

HQF-Net: A Hybrid Quantum-Classical Multi-Scale Fusion Network for Remote Sensing Image Segmentation

Supplementary Material

1. Quantum Operations Used in HQF-Net

HQF-Net employs parameterized quantum circuits to transform compact classical feature representations in the bottleneck and skip pathways. A qubit state can be represented as a superposition of basis states, allowing compressed feature responses to be encoded in a higher-dimensional Hilbert space. Parameterized rotation gates, such as $R_Y(\theta)$, are used to introduce learnable transformations, while CNOT-based entanglement operations model interactions between qubits. In HQF-Net, these quantum operations are not used as a standalone quantum model; instead, they function as feature enrichment modules embedded within a hybrid encoder–decoder segmentation pipeline. This design enables the network to capture structured correlations in compressed latent features while preserving compatibility with classical convolutional processing.

To clarify the role of these quantum modules in HQF-Net, we briefly review the quantum concepts that underpin their design. Specifically, qubit state representations provide the basis for encoding compressed classical features, parameterized quantum gates define learnable transformations over these encoded states, and entanglement enables the modeling of correlations between latent feature components. These concepts form the foundation of the quantum operations used in both the QMoE bottleneck and QSkip pathways.

1.1. Qubit & Hilbert Space Representation

The fundamental unit of quantum information is the qubit, analogous to a classical bit but with the unique property of being able to exist in a superposition [5, 7] of states. A qubit’s state $|\psi\rangle$ is represented as a linear combination of the basis states $|0\rangle$ and $|1\rangle$:

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$$

where α and β are complex probability amplitudes satisfying the normalization condition $|\alpha|^2 + |\beta|^2 = 1$. The ability of qubits to exist in superposition enables quantum systems to encode and process information in exponentially larger representational spaces than classical systems.

A system of Q qubits is described in a tensor product Hilbert space, $\mathcal{H}_{2^Q} = \mathcal{H}_2^{\otimes Q}$, whose dimensionality grows exponentially with Q . For instance, a two-qubit system can exist in a superposition of the four states $|00\rangle, |01\rangle, |10\rangle, |11\rangle$, providing a powerful basis for quantum computation.

1.2. Quantum Gates as Unitary Rotations

Quantum gates are the operations that manipulate qubits. These gates perform unitary transformations that evolve a qubit’s state, and they can be visualized as rotations on the Bloch sphere [4]. The primary gates used in our approach are:

1. **Hadamard Gate (H):** This gate [13] creates superposition by transforming a basis state into an equal superposition of $|0\rangle$ and $|1\rangle$.

$$H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}; \quad H|0\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$$

2. **Rotation Gates:** These parameterized gates [3] rotate the state vector around an axis of the Bloch sphere by an angle θ . These gates are used to introduce learnable transformations and enable flexible feature encoding.

$$R_Y(\theta) = \begin{pmatrix} \cos(\theta/2) & -\sin(\theta/2) \\ \sin(\theta/2) & \cos(\theta/2) \end{pmatrix}$$

The use of these gates [14] allows quantum models to perform transformations on feature representations in ways that may capture intricate interactions that are difficult to model efficiently with standard classical operations.

1.3. Entanglement and Multi-Qubit Gates

Entanglement [6] is a unique quantum phenomenon that describes correlations between qubits that cannot be separated into individual qubit states. These correlations enable quantum circuits to model complex feature interactions in a different representational space from classical systems. The **Controlled-NOT (CNOT)** gate [15] is commonly used to create entanglement between qubits:

$$\text{CNOT}|10\rangle = |11\rangle; \quad \text{CNOT}|01\rangle = |01\rangle$$

By applying a Hadamard gate followed by a CNOT gate, we can create the canonical entangled Bell state:

$$|\Phi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$$

In our model, entanglement is used within quantum circuits at key points in the U-Net architecture, such as the bottleneck, to enhance feature extraction and representation learning. These quantum circuits enable the model to capture high-level correlations across multi-resolution features, which can be beneficial for segmentation tasks.

1.4. Quantum Approaches in Deep Learning

Quantum models have recently been integrated into classical deep learning architectures, where quantum circuits are placed at strategic locations [11, 12], such as the bottleneck layer [2, 8], to enhance feature extraction. The quantum bottleneck approach has shown promise by incorporating quantum circuits at the highest level of feature compression, thereby facilitating better representation learning. Li et al. [9] applied a hybrid model combining classical feature extractors with quantum classifiers, demonstrating that quantum circuits can complement classical architectures in image classification tasks.

