

000
001
002
003
004
005
006
007
008
009
010
011
012
013
014
015
016
017
018
019
020
021
022
023
024
025
026
027
028
029
030
031
032
033
034
035
036
037
038
039
040
041
042
043
044
045
046
047
048
049
050
051
052
053

KANO: KOLMOGOROV-ARNOLD NEURAL OPERATOR

Anonymous authors

Paper under double-blind review

ABSTRACT

We introduce *Kolmogorov-Arnold Neural Operator* (KANO), a dual-domain neural operator jointly parameterized by both spectral and spatial bases with intrinsic symbolic interpretability. We theoretically demonstrate that KANO overcomes the pure-spectral bottleneck of Fourier Neural Operator (FNO): KANO remains expressive over a generic position-dependent dynamics (variable coefficient PDEs) for any physical input, whereas FNO stays practical only to spectrally sparse operators and strictly imposes fast-decaying input Fourier tail. We verify our claims empirically on position-dependent differential operators, for which KANO robustly generalizes but FNO fails to. In the quantum Hamiltonian learning benchmark, KANO reconstructs ground-truth Hamiltonians in closed-form symbolic representations accurate to the fourth decimal place in coefficients and attains $\approx 6 \times 10^{-6}$ state infidelity from projective measurement data, substantially outperforming that of the FNO trained with ideal full wave function data, $\approx 1.5 \times 10^{-2}$, by orders of magnitude.

1 INTRODUCTION

Contemporary science and engineering increasingly operate in regimes where the effective dimensionality and complexity of phenomena and data overwhelm human-designed calibrations and approximations. This motivates data-centric modeling of governing dynamics from observations (Karniadakis et al., 2021; Wang et al., 2023; Carleo et al., 2019). For a learned model to be constituted as a scientific law, it should first *generalize* universally over a well-defined domain, and also should be *interpretable* so that the learned representations can be extracted and reused for verification, testing, and downstream simulation. [Mathematically, physical dynamics are generalized as operators as they are often formalized through PDEs](#) (Courant et al., 1963; Evans, 2022). A large and practically important subclass consists of variable coefficient PDEs, in which at least one term has a coefficient that varies by its variables (Gilbarg et al., 1977); we define physical dynamics governed by such PDEs as *position-dependent dynamics*, when one of the variables that varies the coefficient is position. Examples include fluid flow in media with spatially varying viscosity or conductivity (Kundu et al., 2024), and the Schrödinger equation with a potential that is a function of position operators (Sakurai & Napolitano, 2020). Scientific AI such as operator networks (Kovachki et al., 2023; Lu et al., 2021) that efficiently approximate a generic position-dependent dynamics with tractable interpretability are therefore valuable, which we recognize the absence and aim to fill the gap in this work.

An operator network approximates an arbitrary mapping between infinite-dimensional function spaces by first encoding functions into finite latent vectors and then learning the latent-to-latent map that represents the target operator (Lanthaler et al., 2022). DeepONet of Lu et al. (2021; 2019) implements the most general dense operator network where two neural networks learn both encoding and latent mapping directly from data, based on the theoretical foundation laid by Chen & Chen (1995). Fourier Neural Operator (FNO) of Li et al. (2020), on the other hand, hard-codes the encoding as pseudo-spectral projection with its spectrally diagonal kernels. FNO is provably and empirically superior when its hard-coded sparsity is optimal (Li et al., 2020; Kovachki et al., 2021), but this spectral sparsity becomes maladaptive for position-dependent or otherwise spectrally dense dynamics (Koshizuka et al., 2024; Qin et al., 2024). In such cases, the model size required for a target accuracy can grow super-exponentially (Kovachki et al., 2021), and although the universal approximation guarantee still holds, realistic size FNO may only converge on an in-sample mapping that fails outside the training distribution. Numerous variants of FNO attempted to break this spectral bottleneck. Some have broadened spectral coverage by exploiting factorized (Tran et al., 2021) or multi-scale (You et al., 2024) spectral kernels, and others have injected local spatial kernels alongside

054 the original spectral ones (Wen et al., 2022; Liu-Schiaffini et al., 2024; Liu et al., 2025). Yet all prior
 055 works still privilege the spectral basis and cannot achieve optimal sparsity in the spatial basis.
 056

057 In parallel, interpretability has recently pivoted around Kolmogorov–Arnold Network (KAN) (Liu
 058 et al., 2024b;a), whose edges are trainable univariate functions and thus amenable to human inspection.
 059 Several works demonstrate data-driven scientific modeling with KAN: Chiparova & Popov (2025);
 060 Gashi et al. (2025) use KAN for system identification, and Koenig et al. (2024) replace the MLPs
 061 in Neural ODEs (Chen et al., 2018) with KANs, each reporting symbolic recovery of benchmark
 062 equations and parameters. KANs have also been explored within operator networks: Abueidda
 063 et al. (2025) employed KANs instead of MLPs in DeepONet and Xiao et al. (2024); Wang et al.
 064 (2025) augmented FNO with KANs. Despite performance gains however, prior KAN-based operator
 065 networks have not reported symbolic recovery of the learned operator, leaving the avenue of an
 066 interpretable operator network largely unexplored.

