

Article

Continuous-Variable Quantum Fourier Neural Operator for Solving Partial Differential Equations

Paolo Marcandelli ^{1,*} , Stefano Mariani ¹ , Martina Siena ¹  and Stefano Markidis ^{2,*}

¹ Department of Civil and Environmental Engineering, Politecnico di Milano, 20133 Milan, Italy; stefano.mariani@polimi.it (S.M.); martina.siena@polimi.it (M.S.)

² Department of Computational Science and Technology, KTH Royal Institute of Technology, 114 28 Stockholm, Sweden

* Correspondence: paolo.marcandelli@polimi.it (P.M.); markidis@kth.se (S.M.)

Abstract

Fourier Neural Operators have become a central tool for learning solution operators of partial differential equations, but their spectral layers remain entirely classical and rely on digital Fourier processing. In this work, we introduce the Continuous-Variable Quantum Fourier Neural Operator (CV-QFNO), a Gaussian photonic formulation of the FNO spectral layer. The proposed architecture maps the essential operations of Fourier-domain operator learning, Fourier transformation, mode selection, and channel mixing, onto native continuous-variable optical primitives. In this way, the CV-QFNO provides a photonic quantum analogue of the truncated spectral mechanism underlying the classical FNO, while avoiding the compilation overhead and spectral mismatch that arise in qubit-based Quantum FNO constructions. We extended the framework to both one- and two-dimensional operator learning and validated it on standard PDE benchmarks, including Burgers' equation, heat equation, Navier–Stokes dynamics, and Darcy flow. The results show that the proposed model preserves the predictive accuracy, resolution generalisation, and spectral inductive bias of Fourier neural operators while using structurally constrained photonic parameterisation. Since all the experiments were performed as classical simulations, the contribution should be understood as an architectural and algorithmic blueprint for photonic neural operators rather than as a demonstration of quantum computational advantage.

Keywords: Fourier neural operator; Continuous-Variable Quantum Computing; Photonic machine learning; neural operator; partial differential equations

1. Introduction

Partial differential equations (PDEs) are central to the mathematical modelling of physical and engineering systems, but their numerical solution can become computationally demanding in regimes requiring repeated simulations, real-time prediction, uncertainty quantification, or high-dimensional parametric studies. This has motivated the development of learning-based surrogate models, and in particular neural operators, which aim to approximate solution operators between function spaces rather than individual solution instances [1–3]. Among them, the Fourier Neural Operator (FNO) has become a particularly influential architecture for PDE learning, since it parameterises the nonlocal part of the operator in Fourier space, where many PDE dynamics admit a compact spectral representation [2].

The spectral structure of the FNO makes it especially relevant from the perspective of quantum computing. Each FNO layer applies a Fourier transform, performs a learnable



Academic Editor: Giuliano Benenti

Received: 14 May 2026

Revised: 22 June 2026

Accepted: 24 June 2026

Published: 1 July 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and

conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

transformation on a finite set of retained modes, suppresses the remaining modes, and maps the result back to physical space through the inverse Fourier transform. Thus, the FNO spectral layer is not merely a neural network block but a truncated spectral operator. This structure suggests a natural route toward quantum and photonic implementations, where Fourier transforms and mode-wise operations may be embedded directly into the computational substrate rather than treated as external classical routines.

Qubit-based Quantum Fourier Neural Operators have recently explored this direction by combining unary encodings, Hamming-weight-preserving circuits, and quantum Fourier transforms [4–6]. In this setting, data are represented in fixed-Hamming-weight subspaces, and trainable spectral mixing layers are realised through orthogonal transformations implemented by Reconfigurable Beam Splitter (RBS) gates. This formulation has three important advantages. First, the unary-QFT provides a structured quantum implementation of the Fourier transform: while its gate count scales as $O(N_s \log N_s)$, analogously to the classical FFT, its circuit depth can scale as $O(\log N_s)$ under suitable hardware assumptions allowing parallel operations on disjoint qubit pairs. Second, the architecture provides a direct quantum analogue of the Fourier layer by combining Fourier-domain processing with mode-wise trainable transformations. Third, it constrains the trainable spectral weights to orthogonal matrices, reducing the number of parameters with respect to the unconstrained complex tensors used in classical FNOs. Orthogonal parameterisations are also known to improve conditioning, gradient propagation, and training stability in classical and quantum learning models [7–9].

However, the qubit formulation also presents two structural limitations. First, RBS gates are not native operations on standard qubit hardware and must be compiled into elementary one- and two-qubit gates, introducing a substantial entangling-gate overhead. This makes full QFNO circuits difficult to realise on near-term qubit devices. Second, the parallel QFNO construction does not exactly reproduce the truncated spectral action of the classical FNO: while the selected Fourier modes are processed by trainable circuits, the non-selected modes remain present in the quantum state instead of being set to zero [5]. This introduces a systematic spectral mismatch with the classical FNO layer, whose high-frequency modes are explicitly filtered before applying the inverse Fourier transform.

Continuous-variable (CV) quantum computing provides a natural framework to address both issues. In photonic CV systems, information is carried by bosonic modes and processed through Gaussian operations such as beam splitters, phase shifters, squeezers, and loss channels [10–13]. These operations are directly suited to spectral processing: the Cooley–Tukey structure of the Fourier transform can be mapped onto optical networks of beam splitters and phase rotations, while trainable orthogonal mixing can be implemented by passive interferometers. Moreover, unwanted Fourier modes can be physically suppressed by routing them through loss channels with transmissivity $\eta = 0$, thereby recovering the explicit truncation mechanism of the classical FNO.

In this work, we introduce the Continuous-Variable Quantum Fourier Neural Operator (CV-QFNO), a Gaussian photonic formulation of the FNO spectral layer. The construction builds on the continuous-variable quantum Fourier layer introduced in [14], where a full matrix is embedded into the off-diagonal covariance block of a bipartite Gaussian state and Fourier transformations are implemented through optical Gaussian circuits. Here, we extend that framework from a fixed Fourier processing layer to a trainable neural-operator architecture. In the proposed CV-QFNO, Fourier transformation, mode selection, and channel mixing are all mapped onto native CV optical primitives: optical quantum Fourier transforms perform the spectral change of basis, zero-transmissivity loss channels remove non-retained modes, and pyramidal interferometers implement trainable orthogonal mixing across channels.

This formulation should not be interpreted as a claim of quantum computational advantage. All the experiments in this work were performed as classical simulations of the corresponding Gaussian transformations. Rather, the contribution is architectural and algorithmic: we show that continuous-variable photonic formalism provides a structurally faithful implementation of the FNO spectral mechanism, while avoiding the qubit compilation overhead and the spectral mismatch of previous QFNO constructions.

The main contributions of this work are as follows:

- We introduce the CV-QFNO, a Gaussian photonic formulation of the Fourier Neural Operator spectral layer in which Fourier transformation, mode selection, and channel mixing are expressed in terms of native continuous-variable optical operations.
- We show that the parallel qubit QFNO suffers from a systematic spectral mismatch with the classical FNO, since non-selected Fourier modes are left unchanged. The CV-QFNO removes this mismatch by using loss channels with transmissivity $\eta = 0$, thereby recovering the truncated spectral structure of the classical FNO.
- We develop both one- and two-dimensional CV-QFNO architectures. The one-dimensional construction assigns each retained Fourier mode to an independent optical branch, while the two-dimensional extension introduces cross-circuit interferometric coupling to realise channel mixing across spatial Fourier modes.
- We provide a resource comparison between qubit and continuous-variable formulations, highlighting that the qubit implementation requires a large number of compiled entangling gates, whereas the CV construction is naturally expressed through passive linear-optical elements.
- We benchmark the FNO, QFNO, and CV-QFNO on representative PDE operator-learning tasks, including Burgers' equation, heat equation, Navier–Stokes dynamics, and Darcy flow. The results show that the CV-QFNO preserves the accuracy and resolution-generalisation behaviour of the FNO while using a constrained orthogonal spectral parameterisation.

2. Methods

2.1. Classical Fourier Neural Operator

We briefly recall the classical Fourier Neural Operator (FNO) [2], focusing only on the spectral structure that will serve as the reference target for the qubit and continuous-variable formulations developed below. In the neural-operator framework, the goal is to approximate a solution operator

$$\mathcal{G}^\dagger : \mathcal{A} \rightarrow \mathcal{U},$$

mapping input functions to output functions. Given an input $a : D_j \rightarrow \mathbb{R}^{d_a}$, the FNO first lifts it pointwise to a latent field $v_t : D_j \rightarrow \mathbb{R}^{d_v}$, where t indexes the Fourier layer within the stacked FNO architecture, i.e., v_t denotes the latent field before the t -th Fourier block and v_{t+1} the corresponding output. The latent representation is then updated by Fourier layers of the form

$$v_{t+1}(x) = \sigma\left(Wv_t(x) + \mathcal{F}^{-1}(R_\phi \cdot \mathcal{F}(v_t))(x)\right), \quad (1)$$

where W is a learnable pointwise linear map, σ is a nonlinear activation function, and $\mathcal{F}, \mathcal{F}^{-1}$ denote the Fourier and inverse Fourier transforms along the spatial variables. The symbol “ $\cdot$ ” denotes the mode-wise action of the spectral weights on the Fourier coefficients; in discrete form this corresponds to a matrix–vector or tensor contraction over the channel dimension.

We now make Equation (1) explicit in one and two spatial dimensions. This is not intended as a full review of FNOs, but as a preparation for the quantum constructions in Sections 2.2 and 2.3: both architectures operate on matrix or tensor representations of the

latent field, and must reproduce the Fourier transform, the mode-wise channel mixing, and the suppression of non-retained modes.

2.1.1. One-Dimensional Case

We first consider the common one-dimensional structure underlying the Burgers' and heat-equation benchmarks reported in Section 3. Their specific PDE definitions and numerical settings are given there; here we only need the latent representation on which the FNO spectral layer acts. For both 1D benchmarks, the lifting map P promotes the input $a: D_j \rightarrow \mathbb{R}^{d_a}$ to the latent field $v_t: D_j \rightarrow \mathbb{R}^{d_v}$, evaluated at each of the N_s grid points $\{x_j\}_{j=1}^{N_s}$. Collecting these evaluations column-wise defines the latent matrix

$$A = \begin{pmatrix} v_t(x_1) & v_t(x_2) & \cdots & v_t(x_{N_s}) \end{pmatrix} \in \mathbb{R}^{N_c \times N_s}, \quad N_c = d_v, \quad (2)$$

where each column $v_t(x_j) \in \mathbb{R}^{N_c}$ is the channel vector at spatial location x_j . Applying the Fourier transform $\mathcal{F}$ along the spatial direction transforms each row of A , producing

$$\hat{A} = \begin{pmatrix} \hat{v}_t(1) & \hat{v}_t(2) & \cdots & \hat{v}_t(N_s) \end{pmatrix} \in \mathbb{C}^{N_c \times N_s}, \quad (3)$$

where the hat denotes quantities represented in Fourier space. Thus, $\hat{A}$ is the Fourier-space representation of the latent matrix A , obtained by applying $\mathcal{F}$ along the spatial direction, and $\hat{v}_t(k) \in \mathbb{C}^{N_c}$ collects the k -th Fourier coefficient across all channels simultaneously. The learnable operator $R_\phi^{(k)} \in \mathbb{C}^{N_c \times N_c}$ then acts along the channel direction on each retained mode independently,

$$\hat{v}_t(k) \mapsto R_\phi^{(k)} \hat{v}_t(k), \quad k = 1, \dots, K, \quad (4)$$

while all non-retained modes are set to zero. Since $R_\phi^{(k)}$ is an unconstrained matrix in $\mathbb{C}^{N_c \times N_c}$, writing $R_\phi^{(k)} = \text{Re}(R_\phi^{(k)}) + i \text{Im}(R_\phi^{(k)})$, the product $R_\phi^{(k)} \hat{v}_t(k)$ *cross-mixes* the real and imaginary parts of the mode vector:

$$\begin{aligned} \text{Re}(R_\phi^{(k)} \hat{v}_t(k)) &= \text{Re}(R_\phi^{(k)}) \text{Re}(\hat{v}_t(k)) - \text{Im}(R_\phi^{(k)}) \text{Im}(\hat{v}_t(k)), \\ \text{Im}(R_\phi^{(k)} \hat{v}_t(k)) &= \text{Im}(R_\phi^{(k)}) \text{Re}(\hat{v}_t(k)) + \text{Re}(R_\phi^{(k)}) \text{Im}(\hat{v}_t(k)). \end{aligned} \quad (5)$$

The two output components are therefore independent linear combinations of *both* input components, with no structural coupling between $\text{Re}(R_\phi^{(k)})$ and $\text{Im}(R_\phi^{(k)})$. This unconstrained complex parameterisation uses $2N_c^2$ real degrees of freedom per retained mode. The resulting filtered Fourier matrix is

$$\hat{A}' = \begin{pmatrix} R_\phi^{(1)} \hat{v}_t(1) & \cdots & R_\phi^{(K)} \hat{v}_t(K) & 0 & \cdots & 0 \end{pmatrix} \in \mathbb{C}^{N_c \times N_s}, \quad (6)$$

where the last $N_s - K$ columns are zeroed out. Since $\mathcal{F}^{-1}$ acts along the spatial direction, i.e., row-wise across the N_s columns of $\hat{A}'$, it can be applied outside the matrix to recover the output latent matrix,

$$A' = \mathcal{F}^{-1}(\hat{A}') \in \mathbb{R}^{N_c \times N_s}, \quad A' = \begin{pmatrix} v_{t+1}(x_1) & \cdots & v_{t+1}(x_{N_s}) \end{pmatrix}, \quad (7)$$

whose j -th column $v_{t+1}(x_j)$ is the updated channel vector at spatial location x_j , consistent with Equation (1) evaluated at $x \in D_j \subset \mathbb{R}$.

