

EXTENDING FOURIER NEURAL OPERATORS FOR PARAMETERIZED AND COUPLED PDES

Cheng Jing^{1*} Uvini Balasuriya^{1*} Abhishek Verma² Kallol Bera² Shahid Rauf² Kookjin Lee¹

¹Arizona State University, ²Applied Materials Inc.

{cjing6, ubalasure, Kookjin.Lee}@asu.edu,

{AbhishekKumar.Verma, Kallol-Bera, Shahid-Rauf}@amat.com

ABSTRACT

Parameterized and coupled partial differential equations (PDEs) are central to modeling phenomena in science and engineering, yet neural operator methods that address both aspects remain limited. We extend Fourier neural operators (FNOs) with minimal architectural modifications along two directions. For parameterized dynamics, we propose a hypernetwork-based modulation that conditions the operator on physical parameters. For coupled systems, we conduct a systematic exploration of architectural choices, examining how operator components can be adapted to balance shared structure with cross-variable interactions while retaining the efficiency of standard FNOs. Evaluations on benchmark PDEs, including the one-dimensional capacitively coupled plasma equations and the Gray–Scott system, show that our methods achieve up to 55~72% lower errors than strong baselines, demonstrating the effectiveness of principled modulation and systematic design exploration.

1 INTRODUCTION

Numerical simulation has long served as a cornerstone of scientific and engineering inquiry, underpinning advances in areas ranging from fluid dynamics (Ferziger et al., 2019) and climate modeling (Bauer et al., 2015) to material science (Rappaz et al., 2003) and structural analysis (Macneal & Harder, 1985). High-fidelity simulations, while indispensable, often entail substantial computational cost. These costs become especially significant when conducting parameter studies, uncertainty quantification, and real time decision making. As a result, surrogate modeling has emerged as an essential tool to expedite simulation workflows, enabling rapid approximations that preserve first-order fidelity while dramatically reducing computational burden. Recently, deep learning techniques have gained increasing traction in this domain. Methodologies such as physics-informed neural networks (PINNs) (Raissi et al., 2019; Karniadakis et al., 2021) and neural operators (NOs) (Li et al., 2020a;b; Lu et al., 2021) offer flexible, data-driven approximations of parametrized partial differential equations (PDEs), often delivering orders-of-magnitude speedups over traditional solvers while maintaining acceptable accuracy.

Despite significant progress in NO, most existing studies have considered scenarios with both varying initial conditions and parameter regimes, yet they often overlook the particular, yet prevalent case of parameterized dynamics with fixed initial conditions, typified by equations of the form $u_t = f(u; \mu)$ with $u(0) = u_0$. Such settings are ubiquitous in engineering, where model dynamics depend on physical parameters (e.g., forcing terms in inviscid Burgers’ equation (Rewieński, 2003; Carlberg et al., 2013), diffusivity/reaction constants in chemically reacting flows (Buffoni & Willcox, 2010; Lee & Carlberg, 2020), and Reynolds number and viscosity in fluid dynamics (Stabile & Rozza, 2018)) while initial states remain unchanged across simulations. Studying this regime is important for practical applications, for example, to understand how variations in material properties, operating conditions, or control parameters affect system behavior from a fixed, well-defined state.

In this work, we focus specifically on NOs for parameterized coupled PDEs in fixed-initial-condition settings. To facilitate this exploration, we build upon Fourier neural operators (FNOs) (Li et al.,

*Equal Contribution

2021), a method now well-established in the community for its expressive parametric kernel in Fourier space. Our approach is designed to remain general and modular, making it adaptable to both FNO variants and other neural operator frameworks. To incorporate parameterized dynamics effectively, we introduce an architectural modification in which the meta-physics-knowledge encoded in the governing equations is captured by a larger set of shared base parameters, while parameter-specific variations are represented through a smaller set of task-dependent model parameters (with a task being associated a specific physical parameter). To ensure the overall model remains lightweight, we implement this design using a compact hypernetwork, enabling parameterization while retaining a model size comparable to standard, non-parameterized FNOs.

Further, we extend the (parameterized) FNOs to effectively model the coupled systems while being efficient in the model size. While existing efforts (Xiao et al., 2023) introduce specialized decompositions stream to capture cross-field interactions, there is little systematic guidance of the design principles that govern effective cross-variable coupling in neural operators, particularly when aiming to preserve the efficiency and scalability of standard FNOs. In this paper we address that gap: we present a principled design space for extending FNOs to coupled problems, investigate the lift, Fourier layers, and projection operators for identifying effective locations and mechanisms for enabling cross-variable coupling.

In addition, we introduce a new benchmark PDE based on a simplified one-dimensional capacitively coupled plasma (CCP) model. This system is of practical importance in plasma physics and semiconductor manufacturing, and it exhibits rich parameterized dynamics that challenge existing operator-learning methods. Our formulation provides a tractable yet representative setting for evaluating neural operators under coupled and parameterized conditions.

Our contributions are summarized as

- Parametric extensions of FNOs for effectively modeling parameterized dynamics,
- Systematic dissection and adaptation of FNO architecture components for coupled systems,
- Introduction of a novel benchmark problem with many physical parameters describing simplified sheath dynamics governed by plasma physics, and
- Extensive experimentation with the proposed methods, comparisons with baselines, and validation on benchmark problems.

2 TECHNICAL BACKGROUND

2.1 NEURAL OPERATORS

NOs are a type of data-driven surrogate model designed to learn mappings between functions rather than traditional input-output pairs, making them particularly effective for problems governed by PDEs (Lu et al., 2021; Kovachki et al., 2023; Boullé & Townsend, 2024; Azizzadenesheli et al., 2024). By taking discrete representations of continuous functions as input, NOs produce functional representations that can efficiently approximate and simulate complex physical systems.

Let \mathcal{A} and \mathcal{U} be two function spaces, and consider a nonlinear operator $G : \mathcal{A} \rightarrow \mathcal{U}$ mapping between them. Neural operators are a class of machine learning methods designed to construct a surrogate $G_\Theta \approx G$ within a trial space of neural networks. Given input-output pairs $\{(a^i, u^i)\}_{i=1}^n$, where $x^i \in \mathcal{A}$ and $u^i = G(a^i) \in \mathcal{U}$, a neural operator can be trained in the standard supervised manner by solving: $\min_{\Theta} \sum_{i=1}^n L(G_\Theta(a^i), u^i)$, where L denotes a loss term and Θ denotes the parameters of the neural operator G_Θ .

In practice, the functional input data a^i are first discretized; for example, if $a(x)$ and $u(x)$ are functions, the grid representation $\mathbf{a} = [a(x_1), \dots, a(x_m)]^\top \in \mathbb{R}^{m \times d_a}$ serves as input along with an evaluation point (or set of points) $D = \{x_1, \dots, x_m\}$ in the domain of a , and the neural operator learns a mapping $\mathbf{a} \mapsto G_\Theta(\mathbf{a})$ that predicts $G_\Theta(\mathbf{a})(x_p) \approx u(x_p)$ (Lu et al., 2021; Li et al., 2021). With the assumption that the spatial grid where $a(x)$ and $u(x)$ are discretized is fixed, $(D, \mathbf{a}) \mapsto \mathbf{u}$ (under G) produces a vector \mathbf{u} of operator evaluations, and therefore samples $(\mathbf{x}, \mathbf{a}, \mathbf{u})$ form the training data for the NO learning problem.

2.2 FOURIER NEURAL OPERATORS

FNOs are a specific instantiation of NOs that model learnable nonlinear kernel integral operators, with the integrals efficiently evaluated in the Fourier frequency domain. FNOs, $G_\theta : \mathcal{A} \rightarrow \mathcal{U}$, first projects the input data into hidden representations via local transformation, $v_0(x) = P(a(x))$, where P is typically chosen as a linear projector in practice. FNOs then employ the iterative updates of hidden representations $v_\ell \mapsto v_{\ell+1}$ via

$$v_{\ell+1}(x) := \sigma\left(Wv_\ell(x) + (\mathcal{K}(a; \phi)v_\ell)(x)\right), \quad \forall x \in D \quad (1)$$

where W is a linear transformation and \mathcal{K} denotes a kernel integral operator. The kernel integral operator is defined as

$$(\mathcal{K}(a; \phi)v_\ell)(x) := \int_D \kappa(x, y, a(x), a(y); \phi)v_\ell(y)dy, \quad \forall x \in D \quad (2)$$

where ϕ indicates the learnable parameters that characterize the kernel. For efficient computation of Eq. equation 2, FNOs define the kernel integral operator in Fourier space and perform a convolution operation in that space such that

$$(\mathcal{K}(\phi)v_\ell)(x) = \mathcal{F}^{-1}\left(R_\phi \cdot (\mathcal{F}v_\ell)\right)(x), \quad \forall x \in D, \quad (3)$$

where \mathcal{F} and \mathcal{F}^{-1} denote Fourier and inverse Fourier transformations, respectively, and R_ϕ denotes the Fourier transform of function κ . In practice, R_ϕ is parameterized as a learnable tensor. The last hidden representation $v_T(x)$ is then projected back to the target data space via a local transformation, $u(x) = Q(v_T(x))$, where Q is typically modeled as a shallow multi-layer perceptron (MLP).

3 RELATED WORK

Building on the success of FNOs, there have been many variants extending their capabilities. For example, physics-informed neural operator (PINO) (Li et al., 2024) and Geo-FNO (Li et al., 2023) extend FNO with physical consistency or geometric generality.

More closely related to our work are efforts that enhance FNO architectures to improve expressivity and efficiency. Factorized FNO (Tran et al., 2023), U-FNO (Wen et al., 2022), and U-NO (Rahman et al., 2023) propose structural modifications: separable spectral layers or encoder-decoder designs that enable stable training. Our approach follows a similar philosophy of architectural enhancement, but with emphasis on parameterized and coupled systems rather than general-purpose scalability.

Within the FNO family, explicit modeling of physical parameters remains relatively underexplored. The recently proposed HyperFNO (Alesiani et al., 2022) infers the model parameters of FNO via a hypernetwork conditioned on known physical parameters, illustrating the benefit of encoding parametric dependencies. Another complementary study (Subramanian et al., 2023) explores scaling laws and transfer behavior of neural operators under large datasets and compute budgets. Our method shares the goal of parameter-aware modeling but focuses on a compact modulation mechanism that preserves the lightweight character of standard FNOs.

For coupled systems, prior works such as coupled multiwavelet neural operator (CMWNO) (Xiao et al., 2023) and attention-based multiphysics models (Rahman et al., 2024) have introduced cross-field coupling through wavelet or large-scale pre-training. In contrast, our work provides a systematic and minimal architectural extension to FNOs, achieving efficient cross-variable interaction without sacrificing model compactness.

4 METHODS

Our methods seek minimal yet principled architectural modifications to FNOs that enable them to handle parameterized and coupled dynamical systems. We emphasize minimal modification as a design principle: this ensures (i) the resulting models preserve the efficiency and scalability of FNOs and (ii) performance gains can be attributed to structural insights rather than increased capacity.

To this end, we focus on the three core computational components of FNOs: the lift operator (input projection), the Fourier layers (spectral updates), and the projection operator (output mapping). We systematically explore how to adapt each component to encode parameter dependence and to capture cross-variable interactions in coupled systems. Section 4.1 presents the parameterized extension via hypernetwork-based modulation, and Section 4.2 details the coupled extension by introducing the design spaces and spectral-domain coupling.

4.1 EXTENSIONS TO PARAMETRIZED FNO

A compact hypernetwork-based modulation scheme is proposed to extend FNOs for handling physical parameters, $\mu \in \mathbb{R}^{n_\mu}$. A lightweight hypernetwork (Ha et al., 2017) takes the physical parameters μ as input and outputs a set of shifts, $s(x, \mu) = f^{\text{hyper}}(x, \mu)$, where $s(x, \mu) = [s_1(x, \mu), \dots, s_L(x, \mu)]$ corresponds to layer-specific adjustments. These shifts are incorporated as additional biases in the Fourier layers:

$$v_{\ell+1}(x) := \sigma\left(Wv_\ell(x) + (\mathcal{K}(a; \phi)v_\ell)(x) + s_\ell(x, \mu)\right), \quad \forall x \in D \quad (4)$$

We refer to this variant using “hp” (e.g., hpFNO) indicating the hypernetwork-based approach (Figure 1).

The rationale behind this design is twofold. First, modulation has already proven effective in implicit neural representations (INRs) (Sitzmann et al., 2020; Fathony et al., 2020; Dupont et al., 2022) and physics-informed neural networks (PINNs) (Cho et al., 2023), where conditioning on auxiliary variables enables networks to represent families of functions without retraining. Second, unlike a simple input augmentation, modulation operates directly on internal representations, allowing parameter dependence to influence the dynamics at every layer rather than only at the input. Together, these properties make hp-variant a lightweight yet expressive model for parameterized dynamics.

Compared to HyperFNO (Alesiani et al., 2022), which infers a large part of the base FNO through a hypernetwork, our approach adopts a lightweight shift modulation that perturbs intermediate activations while keeping the core weights fixed. Within the considered scenarios in this work, this design achieves comparable flexibility with substantially lower parameter overhead and complexity.

For completeness, we also consider simpler variant that employs input augmentation to encode physical parameters. In this approach, the parameters are concatenated with the function input as $[a(x); \mu(x)]$, where $\mu(x) = \mu$ is constant across the spatial domain. The lift operator then encodes this augmented input into a latent representation, $v_0(x) = P([a(x), \mu(x)])$. This strategy mirrors prior work, for example in the neural ODE community, where augmenting the input has been shown to enhance the expressivity of the learned dynamics (Dupont et al., 2019; Massaroli et al., 2020; Lee & Parish, 2021), and PINNs (Cai et al., 2021; Cho et al., 2024). We refer to this variant using the prefix “p” (e.g., pFNO) to denote parameterization via input augmentation.

While input augmentation is straightforward and effective in some cases, it may not fully capture how physical parameters influence the underlying dynamics. The hypernetwork-based modulation therefore provides a more expressive yet efficient alternative, conditioning the operator throughout its layers rather than only at input. Within our experiments, the hypernetwork-based approach consistently achieves better performance without significantly increasing model complexity.

4.2 EXTENSIONS TO COUPLED FNO

We now turn to extending FNOs for coupled systems, where multiple interdependent variables evolve under shared dynamics. For simplicity of exposition, we focus on the case of two variables and denote each component by superscript, z^\square for a generic variable z , where $\square \in \{\alpha, \beta\}$. With this notation, the operator of interest becomes $(u^\alpha(x), u^\beta(x)) = G_\Theta(a^\alpha(x), a^\beta(x))$. Our design philosophy follows the same principle outlined at the beginning of this section: to retain the minimal and scalable structure of standard FNOs while introducing only the modifications necessary to capture cross-variable interactions. To this end, we introduce coupling in the spectral domain, where the spectral representations of multiple variables are mixed through a lightweight encoder-decoder mechanism. To support this coupling, we explore how the core FNO computational components, namely the lift operator, P , Fourier layers, and projection operator, Q , can be configured, including choices on whether those operators are shared or defined separately across variables.