067 To address these research gaps, we introduce the **Kolmogorov–Arnold Neural Operator (KANO)**,
 068 an interpretable operator network jointly parameterized in both spatial and spectral bases with KAN
 069 sub-networks embedded in a pseudo-differential operator framework (Hörmander, 2007; Kohn &
 070 Nirenberg, 1965). The key insight is to represent each component of the operator in the basis where
 071 it is sparse: differential terms spectrally, localized terms spatially, to achieve the most compact and
 072 tractable representation. Our work offers three main contributions to the scientific AI community.

- 073 • First, we demonstrate the pure-spectral bottleneck of FNO with an illustrative example and the-
 074oretically analyze why FNO cannot converge closely as desired to a generic position-dependent
 075 dynamics (variable coefficient PDEs) with a practical parameter complexity.
- 076 • Second, we propose a novel framework of KANO that is expressive over a generic position-
 077 dependent dynamics with intrinsic symbolic interpretability. We provide theoretical analysis on
 078 KANO’s dual-domain (spatial and spectral) expressivity along with the empirical evidences of
 079 KANO robustly generalizing on unseen input subspace when FNO fails to.
- 080 • Finally we validate the performance of KANO on some synthetic operators and a quantum simula-
 081 tion benchmark. KANO successfully recovered the closed-form formula accurately to the fourth
 082 decimal place in coefficients. Compared to the FNO baseline, KANO used only 0.03% of the model
 083 parameters, but achieved an order lower relative loss ℓ_2 in our synthetic operator benchmarks, and
 084 a four-order lower state infidelity in the quantum Hamiltonian learning benchmark.

085 To the best of our knowledge, our work is the first to demonstrate and quantify the symbolic
 086 recovery via KAN in operator learning. We shift the paradigm from mere *universal approximation*
 087 in operator learning toward the *universal generalization* of an operator network. Different from
 088 DeepOKAN (Abueidda et al., 2025) which replaces MLPs with KANs in DeepONet, our work
 089 achieves generalization over disjoint out-of-distribution subspace via a novel architecture design.

091 2 BACKGROUND

093 2.1 OPERATOR LEARNING AND FOURIER NEURAL OPERATOR

094 Operator learning approximates mapping between infinite-dimensional function spaces, $\mathcal{G} : \mathcal{A} \rightarrow \mathcal{U}$,¹
 095 from function pairs $\{(\mathbf{a}_i \in \mathcal{A}, \mathbf{u}_i = \mathcal{G}(\mathbf{a}_i) \in \mathcal{U})\}_{i=1}^N$ ²(Kovachki et al., 2023; 2024b). An operator
 096 network \mathcal{G}_θ first encodes input \mathbf{a}_i via *encoder* $\mathcal{E}_m : \mathcal{A} \rightarrow \mathbb{C}^m$ into a latent vector, then learns the *latent*
 097 $\mathbf{T}_\theta : \mathbb{C}^m \rightarrow \mathbb{C}^{m'}$ which the output is reconstructed to approximate the label \mathbf{u}_i via *reconstructor*
 098 $\mathcal{R}_{m'} : \mathbb{C}^{m'} \rightarrow \mathcal{U}$: i.e. $\mathcal{G}_\theta = \mathcal{R}_{m'} \circ \mathbf{T}_\theta \circ \mathcal{E}_m$ (Lanthaler et al., 2022). For fixed $(\mathcal{E}_m, \mathcal{R}_{m'})$, we can
 099 define the projection $\mathbf{\Pi}$ of an operator \mathcal{G} as

$$100 \quad \mathbf{\Pi}(\mathcal{G}) = \mathcal{R}_{m'} \circ \widehat{\mathbf{T}} \circ \mathcal{E}_m \quad \text{where} \quad \widehat{\mathbf{T}} \in \arg \min_{\mathbf{T} : \mathbb{C}^m \rightarrow \mathbb{C}^{m'}} \|\mathcal{G} - \mathcal{R}_{m'} \circ \mathbf{T} \circ \mathcal{E}_m\|. \quad (1)$$

103 DeepONet (Lu et al., 2021; 2019) learns \mathcal{E}_m , $\mathcal{R}_{m'}$, and \mathbf{T}_θ all with two sub-networks. FNO (Li
 104 et al., 2020), on the other hand, *hard-codes* \mathcal{E}_m to be the truncated Fourier transform and \mathcal{R}_m to be
 105 its band-limited inverse.

106 ¹ \mathcal{A} and \mathcal{U} are Banach function spaces (e.g., Sobolev spaces) defined on a bounded domain $D \subset \mathbb{R}^d$.

107 ² In practice, each function is sampled on a discretized grid in D and stored as a vector.