By combining classical architectures with quantum circuits [10], hybrid quantum-classical models leverage the strengths of both: the flexibility of classical deep learning for general tasks and the expressiveness of quantum circuits for complex feature interactions. Building on this hybrid paradigm, HQF-Net inserts quantum modules at two complementary locations: the skip pathways for intermediate feature refinement and the bottleneck for compact latent transformation through expert-based routing.

2. Quantum Feature Processing Modules

In HQF-Net, classical feature maps are first projected into a compact latent representation and encoded into quantum states. Structured parameterized circuits are then used to transform these states, and the resulting measurements are projected back to classical feature space. These quantum modules are used in two parts of the architecture: the QSkip pathways and the QMoE bottleneck. The QSkip branch uses a multi-scale enrichment circuit to refine skip features, while the QMoE module routes latent features through specialized quantum experts designed to capture complementary local, global, and directional dependencies.

1. **Enrichment Multi-Scale Circuit.** This circuit is designed to enhance image features by incorporating both global and local feature correlations. The circuit consists of local convolutional layers that apply localized operations on neighboring qubits to extract fine-grained spatial patterns. Then a global mixing layer that entangles all qubits propagates information across the entirety of the input. This design enables the circuit to preserve very detailed, structured information locally while maintaining the overall structure globally, thereby allowing it to capture and represent many different types of patterns in the data. In this architecture, the circuit is utilized in our QSkip module and before the QMoE block.

2. **Localist Expert.** The Localist Circuit captures local and fine-grained features. This circuit is a simple quantum circuit that has entanglements only between neighboring qubits. In effect, this localist circuit mimics a classical two-dimensional separable convolution operator and allows

the capture of small spatial correlations (e.g., an edge, corner, localized texture). Thus its principal benefit is the retrieval of fine detail that would otherwise not be retrieved from any of the deeper encoder layers or from global operations.

3. **Globalist Expert.** The Globalist circuit captures long-distance dependencies as well as the entire range of global patterns across an entire feature map. The Globalist circuit employs no local convolutions, but rather utilizes deep entanglement of qubits and rotation gates. It also creates the ability to model correlations between non-adjacent qubits, allowing the circuit to capture repeating patterns; assigns symmetry to data; and provides a model for capturing long-term structural dependencies. The Globalist circuit complements the Localist expert by encoding relevant global context that cannot be encoded solely through local operations or local connectivity.

4. **Diagonal Expert.** The Diagonal circuit captures structured, directional dependencies that cannot be represented as either a purely local or a purely global circuit. By using parameterized rotation gates followed by CNOT gates on non-adjacent diagonally indexed qubits, this circuit uses a staggered entanglement topology, allowing it to model diagonal or directional relationships among the features in the image, thus creating a structure to help the system learn to identify complex patterns in the features of the image based on directionality and structure.

5. **Two-Qubit Quantum Filter.** The convolution-like operation is implemented using a two-qubit parameterized unitary, as illustrated in Figure 1. This unit serves as a basic local interaction block within the larger circuit designs used in HQF-Net.

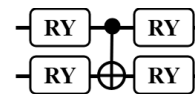


Figure 1. Two-qubit parameterized quantum filter used as a local interaction block in HQF-Net. The unit models pairwise feature interactions and serves as a basic building block within the larger quantum circuit designs.

3. Data Pre-processing and Pipeline

For LandCover.ai and OpenEarthMap, we use random cropping, which samples 224×224 patches from the larger source tiles. For SeasoNet, which has a native patch size of 120×120 , all input images/masks were resized using bilinear/nearest-neighbour interpolation to 224×224 so that they have the same spatial dimensions and can therefore be processed through the encoder and decoder blocks. Finally, each patch is normalised to the range $[0, 1]$.