2.1.2. Two-Dimensional Case

For two-dimensional operator-learning problems, such as the Navier–Stokes and Darcy-flow benchmarks reported in Section 3, the latent field is evaluated on an $N_x \times N_y$

spatial grid. After lifting, we collect the channel values at each spatial location into a tensor $A \in \mathbb{R}^{N_c \times N_x \times N_y}$, with $N_c = d_v$, defined by

$$A_{:,i,j} = v_t(x_i, y_j) \in \mathbb{R}^{N_c}, \quad i = 1, \dots, N_x, \quad j = 1, \dots, N_y. \quad (8)$$

Here the colon in $A_{:,i,j}$ denotes all entries along the channel dimension: for each spatial grid point (x_i, y_j) , the slice $A_{:,i,j}$ is the N_c -dimensional vector collecting all latent channels at that location. Applying the 2D Fourier transform $\mathcal{F}_{2D}$ along both spatial directions transforms each channel independently, producing

$$\hat{A}_{:,k_1,k_2} = \hat{v}_t(k_1, k_2) \in \mathbb{C}^{N_c}, \quad k_1 = 1, \dots, N_x, \quad k_2 = 1, \dots, N_y, \quad (9)$$

where $\hat{v}_t(k_1, k_2)$ collects the (k_1, k_2) -th 2D Fourier coefficients across all channels simultaneously. The learnable operator $R_\phi^{(k_1,k_2)} \in \mathbb{C}^{N_c \times N_c}$ then acts along the channel direction on each retained mode pair independently,

$$\hat{v}_t(k_1, k_2) \mapsto R_\phi^{(k_1,k_2)} \hat{v}_t(k_1, k_2), \quad k_1 = 1, \dots, K_x, \quad k_2 = 1, \dots, K_y, \quad (10)$$

while all non-retained mode pairs are set to zero. As in the 1D case (cf. Equation (5)), the unconstrained complex matrix $R_\phi^{(k_1,k_2)} \in \mathbb{C}^{N_c \times N_c}$ cross-mixes the real and imaginary parts of each retained 2D Fourier mode vector, with $2N_c^2$ real degrees of freedom per retained mode pair. The resulting filtered Fourier tensor is

$$\hat{A}'_{:,k_1,k_2} = \begin{cases} R_\phi^{(k_1,k_2)} \hat{v}_t(k_1, k_2) & k_1 \leq K_x, k_2 \leq K_y, \\ 0 & \text{otherwise,} \end{cases} \quad (11)$$

with $\hat{A}' \in \mathbb{C}^{N_c \times N_x \times N_y}$. Since $\mathcal{F}_{2D}^{-1}$ acts along both spatial directions, it can be applied outside the tensor to recover the output latent field,

$$\begin{aligned} A' &= \mathcal{F}_{2D}^{-1}(\hat{A}') \in \mathbb{R}^{N_c \times N_x \times N_y}, \\ A'_{:,i,j} &= v_{t+1}(x_i, y_j), \end{aligned} \quad (12)$$

whose (i, j) -th slice $v_{t+1}(x_i, y_j)$ is the updated channel vector at spatial location (x_i, y_j) , consistent with Equation (1) evaluated at $(x, y) \in D_j \subset \mathbb{R}^2$.

2.2. Quantum FNO: Qubit Formulation

In this section we review the qubit-based Quantum Fourier Neural Operator introduced in [5], which provides the direct precursor to our continuous-variable formulation. We present the core theoretical ingredients of the architecture, namely data encoding, the unary quantum Fourier transform, and the learnable spectral mixing layer, and we then examine the two structural limitations that motivate the transition to a continuous-variable framework: the hardware compilation overhead inherent to the qubit setting, and the systematic spectral mismatch with the classical FNO introduced by the parallel circuit variant. The notation for the latent representation depends on the spatial dimensionality of the problem: in one dimension the latent field is collected into a matrix $A \in \mathbb{R}^{N_c \times N_s}$, whereas in two dimensions it is collected into a tensor $A \in \mathbb{R}^{N_c \times N_x \times N_y}$. The following treatment, as presented in [5], focuses on the one-dimensional setting, so that the latent representation is the matrix A defined in Equation (2). For an extension to higher-dimensional domains we refer the reader to [6].

2.2.1. Data Encoding

The goal of the encoding step is to embed the latent matrix $A \in \mathbb{R}^{N_c \times N_s}$ defined in Equation (2) into a quantum state, so that subsequent quantum operations reproduce the spectral mixing of Equation (4) at the quantum level. This matrix is encoded as a quantum state in the *unary basis* by constructing the two-register state

$$|A\rangle = \frac{1}{\|A\|} \sum_{i=1}^{N_c} \sum_{j=1}^{N_s} a_{ij} |e_j\rangle |e_i\rangle, \quad (13)$$

where $|e_k\rangle$ denotes the k -th unary basis vector (a computational basis state with a single 1 at position k). This encoding requires $N_c + N_s$ qubits and can be loaded with a butterfly circuit with depth $O(\log N_c + 2N_c \log N_s)$ and gate count $(N_c - 1) + (2N_c - 1)(N_s - 1)$. See [5,15] for more information about unary encoding and circuit depth.

2.2.2. Unary Quantum Fourier Transform

The classical FFT butterfly structure is reproduced quantumly via a circuit composed of single-qubit phase gates and parameterised two-qubit Reconfigurable Beam Splitter (RBS) gates [5]. Applied to the lower (N_s -qubit) register of the encoded state $|A\rangle$ in Equation (13), this *Unary-QFT* implements the row-wise discrete Fourier transform of A entirely within the unary subspace. The full two-register action is

$$\text{QFT} |A\rangle = \frac{1}{\|A\|} \sum_{i=1}^{N_c} \text{QFT} \left(\sum_{j=1}^{N_s} a_{ij} |e_j\rangle \right) |e_i\rangle = \frac{1}{\|A\|} \sum_{i=1}^{N_c} \sum_{j=1}^{N_s} \hat{a}_{ij} |e_j\rangle |e_i\rangle =: |\hat{A}\rangle, \quad (14)$$

where $\hat{a}_{ij}$ are the Fourier coefficients of the i -th row of A , so that the quantum state $|\hat{A}\rangle$ encodes the transformed matrix defined in Equation (3). This is the quantum analogue of applying $\mathcal{F}$ row-wise to A along the spatial direction: the channel index i (upper register) is untouched, while the spatial index j (lower register) is transformed, in direct correspondence with the classical operation $A \mapsto \hat{A}$. The circuit requires $O(N_s \log N_s)$ gates and runs in depth $O(\log N_s)$ when disjoint qubit pairs can be addressed in parallel, as in trapped-ion or cold-atom hardware.

2.2.3. Parallel Quantum Fourier Layer and Its Limitations

The architecture we describe is the Parallel QFL of [5]. After the Unary-QFT has mapped the encoded state $|A\rangle$ to $|\hat{A}\rangle$ (see Equation (14)), the intended operation is to reproduce the filtered spectral matrix of Equation (6). This is done, at the level of the target FNO layer, by applying an independent learnable mixing $R_\phi^{(k)}$ to each of the K retained Fourier modes and setting the remaining $N_s - K$ modes to zero. To this end, K independent quantum circuits are prepared, one for each retained mode $k \in \{1, \dots, K\}$. In circuit k , a pyramidal interferometer W_Q^k (Figure 1) is applied to the upper (N_c -qubit) register, controlled on the k -th qubit of the lower register: the gate fires only when the lower register is in state $|e_k\rangle$, i.e., only for Fourier mode k .

Since each $\text{RBS}(\theta)$ gate restricted to the unary subspace acts as a Givens rotation $G_{ij}(\theta) \in O(N_c)$ and any orthogonal matrix admits a decomposition into at most $\frac{N_c(N_c-1)}{2}$ such rotations [16], the pyramidal circuit with one RBS per pair (i, j) provides a universal parameterisation of $O(N_c)$ [17]. The resulting weight matrices $\{W_Q^k\}_{k=1}^K$ are therefore fully learnable orthogonal operators. This is a deliberate architectural choice: by restricting the spectral mixing to $O(N_c) \subset U(N_c)$ the same real matrix W_Q^k is applied independently to the real and imaginary parts of each Fourier mode vector

$$\text{Re}(W_Q^k \hat{v}_t(k)) = W_Q^k \text{Re}(\hat{v}_t(k)), \quad \text{Im}(W_Q^k \hat{v}_t(k)) = W_Q^k \text{Im}(\hat{v}_t(k)), \quad (15)$$

without the cross-coupling between real and imaginary components that characterises the classical FNO complex weight (cf. Equation (5)). The orthogonality constraint reduces the parameter count per mode from $2N_c^2$ to $\frac{N_c(N_c-1)}{2}$, while improving gradient conditioning and reducing overfitting [7–9].




Figure 1. Pyramidal interferometer on $N_c = 5$ modes implementing an arbitrary 5×5 orthogonal matrix in the unary subspace. Each box represents an $\text{RBS}(\theta_k)$ gate with an independent trainable parameter θ_k . Gates in the same column act on disjoint wire pairs and run in parallel. The circuit uses $\frac{N_c(N_c-1)}{2} = 10$ gates arranged in $2N_c - 3 = 7$ sequential timesteps, with at most $\lfloor N_c/2 \rfloor = 2$ gates per timestep, achieving circuit depth $O(N_c)$.