108 **Fourier Neural Operator (FNO).** Let the domain $D \subset \mathbb{R}^d$ be periodic and write the Fourier
 109 transform \mathcal{F} of function $\mathbf{a}(\mathbf{x})$ as $\hat{\mathbf{a}}(\boldsymbol{\xi})$:
 110

$$111 \quad [\mathcal{F}\mathbf{a}](\boldsymbol{\xi}) = \hat{\mathbf{a}}(\boldsymbol{\xi}) = \int_D \mathbf{a}(\mathbf{x}) e^{-i\boldsymbol{\omega}\mathbf{x}} d\mathbf{x}, \quad \boldsymbol{\omega} = 2\pi\boldsymbol{\xi} \in \mathbb{Z}^d. \quad (2)$$

113 For fixed set of retained modes $\boldsymbol{\xi}_i \in \Xi = \{\boldsymbol{\xi}_1, \dots, \boldsymbol{\xi}_m\} \subset \mathbb{Z}^d$, truncated Fourier transform $\mathcal{F}_m : \mathcal{A} \rightarrow \mathbb{C}^m$ and its band-limited inverse $\mathcal{F}_m^{-1} : \mathbb{C}^m \rightarrow \mathcal{U}$ can be defined as:
 114
 115

$$116 \quad \mathcal{F}_m(\mathbf{a}) = [\hat{\mathbf{a}}(\boldsymbol{\xi}_1), \dots, \hat{\mathbf{a}}(\boldsymbol{\xi}_m)], \quad \mathcal{F}_m^{-1}(\mathcal{F}_m \mathbf{a})(\mathbf{x}) = \sum_{j=1}^m \hat{\mathbf{a}}(\boldsymbol{\xi}_j) e^{2\pi i \boldsymbol{\xi}_j \cdot \mathbf{x}}, \quad (3)$$

119 with a slight abuse of notation. A single Fourier layer \mathcal{L}_{FNO} of FNO is written as:
 120

$$121 \quad \mathcal{L}_{\text{FNO}}(\mathbf{a})(\mathbf{x}) = \sigma \left(\mathcal{F}_m^{-1} \left(\mathbf{R}_{\theta_1}(\boldsymbol{\xi}) \cdot \mathcal{F}_m(\mathbf{a})(\boldsymbol{\xi}) \right) (\mathbf{x}) + \mathbf{W}_{\theta_2} \cdot \mathbf{a}(\mathbf{x}) \right) \quad (4)$$

123 with learnable spectral block-diagonal multiplier $\mathbf{R}(\boldsymbol{\xi})$, parametrized linear transformation \mathbf{W} , and
 124 point-wise nonlinear activation σ . FNO is comprised of iterative \mathcal{L}_{FNO} between lift-up (\mathcal{P}) and
 125 projection (\mathcal{Q}) networks:
 126

$$127 \quad \mathcal{G}_{\theta}^{\text{FNO}}(\mathbf{a}) = \mathcal{Q} \circ \mathcal{L}_{\text{FNO}}^{(\ell)} \circ \dots \circ \mathcal{L}_{\text{FNO}}^{(1)} \circ \mathcal{P}(\mathbf{a}). \quad (5)$$

129 In the perspective of the operator network formulation (1) (Lanthaler et al., 2022), FNO hard-codes
 130 its encoder \mathcal{E}_m as \mathcal{F}_m and reconstructor \mathcal{R}_m as \mathcal{F}_m^{-1} , then learns the latent map \mathbf{T} by its iterative
 131 layers of parametrized linear kernels interleaved by non-linear activations (Kovachki et al., 2021).

132 2.2 KOLMOGOROV-ARNOLD NETWORK

134 KAN (Liu et al., 2024b;a) replaces fixed node activations of traditional MLP with simple sum
 135 operations and train the learnable univariate 1D functions ϕ on edges. With layer width $n_l \rightarrow n_{l+1}$
 136 and input field $\mathbf{x}^{(l)} \rightarrow \mathbf{x}^{(l+1)}$, a KAN layer yields a function matrix $\Phi^{(l)}$ at l^{th} layer as
 137

$$138 \quad \mathbf{x}^{(l+1)} = \Phi^{(l)} \mathbf{x}^{(l)}{}^3, \quad \Phi^{(l)} = [\phi_{q,p}^{(l)}(\cdot)]_{q=1, \dots, n_{l+1}}^{p=1, \dots, n_l}, \quad x_q^{(l+1)} = \sum_{i=1}^{n_l} \phi_{q,p}^{(l)}(x_p^{(l)}), \quad (6)$$

140 so each output channel is a sum of edgewise transforms of the inputs (Liu et al., 2024b;a). In the
 141 original KAN each edge function is a spline expansion
 142

$$143 \quad \phi_{q,p}^{(\ell)}(t) = c_{q,p,0}^{(\ell)} b(t) + \sum_{i=1}^g c_{q,p,i}^{(\ell)} B_i(t), \quad (7)$$

