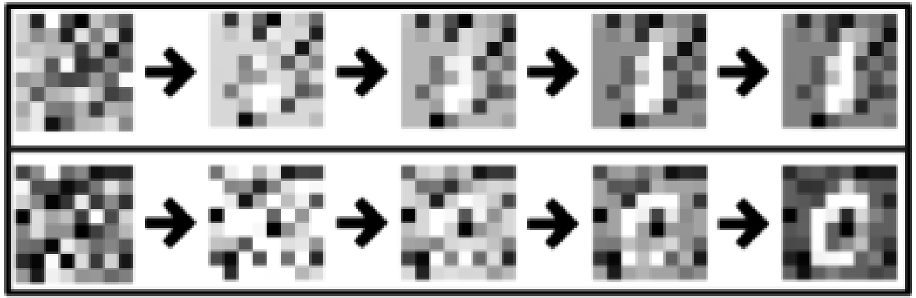

Quantum Few-Shot Learning

- Few-shot learning (FSL) is designed to address supervised learning challenges with a very limited number of training examples: **support set**.
  - The support set consists of a small number of labeled examples, encompassing ***n*** **classes, each with *k* examples**: called ***n-way, k-shot***.
  - We consider quantum FSL to train quantum neural networks (QNNs)

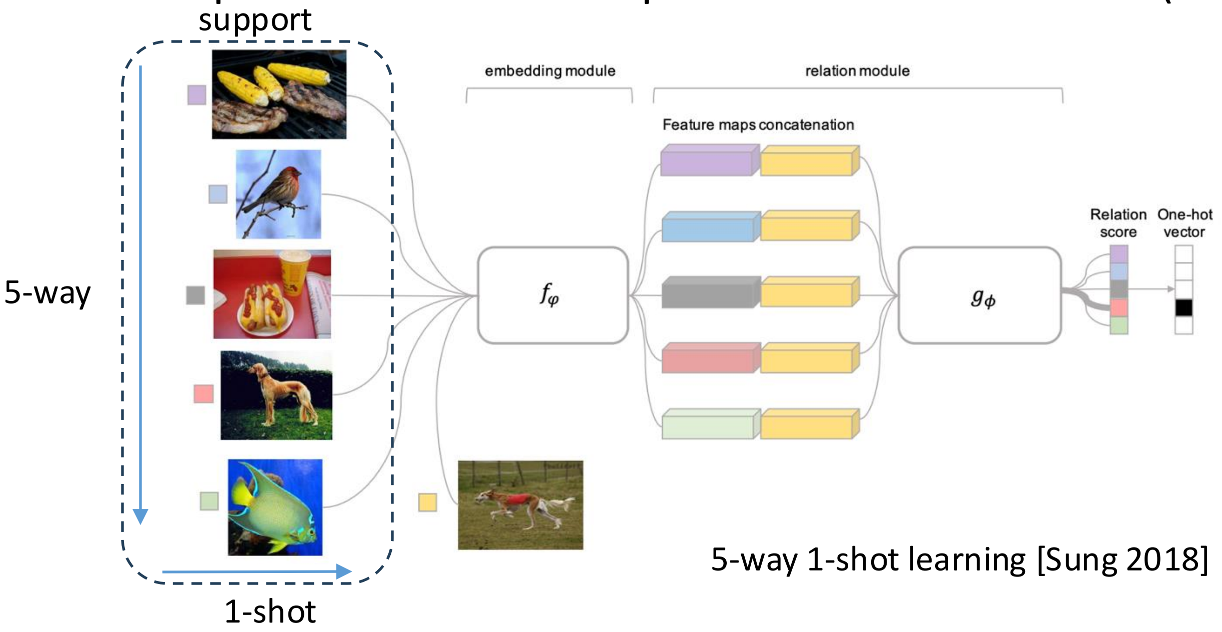

5-way 1-shot learning [Sung 2018]

Method I: Generation Inference

- **QDDM-Based Label-Guided Generation Inference (LGGI)**
  - The primary challenge of QFSL is the limited availability of training data. Thus, expanding the training dataset can significantly enhance the performance of QFL.
  - A small amount of few-shot data is used to train the QDDM. Then, the QDDM is employed to expand the training dataset for QNN.
  - This expanded dataset is then used to train the QNN, which in turn improves its inference accuracy on real data.