Figure 1: The proposed hpFNO architecture. The hypernetwork produces a parameter-dependent ‘shift’ term, $s(\mu, x)$, which acts as a parameter-dependent bias term in each Fourier layer.

Lift operator, P For the local transform, $(Pa)(x) = v_0(x) \in \mathbb{R}^{d_v}$, we consider two main design choices that differ in whether the operator is shared or individual across the two variables.

Ⓜ (Shared): A single operator $P : \mathbb{R}^1 \rightarrow \mathbb{R}^{d_v}$ is shared across the two input components, so that both variables are mapped into the latent space using the same transformation.

Ⓝ (Separate): Two individual mappings, $P = (P^\alpha, P^\beta)$, allow variable-specific latent representation.

Fourier layers (spectral coupling) The main modification enabling cross-variable interaction occurs in the Fourier layers, which consist of two components: the point-wise linear map $Wv_\ell(x)$ (Ⓚ) / (Ⓛ) and the global spectral convolution $\mathcal{K}v_\ell(x)$ (Ⓞ).

Ⓚ / Ⓛ: The point-wise linear map can be implemented either as a single operator shared by both variables (Ⓚ) or as two separate operators (W^α and W^β) (Ⓛ).

Ⓞ For coupled systems, we define the global spectral convolution to perform coupling only in Fourier space. Specifically, each variable is first transformed independently, $\tilde{v}^\square(k) = \mathcal{F}v_\ell^\square(k)$, for $\square = \{\alpha, \beta\}$. Their spectral representations are then combined through a shallow encoder network $f^{\text{enc}}(\tilde{v}^\alpha(k), \tilde{v}^\beta(k))$, followed by the usual mode-wise kernel multiplication $\tilde{\tilde{v}}(k) = R_\phi(k)\tilde{v}(k)$. The result is decomposed into variable-specific coefficients using another shallow network, $(\tilde{\tilde{v}}^\alpha(k), \tilde{\tilde{v}}^\beta(k)) = f^{\text{dec}}(\tilde{\tilde{v}}(k))$, and finally mapped back to the data space via the inverse Fourier transform, $v_{\ell+1}^\square(k) = \mathcal{F}^{-1}(\tilde{\tilde{v}}^\square(k))$, for $\square = \alpha, \beta$.

A key design choice here is that coupling is introduced only in Fourier space, after the individual Fourier transforms. This has several benefits compared to approaches that use entirely separate convolution operators and exchange hidden representations across variables. First, it preserves the core structure of FNOs, keeping the model lightweight and scalable. Second, the Fourier domain naturally captures long-range correlations, making it a well-suited location for cross-variable mixing. Finally, by avoiding repeated exchanges of hidden states in the spatial domain, this design reduces computational overhead while maintaining symmetry between the two variables.

Projection operator, Q The local transform, $(Qv_T)(x) = u(x)$, resembles with that of the lift operator P and, thus, similar choices apply. We focus here on the single-channel output format for each variable, which is the natural setting for coupled systems.

Ⓜ (Shared): A single projection is shared by both outputs, $Q : \mathbb{R}^{d_v} \rightarrow \mathbb{R}^1$.

Ⓝ (Separate): Two operators are separately defined for each output component, $Q = (Q^\alpha, Q^\beta)$, with $Q^\square : \mathbb{R}^{d_v} \rightarrow \mathbb{R}^1$ for $\square \in \alpha, \beta$.

To refine the design space under Ⓝ, we adopt the adaptive basis viewpoint (Cyr et al., 2020), which interprets the last hidden activation as a set of adaptive basis functions, $\Psi(a(x)) = [\psi_1(a(x)), \dots, \psi_{n_q}(a(x))]^\top$, and the weight of the output layer as coefficients, $\Xi = [\xi_1, \dots, \xi_{n_q}]^\top$: $u(x) = \sum_{i=1}^{n_q} \xi_i \psi_i(x)$. In practice, Q is parameterized as a shallow MLP, $(Qv_T)(x) = W_2 \sigma(W_1 v_T(x) + b_1) + b_2$ and the adaptive basis corresponds to $\Psi(a(x)) = \sigma(W_1 v_T(x) + b_1)$

and the coefficients to $\Xi = W_2$. The shared/separate choices are then realized by sharing or splitting the weight matrices and biases.

From this perspective, Q can differ in whether the basis and the coefficients are shared or separate: $\textcircled{Q2a}$ Shared basis and separate coefficients, $\textcircled{Q2b}$ Separate basis and shared coefficients, and $\textcircled{Q2c}$ Separate basis and separate coefficients.

4.3 PUTTING ALL TOGETHER

Based on our experimentation (presented in Appendix D.1.1) with the different combinations of the design choices (further elaborated in Appendix A), we define our model, extended FNOs — FNO_x , as a realization of the combination: $\textcircled{P1} + \textcircled{L2} + \textcircled{G3} + \textcircled{Q2c}$, which employs the shared lift operator, the separate point-wise linear maps, the proposed coupled global spectral convolution layer, and the separate sets of basis functions and separate sets of coefficients in the projection operator.

5 EXPERIMENTAL RESULTS

In this section, we present experimental results comparing the proposed methods with a range of baseline models. All methods are implemented in PYTHON with PYTORCH (Paszke et al., 2019), building upon the original implementation of FNOs (Li et al., 2021) and extending it to handle coupled and parameterized settings. Each experiment is repeated five times with different random seeds, and all computations are performed on an NVIDIA A100 GPU with 80GB memory.

5.1 SETUP

Tasks and training We consider time-dependent PDEs that generate solution trajectories $\{u(x, t; \mu)\}_{t=0}^T$ with $u(x, 0; \mu) = u(0)$. NOs are formulated as one-step evolution maps that approximate the discrete-time flow of the underlying PDE. Specifically, given a finite history of past solution states $\{u(t - \tau)\}_{\tau=0}^{T_{\text{in}}-1} \in \mathcal{A}$, the operator predicts the next state $u(t+1) \in \mathcal{U}: G_{\Theta} : \mathcal{A} \rightarrow \mathcal{U}$. Multi-step forecasts are obtained autoregressively by recursively applying the same operator. During training and testing, the initial input window $(u(T_{\text{in}} - 1), \dots, u(0))$ is assumed to be available, while subsequent rollouts incorporate predicted states back into the input window.

Baselines As baselines of comparisons, we consider methods inherited from Xiao et al. (2023):

- FNO_c : FNOs taking inputs as a vertical concatenation of the discretized input data,
- CFNO : Two separate FNOs sharing exchanging hidden representation in Fourier layers,
- MWT_c : MWTs (Gupta et al., 2021) taking a vertically concatenated input data,
- CMWNO : Multiwavelet NOs specifically designed for coupled systems (Xiao et al., 2023),

a FNO baseline considering parametric extension via hypernetworks (Alesiani et al., 2022):

- HyperFNO_c : FNOs whose model parameters of the lift/projection operators (P/Q), the pointwise linear map, W , and the kernel weight R are inferred from hypernetworks,

other relevant baselines, DeepONets (Lu et al., 2021) and U-Nets (Ronneberger et al., 2015), modified to handle multi input:

- DON_c : DeepONets extended to take a vertical concatenation of the discretized input data,
- U-Net_c : U-Nets taking two-channel input,

and, finally, one variant of a certain combinations of choices from the described design choices:

- FNO_m : the lift operator projects the two-channel input (a^α, a^β) jointly into a shared latent space, i.e., $P : \mathbb{R}^2 \rightarrow \mathbb{R}^{d_v}$ with $v_0(x) = P([a^\alpha(x), a^\beta(x)])$.

We refer to Appendix B.1 for the detailed description of the baselines.

Performance evaluation We evaluate performance using the normalized root mean square error (nRMSE), averaged over all test trajectories:

$$\text{nRMSE} = \frac{1}{n_{\text{test}}} \sum_{i=1}^{n_{\text{test}}} \left\| u_{0:T}^{(i)} - \tilde{u}_{0:T}^{(i)} \right\|_2 / \left\| u_{0:T}^{(i)} \right\|_2 \quad (5)$$

where u and \tilde{u} denote the ground-truth and predicted trajectories. For coupled systems, nRMSE is computed separately for each variable, and the final error is the sum: $\text{err} = \text{nRMSE}^\alpha + \text{nRMSE}^\beta$. For all methods, reported errors are computed on normalized fields; given the scale-invariant nature of the metric, the results are expected to be similar under the original physical scaling.


5.2 1-DIMENSIONAL CAPACITIVELY COUPLED PLASMA FLUID MODEL

In this section, we look at a 1-dimensional capacitively coupled plasma (CCP) model as an example to demonstrate the performance of the proposed model in a parameteric setting.

The simplified CCP model describes a low-temperature plasma that is generated between two electrodes when an alternating voltage is applied, operating in a manner similar to a capacitor. In this system, the oscillating electric field accelerates electrons, which then collide with neutral gas atoms or molecules, producing ions and additional electrons that sustain the plasma. CCPs are widely used in microelectronics manufacturing, including etching and thin-film deposition for semiconductor devices, as well as in materials processing and surface engineering (Lieberman & Lichtenberg, 1994; Rauf et al., 2023).

The governing equation is defined as follows:

$$\begin{aligned} \partial_t n_e &= -\partial_x \Gamma_e + R, & (\text{Electron continuity equation}) \\ \partial_{xx} \phi &= -e/\epsilon_0(n_e - n_{i0}), & (\text{Poisson equation}) \end{aligned} \quad (6)$$



where $n_e(x, t)$ and $\phi(x, t)$ denote electron density and electric potential, which are the solutions of the equations, and ∂_{\square} refers to a partial derivative with respect to \square . Γ_e is electron flux, defined as diffusion flux and drift flux such as $\Gamma_e = -D\partial_x n_e - \mu n_e \partial_x \phi$ with electron diffusion coefficient D and electron mobility coefficient μ , and ion density is defined as $n_{i0} = R_0(x_2 - x_1)\sqrt{m_i/eT_e}$. R and R_0 denote reaction rate and coefficient, respectively. The ion density is assumed constant in this model.

Figure 2: The spatial domain and boundary conditions.

Physical parameters There are several physical parameters characterizing a domain geometry, boundary conditions (BCs), and governing physical dynamics (See Table 5 in Appendix). Among them, in this study, we consider fixed parameters for geometry (illustrated in Figure 2), but varying parameters that define the boundary conditions and the dynamics. In the BCs, the zero boundary condition is given to the electron density (at $x = 0, L$). The electric potential satisfies a Dirichlet condition (at $x = 0$), and a time-periodic Dirichlet condition at $x = L$, $V(t) = V_0 \sin(2\pi t)$. That is, the potential at the right boundary oscillates sinusoidally with amplitude V_0 (driving voltage) and frequency f and we vary V_0 in the experiments. This models the externally applied radio-frequency forcing that sustains the discharge in CCP. Along with V_0 , we consider two other physical parameters that governs the dynamics: reaction coefficient, R_0 , and ion mass, m_i . Figure 3 depicts solution snapshots of $n_e(x, t)$ and $\phi(x, t)$, simulated over a single cycle, configured with $V_0 = 100$, $R_0 = 2.7 \times 10^{20}$, and $m_i = 40 \times 1.67 \times 10^{-27}$.




Figure 3: An example solution snapshot of electron density and electric potential.

5.2.1 DATA SETUP

Data generation We collect solution snapshots by numerically simulating the governing equations using a finite-difference solver. The spatial domain is discretized into 128 cells, and the time domain into 100,000 steps per cycle, with the cycle determined by the boundary driving frequency of 13.56 MHz. A fixed initial condition is used to focus on scenarios where capturing physical dynamics is critical. We then subsample every 1000th snapshot, yielding 100 temporal indices per solution. Physical parameters are sampled on an equidistant grid, which contains 100 elements: $V_0 \in [100, 300]$, $R_0 \in [2.7 \times 10^{19}, 2.7 \times 10^{20}]$, and $m_i \in [1.67 \times 10^{-26}, 6.68 \times 10^{-26}]$. This setup yields 100 trajectories, each corresponding to one parameter configuration, which are then randomly split into training and test sets with a 9:1 ratio.

5.2.2 NUMERICAL RESULTS

1-dimensional parameter space We perform three separate experiments, each on a dataset obtained by varying a single physical parameter: R_0 , V_0 , or m_i . Table 1 summarizes the results, reporting prediction errors on the test set for the proposed methods and the baseline models. The proposed FNO_x architecture (FNO_x , without the parametric extension) provides notable improvement over the baseline methods. The overall second best model in this set of experiments is CWMNO and, compared to its performance, FNO_x achieve the prediction accuracy improved up to 38% (in the varying reaction rate scenario). The parametric extensions, pFNO_x and hpFNO_x , further improve the performance, achieving up to 55% improvement over the best performing non-parameterized baselines. Compared to hyperFNO , which is also informed explicitly by the physical parameters through hypernetworks, the proposed variants, particularly, hpFNO_x demonstrates significant improvements (more than 40% for all cases).

Table 1: Performance comparisons between the proposed methods and the baseline methods. All models are tested on three parameters (R_0 , V_0 and m_i) individually. The performance is measured in the relative ℓ^2 -error (nRMSE); the reported numerical values refer to the mean (\pm std. dev.).

Model	Reaction rate	Driving voltage	Ion mass
FNO_m	0.0403 (\pm 0.0012)	0.0791 (\pm 0.0243)	0.0363 (\pm 0.0067)
FNO_c	0.0375 (\pm 0.0055)	0.0873 (\pm 0.0200)	0.0299 (\pm 0.0030)
CFNO	0.0315 (\pm 0.0093)	0.0428 (\pm 0.0064)	0.0333 (\pm 0.0076)
HyperFNO_c	0.0278 (\pm 0.0117)	0.0355 (\pm 0.0160)	0.0253 (\pm 0.0088)
MWT_c	0.0409 (\pm 0.0032)	0.0639 (\pm 0.0038)	0.0403 (\pm 0.0040)
CWMNO	0.0312 (\pm 0.0023)	0.0526 (\pm 0.0521)	0.0241 (\pm 0.0036)
DON_c	0.0844 (\pm 0.0058)	0.2147 (\pm 0.0414)	0.1035 (\pm 0.0174)
U-NET $_c$	0.1084 (\pm 0.0345)	0.0844 (\pm 0.0354)	0.1719 (\pm 0.1022)
FNO_x	0.0193 (\pm 0.0059)	0.0345 (\pm 0.0108)	0.0212 (\pm 0.0062)
pFNO_x	0.0194 (\pm 0.0075)	0.0278 (\pm 0.0053)	0.0142 (\pm 0.0021)
hpFNO_x	0.0154 (\pm 0.0029)	0.0192 (\pm 0.0040)	0.0128 (\pm 0.0017)

Computational/memory aspects Since prediction accuracy alone does not provide a complete picture of the experimental results, Figure 4 presents additional information under the varying driving voltage scenario, including model size (4a) and computational timing (4b), alongside the models’ prediction accuracy. The model size is measured in the number of trainable model parameters and the timing reports a per-epoch training time measured and averaged across the total number of epochs. The both figures provide the same observation: the proposed methods FNO_x , pFNO_x , and hpFNO_x achieve improvement in accuracy without sacrificing the model size (i.e., memory footprint) and the computational wall time too much. Additionally, Figure 4c shows the test loss trajectories over the training epoch. Along with the computational timing per epoch (Figure 4b), the loss trajectories provide a complete picture on models’ computational requirements in achieving the presented accuracy.