4. Model Architectural Discussion

HQF-Net combines semantic-guided fusion, multi-scale quantum enrichment, and expert-based quantum routing within a U-Net-style encoder-decoder framework. DMCAF injects semantically informed context into encoder features while preserving spatial fidelity through sparse deformable sampling. The bottleneck compresses the high-dimensional encoder output into a compact latent representation, which is enriched through structured quantum transformations before adaptive routing across specialized experts. In parallel, quantum-enhanced skip pathways refine intermediate features passed to the decoder. Together, these components allow HQF-Net to combine local detail preservation, global context modeling, and adaptive feature specialization for remote sensing segmentation.

5. Qualitative Results for OpenEarthMap and SeasoNet

For OpenEarthMap scenes shown in Fig. 3, which contain complex urban layouts and heterogeneous land-cover patterns, HQF-Net demonstrates improved delineation of buildings, roads, and surrounding land areas. Baseline methods sometimes produce fragmented predictions or confuse neighboring classes, particularly in dense urban regions. HQF-Net generates smoother and more complete segmentation maps while maintaining sharp structural boundaries. From the visual comparisons, conventional architectures such as U-Net tend to produce coarser segmentation masks and occasionally miss small structures or thin road segments. Transformer-based models such as SegFormer, DC-Swin, and UNetFormer improve the overall structural representation; however, they still produce fragmented predictions or minor misclassifications around object boundaries and densely packed regions.

Figure 2 presents qualitative comparisons between HQF-Net and several baseline models on the SeasoNet dataset. This dataset contains multi-temporal agricultural scenes with complex vegetation patterns and subtle class transitions, making accurate segmentation particularly challenging. From the visual results, DeepLabv3-based models with different backbones (DenseNet121, ResNet-50, ConvNeXt-Small, and Swin-Tiny) generally capture the large land-cover regions but often produce coarse boundaries and slight inconsistencies between adjacent classes. Similarly, UPerNet and SegFormer provide improved spatial representation but occasionally introduce minor boundary artifacts or fragmented regions, especially around irregular vegetation patterns and water bodies.

Figure 4 shows representative qualitative segmentation results of HQF-Net on the OpenEarthMap and SeasoNet datasets. The examples demonstrate challenging urban and semi-urban scenes with complex spatial layouts, including

dense buildings, roads, and mixed land cover.

Overall, the qualitative results indicate that HQF-Net provides more stable and spatially coherent predictions on the SeasoNet dataset, effectively capturing both large homogeneous regions and finer structural details in complex agricultural environments.

6. Quantum Circuit Implementation and Simulation Environment

All quantum components of our architecture were designed and simulated using **PennyLane** [1], a widely used quantum differentiable programming library. PennyLane’s seamless integration with **PyTorch** was instrumental, allowing the quantum circuits to be inserted directly into the classical deep learning model as trainable layers. This enabled fully end-to-end training of the entire hybrid system using standard backpropagation.

All quantum simulations were run on classical hardware using PennyLane’s high-performance backends in order to ensure computational feasibility and to speed up our experiments. We used the `lightning.gpu` device for GPU execution. These backends provide much greater speed when compared to pure Python. Furthermore, in order to speed up the training process, we used the adjoint differentiation method to compute the gradients of the trainable parameters in the quantum circuits. Compared to the standard parameter-shift rule, adjoint differentiation is substantially quicker for simulation, as it computes all of the gradients with a constant, small number of circuit executions, allowing backpropagation through the quantum layers to be very efficient.

7. Limitations and Future Directions

While our proposed hybrid quantum architectures show competitive performance, this study has several limitations that highlight key challenges and pave the way for future research in hybrid quantum-classical computer vision.

7.1. Computational Cost of Quantum Simulation

The most significant practical limitation of this work is the computational cost associated with simulating quantum circuits. Although our final architectures were designed to reduce the number of quantum circuit evaluations per forward pass, the training time was still substantially longer than that of a purely classical model. The overhead of initializing the quantum simulator, constructing the circuit, and performing state vector simulation, even on a high-performance backend like `lightning.gpu`, remains a major bottleneck. This underscores the critical need for either more optimized classical simulators or, ultimately, fault-tolerant quantum hardware to make the training of such deep hybrid models truly scalable.




Figure 2. Qualitative segmentation results on the SeasoNet dataset showing original images, ground-truth masks, and comparisons with other models.