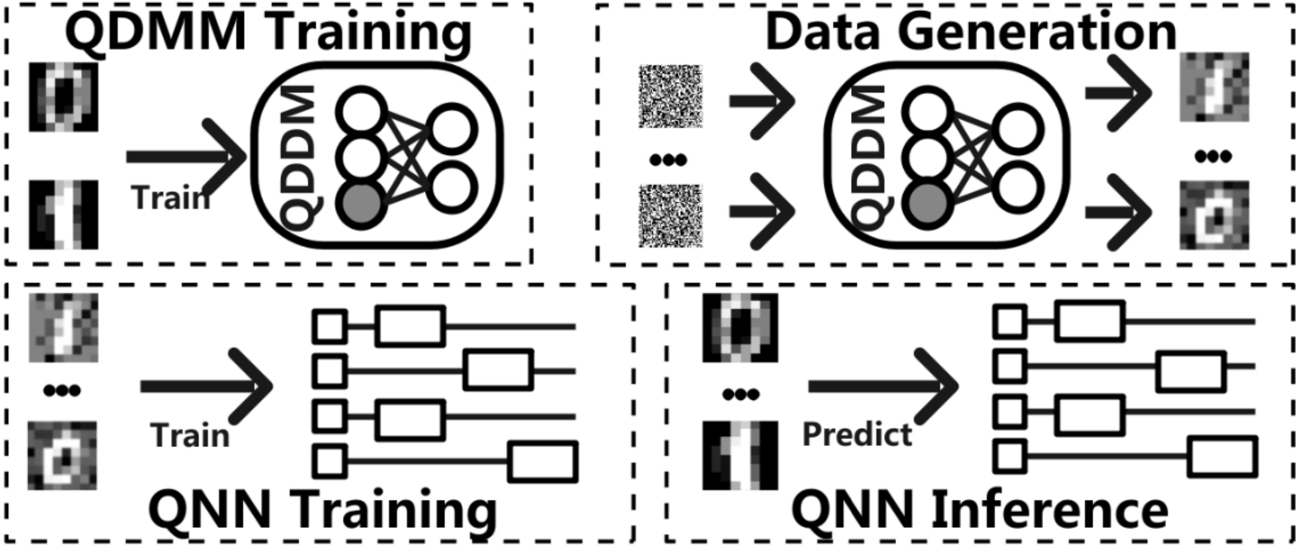

Method II: Diffusion Inference

- **QDM-Based Label-Guided Noise Addition Inference (LGNAI)**
  - The learning objective of the QDDM relies on using a noise predictor to estimate the noise in noisy data compared to the actual noise.
  - The noise predictor's estimation is guided by a label. By using the correct label for guidance, the error between the predicted noise and the actual noise may be minimized.