Study on varying T_{in} Although the true dynamics depend on physical parameters, all non-parametric NOs in Table 1 achieve reasonable accuracy. This is because the NOs take as input a window of past snapshots, characterized by T_{in} , akin to delay embedding in dynamical systems (Takens, 2006). By consuming multiple consecutive time steps, the NOs implicitly encode information about system parameters through temporal correlations, enabling them to infer dynamics.

Thus, the benefits of using the parameterized version of FNO_x can be more pronounced when T_{in} becomes smaller. To see this, we extend the experimentation for varying $T_{\text{in}} = \{10, 5, 2, 1\}$.¹ Table 2 reports the results comparing the performance of FNO_c and FNO_x and their two parametric variants (pFNO_c , hpFNO_c), and (pFNO_x , hpFNO_x), respectively. As T_{in} decreases, the overall model performance, which is expected. For $T_{\text{in}} = \{10, 5\}$, the improvement achieved by hpFNO_x over FNO_x over 20%. For $T_{\text{in}} = \{2, 1\}$, while the hp variants maintain the accuracy level around

¹For $T_{\text{in}} = \{2, 1\}$, training instability was observed across some random seeds; due to this, we trained models with 10 random seeds and report the average over the five runs with the higher test prediction accuracy.

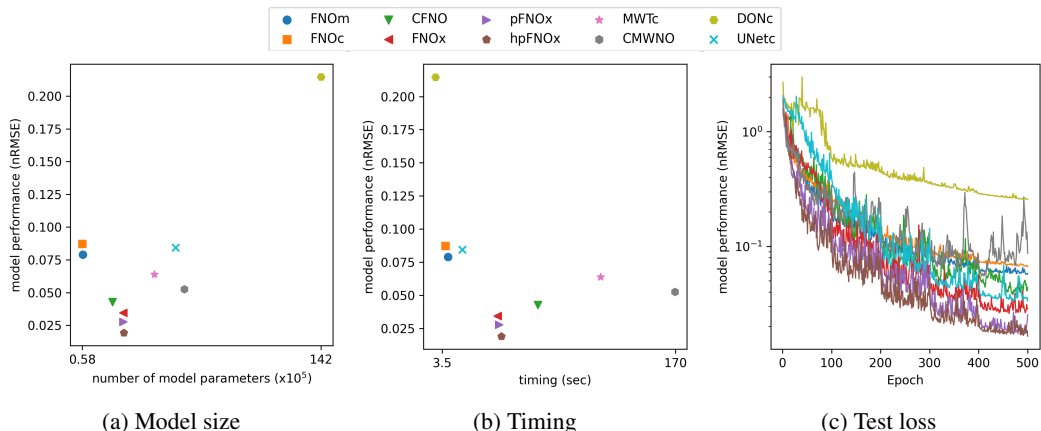


Figure 4: [1D CCP] Plots depict the number of model parameters versus the model performance (left), computational timing versus the model performance (middle), and the test loss trajectories over epochs (right). These results are from the varying driving voltage scenario.

13~16% relative errors. Considering that only initial conditions are available in most practical setups, maintaining this level of performance could be considered as important step forward for FNOs.

Table 2: Performance comparisons of non-parameterized (FNO_c , FNO_x) and parameterized ($pFNO_c$, $hpFNO_c$, $pFNO_x$, $hpFNO_x$) models for varying $T_{in} = \{10, 5, 2, 1\}$. All models are tested on the reaction rate case. The performance is measured in the relative ℓ^2 -error; the reported numerical values refer to the mean (\pm std. dev.). \dagger indicates that the models fail to learn the dynamics and \ddagger indicates that the model performance are averaged from the best 5 runs out of 10 total runs.

Model	$T_{in} = 10$	$T_{in} = 5$	$T_{in} = 2$	$T_{in} = 1$
FNO_c	0.0375 (± 0.0055)	0.1324 (± 0.0304)	1.0048 \dagger (± 0.6356)	1.4484 \dagger (± 0.3128)
$pFNO_c$	0.0334 (± 0.0101)	0.0923 (± 0.0276)	0.2151 \ddagger (± 0.0374)	0.3757 \ddagger (± 0.0679)
$hpFNO_c$	0.0196 (± 0.0021)	0.0804 (± 0.0237)	0.1515 \ddagger (± 0.0070)	0.1609 \ddagger (± 0.0043)
FNO_x	0.0193 (± 0.0059)	0.0406 (± 0.0093)	1.2832 \dagger (± 0.4298)	1.8143 \dagger (± 0.0558)
$pFNO_x$	0.0194 (± 0.0075)	0.0464 (± 0.0158)	0.1640 \ddagger (± 0.0155)	0.2522 \ddagger (± 0.1376)
$hpFNO_x$	0.0154 (± 0.0029)	0.0317 (± 0.0022)	0.1324 \ddagger (± 0.0242)	0.1372 \ddagger (± 0.0100)

5.3 GRAY-SCOTT EQUATIONS

As the second benchmark problem, we consider the Gray-Scott equations, which form a system of coupled reaction-diffusion equations describing the spatio-temporal dynamics of interacting chemical species. We consider a parameterized dynamics of this equation, where the parameters are diffusion coefficients and feed rate. The task of the NOs are performing predictions of time-dependent PDE solutions in the autoregressive manner, and we follow the same training/testing protocols shown in the first benchmark problem. We refer to Appendix E for details (the equations, the experimental setups, and additional results).

Table 3 summarizes the performance of the proposed and baseline models on the Gray-Scott benchmark, where the task is to predict the spatiotemporal evolution of two reacting chemical species under varying feed-rate conditions. The proposed

Table 3: Performance comparisons: The performance is measured in the relative ℓ^2 -error; the reported numerical values refer to the mean (\pm standard deviation)

Model	Performance
FNO_m	0.0116 (± 0.0026)
FNO_c	0.0097 (± 0.0030)
CFNO	0.0161 (± 0.0045)
MWT_c	0.0092 (± 0.0021)
CMWNO	0.0293 (± 0.0078)
DON_c	0.0138 (± 0.0020)
U-Net $_c$	0.0159 (± 0.0071)
FNO_x	0.0075 (± 0.0016)
$pFNO_x$	0.0041 (± 0.0010)
$hpFNO_x$	0.0022 (± 0.0006)

variants consistently outperform all baseline methods, achieving relative ℓ^2 -errors of 0.0075, 0.0041, and 0.0022 for FNO_x , pFNO_x , and hpFNO_x , respectively. Among the baselines, the multiwavelet-based model (MWT_c) performs best with 0.0092 error, yet the parameterized versions of our model reduce this by more than half. These results indicate that incorporating parameter conditioning and lightweight coupling in the Fourier domain substantially improves predictive accuracy while maintaining the compactness of the original FNO architecture. Overall, the hpFNO_x variant achieves up to 72% lower error compared to the strongest baseline, demonstrating effective generalization across parameterized reaction-diffusion dynamics.

6 CONCLUSION

This work has investigated extensions of Fourier neural operators (FNOs) for modeling parameterized and coupled time-dependent partial differential equations. To enhance the capability of FNOs in handling parameterized dynamics, we have proposed a novel modulation-based architectural modification that incorporates parameter information in a lightweight yet effective manner. For coupled systems, we have defined design spaces of FNOs and conducted an in-depth systematic investigation to identify economic architectures that preserve efficiency while capturing inter-component interactions. In addition, we have introduced a novel benchmark problem on plasma physics, which introduces rich parameterized dynamics and contributes a valuable testbed to the community. Experimental results have demonstrated that our proposed approaches achieve significant improvements in predictive performance, while maintaining model size and runtime efficiency comparable to standard FNOs.

7 ETHICS STATEMENT

This study focuses on methodological advances in machine learning for computational physics problems, and we do not anticipate any significant ethical concerns arising from the proposed approaches.

8 REPRODUCIBILITY STATEMENT

To support reproducibility, we provide detailed descriptions of the proposed methods, including model architectures, algorithms, loss definitions, hyperparameters, training setups, and baseline comparisons. All experiments are described in the main text and appendix, and upon acceptance we will release the full source code to further facilitate reproduction of our results.

9 ACKNOWLEDGEMENT

The authors acknowledge funding support from Applied Materials Inc. K. Lee acknowledges partial support from the U.S. National Science Foundation under grant IIS 2338909. K. Lee also acknowledges Research Computing Jennewein et al. (2023) at Arizona State University for providing HPC resources that have contributed to the partial research results reported within this paper.

REFERENCES

- Francesco Alesiani, Makoto Takamoto, and Mathias Niepert. HyperFNO: Improving the generalization behavior of Fourier neural operators. In *NeurIPS 2022 Workshop on Machine Learning and Physical Sciences*, 2022.
- Kamyar Aizzadenesheli, Nikola Kovachki, Zongyi Li, Miguel Liu-Schiaffini, Jean Kossaifi, and Anima Anandkumar. Neural operators for accelerating scientific simulations and design. *Nature Reviews Physics*, 6(5):320–328, 2024.
- Peter Bauer, Alan Thorpe, and Gilbert Brunet. The quiet revolution of numerical weather prediction. *Nature*, 525(7567):47–55, 2015.
- Nicolas Boullé and Alex Townsend. A mathematical guide to operator learning. In *Handbook of Numerical Analysis*, volume 25, pp. 83–125. Elsevier, 2024.

- Marcelo Buffoni and Karen Willcox. Projection-based model reduction for reacting flows. In *40th Fluid Dynamics Conference and Exhibit*, pp. 5008, 2010.
- Shengze Cai, Zhicheng Wang, Sifan Wang, Paris Perdikaris, and George Em Karniadakis. Physics-informed neural networks for heat transfer problems. *Journal of Heat Transfer*, 143(6):060801, 2021.
- Kevin Carlberg, Charbel Farhat, Julien Cortial, and David Amsallem. The GNAT method for nonlinear model reduction: Effective implementation and application to computational fluid dynamics and turbulent flows. *Journal of Computational Physics*, 242:623–647, 2013.
- Woojin Cho, Kookjin Lee, Donsub Rim, and Noseong Park. Hypernetwork-based meta-learning for low-rank physics-informed neural networks. *Advances in Neural Information Processing Systems*, 36:11219–11231, 2023.
- Woojin Cho, Minju Jo, Haksoo Lim, Kookjin Lee, Dongeun Lee, Sanghyun Hong, and Noseong Park. Parameterized physics-informed neural networks for parameterized PDEs. In *International Conference on Machine Learning*, pp. 8510–8533. PMLR, 2024.
- Eric C Cyr, Mamikon A Gulian, Ravi G Patel, Mauro Perego, and Nathaniel A Trask. Robust training and initialization of deep neural networks: An adaptive basis viewpoint. In *Mathematical and Scientific Machine Learning*, pp. 512–536. PMLR, 2020.
- Tobin A Driscoll, Nicholas Hale, and Lloyd N Trefethen. *Chebfun guide*, 2014.
- Emilien Dupont, Arnaud Doucet, and Yee Whye Teh. Augmented neural odes. *Advances in neural information processing systems*, 32, 2019.
- Emilien Dupont, Hyunjik Kim, SM Ali Eslami, Danilo Jimenez Rezende, and Dan Rosenbaum. From data to functa: Your data point is a function and you can treat it like one. In *International Conference on Machine Learning*, pp. 5694–5725. PMLR, 2022.
- Rizal Fathony, Anit Kumar Sahu, Devin Willmott, and J Zico Kolter. Multiplicative filter networks. In *International conference on learning representations*, 2020.
- Joel H Ferziger, Milovan Perić, and Robert L Street. *Computational Methods for Fluid Dynamics*. Springer, 2019.
- Gaurav Gupta, Xiongye Xiao, and Paul Bogdan. Multiwavelet-based operator learning for differential equations. *Advances in neural information processing systems*, 34:24048–24062, 2021.
- David Ha, Andrew M Dai, and Quoc V Le. Hypernetworks. In *International Conference on Learning Representations*, 2017.
- Maximilian Herde, Bogdan Raonic, Tobias Rohner, Roger Käppeli, Roberto Molinaro, Emmanuel de Bézenac, and Siddhartha Mishra. Poseidon: Efficient foundation models for PDEs. *Advances in Neural Information Processing Systems*, 37:72525–72624, 2024.
- Benjamin Holzsuh, Qiang Liu, Georg Kohl, and Nils Thuerey. PDE-Transformer: Efficient and versatile transformers for physics simulations. *arXiv preprint arXiv:2505.24717*, 2025.
- Douglas M Jennewein, Johnathan Lee, Chris Kurtz, William Dizon, Ian Shaeffer, Alan Chapman, Alejandro Chiquete, Josh Burks, Amber Carlson, Natalie Mason, et al. The Sol supercomputer at Arizona State University. In *Practice and Experience in Advanced Research Computing 2023: Computing for the Common Good*, pp. 296–301. 2023.
- George Em Karniadakis, Ioannis G Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3(6):422–440, 2021.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *Proceedings of the 3rd International Conference on Learning Representations (ICLR 2015)*, 2015. URL <https://arxiv.org/abs/1412.6980>.

- Nikola Kovachki, Zongyi Li, Burigede Liu, Kamyar Azizzadenesheli, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Neural operator: Learning maps between function spaces with applications to PDEs. *Journal of Machine Learning Research*, 24(89):1–97, 2023.
- Kookjin Lee and Kevin T Carlberg. Model reduction of dynamical systems on nonlinear manifolds using deep convolutional autoencoders. *Journal of Computational Physics*, 404:108973, 2020.
- Kookjin Lee and Eric J Parish. Parameterized neural ordinary differential equations: Applications to computational physics problems. *Proceedings of the Royal Society A*, 477(2253):20210162, 2021.
- Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Neural operator: Graph kernel network for partial differential equations. *arXiv preprint arXiv:2003.03485*, 2020a.
- Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Andrew Stuart, Kaushik Bhattacharya, and Anima Anandkumar. Multipole graph neural operator for parametric partial differential equations. *Advances in Neural Information Processing Systems*, 33:6755–6766, 2020b.
- Zongyi Li, Nikola Borislavov Kovachki, Kamyar Azizzadenesheli, Kaushik Bhattacharya, Andrew Stuart, Anima Anandkumar, et al. Fourier neural operator for parametric partial differential equations. In *International Conference on Learning Representations*, 2021.
- Zongyi Li, Nikola Kovachki, Chris Choy, Boyi Li, Jean Kossaifi, Shourya Otta, Mohammad Amin Nabian, Maximilian Stadler, Christian Hundt, Kamyar Azizzadenesheli, et al. Geometry-informed neural operator for large-scale 3d pdes. *Advances in Neural Information Processing Systems*, 36:35836–35854, 2023.
- Zongyi Li, Hongkai Zheng, Nikola Kovachki, David Jin, Haoxuan Chen, Burigede Liu, Kamyar Azizzadenesheli, and Anima Anandkumar. Physics-informed neural operator for learning partial differential equations. *ACM/IMS Journal of Data Science*, 1(3):1–27, 2024.
- Michael A Lieberman and Allan J Lichtenberg. Principles of Plasma Discharges and Materials Processing. *MRS Bulletin*, 30(12):899–901, 1994.
- Lu Lu, Pengzhan Jin, Guofei Pang, Zhongqiang Zhang, and George Em Karniadakis. Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators. *Nature machine intelligence*, 3(3):218–229, 2021.
- Richard H Macneal and Robert L Harder. A proposed standard set of problems to test finite element accuracy. *Finite elements in analysis and design*, 1(1):3–20, 1985.
- Stefano Massaroli, Michael Poli, Jinkyoo Park, Atsushi Yamashita, and Hajime Asama. Dissecting neural odes. *Advances in neural information processing systems*, 33:3952–3963, 2020.
- Vinod Nair and Geoffrey E Hinton. Rectified linear units improve restricted Boltzmann machines. In *Proceedings of the 27th international conference on machine learning (ICML-10)*, pp. 807–814, 2010.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. PyTorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.
- Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. Film: Visual reasoning with a general conditioning layer. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018.
- Md Ashiqur Rahman, Zachary E Ross, and Kamyar Azizzadenesheli. U-NO: U-shaped neural operators. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856. URL <https://openreview.net/forum?id=j3oQF9coJd>.

