The Quantum Data Loading Bottleneck

Alexandru Ioniță
AI Multimedia Lab
National University of Science and Technology
POLITEHNICA Bucharest
Bucharest, Romania
aionita1204@upb.ro

Bogdan Ionescu
AI Multimedia Lab
National University of Science and Technology
POLITEHNICA Bucharest
Bucharest, Romania
bogdan.ionescu@upb.ro

***Abstract*—In this work, we focus on the Flexible Representation of Quantum Images (FRQI), a promising encoding scheme for grayscale image data. We characterize its qubit requirements and circuit depth under both idealized conditions and realistic hardware constraints, revealing the severe limitations posed by current quantum devices. These findings reflect the broader issue of quantum data loading, a fundamental bottleneck in many applications, particularly in quantum machine learning.**

***Index Terms*—quantum encoding, quantum machine learning**

I. Introduction

Quantum computing offers fundamentally new modes of information processing by exploiting superposition and entanglement. These potentially allow quantum algorithms to manipulate information in exponentially large vector spaces using relatively few qubits, promising advantages over classical computation for certain tasks [1]. However, encoding classical data into quantum states requires significant resource overhead and is severely constrained by the limitations of current quantum hardware [2].

This challenge is especially pronounced for image data, given its high dimensionality and inherent spatial structure. This article presents the data loading bottleneck using the example of the Flexible Representation of Quantum Images (FRQI) encoding scheme [3] and discusses potential strategies for achieving meaningful near-term progress in Quantum Machine Learning (QML) despite these limitations.

II. Quantum information

A qubit is best viewed as a vector in a two-dimensional complex vector space, spanned by the classically possible states $|0\rangle$ and $|1\rangle$. When measured, the squared modulus of the projection onto each basis state gives the corresponding probability of observing each classical state, 0 or 1.

This framework naturally extends to multi-qubit systems: the state of an n-qubit register is a vector in a 2^n-dimensional complex Hilbert space, spanned by all 2^n classical bit strings. A general pure quantum state can be expressed as:

$$|\psi\rangle = \sum_{i=0}^{2^n-1} c_i|i\rangle, \tag{1}$$

with complex amplitudes c_i which satisfy $\sum_i |c_i|^2 = 1$, enabling interference and entanglement [1].

Operations on quantum registers, implemented as sequences of quantum gates, correspond to unitary transformations on this vector space. Specifically, an operation on n qubits is described by a $2^n \times 2^n$ unitary matrix. This exponential scaling of the underlying vector space quickly overwhelms classical resources [1, 2]. For instance, even the most powerful supercomputers, with roughly 10 petabytes of RAM, can simulate only about 50 qubits without employing approximations. In this light, state-of-the-art noisy intermediate-scale quantum (NISQ) processors with a few hundred qubits already represent a remarkable leap in computational capacity.

III. Quantum image encoding

Quantum algorithms rely on encoding classical data into quantum states, with performance heavily influenced by the chosen encoding scheme. Superposition enables compact representation using relatively few qubits, allowing quantum parallelism. However, this comes at the cost of increased embedding complexity [1]. In this work, we focus on image data, which is particularly challenging due to its high dimensionality, spatial structure, and color information, making efficient quantum encoding non-trivial. In the following, we focus on the Flexible Representation of Quantum Images (FRQI), as it offers greater practicality for near-term quantum hardware [4]. However, the challenges discussed later are broadly applicable to other quantum encoding schemes as well.

FRQI represents grayscale images by encoding spatial position into computational basis states $|i\rangle$, where each i corresponds to the binary representation of a pixel location. The normalized pixel intensity is embedded into the amplitude of a single, shared color qubit, parameterized as an angle: $\cos\theta_i|0\rangle + \sin\theta_i|1\rangle$. For an image of size $2^n \times 2^n$, the full expression of the FRQI representation is:

$$|I\rangle = \frac{1}{2^n} \sum_{i=0}^{2^{2n}-1} (\cos\theta_i|0\rangle + \sin\theta_i|1\rangle) \otimes |i\rangle, \tag{2}$$

where $\otimes$ is the tensor product. FRQI requires $2n$ qubits to represent position, spanning a 2^{2n}-dimensional Hilbert space, and one additional qubit for intensity, yielding a total of $2n+1$ qubits [3]. To illustrate the exponential memory efficiency, $n = 12$ corresponds to a 4096×4096 pixel image, which can be encoded using only 25 qubits. However, this comes at

2nd Workshop on Artificial Intelligence for Multimedia,
AI Multimedia Lab, CAMPUS Research Institute, POLITEHNICA Bucharest,
Bucharest, Romania, 15 July, 2025

the cost of circuit depth: the ideal depth scales as $O(N \cdot n)$, where $N = 2^{2n}$ is the total number of pixels. This depth results from the need of a multi-controlled rotation gate for each pixel, conditioned on the pixel index $|i\rangle$. Because there is a single pixel intensity qubit, these operations can not be performed in parallel. As noted earlier, this trade-off is general: embedding schemes that achieve exponential memory compression typically come at the cost of exponential gate depth with respect to the number of qubits [4].