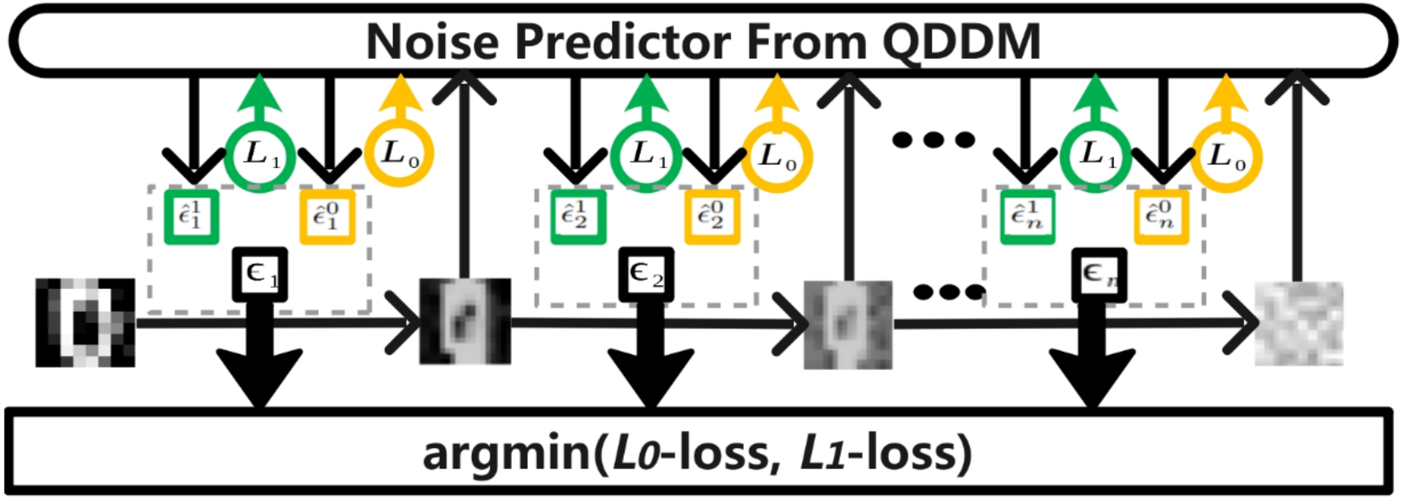

Method III: Denoising Inference

- **QDDM-Based Label-Guided Denoising Inference (LGDI)**
  - During the denoising phase of the QDDM, the noise predictor is utilized to estimate the noise present in the noisy data, which is subsequently subtracted.
  - The noise prediction is guided by labels. Consequently, the final generated images vary according to the guidance provided by different labels.
  - The data generated under the guidance of the true label may be close to the original data.

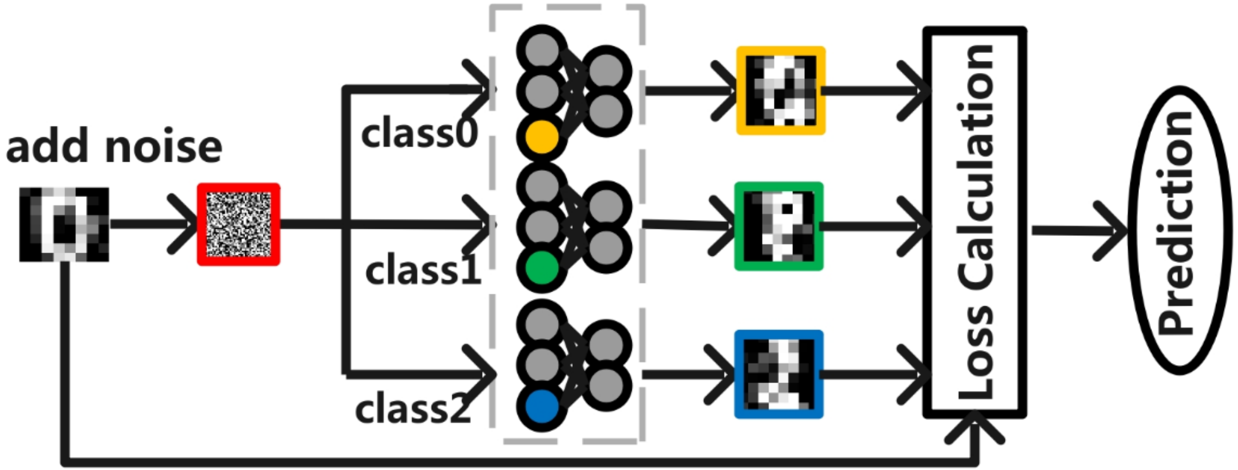

QDDM Training

- We use Adam for 10,000 epochs; labels are encoded with RX gate
- The training loss reflects the discrepancy between the noises predicted by the noise predictor and the actual noises during the denoising phase.