While the restriction to $O(N_c)$ can in principle be relaxed to $U(N_c)$ to recover the cross-mixing structure of the classical FNO (cf. Equation (5)), we retain the real orthogonal parameterisation for the training-efficiency reasons discussed in Section 2.4. The effective action of circuit k is

$$|\hat{A}\rangle \mapsto \frac{1}{\|A\|} \sum_{i=1}^{N_c} \sum_{j=1}^{N_s} \tilde{a}_{ij}^{(k)} |e_j\rangle |e_i\rangle, \quad \tilde{a}_{ij}^{(k)} = \begin{cases} (W_Q^k \hat{v}_t(k))_i & j = k, \\ (\hat{v}_t(j))_i & j \neq k. \end{cases} \quad (16)$$

After applying the Unary-IQFT to the lower register, the output of circuit k is

$$\text{output}_k = \mathcal{F}^{-1}[\hat{v}_t(1), \dots, W_Q^k \hat{v}_t(k), \dots, \hat{v}_t(N_s)], \quad (17)$$

which modifies only mode k but carries the full Fourier content of all the other modes $j \neq k$. Summing the K circuit outputs therefore does not reproduce the classical FNO target $A^* = \mathcal{F}^{-1}[W_Q^1 \hat{v}_t(1), \dots, W_Q^K \hat{v}_t(K), 0, \dots, 0]$: each circuit contributes a spurious background from the non-selected modes, and the sum accumulates K such copies. To see this explicitly, decompose each output by linearity of $\mathcal{F}^{-1}$:

$$\begin{aligned} \text{output}_k &= \mathcal{F}^{-1}[\hat{v}_t(1), \dots, \hat{v}_t(N_s)] + \mathcal{F}^{-1}[0, \dots, (W_Q^k - I) \hat{v}_t(k), \dots, 0] \\ &= A + \mathcal{F}^{-1}[\Delta \hat{A}^{(k)}], \end{aligned} \quad (18)$$

where $\Delta \hat{A}^{(k)} \in \mathbb{C}^{N_c \times N_s}$ has $(W_Q^k - I) \hat{v}_t(k)$ in column k and zero elsewhere. Sum over all K circuits and apply linearity of $\mathcal{F}^{-1}$ to collect the $\Delta \hat{A}^{(k)}$ terms:

$$\begin{aligned} \sum_{k=1}^K \text{output}_k &= K \cdot A + \mathcal{F}^{-1} \left[\sum_{k=1}^K \Delta \hat{A}^{(k)} \right] \\ &= K \cdot A + \mathcal{F}^{-1}[W_Q^1 \hat{v}_t(1), \dots, W_Q^K \hat{v}_t(K), 0, \dots, 0] - \mathcal{F}^{-1}[\hat{v}_t(1), \dots, \hat{v}_t(K), 0, \dots, 0] \\ &= A^* + (K - 1) \cdot A + \mathcal{F}^{-1}[0, \dots, 0, \hat{v}_t(K + 1), \dots, \hat{v}_t(N_s)], \end{aligned} \quad (19)$$

where $A^* = \mathcal{F}^{-1}[W_Q^1 \hat{v}_t(1), \dots, W_Q^K \hat{v}_t(K), 0, \dots, 0]$ is the classical FNO target defined in Equation (7), and we use $A = A_K + \mathcal{F}^{-1}[0, \dots, 0, \hat{v}_t(K+1), \dots, \hat{v}_t(N_s)]$ with $A_K = \mathcal{F}^{-1}[\hat{v}_t(1), \dots, \hat{v}_t(K), 0, \dots, 0]$. The deviation from A^* has two distinct sources: $(K-1) \cdot A$, arising from the $K-1$ spurious background copies accumulated across circuits, and the residual contribution of the non-selected modes $k > K$, which are never zeroed.

Note that if in each circuit k all modes $j \neq k$ were set to zero before the IQFT, the output of circuit k would reduce to $\mathcal{F}^{-1}[0, \dots, W_Q^k \hat{v}_t(k), \dots, 0]$ and the sum over K circuits would give exactly A^* by linearity of $\mathcal{F}^{-1}$. This is precisely the fix introduced in Section 2.3: routing the non-selected modes through loss channels with transmissivity $\eta = 0$ erases their contribution before summation, restoring exact equivalence with Equations (7) and (12).

In terms of circuit resources, the K pyramidal interferometers run on independent registers and therefore contribute a total of $K \cdot \frac{N_c(N_c-1)}{2}$ RBS gates, equivalently $K \cdot N_c(N_c-1)$ CZ gates, at an overall depth of $2N_c - 3 = O(N_c)$, since all K circuits execute in parallel.

2.2.4. RBS Gate and Hardware Overhead

Throughout the whole QFNO implementation, the elementary two-qubit operation is the RBS gate. This gate is not native to any current qubit platform and must be compiled into hardware primitives, incurring a substantial overhead that places the full QFNO beyond near-term reach and motivates the continuous-variable reformulation of Section 2.3. The RBS gate at angle θ acts on two qubits as

$$\text{RBS}(\theta) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos \theta & \sin \theta & 0 \\ 0 & -\sin \theta & \cos \theta & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad (20)$$

which, restricted to the unary subspace $\{|01\rangle, |10\rangle\}$, implements a 2×2 rotation by θ . On standard qubit hardware there is no native RBS gate, so it must be compiled into elementary operations. A minimal decomposition is

$$\text{RBS}(\theta) = (H \otimes H) \text{CZ} (R_Y(\theta) \otimes R_Y(-\theta)) \text{CZ} (H \otimes H), \quad (21)$$

requiring 2 entangling (CZ) gates and 6 single-qubit gates per RBS. Here H denotes the single-qubit Hadamard gate, which maps computational basis states to equal superpositions, while CZ denotes the controlled-Z entangling gate, which applies a phase flip when both qubits are in state $|1\rangle$ [18]. Table 1 summarises the depth and CZ-gate count for each phase of a single Parallel QFNO layer.

For the representative choice $N_c = N_s = 32$, $K = 4$ the total number of CZ gates is 8256 per layer at a sequential circuit depth of 396. Current superconducting and trapped-ion processors achieve two-qubit gate fidelities of 99–99.9% [19]. At 99% fidelity the overall circuit fidelity per layer is

$$0.99^{8256} \approx 10^{-36},$$

and even at the optimistic 99.9% level it falls to $0.999^{8256} \approx 2.6 \times 10^{-4}$, rendering the output indistinguishable from noise for any practically relevant (N_c, N_s) . This positions the full QFNO squarely beyond near-term NISQ capabilities, motivating the shift to a photonic continuous-variable architecture in which the analogous operations are implemented by passive beam splitters, linear-optics elements that realise exact unitary rotations on the field quadratures without any qubit–qubit entangling interaction. The dominant noise source in photonic platforms is propagation loss [20], which is qualitatively different from and typically far smaller than the coherent error rates of two-qubit gates on current qubit hardware.

Table 1. Circuit resources per phase of a single Parallel QFNO layer. Each RBS gate compiles to 2 CZ gates (Equation (21)). Encoding counts from [15]; QFT/IQFT butterfly from [5].

Phase	Depth	RBS Gates	CZ Gates
Encoding	$\log_2 N_c + 2N_c \log_2 N_s$	$(N_c - 1) + (2N_c - 1)(N_s - 1)$	$2[(N_c - 1) + (2N_c - 1)(N_s - 1)]$
Unary-QFT	$\log_2 N_s$	$\frac{N_s}{2} \log_2 N_s$	$N_s \log_2 N_s$
QFL (K parallel circuits)	$2N_c - 3$	$K \cdot \frac{N_c(N_c-1)}{2}$	$K \cdot N_c(N_c - 1)$
Unary-IQFT	$\log_2 N_s$	$\frac{N_s}{2} \log_2 N_s$	$N_s \log_2 N_s$
Total	$O(N_c \log N_s)$	$O(KN_c^2 + N_c N_s + N_s \log N_s)$	

In summary, the qubit QFNO has two fundamental limitations: a spectral mismatch with the classical FNO (Equation (19)) and a prohibitive hardware overhead on qubit platforms, despite the RBS gate being natively a linear-optical element. These are the main motivations for the CV-QFNO introduced in Section 2.3.

2.3. Continuous-Variable Quantum Fourier Neural Operator

We now introduce the Continuous-Variable Quantum Fourier Neural Operator (CV-QFNO), which constitutes the main architectural contribution of this work. The proposed model reformulates the spectral layer of the Fourier Neural Operator within a continuous-variable photonic framework, using the qubit-based QFNO reviewed in Section 2.2 as its direct conceptual precursor.

The qubit-based QFNO encodes data in the unary subspace and implements the learnable spectral mixing through orthogonal rotations realised by RBS gates. Since an RBS gate is the discrete-variable analogue of a linear-optical beam splitter, the structure of the QFNO already suggests a natural connection with photonic hardware. Moving to a full CV formulation [11] makes this connection explicit: the Fourier transform, the trainable orthogonal mixing, and the suppression of unwanted modes can all be expressed directly in terms of Gaussian photonic operations.

This reformulation is not only a change of computational substrate but also addresses a key limitation of the qubit construction. As shown in Section 2.2, the parallel QFNO architecture leaves the non-selected Fourier modes unchanged, rather than setting them to zero. This produces a systematic mismatch with the classical FNO spectral layer (cf. Equation (19)), whose action is explicitly truncated in Fourier space. In the CV setting, by contrast, unused modes can be erased directly by routing them through loss channels with transmissivity $\eta = 0$. Physically, this corresponds to sending the corresponding optical modes to a beam dump; mathematically, it removes their correlations with the active register. As a result, the CV-QFNO can reproduce the truncated spectral operator of the classical FNO exactly on the retained modes.

2.3.1. Bipartite Gaussian Encoding and CV-QFT

We recall only the elements of the continuous-variable Fourier layer that are needed to construct the CV-QFNO. The full derivation of the bipartite Gaussian encoding, the TMS-SVD loading procedure, and the Cooley–Tukey optical QFT is given in [14].

The latent matrix $A \in \mathbb{R}^{N_c \times N_s}$ introduced in Equation (2) is encoded into a bipartite Gaussian state with two registers. The first register, denoted by r_1 , contains N_c bosonic modes associated with the channel dimension, while the second register, denoted by r_2 , contains N_s bosonic modes associated with the spatial dimension. We write x_{r_1}, p_{r_1} for the vectors of position and momentum quadratures of the modes in register r_1 , and x_{r_2}, p_{r_2} for the corresponding quadrature vectors of register r_2 .

Using the TMS-SVD encoding of [14], the singular values of A are loaded by independent two-mode squeezing gates, while the left and right singular vectors are implemented by passive interferometers acting on the two registers. The resulting bipartite Gaussian state encodes the latent matrix in the inter-register covariance blocks as

$$\sigma_{x_{r_1} x_{r_2}} = A, \quad \sigma_{p_{r_1} p_{r_2}} = -A, \quad \sigma_{x_{r_1} p_{r_2}} = \sigma_{p_{r_1} x_{r_2}} = 0. \quad (22)$$

Here $\sigma_{x_{r_1} x_{r_2}}$ denotes the block of the covariance matrix collecting all second-order correlations between the x -quadratures of registers r_1 and r_2 . Analogously, $\sigma_{p_{r_1} p_{r_2}}$ collects the inter-register p -quadrature correlations, while $\sigma_{x_{r_1} p_{r_2}}$ and $\sigma_{p_{r_1} x_{r_2}}$ are the mixed inter-register blocks. Thus, the classical latent matrix A is represented exactly as the x -quadrature cross-correlation block of the bipartite covariance matrix.

On this representation, Fourier transformations are implemented by applying Cooley–Tukey optical QFT circuits to the corresponding registers. Each butterfly of the classical FFT is mapped to Gaussian photonic operations, namely phase rotations and 50:50 beam splitters, as shown in [14]. Applying the QFT to the spatial register r_2 implements the Fourier matrix F_{N_s} on the N_s spatial modes. Since the latent matrix is stored in the cross-covariance block $\sigma_{x_{r_1} x_{r_2}} = A$, this operation transforms all rows of A simultaneously:

$$\sigma'_{x_{r_1} x_{r_2}} = \sigma_{x_{r_1} x_{r_2}} F_{N_s}^\top = A F_{N_s}^\top = \hat{A}. \quad (23)$$

This is the one-dimensional Fourier representation used in the 1D CV-QFNO construction, where the Fourier transform acts along the spatial direction while the channel register is left unchanged.

Analogously, applying the QFT to the channel register r_1 implements F_{N_c} on the N_c channel modes, giving

$$\sigma'_{x_{r_1} x_{r_2}} = F_{N_c} \sigma_{x_{r_1} x_{r_2}} = F_{N_c} A. \quad (24)$$

When QFTs are applied simultaneously to both registers, the two-dimensional Fourier transform is obtained:

$$\hat{A} = F_{N_c} A F_{N_s}^\top. \quad (25)$$

Since the covariance matrix is real, the complex matrix $\hat{A}$ cannot be stored in a single covariance block. Instead, its real and imaginary parts are distributed across two inter-register blocks,

$$\begin{aligned} \sigma'_{x_{r_1} x_{r_2}} &= \text{Re}(F_{N_c} A F_{N_s}^\top) = \text{Re}(\hat{A}), \\ \sigma'_{x_{r_1} p_{r_2}} &= \text{Im}(F_{N_c} A F_{N_s}^\top) = \text{Im}(\hat{A}). \end{aligned} \quad (26)$$

Equivalently,

$$\hat{A} = \sigma'_{x_{r_1} x_{r_2}} + i \sigma'_{x_{r_1} p_{r_2}}. \quad (27)$$

Thus, the CV-QFT provides direct access to the complex Fourier coefficients through covariance blocks of the Gaussian state. For the two-register transform, the optical QFT stage has gate count $O(N_c \log N_c + N_s \log N_s)$ and depth $O(\log(N_c N_s))$, excluding the classical SVD preprocessing and encoding costs discussed in [14].