146 with learnable coefficients for a fixed base 1D function $b(\cdot)$ and 1D B-spline basis $\{B_i\}$. Because
 147 every $\phi_{q,p}$ is a 1D curve, KANs are directly inspectable and amenable to visualization followed by
 148 symbolic regression. On expressivity, Wang et al. (2024) theoretically prove that KANs match MLPs
 149 up to constant depth and width factors; empirically, with appropriate optimization recipes, KANs and
 150 MLPs exhibit comparable scaling on PDE and operator benchmarks (Shukla et al., 2024). Hence,
 151 swapping a latent MLP for a KAN preserves expressivity while enabling symbolic readout.
 152

153 3 THEORETICAL ANALYSIS ON FNO'S PURE-SPECTRAL BOTTLENECK

155 This section first illustrates the pure-spectral bottleneck of FNO. Then we provide a theoretical
 156 analysis and prove that FNO suffers the curse of dimensionality for position-dependent dynamics.
 157 **FNO is proven to have the universal approximation guarantee over any arbitrary non-linear Lipschitz**
 158 **operator (Kovachki et al., 2021; Lanthaler et al., 2025).** This section does not disprove the universal
 159 approximation ability of FNO; it illustrates the limitation on the *generalization ability* of FNO
 160 stemming from its pure-spectral bottleneck on spectrally dense operators.

161 ³Akin to matrix-vector multiplication but follows the third equation Eq. 6 instead of row-vector inner product.

162 3.1 THE PURE-SPECTRAL BOTTLENECK OF FNO
163164 We consider the 1D quantum harmonic oscillator Hamiltonian as an example:
165

$$\mathcal{H}a(x) = -\partial_{xx}a(x) + x^2 \cdot a(x). \quad (8)$$

166 Multiplication and differentiation have a dual relationship under the Fourier transform:
167

$$\mathcal{F}[(-\partial_{xx}a)](\xi) = \xi^2 \cdot \hat{a}(\xi), \quad \mathcal{F}[(x^2 \cdot a)](\xi) = -\partial_{\xi\xi}\hat{a}(\xi). \quad (9)$$

168 In spectral domain, the spatial differential ∂_{xx} is a spectral multiplier ξ^2 , whereas the spatial multiplier
169 x^2 becomes a spectral differential $\partial_{\xi\xi}$. Consider a truncated polynomial basis $\{1, x, x^2, \dots, x^{n-1}\}$
170 and a truncated Fourier basis $e_k(\theta) = e^{ik\theta}$, $k = 0, \dots, n-1$, on a periodic domain. In the spatial
171 (polynomial) basis, the map $a(x) \mapsto x^2 \cdot a(x)$ acts as a two-step up-shift sparse matrix
172

$$\mathbf{S}_n^{(2)} := \begin{bmatrix} 0 & 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \ddots & 1 \\ 0 & 0 & 0 & \cdots & 0 & 0 \\ 0 & 0 & 0 & \cdots & 0 & 0 \end{bmatrix}, \quad (10)$$

173 while in the spectral (Fourier) basis it is a dense Toeplitz matrix (Morrison, 1995)
174

$$\mathbf{T}_n[x^2] := \begin{bmatrix} c_0 & c_{-1} & c_{-2} & \cdots & c_{-n+1} \\ c_1 & c_0 & c_{-1} & \cdots & c_{-n+2} \\ c_2 & c_1 & c_0 & \cdots & c_{-n+3} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ c_{n-1} & c_{n-2} & c_{n-3} & \cdots & c_0 \end{bmatrix}, \quad c_m = \frac{1}{2\pi} \int_0^{2\pi} \theta^2 e^{-im\theta} d\theta. \quad (11)$$

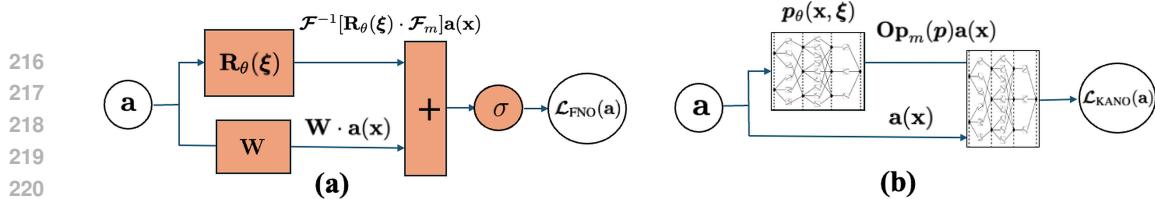
175 Thus each term in \mathcal{H} is sparse in one basis and dense in the other (Morrison, 1995).
176177 An FNO layer \mathcal{L}_{FNO} (4) can easily parametrize $-\partial_{xx}$ by taking $\mathbf{R}(\xi) \approx \xi^2$. However, approximating
178 the dense off-diagonals in $\mathbf{T}_n[x^2]$ to parametrize x^2 must rely on the non-linear activation $\sigma(\cdot)$ since
179 $\mathbf{R}(\xi)$ and \mathbf{W} are spectrally diagonal and hence incapable of mixing modes. Let $\mathbf{z}(\mathbf{u})$ denote the
180 pre-activation for input \mathbf{u} , then the Jacobian of \mathcal{L}_{FNO} at \mathbf{u} gives the first-order approximation of the
181 learned map and its Fourier transform reveals the spectral off-diagonals of itself as
182