- Md Ashiqur Rahman, Robert Joseph George, Mogab Elleithy, Daniel Leibovici, Zongyi Li, Boris Bonev, Colin White, Julius Berner, Raymond A Yeh, Jean Kossaifi, et al. Pretraining codomain attention neural operators for solving multiphysics PDEs. *Advances in Neural Information Processing Systems*, 37:104035–104064, 2024.
- M Raissi, P Perdikaris, and GE Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.
- Michel Rappaz, Michel Bellet, Michel O Deville, and Ray Snyder. *Numerical modeling in materials science and engineering*, volume 20. Springer, 2003.
- Shahid Rauf, Kallol Bera, Jason Kenney, and Prashanth Kothnur. Modeling of plasma processing reactors: review and perspective. *Journal of Micro/Nanopatterning, Materials, and Metrology*, 22(4):041503–041503, 2023.
- Michał Jerzy Rewieński. *A trajectory piecewise-linear approach to model order reduction of nonlinear dynamical systems*. PhD thesis, Massachusetts Institute of Technology, Department of Electrical Engineering . . . , 2003.
- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pp. 234–241. Springer, 2015.
- Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. *Advances in neural information processing systems*, 33:7462–7473, 2020.
- Giovanni Stabile and Gianluigi Rozza. Finite volume POD-Galerkin stabilised reduced order methods for the parametrised incompressible Navier–Stokes equations. *Computers & Fluids*, 173: 273–284, 2018.
- Shashank Subramanian, Peter Harrington, Kurt Keutzer, Wahid Bhimji, Dmitriy Morozov, Michael W Mahoney, and Amir Gholami. Towards foundation models for scientific machine learning: Characterizing scaling and transfer behavior. *Advances in Neural Information Processing Systems*, 36:71242–71262, 2023.
- Floris Takens. Detecting strange attractors in turbulence. In *Dynamical Systems and Turbulence, Warwick 1980: proceedings of a symposium held at the University of Warwick 1979/80*, pp. 366–381. Springer, 2006.
- Alasdair Tran, Alexander Mathews, Lexing Xie, and Cheng Soon Ong. Factorized Fourier neural operators. In *The Eleventh International Conference on Learning Representations*, 2023.
- Gege Wen, Zongyi Li, Kamyar Azizzadenesheli, Anima Anandkumar, and Sally M Benson. U-fno—An enhanced Fourier neural operator-based deep-learning model for multiphase flow. *Advances in Water Resources*, 163:104180, 2022.
- Xiongye Xiao, Defu Cao, Ruochen Yang, Gaurav Gupta, Gengshuo Liu, Chenzhong Yin, Radu Balan, and Paul Bogdan. Coupled multiwavelet neural operator learning for coupled partial differential equations. In *International Conference on Learning Representations*, 2023.

A DESIGN SPACES OF FNOS

A.1 DESCRIPTION ON EACH COMPONENT

Data space

Ⓐ Separate single-channel format. Each variable is treated as an independent function, $a(x) = (a^\alpha(x), a^\beta(x))$, and discretized separately as $\mathbf{a}^\alpha, \mathbf{a}^\beta \in \mathbb{R}^m$. This representation keeps the two components distinct throughout the architecture.

Ⓑ Two-channel format. The input is represented as a vector-valued function, $a : \mathcal{X} \rightarrow \mathbb{R}^2$, $x \mapsto (a^\alpha(x), a^\beta(x))$, which in a discretized form becomes $\mathbf{a} = [\mathbf{a}^\alpha, \mathbf{a}^\beta] \in \mathbb{R}^{m \times 2}$. The output is represented analogously as $u(x) = [u^\alpha(x), u^\beta(x)]^\top \in \mathbb{R}^2$ and the discretized form $\mathbf{u} = [\mathbf{u}^\alpha, \mathbf{u}^\beta] \in \mathbb{R}^{m \times 2}$.

Lift operator, P

Ⓘ The lift operator $P : \mathbb{R}^1 \rightarrow \mathbb{R}^{d_v}$ and shared across the two input components, so that both variables are mapped into the latent space using the same transformation.

Ⓜ Alternatively, the lift operator can consist of two separate mappings, $P = (P^\alpha, P^\beta)$, each specific to one input component. This allows the two variables to have independent latent representations from the beginning.

Ⓝ Compatible with Ⓜ, the lift operator is defined as $P : \mathbb{R}^2 \rightarrow \mathbb{R}^{d_v}$, mapping a two-channel input to a hidden state.

Fourier layers

Ⓛ The point-wise linear map is a single operator shared by the both variables.

Ⓚ The point-wise linear map is defined as two separate operators (W^α, W^β).

Ⓒ For coupled systems, we define the global spectral convolution to perform coupling only in Fourier space. Specifically, each variable is first transformed independently, $\tilde{v}^\square(k) = \mathcal{F}v_\ell^\square(k)$, for $\square = \{\alpha, \beta\}$. Their spectral representations are then combined through a shallow encoder network $f^{\text{enc}}(\tilde{v}^\alpha(k), \tilde{v}^\beta(k))$, followed by the usual mode-wise kernel multiplication $\tilde{\tilde{v}}(k) = R_\phi(k)\tilde{v}(k)$. The result is decomposed into variable-specific coefficients using another shallow network, $(\tilde{\tilde{v}}^\alpha(k), \tilde{\tilde{v}}^\beta(k)) = f^{\text{dec}}(\tilde{\tilde{v}}(k))$, and finally mapped back to the data space via the inverse Fourier transform, $v_{\ell+1}^\square(k) = \mathcal{F}^{-1}(\tilde{\tilde{v}}^\square(k))$, for $\square = \alpha, \beta$.

Ⓓ The standard global spectral convolution performs the Fourier transform, mode-wise kernel multiplication, followed by the inverse Fourier transform: $\mathcal{F}^{-1}(R_\phi \cdot (\mathcal{F}v_\ell))(x)$.

Projection operator, Q

Ⓞ The projection operator is shared across both output components, $Q : \mathbb{R}^{d_v} \rightarrow \mathbb{R}^1$.

Ⓟ The projection operator is defined separately for each output component, $Q = (Q^\alpha, Q^\beta)$, with $Q^\square : \mathbb{R}^{d_v} \rightarrow \mathbb{R}^1$ for $\square \in \alpha, \beta$. From the adaptive-basis view perspective, Q can differ in whether the basis and the coefficients are shared or separate across the two variables:

- Ⓟ_a Shared basis and separate coefficients, and
- Ⓟ_b Separate basis and shared coefficients, and
- Ⓟ_c Separate basis and separate coefficients.

Ⓠ Similar to Ⓝ, this projection operator is defined as $Q : \mathbb{R}^{d_v} \rightarrow \mathbb{R}^2$, mapping a hidden state to a two-channel output.

A.2 SUMMARY OF THE COMPONENTS AND RESULTING MODELS

Table 4 summarizes the architectural components and configuration options discussed in Sections 4.1-4.2 and Appendix A.1, highlighting how shared or separate design choices can be applied to handle parameterized dynamics and multi-variable (coupled) systems.

Table 4: Summary of architectural components and configurations explored in FNO_x.

Component	Design choice / notation	Purpose and interpretation
Parameterization	Input augmentation (pFNO) / Hypernetwork-based modulation (hpFNO)	Introduces explicit dependence on the physical parameter μ . In the hypernetwork-based approach, a small network outputs layer-wise shifts $s_\ell(x, \mu)$ that modulate activations without regenerating full weights.
Lift operator, P	Shared (P_1) / Separate (P_2) across variables	Projects input fields to latent space. In multi-variable settings, a shared lift operator uses the same transformation for all variables, while a separate one allows variable-specific latent representations.
Point-wise linear map, W	Shared (L_1) / Separate (L_2) across variables	Controls local feature transformations in each Fourier layer. Sharing maintains parameter efficiency, while separating allows each variable to evolve differently in latent space.
Global spectral convolution, G	Standard (G_2) / Coupled (G)	Defines non-local interactions in Fourier space. The coupled variant mixes spectral representations of multiple variables via an encoder-decoder, enabling cross-variable interaction within a single operator.
Projection operator, Q	Shared (Q_1) / Separate basis and coefficients ($Q_{2a,b,c}$)	Maps latent states back to the physical output. In coupled systems, using separate bases and coefficients lets each variable form its own representation while improving expressivity.

Now we list realizations of FNOs obtained from some combinations of the above design choices.

- FNO_m: $\textcircled{2} + \textcircled{3} + \textcircled{1} + \textcircled{2} + \textcircled{3}$ - a two-channel input is processed by a standard FNO to produce a two-channel output.
- FNO_c: $\textcircled{1} + \textcircled{1} + \textcircled{1} + \textcircled{2} + \textcircled{1}$ - two single-channel variables are concatenated in a discrete representation, i.e., $[\mathbf{a}^{\alpha\top}, \mathbf{a}^{\beta\top}]^\top \in \mathbb{R}^{2m}$; the concatenated input is processed by a standard FNO to produce a single-channel vertically-concatenated output, i.e., $[\mathbf{u}^{\alpha\top}, \mathbf{u}^{\beta\top}]^\top \in \mathbb{R}^{2m}$.
- FNO_x: This is a class of the proposed models, which can be realized with the combinations of $\textcircled{1} + \textcircled{1}/\textcircled{2} + \textcircled{1}/\textcircled{2} + \textcircled{G} + \textcircled{1}/\textcircled{2}$. In the experimentation, with FNO_x, we refer to $\textcircled{1} + \textcircled{2} + \textcircled{2} + \textcircled{G} + \textcircled{2c}$ unless otherwise stated.

A.3 VISUAL COMPARISONS OF ARCHITECTURES

We provide a set of visual representations of the proposed architectures to illustrate the modifications made in this work and highlight the difference from the original FNO architecture (Figure 5).

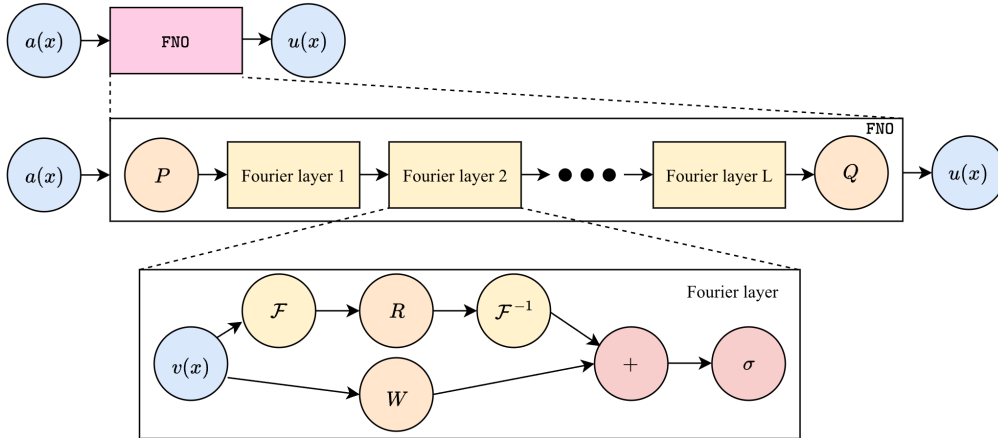


Figure 5: The original FNO architecture

Parametric extension In the first set of illustrations, we present the visual representations of the parametric extension of FNOs: pFNO and hpFNO. Figure 6 depicts the pFNO architecture, which

augments the physical parameters with the input function, which is then used as the input to the FNO layers.

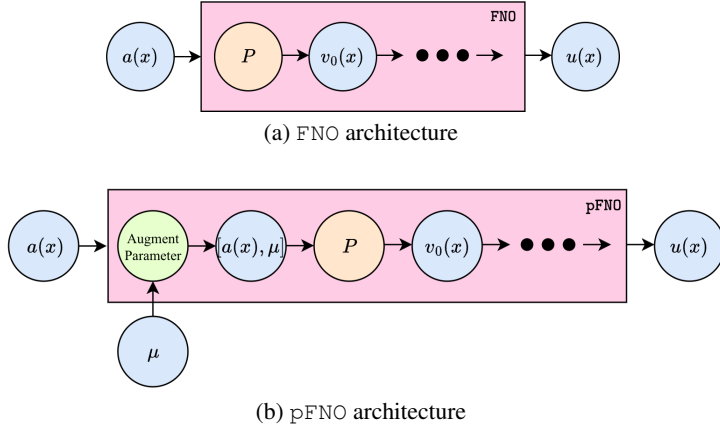


Figure 6: The abstract representations of (a) FNO and (b) pFNO architectures are depicted.

Figure 7 illustrates the additional computational flow from the physical parameter to the shift function $s(\mu, x)$, which then acts as the bias in the Fourier layer.




Figure 7: The proposed h_p FNO architecture. The hypernetwork produces a parameter-dependent ‘shift’ term, $s(\mu, x)$, which acts as a parameter-dependent bias term in each Fourier layer.

Coupling extension

The next set of visual representations explains how the proposed architecture models coupled physical behavior. Figure 8 illustrates the design spaces for the lift operator P and the projection operator Q . These two components determine how information is encoded and decoded across variables. In multi-variable systems, P controls whether the input fields share the same latent representation or retain variable-specific mappings, while Q defines how the learned latent features are projected back into each physical field. Figure 8b further illustrates three possible configurations of Q (i.e., shared basis and coefficients, shared basis with separate coefficients), and fully separate basis and coefficients. The light orange blocks indicate shared operators, whereas the darker orange blocks indicate variable-specific operators.

Figure 9 presents a graphical overview of the proposed coupled extension of the FNO architecture. The coupling mechanism is introduced through an encoder-decoder pair that mixes the spectral representations of multiple variables directly in Fourier space. The encoder aggregates information from the transformed variables, and the decoder distributes the coupled representation back to each channel, enabling cross-variable interaction without substantially increasing model size. This design preserves the efficiency of the original FNO while introducing controlled coupling through a single shared operator.