This covariance-block representation also determines how channel mixing is implemented in the CV-QFNO. A retained Fourier mode k is represented by its real and imaginary parts across the two covariance blocks above. Applying a passive interferometer $W^{(k)} \in O(N_c)$ to the channel register therefore acts on both parts in parallel,

$$\text{Re}(\hat{a}^{(k)}) \mapsto W^{(k)} \text{Re}(\hat{a}^{(k)}), \quad \text{Im}(\hat{a}^{(k)}) \mapsto W^{(k)} \text{Im}(\hat{a}^{(k)}). \quad (28)$$

Consequently, as in the qubit QFNO of Section 2.2, the CV-QFNO uses real orthogonal spectral mixing matrices rather than unconstrained complex matrices. This reduces the number of trainable parameters from $2N_c^2$ to $\frac{N_c(N_c-1)}{2}$ per retained mode and provides a direct passive-interferometric implementation of the mixing layer. In principle, this restriction could be relaxed to unitary mixing in a more general photonic architecture; here we retain the orthogonal parameterisation for consistency with the qubit formulation and for the training-efficiency reasons discussed in Section 2.4.

The novel contribution of the CV-QFNO lies in faithfully reproducing the action of R_ϕ in Equation (1): after the Fourier transform, each retained mode k must be mixed along the channel dimension by an independent learnable weight matrix, while all non-selected modes are set to zero. The local W branch of Equation (1) remains classical and is unchanged; it is the global convolution branch that requires a dedicated optical implementation. We now describe how this is achieved first for a 1D input and then for the 2D case.

2.3.2. 1D Architecture: K Parallel Circuits

In a discrete-variable setting, the one- or two-dimensional Quantum Fourier Layer could in principle be implemented through controlled operations on the upper (channel) register conditioned on the lower (spatial) register. To avoid the resulting CNOT-type overhead one may instead adopt the Parallel QFNO strategy of Section 2.3, which directly inspires the present CV construction. However, the CV Gaussian formalism only admits passive linear-optics operations, namely beam splitters and phase shifters, which act jointly on all modes and therefore cannot implement this type of conditional logic. As a result, sequential mode-selective weight application within a single circuit is not available in this framework. To overcome this limitation, we follow the intuition of the Parallel QFNO and run K independent optical circuits in parallel, one for each retained Fourier mode k . Each circuit receives the same latent matrix A and proceeds as follows; a schematic representation is provided in Figure 2.

As shown in Section 2.3, A is loaded into a $(N_c + N_s)$ -mode Gaussian state via TMS-SVD (Equation (22)), setting $\sigma_{x_{r_1} x_{r_2}} = A$ and $\sigma_{x_{r_1} p_{r_2}} = 0$. Then, QFT_c is applied to the lower N_s -mode register and real and imaginary parts of $\hat{A}$ distribute across two covariance blocks:

$$\sigma'_{x_{r_1} x_{r_2}} = \text{Re}(\hat{A}), \quad \sigma'_{x_{r_1} p_{r_2}} = \text{Im}(\hat{A}), \quad (29)$$

so that $\hat{A} = \sigma'_{x_{r_1} x_{r_2}} + i \sigma'_{x_{r_1} p_{r_2}}$, the CV analogue of Equation (3). An orthogonal interferometer $O(\theta_k) \in O(N_c)$ acting on the upper register applies the same real matrix to both blocks simultaneously (cf. Equation (28)), which is the exact photonic analogue of the FNO applying a weight matrix to both the real and imaginary parts of each complex Fourier mode vector. For notational compactness, all subsequent operations in this section are written directly on the full complex matrix $\hat{A}$, with the understanding that each step acts equivalently on both covariance blocks. At this point, the first K parallel circuits are assigned independent trainable pyramidal interferometers $O(\theta_k)$, with one interferometer for each retained Fourier mode k .

We now focus on a generic k -th circuit; as will be shown below, the full spectral output is then recovered by summing the contributions of all K circuits through the linearity of the inverse Fourier transform. In the k -th circuit, absent further constraints, applying $O(\theta_k)$ to the upper N_c -mode register would mix the channel dimension across all N_s Fourier columns simultaneously, yielding

$$\hat{A}'^{(k)} = \begin{pmatrix} O(\theta_k) \hat{v}_t(1) & \cdots & O(\theta_k) \hat{v}_t(N_s) \end{pmatrix} = O(\theta_k) \hat{A}, \quad (30)$$

which does not reproduce the mode-selective operation of Equation (6).




Figure 2. CV-QFNO spectral layer, branch $k = 1$ (1D case). The latent matrix $A \in \mathbb{R}^{N_c \times N_s}$ is encoded via TMS-SVD into the off-diagonal covariance block $\sigma_{x_{r_1} x_{r_2}}$, with the n -mode upper register (r_1) and the m -mode lower register (r_2). A row-wise QFT_c on the lower register places $\hat{A}$ in that block. The pyramidal interferometer $O(\theta_1)$ is applied to the upper register; loss channels $\eta = 0$ are applied to all lower-register modes $j \neq 1$, leaving mode $k = 1$ active. Each of the remaining $K - 1$ branches is structurally identical, with weight $O(\theta_k)$ and mode k left free. An IQFT_c followed by homodyne measurement yields the output contribution of mode $k = 1$.

To isolate the k -th mode, all lower-register modes $j \neq k$ are routed through loss channels with transmissivity $\eta = 0$, which replace those modes with vacuum and erase their correlations with the upper register. As a result, only the k -th column of $\hat{A}$ survives across the active covariance blocks, and the k -th parallel circuit produces the filtered Fourier contribution

$$\hat{A}'^{(k)} = \begin{pmatrix} 0 & \cdots & 0 & O(\theta_k) \hat{v}_t(k) & 0 & \cdots & 0 \end{pmatrix} \in \mathbb{C}^{N_c \times N_s}, \quad (31)$$

where the only nonzero column is the k -th one, thereby realising Equation (4) exactly. The inverse Fourier transform is then applied independently within each parallel circuit. In the k -th branch this produces

$$A'^{(k)} = \mathcal{F}^{-1}(\hat{A}'^{(k)}) \in \mathbb{R}^{N_c \times N_s}, \quad (32)$$

and homodyne measurement returns the corresponding contribution to the latent output in the physical domain. The K branches are combined only at this stage, by summing their measured outputs:

$$A' = \sum_{k=1}^K A'^{(k)} = \sum_{k=1}^K \mathcal{F}^{-1}(\hat{A}'^{(k)}). \quad (33)$$

Using the linearity of $\mathcal{F}^{-1}$, one obtains

$$A' = \mathcal{F}^{-1}\left(\sum_{k=1}^K \hat{A}'^{(k)}\right) = \mathcal{F}^{-1}\left(\begin{pmatrix} O(\theta_1) \hat{v}_t(1) & \cdots & O(\theta_K) \hat{v}_t(K) & 0 & \cdots & 0 \end{pmatrix}\right), \quad (34)$$

thereby recovering exactly the filtered spectral structure of Equation (7).

2.3.3. Optical Resource Count

Since the CV-QFNO is built entirely from beam splitters and phase shifters, except for the two-mode squeezing gates used in the encoding step, it incurs no entangling-gate

compilation overhead. Table 2 summarises the circuit depth and beam-splitter count associated with each phase of a single CV-QFNO layer.

Table 2. Optical resources per phase of a single CV-QFNO layer (K parallel circuits). BS denotes a passive beam splitter; the encoding step additionally requires $r = \min(N_c, N_s)$ two-mode squeezers (depth 1). Interferometer counts use the Clements rectangular decomposition [21].

Phase	Depth	Beam Splitters
TMS-SVD encoding (U, Σ, V)	$1 + \max(N_c, N_s)$	$\frac{N_c(N_c-1)}{2} + \frac{N_s(N_s-1)}{2}$
CV-QFT _c	$\log_2 N_s$	$\frac{N_s}{2} \log_2 N_s$
Loss channels (K circuits)	1	$K(N_s - 1)$
CV-QFL (K parallel pyramidal)	$2N_c - 3$	$K \cdot \frac{N_c(N_c-1)}{2}$
CV-IQFT _c	$\log_2 N_s$	$\frac{N_s}{2} \log_2 N_s$
Total	$O(\max(N_c, N_s))$	$O(KN_c^2 + N_c^2 + N_s^2)$

For the representative choice $N_c = N_s = 32$ and $K = 4$, the above expressions yield a total of 3260 passive beam splitters and 32 two-mode squeezers, at an overall sequential depth of 105. By comparison, the corresponding qubit QFNO layer requires 8256 CZ gates at depth 396 (Table 1). Since beam splitters are low-loss passive optical elements [20], the fidelity argument that renders the qubit implementation impractical does not carry over to the CV setting, making the CV-QFNO a substantially more realistic photonic architecture. However, we note that the TMS-SVD encoding requires a classical singular value decomposition of the input matrix prior to each forward pass, introducing an additional preprocessing cost of order $O(\min(N_c, N_s)^2 \max(N_c, N_s))$, which is absent in both the qubit and classical FNO pipelines and should therefore be accounted for in any end-to-end latency comparison.

At the same time, it is important to acknowledge the practical limitations of the present construction. Although the CV-QFNO removes the severe entangling-gate overhead of the qubit formulation, a full hardware realisation of the architecture considered here remains far beyond current photonic quantum technology, especially because the model requires multiple parallel optical branches and precise control of Gaussian operations, loss channels, and readout. In addition, quantum machine learning training remains hybrid, with parameter optimisation performed classically, so that any potential advantage at the circuit level must be weighed against the substantial optimisation overhead of the full training loop. For these reasons, the CV-QFNO should presently be regarded primarily as a theoretical and architectural framework, useful for showing how a Continuous-Variable Quantum Fourier Neural Operator could be constructed in a manner faithful to the spectral structure of the classical FNO. For classical PDE problems on classical hardware, the standard FNO remains by far the most practical and effective choice, and the main value of the CV formulation lies not in immediate applicability, but in clarifying what a coherent photonic quantum analogue of the FNO would look like.

2.3.4. 2D Architecture: $K_x K_y$ Parallel Branches

In the two-dimensional setting, the latent representation is the tensor $A \in \mathbb{R}^{N_c \times N_x \times N_y}$ introduced in Equation (8), where N_c denotes the number of latent channels, while N_x and N_y denote the number of sampled grid points along the two spatial directions. Thus, $A_{l,:,:} \in \mathbb{R}^{N_x \times N_y}$ denotes the two-dimensional spatial field associated with the l -th latent channel. Unlike the 1D case, where the latent representation is a matrix with one channel axis and one spatial axis, the 2D tensor has a three-index structure (channel, x -spatial, and y -spatial). This structure does not admit, within the present Gaussian construction, a direct

bipartite encoding that preserves the desired separation between channel mixing and the two spatial Fourier directions.