$$\mathcal{F}[\mathbf{J}(\mathbf{u})](\xi, \xi') = (\mathcal{F}[\sigma'(\mathbf{z}(\mathbf{u}, \cdot))] [\xi - \xi']) \cdot (\mathbf{W} + \mathbf{R}_\theta(\xi')). \quad (12)$$

183 Therefore, *all* off-diagonals arise from the spectrum of the \mathbf{u} -dependent gate $\sigma'(\mathbf{z}(\mathbf{u}, \cdot))$: FNO's
184 nonlinearity *can* create off-diagonals, but they are tied to the input distribution of \mathbf{u} . This is the pure-
185 spectral bottleneck of FNO: *spectral off-diagonals of a learned FNO are tied to the train subspace,*
186 *hence FNO can converge only on the in-sample mapping that fails outside the train distribution*⁴. See
187 Appendix D for further detailed discussion expanding to the arbitrary higher order contribution and
188 deep layered FNO.
189190 3.2 FNO SUFFERS CURSE OF DIMENSIONALITY ON POSITION-DEPENDENT DYNAMICS
191192 As explained previously, *position operator*, $a(x) \mapsto x \cdot a(x)$, is a highly dense Toeplitz map in the
193 spectral basis (Morrison, 1995). Based on the Remark 21 & 22 of Kovachki et al. (2021), we prove
194 that any position-dependent dynamics induces super-exponential scaling in FNO size by the desired
195 error bound: *FNO cannot converge closely as desired on a generic position-dependent dynamics with*
196 *practical model size, hence can only overfit on the in-sample mapping*. We provide Lemma 1, that a
197 single position operator already spreads the input spectra too much for FNO to stay practical, and
198 Theorem 1, expanding Lemma 1 to an arbitrary composition of position operators.
199200 Following from the operator network formulation (1), the error estimate of an operator network \mathcal{G}_θ
201 approximating the ground-truth operator \mathcal{G} in an operator norm is bounded as
202

$$\|\mathcal{G} - \mathcal{G}_\theta\| \leq \underbrace{\|\mathcal{G} - \mathcal{R}_{m'} \circ \widehat{\mathbf{T}} \circ \mathcal{E}_m\|}_{\text{projection error: } \epsilon_{\text{proj}}} + \underbrace{\|\mathcal{R}_{m'} \circ (\widehat{\mathbf{T}} - \mathbf{T}_\theta) \circ \mathcal{E}_m\|}_{\text{latent network error: } \epsilon_{\text{net}}}, \quad (13)$$

203 by the triangle inequality. Latent network error ϵ_{net} follows the well-established scaling law of
204 conventional neural networks (Hornik et al., 1989; Cybenko, 1989). Therefore, whether an operator
205 network is efficient in model and sample size to achieve the desired accuracy hinges on the scalability
206 of the projection error ϵ_{proj} (Lanthaler et al., 2022; Kovachki et al., 2021; 2024a).
207208 ⁴This issue of out-of-distribution fragility from underspecification is well studied by D'Amour et al. (2022)

Figure 1: (a) \mathcal{L}_{FNO} architecture. (b) $\mathcal{L}_{\text{KANO}}$ architecture.

Reviewing Remark 21 & 22 of Kovachki et al. (2021) ϵ_{proj} of FNO is governed by the Fourier tail, the sum of Fourier coefficients outside the retained spectrum Ξ truncated by width m (Gottlieb & Orszag, 1977; Trefethen, 2000): to achieve the desired ϵ_{proj} with practical m , both input and output Fourier tails must decay algebraically or faster. However, even in the optimal case of the band-limited input, if the ground-truth operator is spectrally dense to spread out the input spectra and induce algebraic or slower decay in output Fourier tail, m must scale at least polynomially to suppress ϵ_{proj} : $m \sim \mathcal{O}(\epsilon_{\text{proj}}^{-1/s})$ where s is a geometric constant. Meanwhile, as the latent mapping would be also dense, the size of the latent network, \mathcal{N}_{net} , follows the canonical polynomial neural scaling (Yarotsky, 2017; De Ryck et al., 2021) by the desired ϵ_{net} with its width m^d (d is the input domain dimension) as the exponent: $\mathcal{N}_{\text{net}} \sim \mathcal{O}(\epsilon_{\text{net}}^{-m^d})$. Consequently, this results as the super-exponential scaling in the latent network size $\mathcal{N}_{\text{net}} \sim \mathcal{O}(\epsilon_{\text{net}}^{-\epsilon_{\text{proj}}^{-d/s}})$ even with the optimal band-limited input: *scaling width m to suppress ϵ_{proj} explodes \mathcal{N}_{net} to achieve the desired ϵ_{net} for a generic dense operator.*