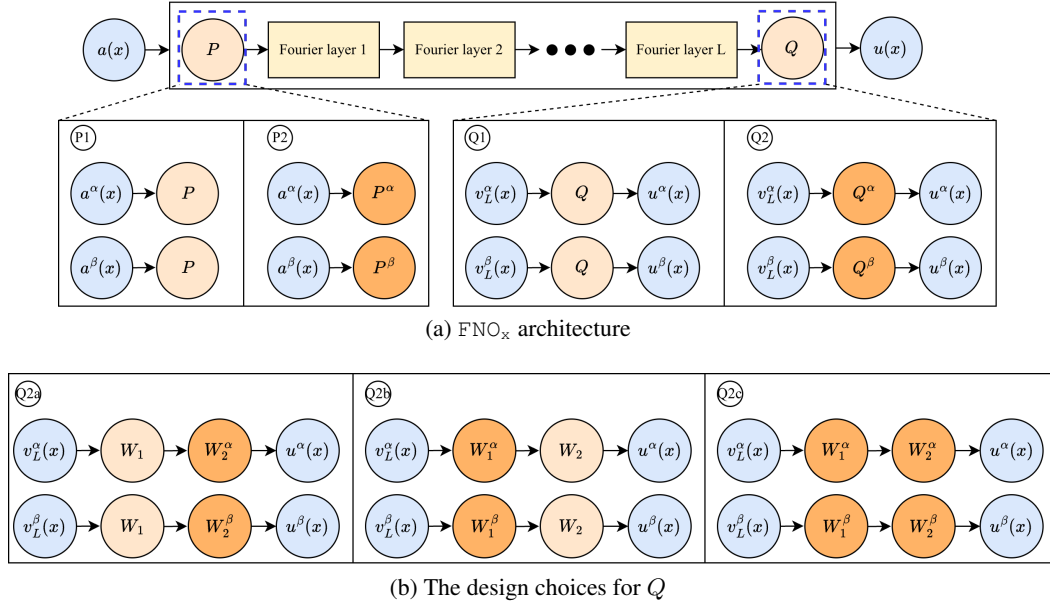



Figure 8: The standard and the proposed FNO architecture for modeling coupled behavior.

Figure 9: The proposed coupled extension of FNO architecture. The \oplus symbol indicates the concatenation operation.

B IMPLEMENTATION DETAILS

B.1 DESCRIPTIONS AND SPECIFICATIONS OF THE BASELINES AND THE PROPOSED MODELS

FNO_m / FNO_c / FNO_x For all FNO-based baselines, we consider the same hyper-parameter 4 Fourier layers with the modes and width ($k = 12$ and $d_v = 20$). P is parameterized as a linear layer and Q is parameterized as a shallow MLP with one hidden layer. All nonlinearity functions are ReLU (Nair & Hinton, 2010). This fixed hyperparameter setup is to control the representational capacity contributed by these hyperparameters. Consequently, any observed performance gains can be attributed to the architectural modifications introduced in each variant rather than to differences in model capacity.

For defining the FNO_x , we consider the combination, $\text{P1} + \text{Q2} + \text{Q2c}$, which employs the shared lift operator, the separate point-wise linear maps, the proposed coupled global spectral convolution layer, and the separate sets of basis functions and separate sets of coefficients in the projection operator. The encoder/decoder, $f^{\text{enc}}/f^{\text{dec}}$, in the Fourier layers are set to simple linear layers. This choice is obtained from the ablation study.

p FNO_c / p FNO_x These parameterized variant takes the standard input (i.e., the input to FNO_c and FNO_x), but augmented with the physical parameters, μ .

hpFNO_c / hpFNO_x These variants require one additional component, the hypernetwork, f^{hyper} . In our implementation, we consider a simple linear layer to parameterize the hypernetwork, following the design used in (Dupont et al., 2022), and we model the hypernetwork to take (x, μ) along with the discrete representation of the windowed history, $\{u(t - \tau)\}_{\tau=0}^{T_{\text{in}}-1}$. Thus, the hypernetwork becomes a local transform similar to the lift operator.

HyperFNO_c HyperFNO (Alesiani et al., 2022) integrates a base FNO with a hyper-network; the hyper-network takes the PDE parameter vector μ as input and generates FNO’s weights. For each Fourier layer ℓ , the hypernetwork generates a subset of the FNO parameters such as the pointwise linear map (W_ℓ), the spectral convolution weights (R_ℓ), and the lift/projection operators (P and Q), according to

$$W_\ell = h_W(\mu; \theta_W), \quad R_\ell = h_R(\mu; \theta_R), \quad P = h_P(\mu; \theta_P), \quad Q = h_Q(\mu; \theta_Q), \quad (7)$$

where $h_\square(\mu; \theta_\square)$ for $\square = \{W, R, P, Q\}$ denotes a hyper-network with its own model parameters θ_\square .

To control model size, two variants are proposed: the Addition form, which scales existing weights by μ -dependent factors, and the Taylor form, which further factorizes the modulation to reduce parameter count.

In the Addition form, the hypernetwork modulates the base FNO weights through low-rank parameterization:

$$\begin{aligned} h_R(\theta_R^\ell, \mu) &= R_0^\ell + \left(V_R^\ell \text{diag}(h_{R,\mu}(\mu)) V_R^{\ell\top} \right) \odot_{\text{row,col}} R_1^\ell, \\ h_W(\theta_W^\ell, \mu) &= W_0^\ell + (U_W^\ell h_{W,\mu}(\mu)) \odot_{\text{row}} W_1^\ell, \end{aligned} \quad (8)$$

where \odot_{row} and $\odot_{\text{row,col}}$ denote row-wise and element-wise modulation, respectively, and $h_{\square,\mu}(\mu)$ refers to a shallow MLP. In our experimentation, we use a two-layer MLP with width 128.

In the Taylor form, the modulation is expressed through a first-order expansion around the base parameters:

$$h_R(\theta_R^\ell, \mu) = R_0^\ell \odot \left(I + V_R^\ell \text{diag}(h_{R,\mu}(\mu)) V_R^{\ell\top} \right), \quad h_W(\theta_W^\ell, \mu) = W_0^\ell \odot \left(I + U_W^\ell h_{W,\mu}(\mu) \right), \quad (9)$$

where \odot denotes element-wise multiplication.

In the experiments described in the main body, we consider the Taylor variant; in the Appendix, we provide the results of both variants. To differentiate those two methods, we denote them by HyperFNO_c-T and HyperFNO_c-A.

MWT_c / CMWNO The multiwavelet-based operator learning framework (MWT) has been introduced as an approach to leverage multiwavelet expansions for differential equations. By decomposing solution operators into multiwavelet bases, MWT captures both global structures and localized variations in a flexible and efficient way, capable of handling multi-scale dynamics and heterogeneous patterns. The coupled extension of this has been introduced in the sequel (Xiao et al., 2023), which is named as MWT_c and is designed to take vertically concatenated input variables.

Extending this idea, the coupled multiwavelet neural operator (CMWNO) adapts the MWT framework to systems of coupled partial differential equations. CMWNO introduces a strategy that decouples the operator kernels across different wavelet components during decomposition and reconstruction. This enables the model to efficiently learn and represent the interactions between multiple coupled fields while keeping computation tractable. Computationally, CMWNO consists of two distinct operators that communicate by exchanging hidden representations. A randomized scheduling mechanism, referred to as the ‘dice’ algorithm, determines the order in which these operators share information, enabling flexible and efficient coordination between the coupled components.

For both MWT and CMWNO, we use the implementation available from the Github repository². The scaling parameter α is set to 12, the channel multiplier c is set to 16, the wavelet polynomial order c is set to 4, and the base polynomial basis is set to the ‘Legendre’ polynomials. The dice strategy

²<https://github.com/joshuaxiao98/CMWNO/>

is utilized as follows: for each batch we throw a random number from a uniform distribution $[0,1]$ and if the sampled random number is below 0.5, the second operator takes the first operator’s hidden representation through the roll-out. On the opposite case, the first operator takes the first operator’s hidden representation through the roll-out.

CFNO CFNO mimics the operations of CMWNO, i.e., two separate operators enabling coupling effect via sharing a hidden representation of one operator to another. For CFNO, the backbone is FNO instead of CFNO. Also, the same dice approach described above is utilized.

Same as other FNO-based baselines, we consider the same hyper-parameter 4 Fourier layers with the modes and width ($k = 12$ and $d_v = 20$). P is parameterized as a linear layer and Q is parameterized as a shallow MLP with one hidden layer. All nonlinearity functions are ReLU. For CFNOs, we have two separate realizations of FNOs and in the third layer of CFNOs, the hidden representation from one CFNO is passed to another CFNO, where the direction is dictated by the dice approach.

DeepONets For the implementation of DeepONets (DONs), we base our implementation on a publicly available PyTorch version.³ We modify the DONs to take multivariate input and produce multivariate output. We explore the combination of a (shared/separate) trunk net and a (shared/separate) branch net, and find that two separate branch networks and a shared trunk network perform the best. Each branch network takes the same input, which is concatenated time-windowed snapshots; then one branch network produces coefficients for the electron density and another branch network produces coefficients for the electric potential. That is,

$$\begin{aligned} n_e(x, t + 1) &= \sum_{i=1}^p f_i^{\text{branch}, n_e}(\bar{a}_{\leq t}) f_i^{\text{trunk}, n_e}(x), \\ \phi(x, t + 1) &= \sum_{i=1}^p f_i^{\text{branch}, \phi}(\bar{a}_{\leq t}) f_i^{\text{trunk}, \phi}(x) \end{aligned} \quad (10)$$

where f^{branch, n_e} and $f^{\text{branch}, \phi}$ denote two separate MLPs, taking time-windowed input

$$\bar{a}_{\leq t} = \begin{bmatrix} n_e(x, t) & \cdots & n_e(x, t - T_{\text{in}} + 1) \\ \phi(x, t) & \cdots & \phi(x, t - T_{\text{in}} + 1) \end{bmatrix}, \quad (11)$$

where T_{in} decides the window length. The trunk net is modeled as a separate network, producing two separate sets of basis such that

$$[f^{\text{trunk}, n_e}(x), f^{\text{trunk}, \phi}(x)] = [W^{n_e} f^{\text{trunk}}(x) + b^{n_e}, W^{\phi} f^{\text{trunk}}(x) + b^{n_{\phi}}], \quad (12)$$

where f^{trunk} denote a shared MLP, and (W^{n_e}, b^{n_e}) and $(W^{\phi}, b^{n_{\phi}})$ denote a set of model parameters specific to the electron density and electric potential, respectively.

For each branch networks, we use 4 hidden layers with 1024 neurons in each layer. For the trunk network, we use 3 hidden layers with 1024 neurons. For both network types, we consider ReLU for nonlinear activation.

U-Net For the implementation of U-Nets, we base our implementation on the code base we use for DeepONets. We modify the U-Net to take 1-dimensional and 2 channel inputs, and produce the output with the same dimensional specification. The U-Net architecture we consider has an encoder-decoder structure with skip connections.

The encoder progressively reduces the spatial resolution while increasing the feature dimension. It begins with an initial block (DoubleConv \rightarrow ReLU), followed by four downsampling stages (MaxPool1d \rightarrow DoubleConv \rightarrow ReLU), where each stage halves the sequence length and expands the channel dimension. The spatial resolution is halved at every downsampling and the channel dimension is updated (32 \rightarrow 64 \rightarrow 64 \rightarrow 128 \rightarrow 128).

The decoder mirrors this process with four upsampling stages (Upsample \rightarrow Concat (skip-connection) \rightarrow DoubleConv \rightarrow ReLU). At each stage, the feature map is upsampled by a factor of two, concatenated with the corresponding encoder feature (skip connection), and refined by a double convolution block. Finally, the output layer applies 1D convolution to project the features to the desired number of output channels.

³This is accessible through <https://github.com/camlab-ethz/ConvolutionalNeuralOperator>.

B.2 TRAINING DETAILS

All FNO variants and wavelet-based operators are trained for 500 epochs using the Adam optimizer (Kingma & Ba, 2015) with a learning rate of 0.0025, weight decay 0.0001, and mini-batches of size 10. For U-Nets, we consider the same hyperparameters, except for the total training epochs, which is set to 1,000, which is to mitigate the effect of the increased model sizes. For DeepONets, we again consider the same hyperparameters, except for the learning rate, which is set to 0.0001 and the total training epochs, which is set to 2,500 which is to compensate the decreased learning rate. For all baseline models, we repeat 5 different runs for varying random seeds. For all methods, we apply a linearly decaying temporal weighting to the rollout loss, emphasizing short-horizon prediction accuracy while progressively reducing the contribution of longer-horizon errors to mitigate compounding autoregressive effects.

B.3 ORTHOGONALITY LOSS

To measure the orthogonality of the learned basis functions, $\Psi(x) = [\psi_1(x), \dots, \psi_{n_q}(x)]$, we define a discrete inner product over the spatial domain. For a spatial domain D and a measure $\rho(x)$, the continuous inner product is $\langle \psi_i, \psi_j \rangle = \int_D \psi_i(x) \psi_j(x) d\rho(x)$. In practice, we approximate this integral using numerical quadrature. Specifically, we employ Gauss-type quadrature rules, which select quadrature nodes $x_k \in D$ and associated weights w_k , yielding $\langle \psi_i, \psi_j \rangle \approx \sum_{m=1}^M w_k \psi_i(x_k) \psi_j(x_k)$.

With this discrete formulation, we compute the Gram matrix, $[G]_{ij} = \langle \psi_i, \psi_j \rangle$, for $i, j = 1, \dots, n_q$ and quantify deviation from orthonormality via $\mathcal{L}_{\text{ortho}} = \|G - I_{n_q}\|_F^2$, leading to the ultimate loss $\mathcal{L} = \mathcal{L}_{\text{mse}} + \lambda_{\text{ortho}} \mathcal{L}_{\text{ortho}}$. This approach allows for a principled enforcement of orthonormality in the learned basis while ensuring consistency with the underlying continuous inner product and measure. For λ_{ortho} , we consider $\{0.1, 0.01, 0.001\}$.