For this reason, we allocate N_c independent optical circuits, one for each channel l , and we encode in the l -th circuit the spatial slice $A_{l,:} \in \mathbb{R}^{N_x \times N_y}$ via TMS-SVD into a bipartite Gaussian state. In this 2D construction, the two registers are denoted by r_x and r_y : r_x contains N_x bosonic modes associated with the grid points along the x -direction, while r_y contains N_y bosonic modes associated with the grid points along the y -direction. By Equation (22), the x -quadrature inter-register covariance block in the l -th circuit satisfies

$$\sigma_{x_{r_x} x_{r_y}}^{(l)} = A_{l,:}, \quad l = 1, \dots, N_c. \quad (35)$$

A two-dimensional CV-QFT is then applied independently within each circuit by acting simultaneously with QFT_x on register r_x and QFT_y on register r_y , as in Equation (26). Because the 2D DFT result $\hat{A}_{l,:} = F_{N_x} A_{l,:} F_{N_y}^\top \in \mathbb{C}^{N_x \times N_y}$ is complex, real and imaginary parts are distributed across two covariance blocks:

$$\sigma_{x_{r_x} x_{r_y}}^{(l)} \mapsto \text{Re}(\hat{A}_{l,:}), \quad \sigma_{x_{r_x} p_{r_y}}^{(l)} \mapsto \text{Im}(\hat{A}_{l,:}), \quad (36)$$

so that $\hat{A}_{l,k_1,k_2} = (\sigma_{x_{r_x} x_{r_y}}^{(l)} + i \sigma_{x_{r_x} p_{r_y}}^{(l)})_{k_1,k_2}$ (cf. Equation (27)). The trainable pyramidal interferometer $O(\theta_{k_1,k_2}) \in O(N_c)$ acting across the N_c circuits applies the same real orthogonal matrix simultaneously to both covariance blocks,

$$\text{Re}(\hat{v}_t(k_1, k_2)) \mapsto O(\theta_{k_1,k_2}) \text{Re}(\hat{v}_t(k_1, k_2)), \quad \text{Im}(\hat{v}_t(k_1, k_2)) \mapsto O(\theta_{k_1,k_2}) \text{Im}(\hat{v}_t(k_1, k_2)), \quad (37)$$

the two-dimensional analogue of Equation (28). For notational compactness, the remainder of this section writes all operations on the full complex quantities $\hat{A}_{l,:}$ and $\hat{v}_t(k_1, k_2)$, with the understanding that each step acts equivalently on both covariance blocks.

We now focus on a fixed retained mode pair (k_1, k_2) with $k_1 \in \{1, \dots, K_x\}$ and $k_2 \in \{1, \dots, K_y\}$; the full output is then recovered by summing the contributions of all the retained mode pairs, as we have previously done in the one-dimensional problem. To isolate the pair (k_1, k_2) , all r_x -modes $i \neq k_1$ and all r_y -modes $j \neq k_2$ are routed through loss channels with transmissivity $\eta = 0$, see Figure 3. This erases all correlations except those associated with the selected Fourier entry, so that in the l -th circuit only the scalar coefficient $\hat{A}_{l,k_1,k_2}$ survives. Collecting these coefficients across the N_c circuits defines the channel vector

$$\hat{v}_t(k_1, k_2) = (\hat{A}_{1,k_1,k_2}, \dots, \hat{A}_{N_c,k_1,k_2})^\top \in \mathbb{C}^{N_c}. \quad (38)$$

A trainable pyramidal interferometer $O(\theta_{k_1,k_2}) \in O(N_c)$ then acts across the N_c circuits on the first register (r_x), connecting the k_1 -th mode wire of each circuit into a shared beam-splitter network. Because r_x carries the row index of each covariance block, this action corresponds to a left-multiplication and implements the channel-mixing operation of Equation (10):

$$\hat{v}_t(k_1, k_2) \mapsto O(\theta_{k_1,k_2}) \hat{v}_t(k_1, k_2). \quad (39)$$

Equivalently, the branch associated with the retained mode pair (k_1, k_2) produces the filtered Fourier contribution

$$\hat{A}'_{l,i,j}(k_1, k_2) = \begin{cases} [O(\theta_{k_1,k_2}) \hat{v}_t(k_1, k_2)]_l, & (i, j) = (k_1, k_2), \\ 0, & \text{otherwise,} \end{cases} \quad (40)$$

that is, a tensor whose only nonzero spatial-frequency entry is the selected mode pair (k_1, k_2) .




Figure 3. CV-QFNO spectral layer: mode pair (k_1, k_2) (2D case). N_c optical circuits run in parallel, one per channel l . Each circuit encodes the 2D slice $A_{l,:}$ via TMS-SVD and applies a 2D CV-QFT on both registers. The pyramidal interferometer $O(\theta_{k_1, k_2})$ connects the k_1 -th upper-register wire across all N_c circuits, mixing the channel dimension for all column modes. Loss channels $\eta = 0$ on upper-register modes $i \neq k_1$ and lower-register modes $j \neq k_2$ then isolate the target mode pair, realising Equation (10). Homodyne measurement yields the output contribution of mode pair (k_1, k_2) .

The two-dimensional inverse Fourier transform is then applied independently within each branch, yielding the corresponding contribution in the physical domain. After homodyne measurement, the full latent output is recovered by summing the contributions of all retained mode pairs:

$$A' = \sum_{k_1=1}^{K_x} \sum_{k_2=1}^{K_y} \mathcal{F}_{2D}^{-1}(\hat{A}'^{(k_1, k_2)}). \quad (41)$$

By linearity of $\mathcal{F}_{2D}^{-1}$, this can be written as

$$A' = \mathcal{F}_{2D}^{-1} \left(\sum_{k_1=1}^{K_x} \sum_{k_2=1}^{K_y} \hat{A}'^{(k_1, k_2)} \right), \quad (42)$$

which exactly reconstructs the filtered spectral tensor of Equation (12). A schematic representation of a single mode-pair branch is shown in Figure 3.

2.3.5. Optical Resource Count (2D)

The 2D layer comprises $N_c \times K_x K_y$ parallel circuit instances, one per channel per retained mode pair, plus $K_x K_y$ cross-circuit interferometers for channel mixing. Table 3 summarises the resource budget.

Compared with the 1D case (Table 2), the sequential depth grows only modestly: the 2D QFT replaces a single-register transform with a parallel two-register transform (depth $\max(\log_2 N_x, \log_2 N_y)$ instead of $\log_2 N_s$), leaving the overall depth dominated by the encoding and channel-mixing stages. However, the beam-splitter count scales as $O(N_c K_x K_y (N_x^2 + N_y^2))$, reflecting the N_c parallel circuit copies required by the 2D architecture; this additional cost is a direct consequence of the three-index structure of the 2D latent tensor and is shared by any photonic implementation of the 2D operator.

Table 3. Optical resources per phase of a single CV-QFNO layer in the 2D case ($K_x K_y$ retained mode pairs, N_c channels). Depths of phases acting on disjoint registers are computed in parallel; the sequential circuit depth is dominated by the encoding and the channel mixing interferometer. Interferometer counts use the Clements rectangular decomposition [21].

Phase	Depth	Beam Splitters (Total)
TMS-SVD encoding (N_c circuits per branch)	$1 + \max(N_x, N_y)$	$N_c K_x K_y \left[\frac{N_x(N_x-1)}{2} + \frac{N_y(N_y-1)}{2} \right]$
2D CV-QFT (both registers in parallel)	$\max(\log_2 N_x, \log_2 N_y)$	$N_c K_x K_y \left[\frac{N_x}{2} \log_2 N_x + \frac{N_y}{2} \log_2 N_y \right]$
Loss channels (r_x and r_y)	1	$N_c K_x K_y (N_x + N_y - 2)$
Channel mixing $O(\theta_{k_1, k_2})$ ($K_x K_y$ interferometers)	$2N_c - 3$	$K_x K_y \cdot \frac{N_c(N_c-1)}{2}$
2D CV-IQFT (both registers in parallel)	$\max(\log_2 N_x, \log_2 N_y)$	$N_c K_x K_y \left[\frac{N_x}{2} \log_2 N_x + \frac{N_y}{2} \log_2 N_y \right]$
Total	$O(\max(N_x, N_y, N_c))$	$O\left(N_c K_x K_y (N_x^2 + N_y^2) + K_x K_y N_c^2\right)$

2.4. Simulation Setup and Orthogonal Parameterisation

The FNO, QFNO, and the proposed CV-QFNO are algorithmically distinct architectures rooted in different computational substrates and physical principles: classical tensor operations, qubit circuits restricted to a structured Hilbert subspace, and Gaussian photonic modes, respectively. Despite these differences, all three models admit exact classical linear-algebraic representations. The quantum gates and optical transformations appearing in the QFNO and CV-QFNO can therefore be replaced by their corresponding matrix actions, allowing the full architectures to be simulated and trained in PyTorch 2.7.0 with CUDA 12.6 without introducing an additional approximation of the underlying circuit model. This is the standard validation strategy in quantum-inspired machine learning prior to the availability of hardware capable of running the complete architecture at the required scale [4,5,15], and it was the approach adopted in this work.

In the numerical experiments, the three models shared the same overall neural-operator structure, consisting of a lifting map, four spectral blocks, and a final projection. They differed only in the spectral convolution layer, where the quantum-inspired structure enters. In the classical FNO, each retained Fourier mode is mixed across the channel dimension by an unconstrained complex matrix $R_\phi^{(k)} \in \mathbb{C}^{N_c \times N_c}$. In the QFNO and CV-QFNO, this operation is replaced by an orthogonal matrix $W^{(k)} \in O(N_c)$, reflecting the unitary or passive-interferometric origin of the corresponding quantum transformation. Earlier quantum-inspired networks typically construct such orthogonal matrices as ordered products of $M = N_c(N_c - 1)/2$ Givens rotations,

$$W = G_{i_1 j_1}(\theta_1) G_{i_2 j_2}(\theta_2) \cdots G_{i_M j_M}(\theta_M),$$

thereby mirroring the gate-by-gate structure of RBS circuits or optical beam-splitter networks [4,15,16]. Although physically natural, this representation is less convenient for classical training, since the matrix W must be reconstructed at every forward pass through a sequential product of $O(N_c^2)$ rotation factors.

In this work, both the QFNO and the CV-QFNO instead used a matrix-exponential parameterisation of the orthogonal spectral-mixing matrices. For each retained mode, we introduced a skew-symmetric generator $A \in \mathbb{R}^{N_c \times N_c}$, with $A^\top = -A$, whose independent upper-triangular entries were the trainable parameters:

$$A_{ij} = \theta_{ij}, \quad A_{ji} = -\theta_{ij} \quad \text{for } i < j, \quad A_{ii} = 0, \quad W^{(k)} = \exp(A). \quad (43)$$

This guaranteed orthogonality by construction, since $\exp(A)^\top = \exp(A^\top) = \exp(-A) = \exp(A)^{-1}$. The parameterisation used exactly $M = N_c(N_c - 1)/2$ free parameters, the same number as the Givens decomposition, but avoided the explicit sequential product of rotations during training. In PyTorch 2.7.0, the exponential is evaluated by `torch.linalg.matrix_exp`, which uses scaling-and-squaring methods with Padé approximants [22], at computational cost $O(N_c^3)$. This cost belongs only to the classical simulation: on a physical CV photonic device, a trained orthogonal transformation would be compiled into beam-splitter and phase-shifter settings, for instance through the Clements decomposition [21], and executed by passive propagation through the interferometric network rather than by explicit matrix exponentiation.

The remaining distinction between the QFNO and the CV-QFNO concerns the treatment of Fourier modes beyond the retained cutoff. In the classical FNO, non-retained modes are set to zero before applying the inverse Fourier transform. The CV-QFNO reproduces this operation exactly in the simulated model by initialising the output spectrum to zero and writing only the retained, orthogonally mixed modes; this is the classical linear-algebraic counterpart of routing the non-selected optical modes through loss channels with transmissivity $\eta = 0$. The QFNO, by contrast, follows the behaviour of the parallel qubit construction: the full input spectrum is copied, and only the retained modes are overwritten by the mixed values. The modes beyond the cutoff therefore remain unchanged, because they are not explicitly discarded by the qubit circuit. This difference is the source of the systematic spectral contamination discussed in Section 2.2.

The simulations reported here should therefore be interpreted as a classical validation of the proposed architectures, not as a demonstration of quantum computational advantage. Since every quantum or photonic operation is replaced by its matrix equivalent and executed on a GPU, the resulting pipeline is a classical algorithm: it does not realise quantum speedup, exponential state-space compression, or genuine quantum parallelism. Its purpose is architectural. The experiments show that the CV formulation provides a spectrally faithful and hardware-compatible quantum analogue of the FNO layer, with a native mechanism for exact mode suppression and a direct mapping of Fourier and mixing operations onto Gaussian photonic primitives. Any practical computational advantage would require deployment on native photonic hardware, where these transformations are carried out by the physical propagation of light rather than by explicit matrix multiplication, or the use of non-Gaussian resources leading to regimes not efficiently captured by classical Gaussian emulation. These perspectives are discussed further in the conclusion.