Lemma 1 (Position operator elongates Fourier tail). *A single position operator, spatial multiplier by x , induces algebraic decay in output Fourier tail when the input is band-limited.*

Sketch of proof. Position operator is kernel $\hat{x}(\xi) \propto 1/\xi$ in spectral basis. Hence, every mode outside the input spectrum picks up a coefficient of size $\sim 1/|\xi|$, ending up as $|\hat{v}(\xi)| \gtrsim 1/|\xi|$ in the output spectrum. See Appendix C.2 for restatement and full proof. \square

Theorem 1 (Curse of dimensionality on position operators). *Any arbitrary composition of position operators requires FNO to scale super-exponentially on its model size by the desired accuracy.*

Sketch of proof. Iteratively apply Lemma 1, then any arbitrary composition of position operators induce algebraic or slower decay in output Fourier tail even for the optimal band-limited input. This results in super-exponential scaling of latent network size by the desired error as discussed above (Kovachki et al., 2021). See Appendix C.3 for restatement and full proof. \square

Remark 1. *What is missing from the upper bound analysis of FNO by Kovachki et al. (2021) is the effect of the wide lift-up and projection networks. For the generalization guarantee arguments of this work, the upper bound analysis is still sound, yet it should be clarified that the theoretical analysis on the role of the lift-up and projection networks is yet an open research question. Recent studies by Lanthaler et al. (2025); Lanthaler (2024) provide better view on it with lower bound analysis.*

4 KOLMOGOROV-ARNOLD NEURAL OPERATOR

Motivated by the pure-spectral bottleneck of FNO, we propose the Kolmogorov-Arnold Neural Operator (KANO), an operator network capable of converging closely as desired on a generic position-dependent dynamics with practical model size. We first introduce the KANO architecture, and provide theoretical analysis on its dual-domain expressivity in the following section.

4.1 KANO ARCHITECTURE

KANO utilizes an iterative structure of KANO layers $\mathcal{L}_{\text{KANO}}$ to learn the unknown operator, akin to FNO. However, KANO excludes the wide lift-up and projection networks to maximize tractability since it is known that wide KANs are fragile to symbolic recovery (Noorizadegan et al., 2025):

$$\mathcal{G}_{\theta}^{\text{KANO}} = \mathcal{L}_{\text{KANO}}^{(\ell)} \circ \dots \circ \mathcal{L}_{\text{KANO}}^{(1)}. \quad (14)$$

$$\mathcal{L}_{\text{KANO}}(\mathbf{a})(\mathbf{x}) = \Phi_{\theta_1} \left(\mathcal{F}_m^{-1} [p_{\theta_2}(\mathbf{x}, \xi) * \mathcal{F}_m(\mathbf{a})(\xi)](\mathbf{x}), \mathbf{a}(\mathbf{x}) \right), \quad (15)$$

270 where Φ is a KAN sub-network for learnable non-linear activation; $p(\mathbf{x}, \xi)$ is another KAN sub-
 271 network, a pseudo-differential symbol jointly parametrized by both spatial \mathbf{x} and spectral ξ bases⁵.
 272 Note the “ $*$ ” notation in $\mathcal{L}_{\text{KANO}}$ (15) instead of the block-diagonal multiplication notation “ \cdot ” in
 273 \mathcal{L}_{FNO} (4). The spatial basis \mathbf{x} of the symbol $p(\mathbf{x}, \xi)$ is convolution (differential) in spectral domain
 274 by the dual relationship (9). Therefore, the pseudo-differential symbol calculus of $p(\mathbf{x}, \xi)$ needs
 275 to be done by quantizing on both spatial and spectral domain (Hörmander, 2007), and we choose
 276 Kohn-Nirenberg quantization (Kohn & Nirenberg, 1965) to compute the symbol calculus in $\mathcal{L}_{\text{KANO}}$:
 277

$$\mathcal{F}_m^{-1}[\mathbf{p}(\mathbf{x}, \xi) * \mathcal{F}_m(\mathbf{a})(\xi)](\mathbf{x}) := \left(\frac{h}{L}\right)^d \sum_{\xi \in \Xi} \sum_{\mathbf{y} \in \mathcal{Y}} e^{i(\mathbf{x}-\mathbf{y}) \cdot \xi} \mathbf{p}(\mathbf{x}, \xi) \mathbf{a}(\mathbf{y}), \quad (16)$$