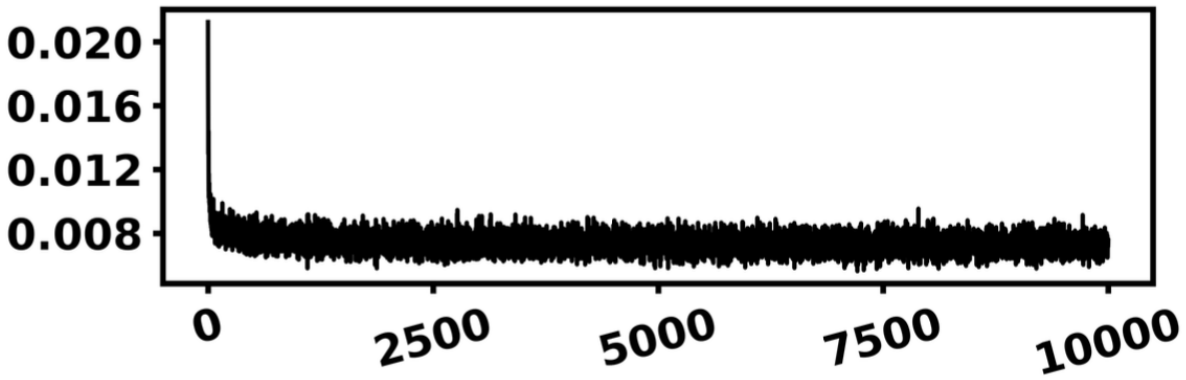

Figure 8: Training Loss Trends during QDDM Model Training.

Performance of Quantum Few-Shot Learning

- QDDM-based FSL offers significant improvement in accuracy

Dataset	Tasks	QDiff-LGDI	QDiff-LGNAI	QDiff-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		$0.754_{\pm 0.015}$	$0.795_{\pm 0.004}$	$0.719_{\pm 0.045}$	$0.574_{\pm 0.060}$	$0.546_{\pm 0.140}$	$0.678_{\pm 0.150}$	$0.666_{\pm 0.120}$

QDDM-based FSL

Impact of Quantum Noise

- We demonstrated robustness against quantum noise on IBM_Almaden quantum processors