All simulations were run on a workstation equipped with an Intel Core i7-14700F processor, 32 GB DDR5 RAM, and an NVIDIA GeForce RTX 4070 GPU with 12 GB of GDDR6X VRAM. PyTorch 2.7.0 training was accelerated on the GPU using CUDA.

3. Results

We evaluated the proposed CV-QFNO against two reference architectures: the classical Fourier Neural Operator (FNO) [2] and the qubit-based Quantum Fourier Neural Operator (QFNO) [5]. The comparison was performed on four standard PDE benchmarks: the one-dimensional Burgers' equation, the one-dimensional heat equation, the two-dimensional Navier–Stokes equation in vorticity form, and the two-dimensional Darcy flow problem. Within each benchmark, all the models are trained using identical data splits, training schedules, and optimisation settings.

The main result of the paper is the formulation of a CV photonic analogue of the FNO spectral layer that faithfully reproduces its truncated Fourier structure. As discussed above, this formulation is particularly natural in CV quantum computing: the Fourier transform, mode filtering, and trainable channel mixing can be implemented directly in

terms of Gaussian photonic operations, such as beam splitters, phase shifters, loss channels, and interferometric networks. In this sense, the CV-QFNO provides a hardware-compatible route to a photonic implementation of the FNO spectral mechanism, while avoiding the qubit-gate compilation overhead of discrete-variable QFNO constructions.

Within this perspective, the numerical benchmarks show that CV-QFNO retains competitive predictive accuracy with respect to both FNO and QFNO, while using substantially fewer trainable parameters than the classical model. This parameter reduction follows from the orthogonal parameterisation of the spectral-mixing matrices induced by the pyramidal interferometers. Since CV-QFNO and QFNO use the same orthogonal channel-mixing structure, their trainable-parameter counts were identical in all the experiments.

3.1. One-Dimensional Experiments

We first considered two one-dimensional operator-learning benchmarks: the viscous Burgers' equation,

$$\partial_t u(x, t) + u(x, t) \partial_x u(x, t) = \nu \partial_{xx} u(x, t), \quad x \in [0, 1], \quad t \in (0, T], \quad (44)$$

and the heat equation,

$$\partial_t u(x, t) = \alpha \partial_{xx} u(x, t), \quad x \in [0, 1], \quad t \in (0, T]. \quad (45)$$

Both problems were posed on the one-dimensional spatial domain $D = [0, 1]$, so $d = 1$, and governed the time evolution of a scalar field $u : D \times [0, T] \rightarrow \mathbb{R}$, where x was the spatial coordinate and t was time. In Equation (44), $u(x, t)$ denotes the scalar velocity field, $\nu > 0$ is the viscosity (set to $\nu = 0.1$ in our experiments), while $\partial_t u$, $\partial_x u$, and $\partial_{xx} u$ denote differentiation with respect to time, the first spatial derivative, and the second spatial derivative, respectively. In Equation (45), $u(x, t)$ represents the temperature field and $\alpha > 0$ is the diffusion coefficient (set to $\alpha = 0.01$ and final time $T = 1$).

In both cases, the learning task was formulated as the approximation of a solution operator mapping an initial condition to the solution at a later time. Following the standard FNO setup, the model input was the initial condition concatenated with the spatial coordinate as a positional feature,

$$a(x) = (u(x, 0), x) \in \mathbb{R}^{d_a}, \quad d_a = 2, \quad (46)$$

and the target output was the scalar solution at the final time,

$$u(x, T) \in \mathbb{R}^{d_u}, \quad d_u = 1. \quad (47)$$

Thus, each training sample corresponded to an input–output pair

$$a(\cdot) = (u(\cdot, 0), \cdot) \mapsto u(\cdot, T),$$

and the learned operator approximated

$$\mathcal{G}^\dagger : a(\cdot) \mapsto u(\cdot, T).$$

In the implementation, the input was evaluated on a uniform grid of N_s spatial points and lifted to a latent representation with N_c channels. Therefore, at each Fourier layer, the latent state had the matrix structure $A \in \mathbb{R}^{N_c \times N_s}$ that was introduced in Equation (2).

The Fourier transform was applied along the spatial direction, i.e., across the N_s grid points, as in Equation (23), while the spectral mixing acted along the channel dimension on

each retained Fourier mode. This was precisely the one-dimensional structure used by the classical FNO, the qubit QFNO, and the CV-QFNO.

The two PDEs probed complementary regimes of operator learning. Burgers' equation combines nonlinear advection with viscous diffusion and may develop sharp solution features, making it a nonlinear benchmark for spectral operator learning. The heat equation is linear and purely diffusive; in Fourier space, its coefficients decay exponentially in time, making it a controlled benchmark for evaluating whether the learned spectral layer captures dissipative dynamics.

Both one-dimensional experiments used the same training protocol. We used $N_{\text{train}} = 200$ input–output pairs and $N_{\text{test}} = 20$ test pairs, discretised on a uniform grid of 512 points. All the models were trained for 300 epochs with the Adam optimiser [23], using an initial learning rate of 10^{-3} , reduced by a factor of two every 100 epochs. For both benchmarks, we analysed the models through three complementary diagnostics: validation loss over training, zero-shot resolution generalisation, and spectral accuracy. The first diagnostic compared the optimisation behaviour of the architectures; the second measured whether a model trained at one resolution could be evaluated on unseen grids; and the third assessed whether the learned operator reproduced the relevant Fourier content of the target dynamics.

Figure 4 reports the loss evolution and resolution generalisation for both one-dimensional benchmarks. To ensure that the quantitative comparisons between architectures reflected genuine structural differences, all the one-dimensional experiments were repeated independently with four different random seeds (0, 11, 22, 33).

For Burgers' equation, the classical FNO contained 419,009 trainable parameters, whereas QFNO and CV-QFNO each used 122,561, corresponding to a $3.4\times$ reduction. For the heat equation, the FNO used 28,145 parameters, while the QFNO and CV-QFNO used 9329, giving a $3.0\times$ reduction.

This parameter reduction, however, should not be interpreted as a direct wall-clock speedup in the present classical simulations. In the QFNO and CV-QFNO, the quantum or photonic operations are emulated through classical linear algebra, and a substantial part of the cost comes from constructing the orthogonal spectral mixing matrices, as described in Section 2.4. Therefore, training-time comparisons would mix the effect of fewer trainable parameters with the overhead of simulating the quantum-inspired layer. Despite this smaller parameter count, the two quantum-inspired models retained accuracy comparable to the classical FNO, as confirmed by the converged test errors in Table 4a,b. The significance of this result is primarily architectural rather than computational: it shows that the FNO spectral layer can be replaced by a more constrained orthogonal parameterisation without a substantial loss of predictive performance. Such a constraint may improve optimisation stability and become advantageous in future native quantum or photonic implementations.

The loss curves highlight different optimisation behaviours across the two benchmarks. For Burgers' equation, the FNO converged to a relative ℓ_2 error of $(2.52 \pm 0.19) \times 10^{-2}$ (Table 4a), while the QFNO and CV-QFNO converged to $(2.97 \pm 0.13) \times 10^{-2}$ and $(2.93 \pm 0.16) \times 10^{-2}$, respectively. Despite reaching a slightly lower final test error, the FNO exhibited considerably more unstable training dynamics. Quantifying this over the four independent seeds, the average cross-seed standard deviation of the log test loss, computed epoch-by-epoch across the full training trajectory, was 0.22 for the FNO compared with 0.13 for the QFNO and 0.15 for the CV-QFNO, indicating that the FNO was approximately 70% more sensitive to the random initialisation than the orthogonal models. This proves consistently that the orthogonal spectral-mixing layers produce a substantially smoother and more reproducible optimisation landscape with Burgers' equation.

Table 4. Convergence summary. Relative ℓ_2 test error at convergence for all four benchmarks. Each entry was computed from four independent training runs with different random seeds: for each seed the test loss was averaged over the last 20 epochs to obtain a per-seed converged value, and the reported mean $\pm$ standard deviation was then taken across the four seeds.

Model	Relative ℓ_2 (Mean $\pm$ Std)
(a) Burgers' equation (1D)	
FNO	$(2.52 \pm 0.19) \times 10^{-2}$
QFNO	$(2.97 \pm 0.13) \times 10^{-2}$
CV-QFNO	$(2.93 \pm 0.16) \times 10^{-2}$
(b) Heat equation (1D)	
FNO	$(3.27 \pm 0.10) \times 10^{-2}$
QFNO	$(2.85 \pm 0.17) \times 10^{-2}$
CV-QFNO	$(3.39 \pm 0.11) \times 10^{-2}$
(c) Navier–Stokes (2D)	
FNO	$(4.32 \pm 0.23) \times 10^{-2}$
QFNO	$(4.37 \pm 0.47) \times 10^{-2}$
CV-QFNO	$(4.31 \pm 0.51) \times 10^{-2}$
(d) Darcy flow (2D)	
FNO	$(5.10 \pm 0.59) \times 10^{-2}$
QFNO	$(6.73 \pm 0.53) \times 10^{-2}$
CV-QFNO	$(6.55 \pm 0.68) \times 10^{-2}$

For the heat equation, a qualitatively different behaviour emerged. As shown in Figure 4b, the FNO loss curve rose in the final training epochs, a clear sign of overfitting. This is confirmed by the convergence statistics: the FNO reached a test error of $(3.27 \pm 0.10) \times 10^{-2}$ (Table 4b), higher than its own training loss on all four seeds. By contrast, the QFNO converged to the lowest test error of $(2.85 \pm 0.17) \times 10^{-2}$, and the CV-QFNO to $(3.39 \pm 0.11) \times 10^{-2}$, with both orthogonal models showing stable descent throughout training.

Taken together, both one-dimensional benchmarks indicate that constraining the spectral-mixing matrices to the orthogonal group acted as an implicit regulariser: it stabilised the optimisation trajectory across the seeds with Burgers' equation and prevented overfitting with the heat equation, in both cases without sacrificing predictive accuracy relative to the unconstrained classical FNO.

The resolution-generalisation results, shown in the right panels of Figure 4, confirm that all three architectures preserved the resolution-invariant character expected from neural operators. The models were trained at the base resolution of 512 grid points and evaluated zero-shot at the other resolutions shown, without retraining. Across both the Burgers' and the heat benchmarks, the QFNO and the CV-QFNO remained competitive with the FNO over the full tested range, with no systematic degradation at unseen spatial resolutions.

Figure 5 provides a direct spectral diagnostic of the learned operators. In both benchmarks, the CV-QFNO reproduced the reference spectrum with essentially the same fidelity as the FNO and confined the spectral activation to the retained modes.

As explained in Section 2.4, the experiments were performed as classical PyTorch 2.7.0 simulations of the corresponding CV operations. Therefore, the mode-isolation step was implemented numerically by setting the non-selected Fourier modes to zero, which is the classical linear-algebra equivalent of routing those modes through CV loss channels with transmissivity $\eta = 0$. The resulting spectra show that this simulated loss-channel

mechanism preserved the truncated spectral inductive bias of the classical FNO without introducing spurious spectral artefacts.




Figure 4. Training dynamics and resolution generalisation: 1D benchmarks. Left: test loss evolution over training epochs. Each solid line shows the geometric mean of the relative ℓ_2 loss across four independent runs (seeds 0, 11, 22, 33); the shaded band spans ± 1 geometric standard deviation across the seeds, reflecting run-to-run variability due to random initialisation. Right: relative ℓ_2 test loss as a function of spatial resolution. The models were trained at the base resolution of 512 grid points and evaluated zero-shot at all resolutions shown, without retraining.




Figure 5. Spectral analysis: 1D benchmarks. Power spectra of the predicted solutions and ground truth for FNO, QFNO, and CV-QFNO. The classical FNO and the proposed CV-QFNO exhibited consistent spectral behaviour: both reproduced the relevant low-frequency content and enforced a sharp truncation beyond the retained Fourier modes. In contrast, the QFNO displayed nonzero spectral contributions beyond the cutoff, visible as horizontal channel-dependent bands, due to the non-suppressed Fourier modes inherited from the parallel qubit construction.