7.2. Challenges in Hybrid Model Training and Debugging

Developing and training our hybrid models that are more fully integrated posed challenges beyond those of standard deep learning pipelines. All of these models were able to converge; however, maintaining stable and consistent gradient flow through the quantum layers was non-trivial. This required careful initialization and the use of adjoint differentiation to maintain stable training. However, as quantum circuits become more sophisticated and deeper, it is

likely that they will also face greater optimization difficulties, such as barren plateaus, than what has been studied here.

7.3. Simulation vs. Real-World Hardware

All experiments in this study were conducted on classical computers using quantum simulation. As a result, the findings are approximate representations of how quantum modules would behave under optimal conditions, rather than real quantum hardware. Assessing the performance of these

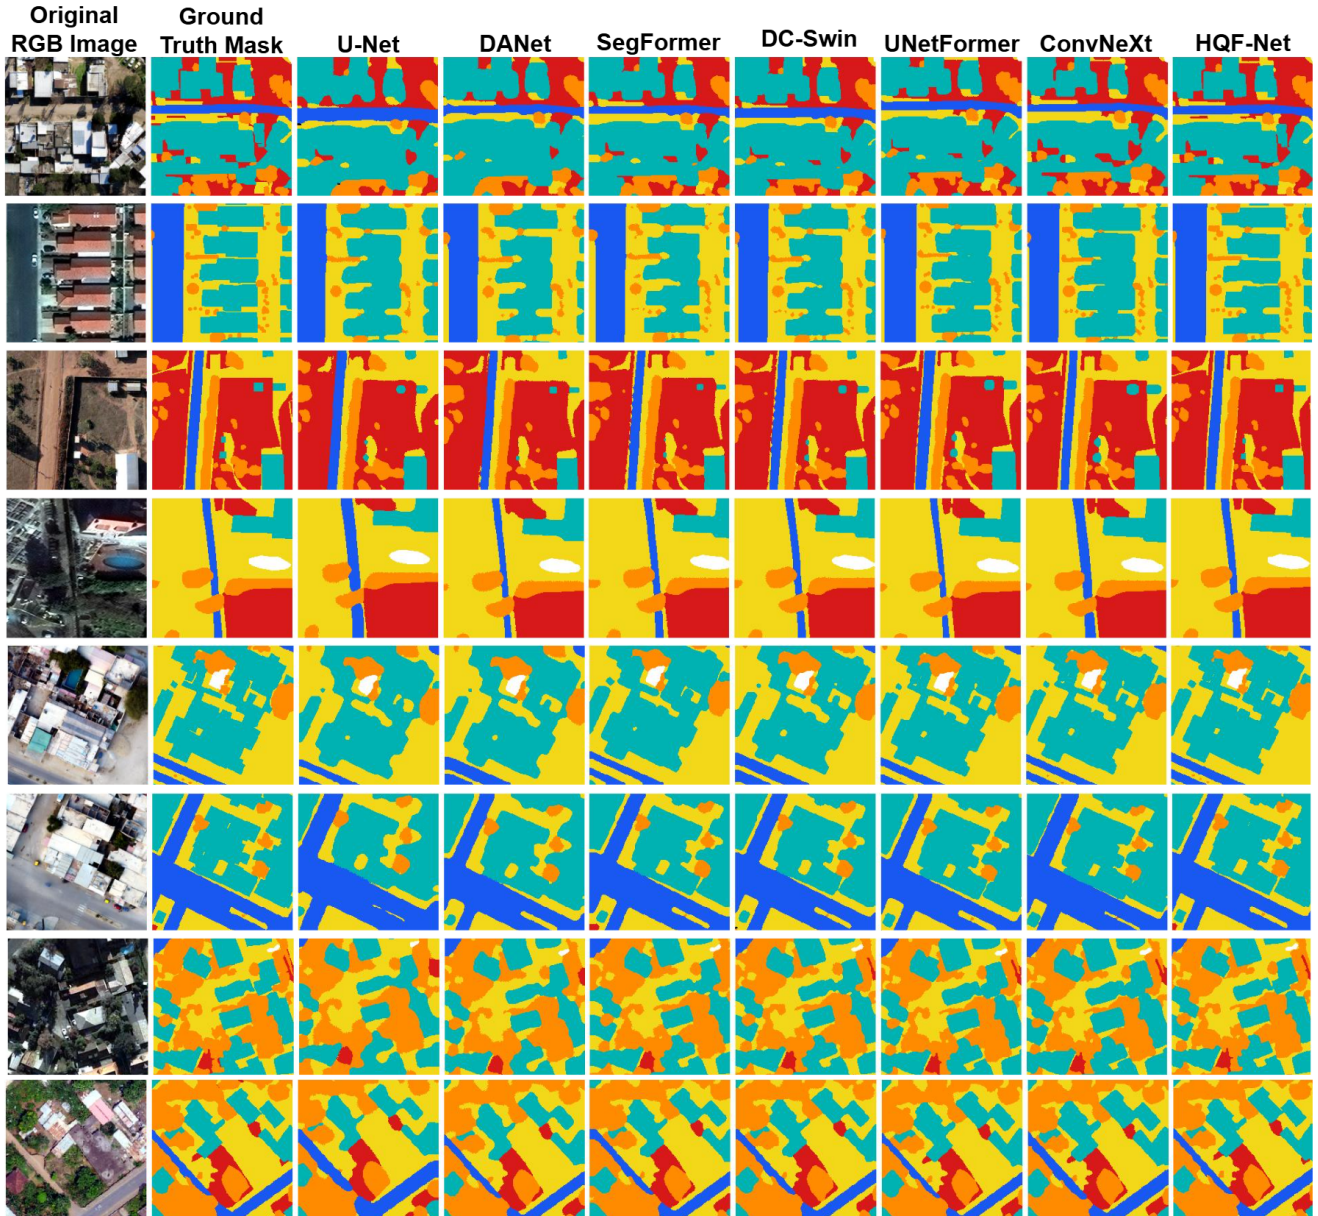


Figure 3. Qualitative segmentation results on the OpenEarthMap dataset showing original images, ground-truth masks, and comparisons with other models.

designs on current NISQ platforms remains an unresolved problem, due to gate errors, qubit decoherence, measurement noise, and limited qubit interconnectivity in today’s physical implementations. Additional research on error mitigation, hardware-aware circuit transpilation, and noise-robust quantum module design would be needed for practical empirical deployment. Therefore, the reported results should be interpreted as an optimistic estimate of the potential behavior of these architectures within the simulated backends.

7.4. Limited Architectural Search

In summary, although we explored several novel architectural designs, the design space for hybrid quantum-classical architectures remains very large. The quantum expert circuits, DINOv3 fusion technique, and U-Net base are merely one point in a very wide design space. More research that examines a larger area will generate other quantum circuit designs, different attention-based algorithms, and various forms of classical backbones resulting in better performance than we have demonstrated in this research paper. This work introduces one promising architecture and evaluates