280 where for a periodic domain $D = (-\frac{L}{2}, \frac{L}{2})^d$, $\mathcal{Y} = \{\mathbf{y}_1, \dots, \mathbf{y}_m\} \subset D$ is a uniform discretization
 281 with spacing h and $\mathbf{x} \in D$ is an evaluation point. We denote Kohn-Nirenberg quantization (16) as an
 282 operator $\mathbf{Op}_m(p) := \mathcal{F}_m^{-1}[\mathbf{p}(\mathbf{x}, \xi) * \mathcal{F}_m]$ defined by the symbol $p(\mathbf{x}, \xi)$. In the operator network
 283 formulation (1) introduced in Section 2.1, KANO’s projection Π_{KANO} is then defined as:

$$\Pi_{\text{KANO}}(\mathcal{G}) := \mathbf{Op}_m(p_{\mathcal{G}}), \quad p_{\mathcal{G}} \in \arg \min_p \|\mathcal{G} - \mathbf{Op}_m(p)\|. \quad (17)$$

287 **Symbolic Interpretability of KANO.** By using compact KANs each for the symbol $p(\mathbf{x}, \xi)$ and
 288 non-linear activation Φ in every KANO layer $\mathcal{L}_{\text{KANO}}$ (15), KANO network $\mathcal{G}_{\theta}^{\text{KANO}}$ (14) is fully
 289 inspectable by visualizing the learned edges of all its KANs, potentially allowing closed-form
 290 symbolic formula of the learned operator with the manual provided by Liu et al. (2024b;a). In
 291 addition, recent endeavors have greatly expanded KAN’s symbolic recovery capacity to non-smooth,
 292 discontinuous targets with high irregularities (Yu et al., 2024; Aghaei, 2024; Lei et al., 2025; Yang
 293 et al., 2025; Shiraishi et al., 2025). All of such advancements are easily and directly applicable in our
 294 KANO framework as well, when facing an operator with high irregularity coefficients.

295 **Remark 2 (Complexity analyses of KANO).** As apparent in Eq. (16), KANO layer must perform
 296 double sum which can be computationally heavy. However, for the target operator class of variable-
 297 coefficient PDEs such as position-dependent dynamics, we show this can be compensated in principle
 298 by the parameter efficiency we prove in the following Section 4.2. See Appendix E.

4.2 KANO’s DUAL-DOMAIN EXPRESSIVITY

300 In contrast to FNO, KANO exploits sparse representations in both the spatial and spectral domains,
 301 hence decoupling the scaling of ϵ_{proj} and ϵ_{net} by never letting the latent map be a dense convolution. For
 302 instance, for the quantum harmonic oscillator in Eq. (8), a KANO layer $\mathcal{L}_{\text{KANO}}$ (15) can parametrize
 303 \mathcal{H} by taking $p(x, \xi) \approx x^2 + \xi^2$, both $-\partial_{xx}$ and x^2 terms are each represented where they are
 304 sparse, both leveraging the shift form $\mathbf{S}_n^{(2)}$ (10). By jointly parameterizing the operator in both
 305 spatial and frequency domains, KANO *cherry-picks* the sparse representation for every term in
 306 position-dependent dynamics, building the right inductive bias well-known to be essential for out-of-
 307 distribution generalization and model efficiency (Goyal & Bengio, 2022; Trask et al., 2018).

308 This dual-domain expressivity of KANO first alleviates the input constraint; we first explain that
 309 ϵ_{proj} of KANO scales practically by its width m for *any* physical input. Then we provide Theorem 2:
 310 as long as the KANO projection (17) of an operator generates smooth symbols KAN can easily
 311 approximate, ϵ_{net} scales practically by compact KAN sub-networks independent of ϵ_{proj} . In conclusion,
 312 *KANO can converge closely as desired to a generic position-dependent dynamics with practical*
 313 *model size using any physical input, robustly generalizing outside the train subspace.*

315 **KANO practically has no input constraint.** According to the quadrature bound from Demanet &
 316 Ying (2011, Thm. 1&2), the error estimate of Kohn-Nirenberg quantization (16) obeys

$$\|\mathcal{G} - \mathbf{Op}_m(p_{\mathcal{G}})\| \leq C B m^{-s}, \quad (18)$$

317 given norm-bound (finite-energy)⁶ input of $\mathcal{A}_B = \{\mathbf{u} : \|\mathbf{u}\| \leq B\}$ where s, C are geometric
 318 constants. Hence KANO width m scales polynomially by the desired ϵ_{proj} given any physical data.

319 ⁵Shin et al. (2022) first employed pseudo-differential operator framework for neural operator. They presumed
 320 the symbol $p(\mathbf{x}, \xi)$ to be separable as $p(\mathbf{x}, \xi) = p_{\mathbf{x}}(\mathbf{x}) \cdot p_{\xi}(\xi)$, and used MLP sub-networks while retaining
 321 the lift-up and projection networks of a generic neural operator architecture (Kovachki et al., 2023).