3.2. Two-Dimensional Experiments

We next considered two two-dimensional operator-learning benchmarks: the incompressible Navier–Stokes equation in vorticity form,

$$\partial_t \omega(x, y, t) + (u(x, y, t) \cdot \nabla) \omega(x, y, t) = \nu \Delta \omega(x, y, t), \quad (x, y) \in [0, 1]^2, \quad t \in (0, T], \quad (48)$$

and the Darcy flow equation,

$$-\nabla \cdot (a(x, y) \nabla u(x, y)) = f(x, y), \quad (x, y) \in [0, 1]^2. \quad (49)$$

Both problems are posed on the two-dimensional spatial domain $D = [0, 1]^2$, so $d = 2$, but they represent different classes of PDE operators. In Equation (48), ω denotes the scalar vorticity, $u = (u_x, u_y)$ is the divergence-free velocity field, $\Delta \omega$ is the spatial Laplacian of the vorticity, and $\nu > 0$ is the kinematic viscosity. The velocity is recovered from the vorticity through the Biot–Savart relation. In Equation (49), $a(x, y)$ is the spatially varying permeability, $u(x, y)$ is the pressure field, ∇u is the pressure gradient, and $f(x, y)$ is a prescribed source term.

We used $N = 30$ trajectories at $\nu = 10^{-3}$ for Navier–Stokes, generated with a pseudo-spectral solver on a 256×256 grid over 40 time steps, and $N = 500$ input–output pairs for Darcy flow, solved on a 421×421 finite-difference grid. The permeability was a thresholded Gaussian random field with spectral density $(4\pi^2 |k|^2 + \tau^2)^{-\alpha}$, $\alpha = 2$, $\tau = 3$, taking values $a \in \{3, 12\}$.

For the Navier–Stokes benchmark, the learning task was formulated as a one-step time-evolution problem. The input consisted of the spatial coordinates concatenated with T_{in} consecutive vorticity snapshots,

$$a(x, y) = (x, y, \omega(x, y, t_1), \dots, \omega(x, y, t_{T_{\text{in}}})) \in \mathbb{R}^{d_a}, \quad d_a = T_{\text{in}} + 2, \quad (50)$$

and the target output was the vorticity at the next time step,

$$u(x, y) = \omega(x, y, t_{T_{\text{in}}+1}) \in \mathbb{R}^{d_u}, \quad d_u = 1. \quad (51)$$

Thus, the learned operator approximated the local-in-time evolution map

$$\mathcal{G}^\dagger : (x, y, \omega(\cdot, t_1), \dots, \omega(\cdot, t_{T_{\text{in}}})) \mapsto \omega(\cdot, t_{T_{\text{in}}+1}).$$

For the Darcy flow, the learning task was instead static. The input was the permeability field concatenated with the spatial coordinates,

$$a(x, y) = (x, y, a(x, y)) \in \mathbb{R}^{d_a}, \quad d_a = 3, \quad (52)$$

and the target output was the scalar pressure solution,

$$u(x, y) \in \mathbb{R}^{d_u}, \quad d_u = 1. \quad (53)$$

The corresponding solution operator was

$$\mathcal{G}^\dagger : a(\cdot, \cdot) \mapsto u(\cdot, \cdot),$$

where $a(x, y)$ denotes the permeability coefficient in Equation (49). This elliptic operator is nonlocal: local changes in the permeability can affect the pressure solution across the whole domain.

In the implementation, both benchmarks were evaluated on a 64×64 uniform grid. After the lifting map, the latent field was represented by the tensor $A \in \mathbb{R}^{N_c \times N_x \times N_y}$ with $N_x = N_y = 64$, as introduced in Equation (8), where N_c was the latent channel dimension. The Fourier transform was applied along the two spatial directions (see Equation (25)) while the spectral mixing acted along the channel dimension for each retained Fourier mode pair (k_1, k_2) . This is the two-dimensional spectral structure used by the classical FNO and reproduced, with constrained orthogonal mixing, by the QFNO and CV-QFNO architectures.

Both datasets were subsampled for training: factor 4 to 64×64 for Navier–Stokes ($N_{\text{train}} = 25$, $N_{\text{test}} = 5$, 100 epochs) and factor 5 to 85×85 for Darcy flow ($N_{\text{train}} = 400$, $N_{\text{test}} = 50$, 150 epochs).

The different dataset sizes reflect the different structure of the two learning problems. In the time-stepping formulation of Navier–Stokes, consecutive vorticity fields are strongly correlated, the dynamics are smooth over short time intervals, and each trajectory contributes multiple input–output pairs to the effective training set. Darcy flow, by contrast, is a static elliptic problem: each input is an independent realisation of a thresholded Gaussian random permeability field $a(x, y) \in \{3, 12\}$ paired with the corresponding pressure solution $u(x, y)$. Since there is no temporal structure to exploit, and since the elliptic solution operator $a \mapsto u$ is nonlocal through the associated Green’s operator, a larger ensemble is required to represent the input distribution adequately.

Both benchmarks were trained with the Adam optimiser, using an initial learning rate of 2×10^{-3} , reduced by a factor of two every 20 epochs. As in the one-dimensional case, all the two-dimensional experiments were repeated across four independent random seeds (0, 11, 22, 33); the loss curves in Figure 6 show the geometric mean with ± 1 geometric standard deviation band, and the converged test errors are reported in Table 4c,d. On Navier–Stokes, all three architectures converged to statistically equivalent accuracy: the FNO reached $(4.32 \pm 0.23) \times 10^{-2}$, the QFNO $(4.37 \pm 0.47) \times 10^{-2}$, and the CV-QFNO $(4.31 \pm 0.51) \times 10^{-2}$, with differences well within the cross-seed variability. The quantum-inspired models achieved this with $4.1 \times$ fewer parameters (223,797 versus 926,517 for the FNO). On the Darcy flow, the FNO converged to $(5.10 \pm 0.59) \times 10^{-2}$, while the QFNO and CV-QFNO reached $(6.73 \pm 0.53) \times 10^{-2}$ and $(6.55 \pm 0.68) \times 10^{-2}$ respectively, again with a $4.1 \times$ parameter reduction (910,977 versus 3,698,817). In this static elliptic setting, the orthogonal constraint did not match the accuracy of the unconstrained FNO; the difference between the CV-QFNO and the QFNO was small and lay within the cross-seed variability. Nevertheless, the orthogonal parameterisation produced a more stable optimisation trajectory: the average cross-seed log-standard deviation of the test loss, computed epoch-by-epoch over the full training run, was 0.154 for the FNO, compared with 0.109 for the QFNO and 0.128 for the CV-QFNO. This pattern mirrors the one observed for Burgers’ equation in the one-dimensional case, suggesting that the orthogonal constraint consistently reduces sensitivity to random initialisation, even when it does not yield a lower final test error.

3.3. Summary of Benchmark Results

Table 4 collects the converged test errors for all the models across the four benchmarks. For each model and benchmark, four independent runs were performed with different random seeds (0, 11, 22, 33); within each run the test loss was averaged over the last 20 epochs to obtain a stable plateau estimate, and the reported mean and standard deviation were computed across these four per-seed values.




Figure 6. Training dynamics and resolution generalisation: 2D benchmarks. Left: test loss evolution over training epochs. Each solid line shows the geometric mean of the relative ℓ_2 loss across four independent runs (seeds 0, 11, 22, 33); the shaded band spans ± 1 geometric standard deviation across seeds. Right: relative ℓ_2 test loss as a function of spatial resolution, evaluated zero-shot without retraining.

Across all the experiments, the QFNO and CV-QFNO used $4.1\times$ fewer trainable parameters than the FNO (1D Burgers': 122,561 vs. 419,009; 1D heat: 9329 vs. 28,145; 2D Navier–Stokes: 223,797 vs. 926,517; 2D Darcy: 910,977 vs. 3,698,817), a direct consequence of restricting the spectral-mixing matrices to the orthogonal group.

For Burgers' equation, all three models reach comparable accuracy, with the FNO attaining a marginally lower mean error. For the heat equation, the orthogonal constraint acted as an implicit regulariser: the FNO exhibited overfitting (its test loss exceeded its training loss on all four seeds), whereas the QFNO achieved the best test error of $(2.85 \pm 0.17) \times 10^{-2}$ and the CV-QFNO remained close to the FNO at $(3.39 \pm 0.11) \times 10^{-2}$. For Navier–Stokes, the three architectures converged to statistically equivalent errors, (4.32 ± 0.23) , (4.37 ± 0.47) , and $(4.31 \pm 0.51) \times 10^{-2}$ for the FNO, QFNO, and CV-QFNO respectively, so the orthogonal parameterisation matched classical accuracy with a $4.1\times$ parameter reduction. For the Darcy flow, the FNO retained a clear accuracy advantage $((5.10 \pm 0.59) \times 10^{-2})$ over the constrained models $((6.73 \pm 0.53)$ and $(6.55 \pm 0.68) \times 10^{-2}$ for the QFNO and CV-QFNO); the difference between the two quantum-inspired models was small relative to the cross-seed variability.

Taken together, the results confirm that the CV-QFNO constitutes a viable photonic analogue of the FNO spectral layer: it matched or closely approached classical accuracy on three out of the four benchmarks, systematically using $4.1\times$ fewer parameters, and it provided a hardware-compatible route to a photonic implementation of the FNO spectral mechanism.

This behaviour can be understood from the spectral properties of the benchmark problems. As discussed in Section 2.2 and as illustrated in Figure 5, the QFNO does not exactly implement the FNO spectral truncation: Fourier modes outside the retained band

are not suppressed but are carried unchanged through the circuit, leading to residual high-frequency activations. In the present benchmarks, however, the solutions are relatively smooth so their Fourier energy decays rapidly with the mode index. As a result, most of the relevant dynamics are already contained in the retained modes, while the unsuppressed modes carry only a small amount of energy. Their effect on the final prediction is therefore limited, and the learnable orthogonal weights can partly compensate for their presence during training.

The $\eta = 0$ loss-channel mechanism of the CV-QFNO removes this mismatch by enforcing the desired spectral truncation by construction, rather than relying on training to reduce the influence of the extra modes. Its advantage is therefore expected to become more visible in regimes where important dynamics extend to higher Fourier modes, such as low-viscosity Burgers' equation, high-Reynolds-number Navier–Stokes flows, or deeper architectures where small spectral errors may accumulate across stacked Fourier blocks. The benchmarks considered here mostly lie in a smooth, spectrally compact regime, which explains why the difference between the CV-QFNO and QFNO remained moderate across all four benchmarks.

4. Discussion

The results presented in this work should be interpreted primarily as an architectural and algorithmic validation of the proposed CV-QFNO framework. Although the construction was motivated by CV quantum computing and by its natural photonic implementation, all the numerical experiments were performed as classical simulations. Every optical or quantum operation was replaced by its matrix-algebraic equivalent and executed on a GPU. The resulting pipeline is therefore a classical algorithm, and no quantum or photonic hardware is involved in the reported results. Consequently, the experiments should not be interpreted as a demonstration of quantum computational advantage.

This distinction is particularly important in the present Gaussian setting. Gaussian CV circuits admit efficient classical descriptions in terms of covariance matrices, and the simulations performed in this work exploited exactly this linear-algebraic structure. The value of the proposed model therefore does not lie in outperforming classical neural operators on conventional hardware but in showing that the spectral mechanism of the FNO admits a coherent CV photonic formulation. In particular, the CV-QFNO identifies a direct optical analogue of the three operations that define the FNO spectral layer: Fourier transformation, mode selection, and channel mixing. Within this formulation, the $\eta = 0$ loss-channel mechanism provides a principled photonic counterpart of the mode-zeroing step of the classical FNO, resolving the spectral mismatch of the qubit-based QFNO.