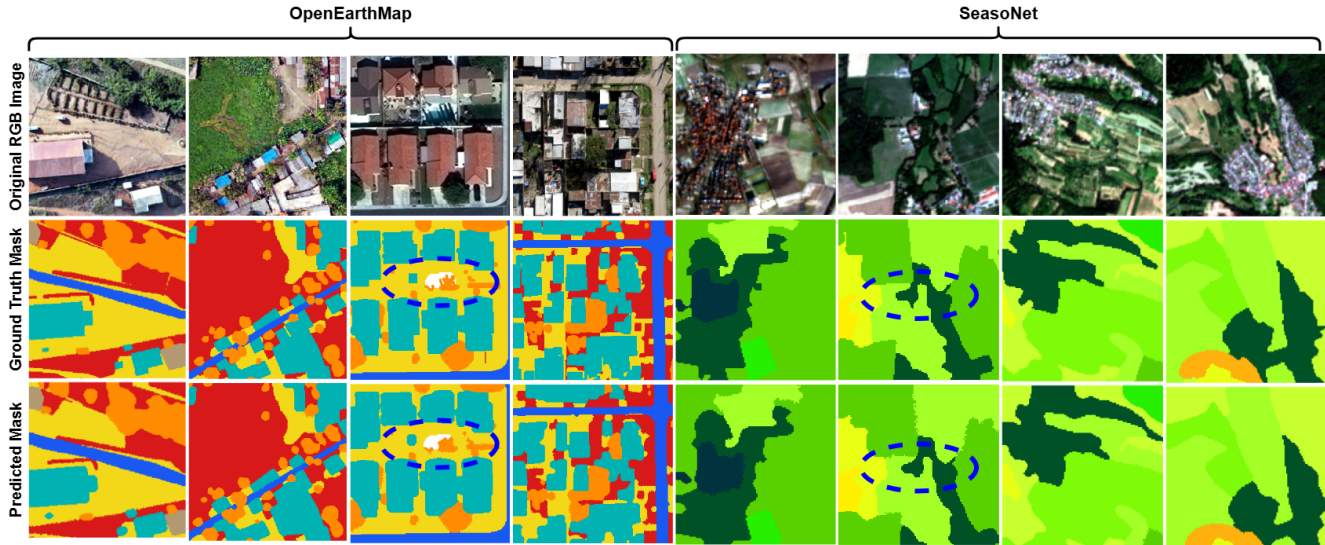


Figure 4. Qualitative segmentation results on the OpenEarthMap and SeasonNet datasets showing original images, ground-truth masks, and HQF-Net predictions.

its capabilities across several settings, but it does not claim to exhaust the broader design space.

References

- [1] Ville Bergholm, Josh Izaac, Maria Schuld, Christian Gogolin, Shahnawaz Ahmed, Vishnu Ajith, M Sohaib Alam, Guillermo Alonso-Linaje, Bharath AkashNarayanan, Ali Asadi, et al. Pennylane: Automatic differentiation of hybrid quantum-classical computations. *arXiv preprint arXiv:1811.04968*, 2018. 3
- [2] Francesca De Falco, Andrea Ceschini, Alessandro Sebastianelli, Bertrand Le Saux, and Massimo Panella. Towards efficient quantum hybrid diffusion models. *arXiv preprint arXiv:2402.16147*, 2024. 2
- [3] Yuxuan Du, Min-Hsiu Hsieh, Tongliang Liu, and Dacheng Tao. Expressive power of parametrized quantum circuits. *Physical Review Research*, 2(3):033125, 2020. 1
- [4] Ian Glendinning. The bloch sphere. In *QIA meeting. Vienna*, 2005. 1
- [5] Lov K Grover. The advantages of superposition. *Science*, 280(5361):228–228, 1998. 1
- [6] Ryszard Horodecki, Paweł Horodecki, Michał Horodecki, and Karol Horodecki. Quantum entanglement. *Reviews of Modern Physics*, 81(2):865–942, 2009. 1
- [7] Ciaran Hughes, Joshua Isaacson, Anastasia Perry, Ranbel F Sun, and Jessica Turner. Introduction to superposition. In *Quantum computing for the quantum curious*, pages 1–5. Springer, 2020. 1
- [8] Md Majedul Islam, Rashik Shahriar Akash, Sayed Mehedi Azim, and Selena He. Qpolypnet: A quantum-inspired deep learning model for polyp segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 969–978, 2025. 2
- [9] Wei Li, Peng-Cheng Chu, Guang-Zhe Liu, Yan-Bing Tian, Tian-Hui Qiu, and Shu-Mei Wang. An image classification algorithm based on hybrid quantum classical convolutional neural network. *Quantum Engineering*, 2022(1):5701479, 2022. 2
- [10] Artur Miroszewski, Jakub Nalepa, Bertrand Le Saux, and Jakub Mielczarek. Quantum machine learning for remote sensing: exploring potential and challenges. *arXiv preprint arXiv:2311.07626*, 2023. 2
- [11] Soronzonbold Otgonbaatar and Mihai Datcu. Classification of remote sensing images with parameterized quantum gates. *IEEE Geoscience and Remote Sensing Letters*, 19:1–5, 2021. 2
- [12] Alessandro Sebastianelli, Daniela Alessandra Zaidenberg, Dario Spiller, Bertrand Le Saux, and Silvia Liberata Ullo. On circuit-based hybrid quantum neural networks for remote sensing imagery classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 15:565–580, 2021. 2
- [13] Dan J Shepherd. On the role of hadamard gates in quantum circuits. *Quantum Information Processing*, 5(3):161–177, 2006. 1
- [14] Colin P Williams. Quantum gates. In *Explorations in quantum computing*, pages 51–122. Springer, 2011. 1
- [15] Hiu Yung Wong. Quantum gate introduction: Not and cnot gates. In *Introduction to Quantum Computing: From a Layperson to a Programmer in 30 Steps*, pages 133–141. Springer, 2023. 1