C 1D CCP: DESCRIPTION

C.1 GOVERNING EQUATIONS

The simplified one-dimensional capacitively coupled plasma physics is described by two equations: electron continuity equation and Poisson equation. The electron continuity equation describes the conservation of particles; electron density changes in time because of electron flows out at the surface and because they can be created/destroyed locally due to the chemical reactions. The Poisson equation relates how the electrostatic potential varies in space to the distribution of charges (i.e., electrons and ions) in the plasma. The governing equations are defined as follows:

$$\begin{aligned} \frac{\partial n_e}{\partial t} &= -\frac{\partial \Gamma_e}{\partial x} + R, & \text{(Electron continuity equation)} \\ \partial_{xx} \phi &= \frac{e}{\epsilon_0} (n_e - n_{i0}), & \text{(Poisson equation)} \end{aligned} \quad (13)$$

where $n_e(x, t)$ and $\phi(x, t)$ denote electron density and electric potential, which are the solutions of the equations, and ∂_{\square} refers to a partial derivative with respect to \square . Γ_e is electron flux, defined as diffusion flux and drift flux such as

$$\Gamma_e = -D \partial_x n_e - \mu n_e \partial_x \phi \quad (14)$$

with electron diffusion coefficient D and electron mobility coefficient μ , which is defined as $D = eT_e/m_e \vartheta_m$. Here, $\mu = e/m_e \vartheta_m$. Here, e , T_e , m_e , and ϑ_m denote elementary charge, electron temperature, electron mass, and collision frequency. The ion density is defined as $n_{i0} = R_0(x_2 - x_1) \sqrt{m_i/eT_e}$, where $R(x)$ and R_0 denote reaction rate and coefficient, respectively. The reaction rate is a spatially dependent quantity such that

$$R(x) = \begin{cases} R_0, & x \in [x_1, x_2] \cup [L - x_2, L - x_1] \\ 0, & \text{otherwise} \end{cases} \quad (15)$$

The boundary conditions are specified as

$$\begin{aligned} n_e &= 0, \phi = 0, & \text{at } x = 0, \\ n_e &= 0, \phi = V(t) & \text{at } x = L, \end{aligned} \quad (16)$$

where $V(t) = V_0 \sin(2\pi t)$. That is, the zero boundary condition is given to the electron density (at $x = 0, L$). The electric potential satisfies a homogeneous Dirichlet condition (at $x = 0$), and a time-periodic Dirichlet condition at $x = L$. That is, the potential at the right boundary oscillates sinusoidally with amplitude V_0 and frequency f .

Table 5 lists up the important parameters, their description, and their values. Other important constants are elementary charge, $e = 1.6 \times 10^{-19}$ (C) and vacuum permittivity $\epsilon_0 = 8.854 \times 10^{-12}$ ($C^2 kg^{-1} m^{-3} s^2$).

Table 5: Input parameter symbols, descriptions, and values

Symbols	Descriptions	Values (unit)
L	Domain Length	0.025 (m)
x_1, x_2	Reaction area	0.005(m), 0.01(m)
f	Driving frequency	13.56 (MHz)
V_0	Driving voltage amplitude	[100, 300] (V)
R_0	Reaction raga coefficients	$[2.7 \times 10^{19}, 2.7 \times 10^{20}] (m^{-3} s^{-1})$
T_e	Electron temperature	3 (eV)
m_i	Ion mass	$[1.67 \times 10^{-26}, 6.68 \times 10^{-26}]$ (kg)
m_e	Electron mass	9.109×10^{-31} (kg)
ϑ_m	Collision frequency	$10^8 (s^{-1})$

C.2 VISUALIZATION OF THE SOLUTION SNAPSHOTS

To provide better understanding on the dynamics of the benchmark problem, we present some collections of solution snapshots for varying input parameters, reaction rate (Figure 10) and driving voltage (Figure 11) in the first RF cycle. All Figures presents the two spatio-temporal fields, the electron density $n_e(x, t)$ and the electric potential $\phi(x, t)$, on a $t - x$ plane, in a heatmap format.

Figure 10 shows the solution snapshots for varying reaction rates $R_0 = \{2.7 \times 10^{19}, 1.485 \times 10^{20}, 2.7 \times 10^{20}\} m^{-3} s^{-1}$. The reaction rate serves as a quantity that determines how fast ionization occurs. The reaction rate drives the plasma density and sustainment in a CCP discharge. The higher the reaction rates, the faster the ionization is, and the higher plasma density is. Reaction rate can be controlled using power in plasma processing systems. Figure 10 shows some example solution snapshots for varying reaction rate; from the left panel to the right panel, the reaction rate increases, which results in higher electron density and thinner sheaths.

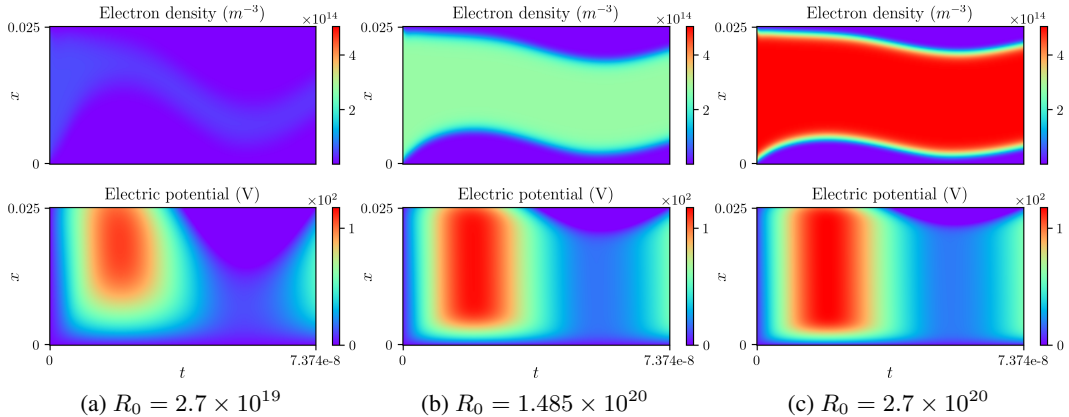


Figure 10: [1D CCP] Solution snapshots presented in heatmap; the solutions are collected for varying reaction rates $R_0 = \{2.7 \times 10^{19}, 1.485 \times 10^{20}, 2.7 \times 10^{20}\}$.

Figure 11 shows the solution snapshots for varying driving voltages $V_0 = \{100, 200, 300\}$ in the first RF cycle. The driving voltage, given as the boundary condition, establishes the time-varying electric potential fields. The higher the driving voltage, the larger is the sheath voltage, resulting in stronger electric fields. The sheath is thicker at higher electric field. Figure 11 shows example solution snapshots for varying driving voltage: (a) $V_0 = 100$, (b) $V_0 = 200$, and (c) $V_0 = 300$.




Figure 11: [1D CCP] Solution snapshots presented in heatmap; the solutions are collected for varying driving voltages $V_0 = \{100, 200, 300\}$.

Likewise, varying the ion-mass also provides variations in the dynamics. However, compared to other two physical parameters, the reaction rate and the driving voltage, changing the ion-mass parameter has slightly less physical impact on the dynamics and thus, the solution snapshots.

D 1D CCP: ADDITIONAL EXPERIMENTAL DETAILS AND RESULTS

D.1 1D CCP RESULTS

In the following sections, we present additional experimental results that are not presented in the main body due to the space limit. The additional experiments include an ablation study on different design choices of FNOs, a study on effectiveness of the proposed hypernetwork and modulation based architecture for the parametric extension, and a study on extrapolation in the parameter space.

D.1.1 ADDITIONAL RESULTS

Continuing from the experimental results presented in the main text (i.e., the “1-dimensional parameter space” experimentation), we present the further detailed information regarding the computations: the model sizes and the computational timings, depicted in Figure 12. This time, the information obtained from the experimentation with the datasets that vary the reaction rate and the ion masses. Again, the model size is measured in the number of trainable model parameters and the timing reports a per-epoch training time measured and averaged across the total number of epochs. We can make the similar observations here: the proposed methods FNO_x , pFNO_x , and hpFNO_x achieve improvement in accuracy without sacrificing the model size and the computational wall time. In case of the ion-mass scenario, we see relatively small improvements compared to other two scenarios (the reaction rates and the driving voltage). This is largely due to the fact that varying the ion-mass provide less changing physical dynamics and taking the windowed history input, $\{u(t - \tau)\}_{\tau=0}^{T_{\text{in}}-1}$, allows other non-parameterized NOs to infer those physical parameters to some extent. Still, the proposed methods FNO_x , pFNO_x , and hpFNO_x improves the performance in that scenario while minimally affecting the model size and the training time.

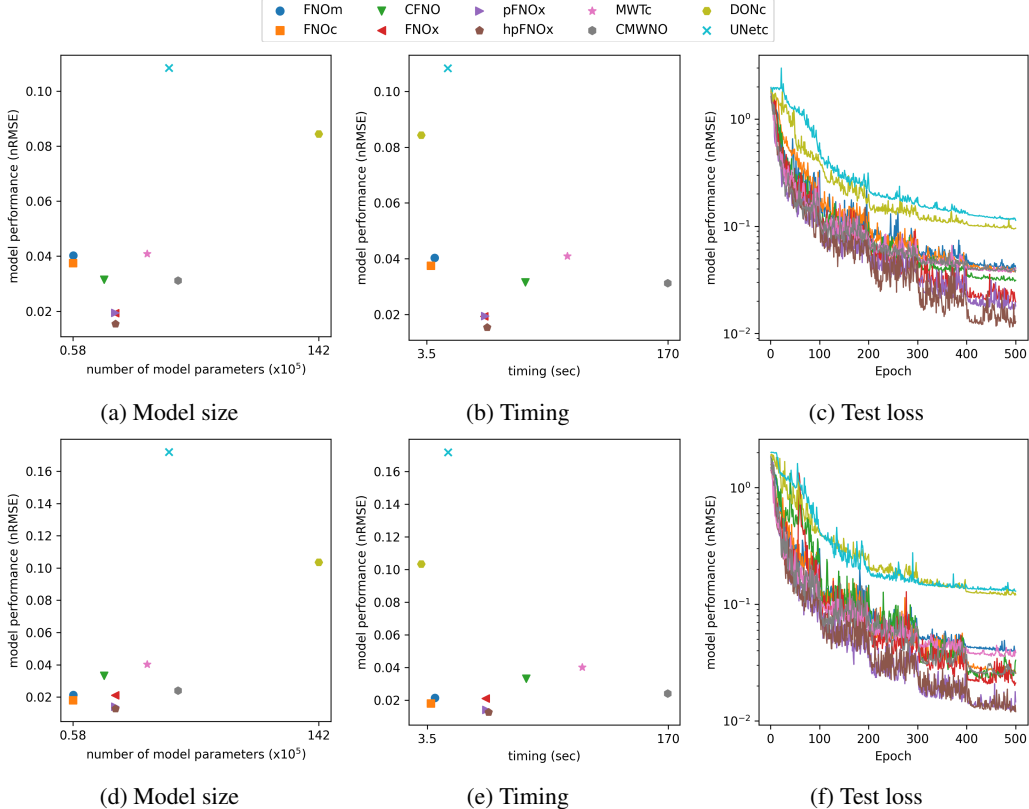


Figure 12: [1D CCP] Plots depict the number of model parameters versus the model performance (left), computational timing versus the model performance (right), for varying reaction rates R_0 (top) and ion masses m_i (bottom)

D.1.2 COUPLED EXTENSION

Ablation on model architectures To investigate the effectiveness of each design component, we perform an ablation study on several different combinations of the design choices. We begin with the *default* combination, $\textcircled{P1} + \textcircled{L1} + \textcircled{Q1}$, which consists of the shared lift/projection operators and the point-wise linear map in the Fourier layers. We first consider replacing the point-wise linear map in the Fourier layer to be a set of separate point-wise linear maps, W^α and W^β . After that, we investigate the projection operator by testing three different combinations: Shared/separate basis functions Ψ and shared/separate coefficients Ξ . Finally, we test the effect of a set of separate lift operators. Table 6 reports the results of this ablation study. The normalization shown in the last

Table 6: Performance comparisons of different realizations of FNO_x . The models are trained for 1-dimensional parameter space, (reaction rate). The performance is measured in the relative ℓ^2 -error (nRMSE); the reported numerical values refer to the mean (\pm std. dev).

Variant	Description	Relative errors
$\textcircled{P1} + \textcircled{L1} + \textcircled{Q1}$	Default setup	0.0341 (\pm 0.0044)
$\textcircled{P1} + \textcircled{L2} + \textcircled{Q1}$	Separate W	0.0281 (\pm 0.0048)
$\textcircled{P1} + \textcircled{L2} + \textcircled{Q2a}$	Separate W and coefficients Ξ	0.0259 (\pm 0.0032)
$\textcircled{P1} + \textcircled{L2} + \textcircled{Q2b}$	Separate W and basis Ψ	0.0315 (\pm 0.0028)
$\textcircled{P1} + \textcircled{L2} + \textcircled{Q2c}$	Separate W and basis Ψ and coefficients Ξ	0.0275 (\pm 0.0062)
$\textcircled{P1} + \textcircled{L2} + \textcircled{Q2b}$ w. norm	Separate W and coefficients Ξ with layer norm	0.0205 (\pm 0.0043)
$\textcircled{P1} + \textcircled{L2} + \textcircled{Q2c}$ w. norm	Separate W and basis Ψ and coefficients Ξ with layer norm	0.0193 (\pm 0.0059)
$\textcircled{P2} + \textcircled{L2} + \textcircled{Q2c}$ w. norm	Separate P and W and basis Ψ and coefficients Ξ with layer norm	0.0204 (\pm 0.0035)




Figure 13: [1D CCP] (a) The model performance measured in nRMSE across varying $\nu_m = \{1, 2, 4\} \times 10^8$. (b) and (c): Solution snapshots collected for varying $\nu_m = \{2 \times 10^8, 4 \times 10^8\}$ while keeping all other parameter values.

entries of the table is layer norm applied to the pre-activation values in the projection operator (i.e., $W_1 v_T(x) + b_1$). From this ablation study, we choose the combination, $\textcircled{1} + \textcircled{2} + \textcircled{2c}$ with layer norm, as FNO_x for all experimentation.

Varying collisional coupling To examine the effect of coupling strength in the 1D CCP system, we vary the electron–neutral collision frequency parameter ν_m across 1×10^8 , 2×10^8 , and 4×10^8 . As summarized in Figure 13a, we observe a general trend of increasing nRMSE with higher ν_m across all models. While the proposed hpFNO_x exhibits a mild degradation at the largest ν_m , it still achieves the lowest overall errors (nRMSE ≈ 0.015 – 0.018).

This behavior is consistent with the underlying plasma physics: as the system becomes more collisional, electron transport is increasingly suppressed. In particular, increasing ν_m strengthens the coupling between charged species and the background gas through enhanced momentum exchange, while simultaneously reducing both the diffusion coefficient $D = eT_e / (m_e \nu_m)$ and the mobility $\mu = e / (m_e \nu_m)$. The reduced transport leads to stiffer solution snapshots with sharper spatial gradients and weaker temporal smoothness (Figures 13b–13c). Consequently, the resulting fields exhibit steeper localized structures that are more challenging for spectral operator networks to approximate.

Taken together, these results suggest that higher-collisional regimes introduce both stronger coupling and increased stiffness, leading to moderately higher errors. We note, however, that it is difficult to fully isolate the influence of coupling strength alone, since changes in ν_m simultaneously affect both the degree of coupling and the stiffness of the governing dynamics, thereby altering the learning difficulty and the numerical conditioning of the operator.

D.1.3 PARAMETRIC EXTENSION

Study on the effectiveness of the proposed model architecture To further investigate the effectiveness of the proposed parametric extension, we compare the performances of the non-parameterized versions, FNO_c and FNO_x against their own two parameterized versions, (pFNO_c, hpFNO_c) and (pFNO_x, hpFNO_x).