The numerical results support this interpretation. Across the one- and two-dimensional benchmarks, the CV-QFNO achieved predictive accuracy comparable to the classical FNO and the qubit-based QFNO while using a substantially smaller number of trainable parameters. It should be emphasised, however, that the observed accuracy differences are entirely attributable to two classical algorithmic choices: the orthogonal constraint on the spectral-mixing matrices and the exact spectral truncation of non-retained Fourier modes. The role of the CV photonic formulation is not to introduce a new algorithmic mechanism but to provide a physically grounded motivation for these constraints, showing that they arise naturally from Gaussian photonic operations. Consequently, the parameter reduction should not be interpreted as a wall-clock speedup: in the present simulations, quantum and photonic operations were emulated through classical linear algebra, and training-time comparisons would mix the effect of a smaller parameter count with the overhead of constructing the orthogonal mixing matrices. The practical relevance of this blueprint lies in regimes that classical simulation cannot access: when the problem

is natively continuous-variable or quantum in nature, the CV-QFNO could operate directly on physical quantum states, with entanglement between registers and non-Gaussian resources enabling spectral interactions outside the efficiently simulable Gaussian sector, potentially opening pathways toward quantum operator learning with no efficient classical counterpart.

The spectral diagnostics further clarify the structural difference between the architectures. The classical FNO and the proposed CV-QFNO both enforce a hard cutoff beyond the retained Fourier modes. By contrast, the parallel qubit-QFNO construction leaves non-selected modes unchanged, which produces residual spectral activations beyond the cutoff. This behaviour is visible in the power spectra and directly reflects the spectral mismatch analysed in Section 2.2. The CV-QFNO removes this mismatch by suppressing non-retained modes through zero-transmissivity loss channels before the inverse Fourier transform.

A second limitation concerns hardware realisability. Although the CV-QFNO removes the entangling-gate compilation overhead of the qubit formulation, a full photonic implementation of the architecture remains beyond current integrated photonic capabilities. The required number of optical modes, parallel branches, Gaussian transformations, loss channels, and readout stages would pose substantial experimental challenges. Moreover, training remains a hybrid process in which the optimisation of the trainable parameters is performed classically. Any advantage at the circuit level would therefore need to be assessed together with the cost of data preparation, readout, and classical optimisation. A full hardware analysis, including squeezing requirements, optical loss budgets, interferometer depth, and measurement noise, is deferred to future work targeting natively quantum or continuous-variable problems, where such a characterisation becomes most meaningful.

For the classical PDE benchmarks considered here, standard neural operators remain the most practical approach. They are accurate, computationally efficient, and directly deployable on conventional hardware. The parameter reduction observed in the CV-QFNO is structurally interesting and reflects the orthogonal parameterisation of the spectral layer, but it does not by itself provide a sufficient practical reason to replace classical FNOs with photonic hardware for these tasks. The main significance of the present work is therefore conceptual: it establishes a faithful CV photonic blueprint for the FNO spectral layer and clarifies how a quantum-photonic operator-learning architecture can reproduce the truncated Fourier inductive bias of the classical model.

5. Conclusions

This work introduces the Continuous-Variable Quantum Fourier Neural Operator (CV-QFNO), a Gaussian photonic analogue of the spectral layer of the classical Fourier Neural Operator. Building on bipartite Gaussian encoding and the optical quantum Fourier transform, the CV-QFNO represents latent fields in the covariance structure of Gaussian states and maps Fourier-domain processing, mode selection, and channel mixing onto native continuous-variable optical operations.

The main architectural contribution is a CV spectral mixing layer that recovers the truncated spectral structure of the classical FNO. In the one-dimensional setting, each retained Fourier mode is processed in an independent optical branch: non-selected modes are removed through loss channels with transmissivity $\eta = 0$, while a pyramidal interferometer applies a learnable orthogonal transformation across the channel dimension. The construction is then extended to two spatial dimensions through a cross-circuit interferometric coupling, providing a photonic generalisation of the same mode-wise principle.

This formulation addresses two limitations of qubit-based QFNOs. First, it avoids the compilation overhead associated with RBS gates, since the corresponding transformations are native linear-optical operations in the CV setting. Second, it removes the spectral

mismatch of the parallel qubit-QFNO architecture, where non-selected Fourier modes remain unchanged, by suppressing them directly through $\eta = 0$ loss channels. As a result, the CV-QFNO restores the hard spectral truncation that defines the classical FNO layer.

We also introduce a matrix-exponential parameterisation of the orthogonal spectral-mixing matrices, which preserves the number of trainable parameters of the Givens construction while avoiding an explicit sequential product of rotations during classical training. Our numerical experiments on Burgers', heat, Navier–Stokes, and Darcy-flow benchmarks showed that the CV-QFNO achieves competitive predictive accuracy with the FNO and QFNO, while using substantially fewer trainable parameters than the classical FNO. The spectral diagnostics confirmed the key structural feature of the CV construction: the CV-QFNO reproduces the hard cutoff of the classical FNO spectrum, whereas the QFNO retains residual activations beyond the selected modes.

Overall, the CV-QFNO provides a coherent Gaussian photonic formulation of Fourier-domain operator learning. Rather than demonstrating immediate computational advantage, the present work defines an architectural blueprint for photonic neural operators.

The most promising future directions concern settings in which the input data are themselves optical or quantum in nature, rather than classical fields already discretised on a digital grid. For natively optical inputs, such as coherent signals generated by interferometers, wavefront sensors, or free-space optical communication systems, a photonic implementation of the CV-QFNO could in principle operate directly on the incoming field, reducing or bypassing the analogue-to-digital conversion and classical encoding overheads. For intrinsically quantum inputs, as in quantum sensing, quantum simulation, or quantum communication, the bipartite CV architecture could act directly on state-level correlations rather than on classical samples extracted from them.

A longer-term route toward computational regimes beyond efficient classical emulation would require moving beyond the Gaussian framework. Non-Gaussian resources, such as cubic-phase operations or photon-number-resolving measurements, could introduce nonlinear spectral interactions outside the efficiently simulable Gaussian sector. Exploring these directions, together with higher-dimensional generalisations of the architecture and applications to more complex time-dependent problems, represents a natural next step for this line of research.

Author Contributions: P.M. conceived the research idea, developed the methodology, implemented the models, carried out all the experiments, and wrote the manuscript. S.M. (Stefano Mariani) and M.S. contributed to the revision of the manuscript, supported the research, and acquired the funding. S.M. (Stefano Markidis) contributed to the scientific discussion and provided guidance and feedback throughout the project. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by ICSC—“National Research Centre in High Performance Computing, Big Data and Quantum Computing”, funded by the European Union—NextGenerationEU.

Data Availability Statement: The code and data supporting the findings of this study are available in the GitHub repository (<https://github.com/PaoloMarcandelli>, accessed on 25 June 2026) of the corresponding author. They are currently private and will be made publicly available upon publication of this article.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Lu, L.; Jin, P.; Pang, G.; Zhang, Z.; Karniadakis, G.E. Learning Nonlinear Operators via DeepONet Based on the Universal Approximation Theorem of Operators. *Nat. Mach. Intell.* **2021**, *3*, 218–229. [[CrossRef](#)]
2. Li, Z.; Kovachki, N.; Aizzadenesheli, K.; Liu, B.; Bhattacharya, K.; Stuart, A.; Anandkumar, A. Fourier Neural Operator for Parametric Partial Differential Equations. *arXiv* **2020**, arXiv:2010.08895.

3. Kovachki, N.; Li, Z.; Liu, B.; Azizzadenesheli, K.; Bhattacharya, K.; Stuart, A.; Anandkumar, A. Neural Operator: Learning Maps Between Function Spaces with Applications to PDEs. *J. Mach. Learn. Res.* **2023**, *24*, 1–97.
4. Kerenidis, I.; Prakash, A. Quantum machine learning with subspace states. *arXiv* **2022**, arXiv:2202.00054. [[CrossRef](#)]
5. Jain, N.; Landman, J.; Mathur, N.; Kerenidis, I. Quantum Fourier networks for solving parametric PDEs. *Quantum Sci. Technol.* **2024**, *9*, 035026. [[CrossRef](#)]
6. Marcandelli, P.; He, Y.; Mariani, S.; Siena, M.; Markidis, S. Partitioned Hybrid Quantum Fourier Neural Operators for Scientific Quantum Machine Learning. In Proceedings of the 2025 IEEE International Conference on Quantum Computing and Engineering (QCE), Albuquerque, NM, USA, 30 August–5 September 2025; pp. 1676–1687.
7. Li, S.; Jia, K.; Wen, Y.; Liu, T.; Tao, D. Orthogonal Deep Neural Networks. *IEEE Trans. Pattern Anal. Mach. Intell.* **2021**, *43*, 1352–1368. [[CrossRef](#)] [[PubMed](#)]
8. Wang, J.; Chen, Y.; Chakraborty, R.; Yu, S.X. Orthogonal Convolutional Neural Networks. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Seattle, WA, USA, 13–19 June 2020; pp. 11502–11512. [[CrossRef](#)]
9. Monbroussou, L.; Mamon, E.Z.; Landman, J.; Grilo, A.B.; Kukla, R.; Kashefi, E. Trainability and Expressivity of Hamming-Weight Preserving Quantum Circuits for Machine Learning. *Quantum* **2025**, *9*, 1745. [[CrossRef](#)]
10. Lloyd, S.; Braunstein, S.L. Quantum Computation over Continuous Variables. *Phys. Rev. Lett.* **1999**, *82*, 1784–1787. [[CrossRef](#)]
11. Weedbrook, C.; Pirandola, S.; García-Patrón, R.; Cerf, N.J.; Ralph, T.C.; Shapiro, J.H.; Lloyd, S. Gaussian quantum information. *Rev. Mod. Phys.* **2012**, *84*, 621–669. [[CrossRef](#)]
12. Lau, H.K.; Pooser, R.; Siopsis, G.; Weedbrook, C. Quantum Machine Learning over Infinite Dimensions. *Phys. Rev. Lett.* **2017**, *118*, 080501. [[CrossRef](#)] [[PubMed](#)]
13. Schuld, M.; Killoran, N. Quantum Machine Learning in Feature Hilbert Spaces. *Phys. Rev. Lett.* **2019**, *122*, 040504. [[CrossRef](#)] [[PubMed](#)]
14. Marcandelli, P.; Mariani, S.; Siena, M.; Markidis, S. A Continuous-Variable Quantum Fourier Layer: Applications to Filtering and PDE Solving. *Algorithms* **2026**, *19*, 370. [[CrossRef](#)]
15. Cherrat, E.A.; Kerenidis, I.; Mathur, N.; Landman, J.; Strahm, M.; Li, Y.Y. Quantum Vision Transformers. *Quantum* **2024**, *8*, 1265. [[CrossRef](#)]
16. Kerenidis, I.; Landman, J.; Mathur, N. Classical and Quantum Algorithms for Orthogonal Neural Networks. *arXiv* **2021**, arXiv:2106.07198. [[CrossRef](#)]
17. Reck, M.; Zeilinger, A.; Bernstein, H.J.; Bertani, P. Experimental realization of any discrete unitary operator. *Phys. Rev. Lett.* **1994**, *73*, 58–61. [[CrossRef](#)] [[PubMed](#)]
18. Nielsen, M.A.; Chuang, I.L. *Quantum Computation and Quantum Information*, 10th ed.; Cambridge University Press: Cambridge, UK, 2010.
19. Krantz, P.; Kjaergaard, M.; Yan, F.; Orlando, T.P.; Gustavsson, S.; Oliver, W.D. A quantum engineer’s guide to superconducting qubits. *Appl. Phys. Rev.* **2019**, *6*, 021318. [[CrossRef](#)]
20. Flamini, F.; Spagnolo, N.; Sciarrino, F. Photonic quantum information processing: A review. *Rep. Prog. Phys.* **2018**, *82*, 016001. [[CrossRef](#)] [[PubMed](#)]
21. Clements, W.R.; Humphreys, P.C.; Metcalf, B.J.; Kolthammer, W.S.; Walmsley, I.A. An Optimal Design for Universal Multiport Interferometers. *Optica* **2016**, *3*, 1460–1465. [[CrossRef](#)]
22. Higham, N.J. The scaling and squaring method for the matrix exponential revisited. *SIAM J. Matrix Anal. Appl.* **2005**, *26*, 1179–1193. [[CrossRef](#)]
23. Kingma, D.P.; Ba, J. Adam: A Method for Stochastic Optimization. *arXiv* **2014**, arXiv:1412.6980.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.