IV. The bottleneck

It is clear from the FRQI example that the data-loading bottleneck arises primarily from the circuit depth overhead. Current quantum hardware is in the Noisy Intermediate-Scale Quantum (NISQ) era, characterized by limited qubit counts and substantial noise. The runtime of quantum algorithms is constrained by the coherence time, the interval during which quantum states retain coherence before decaying into classical mixtures. Furthermore, quantum operations are analog and imperfect, leading to accumulating gate errors. Collectively, these limitations impose a practical upper bound on executable quantum circuits, typically around 1000 native gates, beyond which meaningful quantum output becomes unreliable [2].

The fact that each pixel requires a controlled rotation would restrict data loading to images with at most 1000 pixels. However, realistic implementations must account for additional overhead. Unlike theoretical models, NISQ devices support only a restricted set of native gates (e.g., R_z, $\sqrt{X}$, and CNOT), and exhibit limited qubit connectivity. For example, implementing a single controlled rotation may require up to 32 native gates under standard gate decompositions, reducing the feasible image size to approximately 31 pixels. Furthermore, qubit routing via $SWAP$ operations introduces additional depth overhead when logical qubits are not directly connected, further tightening the constraints [2, 4].

These estimates are not exact but illustrate the scale of the resource gap between theoretical and experimental implementations. Noisy simulations often fail to capture the full behavior of real devices, and quantum hardware access remains limited. A notable hardware demonstration was conducted by Geng et al. [4], who tested an optimized variant of FRQI. Although the theoretical encoding was designed to support images up to 64×64 pixels, experimental results on superconducting quantum hardware showed that only 2×2 pixel images could be reliably reconstructed from quantum encoding. Figure 1 compares the target image, simulated reconstruction, and hardware reconstruction for the failed 4×4 image case.

V. Discussion

In image-based machine learning, datasets are typically large and high-dimensional, often comprising millions of pixels. The previous analysis focused exclusively on quantum data loading, which already presents a significant bottleneck. This issue is particularly acute in quantum machine learning (QML), where the overhead of encoding large datasets limits feasibility on current NISQ hardware [2].

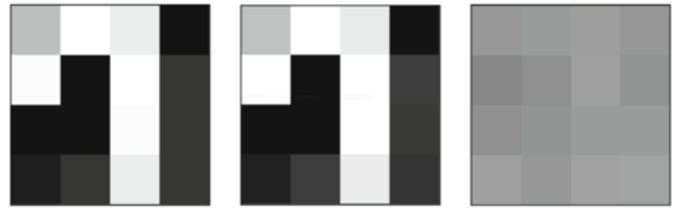

(a) Input image (b) `qasm_simulator` (c) `ibmq_toronto`

Fig. 1: Input image and FRQI reconstructions, both simulated and from quantum hardware. Image credit: [4]

Although quantum devices have shown early utility in simulation and optimization, clear QML utility remains elusive. The gap lies largely in the mismatch between dataset complexity and current hardware constraints. Hybrid quantum-classical approaches aim to reduce quantum resource demands, but they bring new challenges, such as repeated classical–quantum communication and unclear quantum advantage [5].

To advance QML under existing hardware limitations, attention should shift toward small, low-dimensional datasets that remain challenging for classical models. Such tasks avoid the data loading barrier while offering a meaningful testbed for quantum processing. Identifying and targeting these problems is essential to exploring the practical potential of future QML applications, beyond the current data loading bottleneck.

VI. Conclusion

Quantum image encoding schemes such as FRQI demonstrate the promise of exponential compression in qubit usage, but also expose the critical bottleneck of circuit depth. While simulations suggest feasibility for encoding high-resolution images, real quantum hardware imposes severe constraints. A promising direction in the context of QML is to find alternative datasets that are information-rich yet low in size or dimensionality, thereby enabling the exploration of competitive QML algorithms under realistic hardware constraints.

References

[1] Michael A Nielsen and Isaac L Chuang. *Quantum computation and quantum information*. Cambridge university press, 2010.

[2] Christian Bauckhage et al. *Quantum Machine Learning - State of the Art and Future Directions*. Tech. rep. Germany: Federal Office for Information Security, 2022, p. 122.

[3] Marina Lisnichenko and Stanislav Protasov. “Quantum image representation: A review”. In: *Quantum Machine Intelligence* 5.1 (2023), p. 2.

[4] Alexander Geng et al. “Improved FRQI on superconducting processors and its restrictions in the NISQ era”. In: *Quantum Information Processing* 22.2 (2023), p. 104.

[5] Alexandru Ionita and Bogdan Ionescu. “Hybrid Quantum-Classical Neural Networks for Image Classification: Challenges and Perspectives”. In: *ISSCS* (2025).