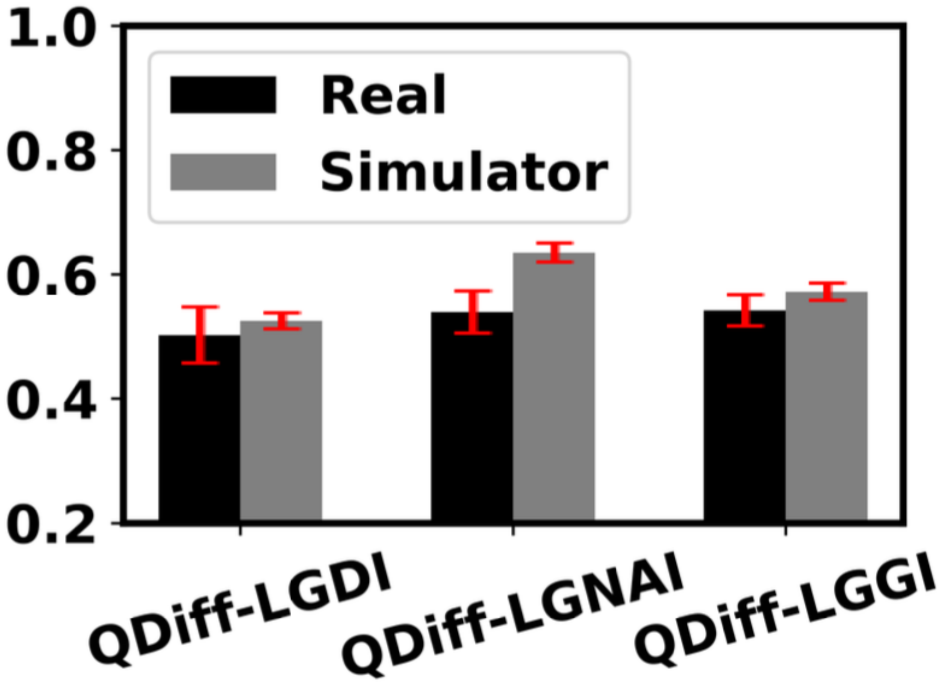

3-way, 1-shot

Impact of QNN Ansatz

- The impact of the selection of QNNs utilized in Qdiff-LGGI
  - Different QNNs have different expressive abilities and different information extraction capabilities for input images.

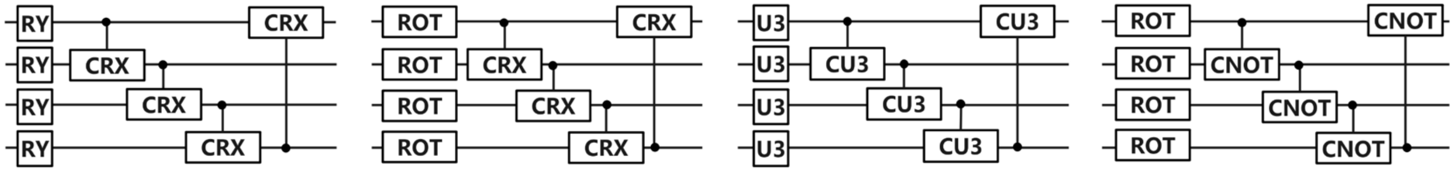

(a)C14 (b)QMLP (c) Quantumnat (d)OPTIC

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	$36 \times n$

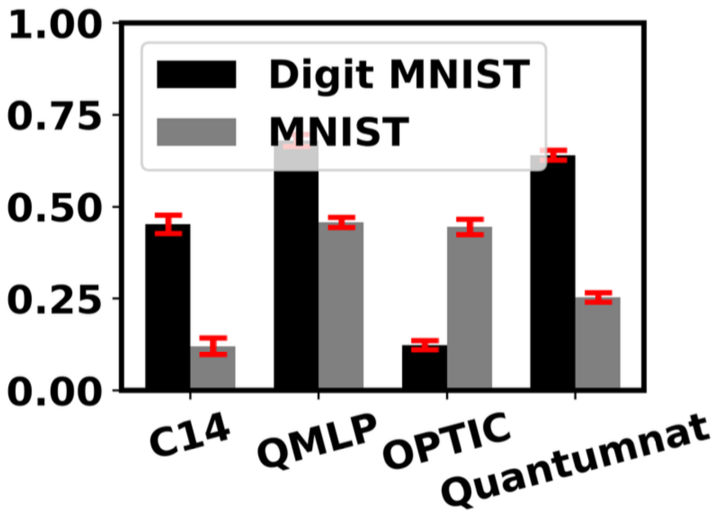

Performance of Quantum Zero-Short Learning

- We utilize the Digit MNIST dataset to train the QDDM and then use Qdiff-based algorithms to complete the inference on the MNIST dataset. The same strategy is employed for experiments on the MNIST dataset.