Table 7 reports the models’ performance measured in all three physical parameter scenarios. Taking the non-parameterized version as the baseline, Table 7 reports the achieved prediction accuracy of each model and their improvements over the non-parameterized versions, presented in the percentage. The results indicate that simply augmenting physical parameters to input provide some improvements in prediction accuracy. However, the new architectural modification (i.e., hyper-network and modulation) brings significantly improved results (upto 68.01% improvement). As a reference point, we also report the performance of two variants of HyperFNO_c–T and –A (Taylor and Addition) as well. The overall performance of HyperFNO_c are in-between those of pFNO_c and hpFNO_c, suggesting that the parameter-efficient modulation technique (i.e., latent shift modulation) provides benefits in the given experimental contexts.

Study on the out-of-distribution samples In the main experiments, we follow the standard FNO testing strategy. In our parameterized scenarios, however, we can explicitly sample parameters that

Table 7: Performance comparisons of non-parameterized (FNO_c , FNO_x) and parameterized ($pFNO_c$, $hpFNO_c$, $pFNO_x$, $hpFNO_x$) models. All models are tested on three scalar parameters (reaction rate, driving voltage, and ion mass) individually. The performance is measured in the relative ℓ^2 -error; the reported numerical values refer to the mean (\pm standard deviation, improvement over non-parameterized version).

Model	Reaction rate	Driving voltage	Ion mass
FNO_c	0.0375 (\pm 0.0055)	0.0873 (\pm 0.0200)	0.0299 (\pm 0.0030)
$pFNO_c$	0.0334 (\pm 0.0101, 10.94 %)	0.0454 (\pm 0.0094, 47.92 %)	0.0284 (\pm 0.0023, 5.08 %)
$hpFNO_c$	0.0196 (\pm 0.0021, 47.57 %)	0.0279 (\pm 0.0016, 68.01 %)	0.0210 (\pm 0.0036, 29.84 %)
FNO_x	0.0193 (\pm 0.0059)	0.0345 (\pm 0.0108)	0.0212 (\pm 0.0062)
$pFNO_x$	0.0194 (\pm 0.0075, -0.51 %)	0.0278 (\pm 0.0053, 19.30 %)	0.0142 (\pm 0.0021, 32.91 %)
$hpFNO_x$	0.0154 (\pm 0.0029, 20.12 %)	0.0192 (\pm 0.0040, 44.40 %)	0.0128 (\pm 0.0017, 39.83 %)
Hyper FNO_c -T	0.0278 (\pm 0.0117)	0.0355 (\pm 0.0160)	0.0253 (\pm 0.0088)
Hyper FNO_c -A	0.0276 (\pm 0.0119)	0.0401 (\pm 0.0143)	0.0250 (\pm 0.0098)

lie entirely outside the training range, which we define as out-of-distribution (OOD) testing. As a benchmark case, we vary the reaction rate R_0 . For training, 101 equidistance values of the reaction rate is sampled from $[2.7 \times 10^{19}, 2.7 \times 10^{20}]$; for OOD testing, the interval is expanded outward by 0.2×10^{19} in both directions. Table 8 reports the performance of all considered models on ID test samples and OOD test samples. Overall, the performance measured on OOD produces increased numbers, but the trend across the compared methods remains the same. The proposed FNO_x model family produce more accurate predictions than the existing methods.

Figure 14 depicts the per-parameter nRMSE measured for the reaction coefficients that are sampled in the OOD region. As the index increases (from 1 to 10, sampled uniformly from $[2.7 \times 10^{20}, 2.9 \times 10^{20}]$), the parameter is farther away from the training distribution. All considered methods show a gradual increase in nRMSE as the chosen parameter departs from the training range. The parameterized versions ($pFNO_c$, $hpFNO_c$) and ($pFNO_x$, $hpFNO_x$) achieve lower errors compared to their non-parameterized counterparts, indicating improved extrapolation capability. This trend is expected for data-driven surrogate models, and the smooth growth of nRMSE suggests that the proposed parameterized extensions remain stable even when evaluated beyond the training domain. We note that this analysis is intended to investigate general extrapolation behavior rather than to claim superior OOD performance compared to large-scale or foundation models (e.g., (Rahman et al., 2024; Herde et al., 2024; Holzschuh et al., 2025)).

Table 8: Performance comparisons of non-parameterized (FNO_c , FNO_x) and parameterized ($pFNO_c$, $hpFNO_c$, $pFNO_x$, $hpFNO_x$) models for OOD samples (reaction rate). The performance is measured in the relative ℓ^2 -error; the reported numerical values refer to the mean (\pm standard deviation, improvement over non-parameterized version).

Model	ID	OOD
FNO_m	0.0403 (\pm 0.0012)	0.0628 (\pm 0.0075)
FNO_c	0.0375 (\pm 0.0055)	0.0673 (\pm 0.0097)
$pFNO_c$	0.0334 (\pm 0.0101)	0.0498 (\pm 0.0069)
$hpFNO_c$	0.0196 (\pm 0.0021)	0.0374 (\pm 0.0051)
FNO_x	0.0193 (\pm 0.0059)	0.0446 (\pm 0.0094)
$pFNO_x$	0.0194 (\pm 0.0075)	0.0343 (\pm 0.0059)
$hpFNO_x$	0.0154 (\pm 0.0029)	0.0303 (\pm 0.0050)

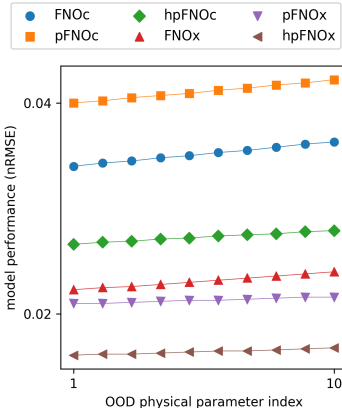


Figure 14: nRMSE over the physical parameter indices.

2-dimensional parameter space Next we extend our investigation to 2D input parameter space, a parameter space spanned by the reaction rate and the driving voltage. For input parameter sampling, we consider the same ranges for each parameter $V_0 \in [100, 300]$, $R_0 \in [2.7 \times 10^{19}, 2.7 \times 10^{20}]$, construct a uniform 2D mesh grid (21×21), and randomly split them into training and test sets with 9:1 ratio. We can make observations that are similar with the ones from the 1D input parameter cases; the proposed architecture FNO_x provides improvement in prediction accuracy and parameterized versions of FNO_c and FNO_x further reduce the prediction errors, leading to the best performance with hpFNO_x . The model size and the computational timing follow the similar trends with the 1D case, only except the slightly increased per-epoch computation time due to the increased number of data instances in the dataset for all methods.

Table 9: Performance comparisons for the 2D input parameter case.

Model	Performance
FNO_m	0.0285 (± 0.0028)
FNO_c	0.0252 (± 0.0026)
CFNO	0.0223 (± 0.0023)
MWT_c	0.0305 (± 0.0028)
CMWNO	0.0298 (± 0.0023)
DON_c	0.0409 (± 0.0077)
U-Net _c	0.0461 (± 0.0160)
FNO_x	0.0163 (± 0.0009)
pFNO_x	0.0140 (± 0.0012)
hpFNO_x	0.0124 (± 0.0011)

Adaptive basis view point In the next set of experiments, we investigate further on the adaptive basis viewpoint given on the last projection operator Q . We first vary the width (that is, the number of basis functions) of the last layer of Q for $n_q = \{8, 16, 32, 64, 128\}$ and measure the prediction performance of hpFNO_x on the varying reaction rate scenario. Taking the $d_{\text{basis}} = 8$ as a baseline, Figure 15 reports the percentage of improvements (shown in the bottom filled markers). Next, we test the idea of having (close to) orthogonal basis functions and repeat the same experiments. An orthonormality condition imposed to the learned basis, $\Psi(a(x))$; we could achieve further improvement (shown in the top filled markers). As d_{basis} becomes small, the effect of the orthogonality enforcement becomes higher, leading to $\sim 10\%$ improvements. We define the orthogonality of the learned basis functions and its computation in Appendix B.3.




Figure 15: Num. of bases v.s. performance improvement.

Study on hyper-network configurations To assess the impact of hypernetwork design, we test several configurations varying in depth and width. Specifically, we consider hypernetworks with 1 and 2 hidden layers (width 256) and 1 hidden layer (width 128), which yield nRMSE values in the range of 0.0171-0.0179 compared to 0.0154 for the linear baseline. Deeper or narrower configurations, such as a 3-layer network (width 256, nRMSE = 0.0200) or 3-layer network (width 32, nRMSE = 0.0230), show a mild degradation in performance. These results indicate that hpFNO_x performance is relatively insensitive to moderate variations in hypernetwork architecture, while deeper or overly compact designs may introduce minor instability. The linear configuration remains the most efficient and accurate choice.

Comparison of modulation techniques In addition to the proposed shift (bias) modulation, we also evaluate variants that modulate both the layer weights and biases based on the physical parameters. We denote this approach by `weight` modulation. The first variant is the weight modulation following the formulation suggested in the HyperFNO paper:

$$h_W(\theta_W^\ell, \mu) = W_0^\ell + (U_W^\ell h_{W,\mu}(\mu)) \odot_{\text{row}} W_1^\ell, \quad (17)$$

where $h_{W,\mu}(\mu)$ denotes the MLP-encoded parameter vector. This configuration allows both $W_\ell(\mu)$ and $s_\ell(x, \mu)$ to vary with the physical parameter μ .

For the next variant, we augment the additive shift modulation with a multiplicative gating mechanism, which is in a similar spirit with Feature-wise linear modulation (FiLM) (Perez et al., 2018). We denote this approach by `scale` modulation. With this augmentation, the Fourier layer becomes:

$$v_{\ell+1}(x) = \sigma\left(\gamma_\ell(x, \mu) \left[W v_\ell(x) + (\mathcal{K}(a; \phi) v_\ell)(x) \right] + s_\ell(x, \mu) \right), \quad \forall x \in D. \quad (18)$$

where the scale factor, $\gamma(x, \mu) = [\gamma_1(x, \mu), \dots, \gamma_L(x, \mu)]$, is inferred from a hypernetwork, along with the shift $s(x, \mu)$.

Table 10: Performance comparisons of different modulation techniques applied to hpFNO_c.

	Reaction term	Driving voltage	Ion mass
Shift only	0.0154 (\pm 0.0029)	0.0279 (\pm 0.0016)	0.0128 (\pm 0.0017)
Weight + shift	0.0157 (\pm 0.0028)	0.0214 (\pm 0.0056)	0.0134 (\pm 0.0036)
Scale + shift	0.5037 (\pm 0.0144)	0.0279 (\pm 0.0636)	0.8550 (\pm 1.8748)
Scale-0.1 + shift	0.0151 (\pm 0.0036)	0.0170 (\pm 0.0029)	0.0137 (\pm 0.0059)
Scale-0.5 + shift	0.0146 (\pm 0.0022)	0.0180 (\pm 0.0031)	0.0157 (\pm 0.0053)

Compared to using a shift term alone, the inclusion of $\gamma_\ell(x, \mu)$ makes the modulation strictly more expressive, as it enables the network to multiplicatively amplify or suppress the latent responses in a parameter-dependent manner. However, multiplicative gating can introduce training instabilities if γ_ℓ varies too aggressively across the parameter space. To mitigate this, we employ a stabilization strategy by parameterizing the gate as

$$\gamma_\ell(x, \mu) = 1 + \eta \tanh(\tilde{\gamma}_\ell(x, \mu)), \quad (19)$$

which keeps the modulation centered near identity at initialization while still allowing smooth, learnable deviations driven by the physical parameter. Here, $\tilde{\gamma}$ becomes the hypernetwork output and η is the hyper-parameter to decide the amplitude of deviation. We denote this variant by `scale- η` , where the suffix indicates the chosen value of η .

Table 10 provides summarized results. Overall, the amplitude-based multiplicative gating generally offers performance improvements, but it introduces several considerations. First, naive multiplicative modulation is unstable, since the unconstrained scaling factors can deviate excessively and cause large errors during training. Second, the stabilized formulation proposed in this work, which applies a scaled tanh function to control the amplitude of the gate, significantly improves robustness but requires searching over the hyperparameter η to determine the appropriate scaling range. Third, the benefits of multiplicative gating are not universal across all settings; for example, in the ion mass benchmark, the configuration with the shift modulation alone achieves the best accuracy. These observations indicate that although multiplicative gating is more expressive in principle and can yield improved performance, its practical effectiveness depends on careful amplitude control and the characteristics of the specific parameterized problems.

In summary, a reasonably strong and stable parameterization can be obtained using the shift modulation alone, which delivers reliable performance across the benchmarks considered in this work. However, the amplitude based multiplicative gating provides additional expressive capacity, and with a more refined search over its scaling hyperparameters, there is room for achieving higher accuracy. This suggests a practical tradeoff: the shift modulation is simple and robust, whereas the amplitude based modulation can offer further performance gains when carefully tuned.

E GRAY-SCOTT EQUATIONS

The Gray-Scott equations form a system of coupled reaction-diffusion equations that describe the spatiotemporal dynamics of interacting chemical species. This model can reproduce a rich variety of natural patterns, including structure resembling bacterial colonies, spiral waves, and coral-like formations. In the system, two state variables u and v each undergo independent diffusion and linear growth/decay, while their interaction is governed by the nonlinear reaction term.

The governing equations are defined as follows:

$$\begin{aligned} \frac{\partial u}{\partial t} &= \epsilon_1 \frac{\partial^2 u}{\partial x^2} + F(1 - u) - \lambda_1 uv^2, \\ \frac{\partial v}{\partial t} &= \epsilon_2 \frac{\partial^2 v}{\partial x^2} - (K + F)v + \lambda_2 uv^2, \end{aligned} \quad (20)$$

where $u(x, t)$ denotes concentration of the feed species, which is continuously supplied to the system and $v(x, t)$ denotes concentration of the reactant species, which is produced and consumed through interaction with u . The diffusion coefficients of u and v are denoted as ϵ_1 and ϵ_2 , respectively. Other

parameters include: F , the feed rate, representing the rate at which species u is introduced into the system from an external source, K , the decay rate of species v , which accounts for natural removal, and λ , the reaction rate constant associated with the nonlinear term $-\lambda uv^2$, which describes the autocatalytic process where one unit of u reacts with two units of v .

E.1 EXPERIMENTATION

In the experiment, we consider the parameterized dynamics realized by varying the diffusion coefficient $\epsilon_1 \in [0.1, 10]$ and varying the feed rate $F \in [0.1, 10]$. We collect 101 samples of equi-distanced coefficient/feed rate values while the other parameters are fixed as $\epsilon_2 = 0.05$, $K = 0.1$ and $(\lambda_1, \lambda_2) = (2, 5)$. For sampling an initial condition, we follow the approach described in (Xiao et al., 2023), sampling $u(x, 0)$ using the smooth random functions in the Chebfun package (Driscoll et al., 2014) and sampling $v(x, 0)$ from Gaussian random field, $v(x, 0) \sim \mathcal{N}(0, 7^4(-2\pi + 7^2 I)^{-2.75})$. For collecting trajectories for constructing training/test datasets, initial value problems are numerically solved using a fourth-order time-integrator, exponential time differencing fourth-order Runge–Kutta (ETDRK4), utilizing the chebfun software. We base the implementation the publicly available Matlab script⁴, which is the implementation accompanied by the research work (Xiao et al., 2023). The IVPs are solved with a spatial resolution, 1024, and then the collected solution snapshots are subsampled in the spatial domain, leading to the final spatial resolution, 128. The time integration is performed for the time interval $[0, 0.1]$ with 31 time steps.