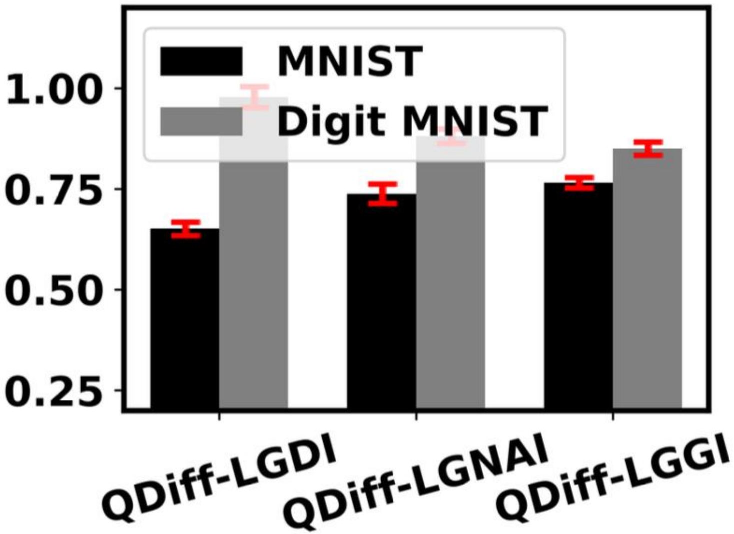

2-way

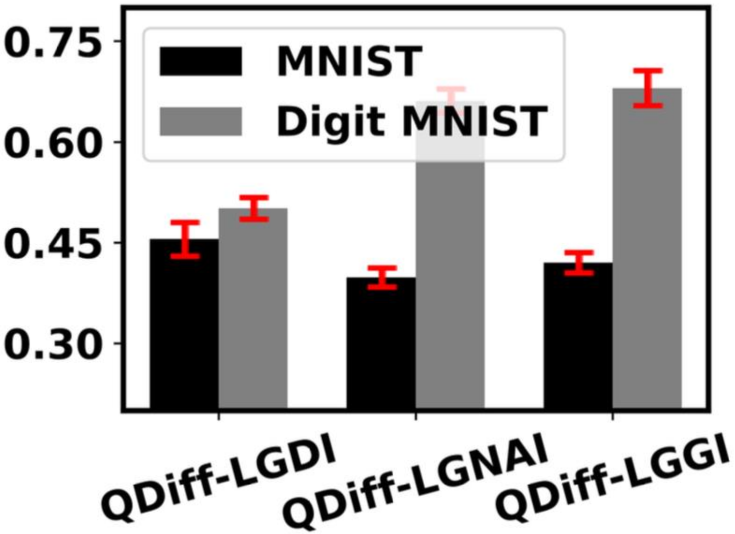

3-way

Summary

- We addressed overviews of generative AI and quantum AI
- We proposed new methods for quantum few-shot learning using quantum diffusion models
  - We introduced 3 different approaches based on QDDM’s generation/diffusion/denoising
  - Our methods demonstrated significant improvement up to 30% gain
  - We also validated that our method has a high resilience to the quantum noise
  - We evaluated different QNN ansatz
  - Zero-shot capability was discussed too
- Future work:
  - Enhancing the capabilities of QDM through improvements in model architecture and optimization techniques, enabling more intricate datasets with diverse and high-dimensional features for diverse real-world applications.
- Questions?
  - koike@merl.com

Impact of Diffusion/Denoising Steps

- With the change of Diffusion, Denoising steps, the performance of different Qdiff-based algorithm on Digit MNIST vs. MNIST varies.

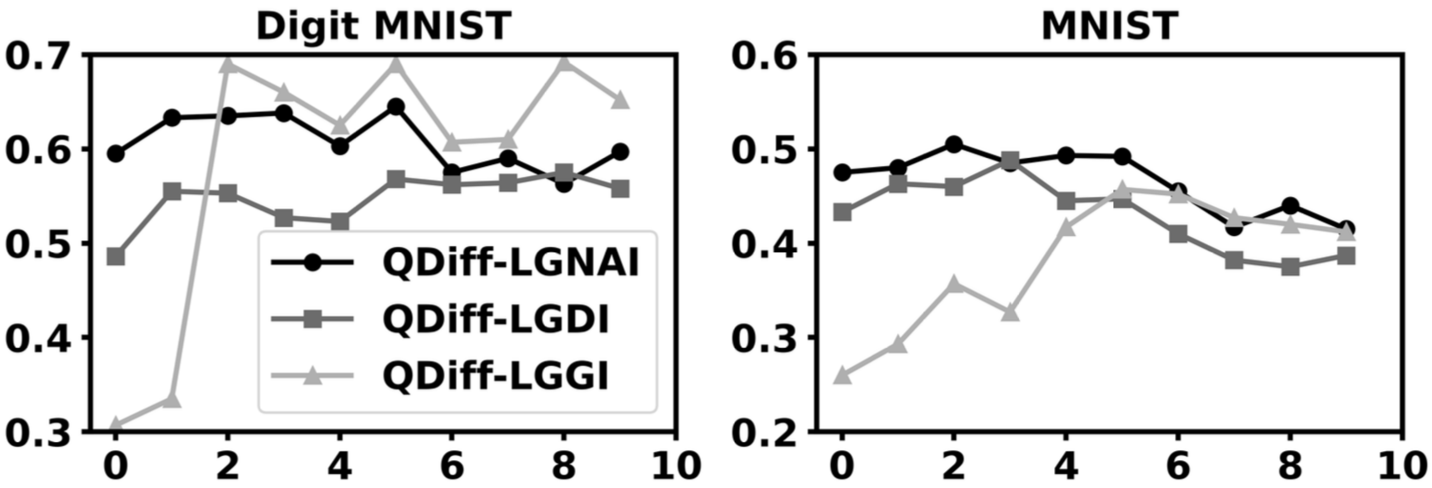

- As the number of steps increases, the performance of Qdiff-LGGI improves.
- Qdiff-LGDI and Qdiff-LGNAI should use a moderate number of steps.

Quantum Neural Network: QNN

- Quantum Neural Network (QNN) performs inference tasks: consisting of
  - a data **encoder** that embeds a classical input x into a quantum state $|x\rangle$
  - a **Variational Quantum Circuit (VQC)** that entangles the quantum state
  - a **measurement layer** that maps the output quantum state to a classical vector.

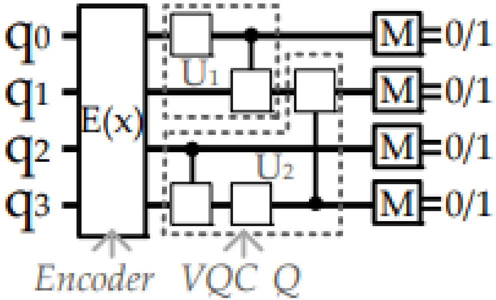

(a) A standard QNN

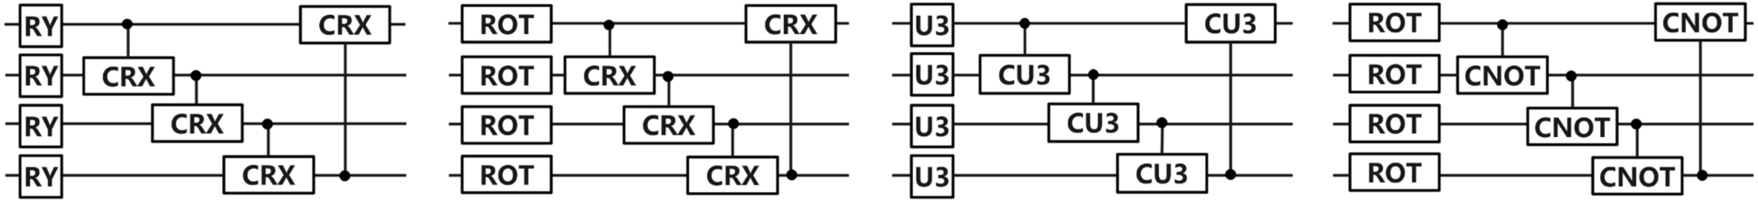

(a)C14 (b)QMLP (c) Quantumnat (d)OPTIC

MITSUBISHI
ELECTRIC

Changes for the Better