Table 11: Performance comparisons between the proposed methods and the baseline methods. All models are tested on a scalar parameter: the diffusion coefficient. The performance is measured in the relative ℓ^2 -error; the reported numerical values refer to the mean (\pm standard deviation)

Model	Diffusion coeff, ϵ_1	Feed rate, F
FNO _m	0.0129 (\pm 0.0032)	0.0116 (\pm 0.0026)
FNO _c	0.0089 (\pm 0.0052)	0.0097 (\pm 0.0030)
CFNO	0.0093 (\pm 0.0016)	0.0161 (\pm 0.0045)
MWT _c	0.0055 (\pm 0.0011)	0.0092 (\pm 0.0021)
CWMNO	0.0048 (\pm 0.0015)	0.0293 (\pm 0.0078)
DON _c	0.0243 (\pm 0.0029)	0.0138 (\pm 0.0020)
U-Net _c	0.0239 (\pm 0.0160)	0.0159 (\pm 0.0071)
FNO _x	0.0039 (\pm 0.0002)	0.0075 (\pm 0.0016)
pFNO _x	0.0027 (\pm 0.0008)	0.0041 (\pm 0.0010)
hpFNO _x	0.0022 (\pm 0.0010)	0.0022 (\pm 0.0006)

Prediction accuracy Table 11 reports the prediction accuracy of all considered method. The results provide similar observations with those of the 1D CCP experimentation; the proposed methods, the parameterized variants of FNO_x achieve the lowest errors compared to existing baselines with 54% (for the diffusion coefficient case) and 72% (for the feed rate case) improvements. The coupled extension of FNO, FNO_x, improves the performance compared to other baselines. This is achieved without significantly scarifying the training/inference time and the memory footprint compared to FNO_m or FNO_c while other baselines require significantly larger-sized models and/or computational times. In the varying feed rate scenario, the approaches leveraging two separate operators struggle further to capture the accurate dynamics, resulting in higher prediction errors.

Computational/memory aspects Figure 16 further provides more information on computing. The figure presents the model sizes (i.e., memory requirements) and the computational timing for per epoch operation measured in seconds. The first row presents the results for the diffusion coefficients and the bottom row presents the results for the feed rate. Again the proposed variants pFNO_x and hpFNO_x serve as an efficient (in both memory and computation) and effective surrogate for modeling the parameterized Gray–Scott equations. Along with these information, the third column

⁴This script is available in <https://github.com/joshuaxiao98/CMWNO/>.

of Figure 16 presents the test loss achieved over the course of the training. The proposed $hpFNO_x$ demonstrates faster convergence to more accurate predictions.

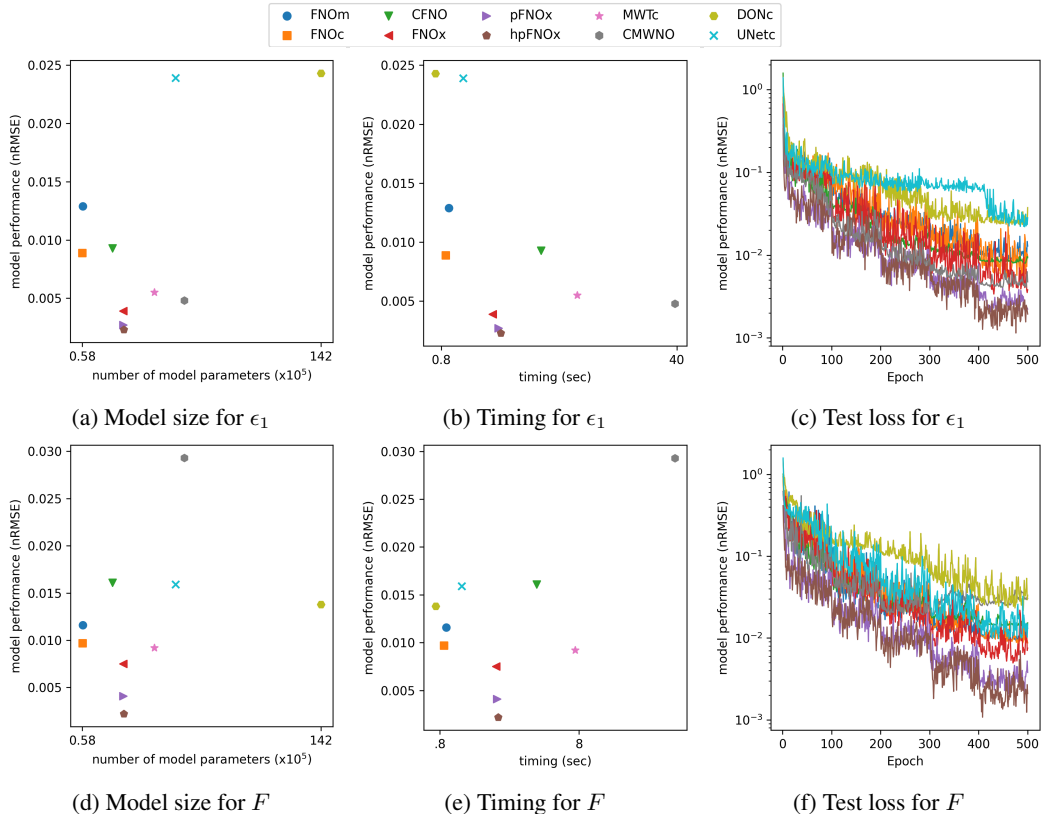


Figure 16: [GS] The results of the experiments for varying diffusion coefficients (top) and varying feed rates (bottom). Plots depict the number of model parameters versus the model performance (left), computational timing versus the model performance (middle), and the test loss trajectories over epochs (right).

Varying spatial resolutions We further conduct additional experiments on the Gray–Scott benchmark of varying diffusion coefficients using datasets generated at three different spatial resolutions: 64, 128, and 256 grid points. All models are trained separately on data from each resolution under identical training configurations and evaluated on corresponding test sets. The results, summarized in Figure 17, show that all considered methods maintain comparable accuracy across resolutions, consistent with the resolution-invariant nature of neural operator formulations. Among them, FNO_x consistently achieves the lowest relative ℓ^2 -errors (approximately 3.3×10^{-3} across all grid sizes). These results confirm that the proposed architecture preserves accuracy under varying spatial discretizations while maintaining the efficiency of operator learning.




Figure 17: [GS] nRMSE over varying resolutions

F 2D CHEMICALLY REACTING FLOW

As a next benchmark, we consider the two-dimensional chemically reacting flow of a premixed H_2 -air flame at constant and uniform pressure Buffoni & Willcox (2010), which is modeled as a coupled system of four PDEs. Each equation governs the temporal and spatial evolution of one of the four thermo-chemical state variables, temperature and the mass fractions of H_2 , O_2 , and H_2O , which are nonlinearly coupled through diffusion, convection, and reaction terms. The governing equations for

the system is expressed as

$$\frac{\partial \mathbf{w}(\mathbf{x}, t; \boldsymbol{\mu})}{\partial t} = \nabla \cdot (\nu \nabla \mathbf{w}(\mathbf{x}, t; \boldsymbol{\mu})) - \mathbf{v} \cdot \nabla \mathbf{w}(\mathbf{x}, t; \boldsymbol{\mu}) + \mathbf{q}(\mathbf{w}(\mathbf{x}, t; \boldsymbol{\mu}); \boldsymbol{\mu}), \quad (21)$$

where $\mathbf{w} = [w_T, w_{\text{H}_2}, w_{\text{O}_2}, w_{\text{H}_2\text{O}}]^\top$ represents the state vector of the four coupled fields.

On the right-hand side, the first term captures molecular diffusion with diffusivity $\nu = 2 \text{ cm}^2/\text{s}$, the second term represents advection by a constant, divergence-free velocity field $\mathbf{v} = [50 \text{ cm/s}, 0]^\top$, and the third term \mathbf{q} accounts for nonlinear chemical reactions coupling all species. The system is therefore strongly coupled: the evolution of each variable depends on the others through \mathbf{q} and, to a lesser degree, through diffusive transport.

The reaction source term follows an Arrhenius-type law:

$$\mathbf{q}(\mathbf{w}; \boldsymbol{\mu}) = [q_T(\mathbf{w}; \boldsymbol{\mu}), q_{\text{H}_2}(\mathbf{w}; \boldsymbol{\mu}), q_{\text{O}_2}(\mathbf{w}; \boldsymbol{\mu}), q_{\text{H}_2\text{O}}(\mathbf{w}; \boldsymbol{\mu})]^\top, \quad (22)$$

where

$$\begin{aligned} q_T(\mathbf{w}; \boldsymbol{\mu}) &= Q q_{\text{H}_2\text{O}}(\mathbf{w}; \boldsymbol{\mu}), \\ q_i(\mathbf{w}; \boldsymbol{\mu}) &= -v_i \left(\frac{W_i}{\rho} \right) \left(\frac{\rho w_{\text{H}_2}}{W_{\text{H}_2}} \right)^{v_{\text{H}_2}} \left(\frac{\rho w_{\text{O}_2}}{W_{\text{O}_2}} \right)^{v_{\text{O}_2}} A e^{-\frac{E}{R w_T}}, \end{aligned} \quad (23)$$

for $i \in \{\text{H}_2, \text{O}_2, \text{H}_2\text{O}\}$.

The stoichiometric coefficients are $(v_{\text{H}_2}, v_{\text{O}_2}, v_{\text{H}_2\text{O}}) = (2, 1, -2)$, and the molecular weights are $(W_{\text{H}_2}, W_{\text{O}_2}, W_{\text{H}_2\text{O}}) = (2.016, 31.9, 18) \text{ g/mol}$. The mixture density is $\rho = 1.39 \times 10^{-3} \text{ g/cm}^3$, the universal gas constant is $R = 8.314 \text{ J/(mol}\cdot\text{K)}$, and the heat of reaction is $Q = 9800 \text{ K}$.

The system depends on two physical parameters, $\boldsymbol{\mu} = (\mu_1, \mu_2) = (A, E)$, corresponding to the pre-exponential factor and activation energy, respectively. Hence, this problem represents a four-variable coupled nonlinear PDE system with two input parameters, capturing the essential thermo-chemical coupling of a reacting H_2 -air mixture.

Figure 18 illustrates the computational geometry and the corresponding boundary conditions, defined as follows:

- Γ_2 : inflow boundary with Dirichlet conditions specified by $T = 950 \text{ K}$ and $(w_{\text{H}_2}, w_{\text{O}_2}, w_{\text{H}_2\text{O}}) = (0.0282, 0.2259, 0)$,
- Γ_1 and Γ_3 : boundaries with Dirichlet conditions $T = 300 \text{ K}$ and $(w_{\text{H}_2}, w_{\text{O}_2}, w_{\text{H}_2\text{O}}) = (0, 0, 0)$,
- Γ_4, Γ_5 , and Γ_6 : boundaries with homogeneous Neumann conditions,




Figure 18: [ADR] The spatial domain and boundary conditions.

and the initial condition is set to $T = 300 \text{ K}$ and $(w_{\text{H}_2}, w_{\text{O}_2}, w_{\text{H}_2\text{O}}) = (0, 0, 0)$, indicating that the domain initially contains no chemical species.

F.1 EXPERIMENTATION

The numerical data are generated using a finite-difference scheme on a uniform grid of 64×32 points (i.e., $4 \times 64 \times 32$, 4 variables). For training and testing NOs, we subsample the horizontal direction to 32, leading to 32×32 for four variables. Time integration employs the second-order backward differentiation formula (BDF2) with a constant time step $\Delta t = 10^{-4} \text{ s}$ and a final simulation time of $t = 0.02 \text{ s}$ (corresponding to $n_t = 200$ time steps). For training and testing NOs, we subsample every 5th snapshot, resulting in 40 temporal indices for each trajectory.

The two model parameters, the pre-exponential factor A and the activation energy E , are sampled uniformly from the rectangular parameter domain $A, E \in [2.3375 \times 10^{12}, 6.5 \times 10^{12}] \times [5.625 \times 10^3, 9 \times 10^3]$. For data collection, we construct an 8×8 uniform grid over this domain, resulting in 64 trajectories, which are then split into 56:8 (train:test).

Figure 19 summarizes the performance and efficiency of the proposed and baseline models. In this experiment, we consider a subset of the baselines considered in other experiments, namely FNO_c ,

MWT_c , and DON_c . Particularly, $CFNO$ and $CMWNO$ are excluded from this comparison, as their architectures are primarily demonstrated for two-variable systems and it is hard to find their extension to higher-order coupling is not explicitly discussed in the original formulations. In the results, we observe that the proposed variants, FNO_x and $hpFNO_x$, achieve the lowest prediction errors while maintaining competitive model size and computational cost. Relative to FNO_c , FNO_x improves accuracy by about 41%, and $hpFNO_x$ further reduces the error by approximately 61%. Compared to the multiwavelet baseline MWT_c , $hpFNO_x$ still provides around 41% lower error while remaining lightweight. DON_c yields nRMSE values roughly an order of magnitude larger than the other models and is therefore omitted from Figure 19 for visual clarity. Overall, the results demonstrate that spectral coupling and parameter conditioning enhance predictive accuracy and efficiency without compromising the compact structure of the original FNO architecture.

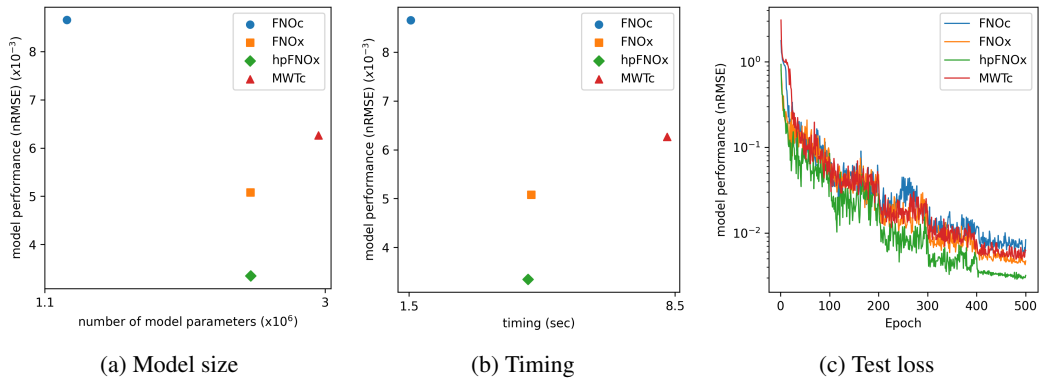


Figure 19: [ADR] The results of the experiments for varying pre-exponential factor and the activation energy (A, E). Plots depict the number of model parameters versus the model performance (left), computational timing versus the model performance (middle), and the test loss trajectories over epochs (right).