322 ⁶Norm here and Equation Eq. 18 is the Sobolev norm

Theorem 2 (KANO stays practical for smooth symbol). *If the KANO projection of an operator \mathcal{G} , $\Pi_{\text{KANO}}(\mathcal{G})$ (17), generates a finite composition of smooth symbols $p_{\mathcal{G}}(\mathbf{x}, \xi)$ and finite-degree non-linearities, the model size of KANO scales polynomially by the desired accuracy ε .*

Sketch of proof. Choosing $m \sim (B/\varepsilon)^{1/s}$ scales projection error down to $\varepsilon/2$ by Eq. (18). A fixed-width KAN then approximates the symbols to accuracy $\varepsilon/2$ with $\mathcal{O}(\varepsilon^{-d/(2s_p)})$ parameters (Wang et al., 2024, Corol. 3.4) (s_p is a geometric constant). The finite-degree non-linearities add only constant-size weights by the activation KAN, so the total parameter count is $\mathcal{O}(\varepsilon^{-d/(2s_p)})$. See Appendix C.4 for restatement and full proof. \square

Corollary 1 (KANO is practical for generic position-dependent dynamics). *For a finite composition of spatial and spectral multipliers of maximum r -differentiable symbols with finite-degree non-linearity, Theorem 2 yields $|\Theta| = \mathcal{O}(\varepsilon^{-d/(2r)})$.*

Remark 3 (Scope of KANO). *Recent studies demonstrate that wide lift-up and projection networks are essential for strong performance on high-dimensional benchmarks (Diab & Al Kobaisi, 2025; Eker, 2024; Liu et al., 2023). In contrast, KANO is designed to prioritize symbolic recovery with robust generalization, complementary to the scope of FNO. Because the core dual-domain expressivity is mathematically agnostic to dimensionality, extending KANO to conventional high-dimensional use cases is a natural and promising direction for future work.*

5 EXPERIMENTAL RESULTS

5.1 SYNTHETIC-OPERATOR GENERALIZATION BENCHMARKS

We benchmark FNO-based models and KANO on three position-dependent operators:

$$\mathcal{G}_1 f = x^2 \cdot f - \partial_{xx} f, \quad \mathcal{G}_2 f = x \cdot \partial_x f + \partial_{xx} f, \quad \mathcal{G}_3 f = f^3 + x \cdot \partial_x f + \partial_{xx} f.$$

Our goal is to quantify and compare the generalization of each model. We train the models *only* with Group A dataset and evaluate them on the *unseen* Group B dataset.

- **Group A (Training families):** Periodic, Chirped Cosine, Sine Beats.
- **Group B (Testing families):** Sinc Pulse, Gaussian \times Hermite, Wave Packet.

For each operator, we generate 2000 train pairs from Group A and 400 test pairs from Group B to evaluate the generalization by comparing the ratio between the average relative ℓ_2 loss over each group (**Loss Test**). We also interpolate the Group A and B function samples in 100 steps, apply ground-truth operators in each step to build the interpolated dataset, and evaluate the loss ratio to that of the Group A samples (**Interpolation Test**). We trained FNO models of 2 layers, 64 width with no mode truncation, and used one-layer KANs of grid 10 cubic B-splines edges for the KANO model. For other FNO variant baselines, we used 2-layer models of similar size to FNO. Lastly, we trained a KANO variant for an ablation study, which we replaced the KAN subnetworks with compact MLPs of 32 hidden width and 2 hidden layers (**KANO_MLP**). See Appendix B.1 for experiment details. We used Adam optimizer and relative ℓ_2 loss for training.

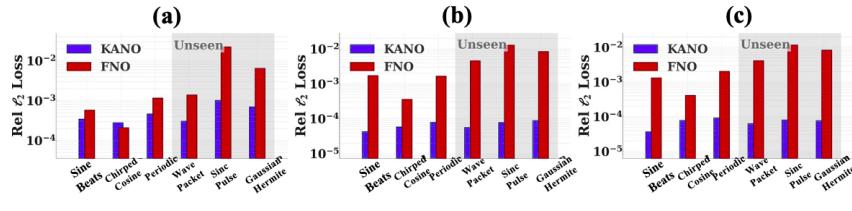
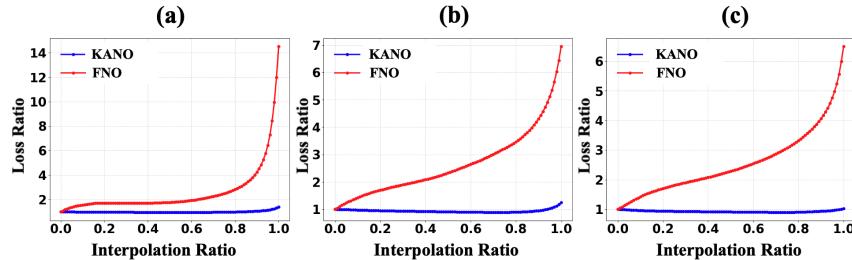


Figure 2: Loss test results. (a) \mathcal{G}_1 (b) \mathcal{G}_2 (c) \mathcal{G}_3 . Note the logarithmic scale.

378 Table 1: Relative ℓ_2 losses ($\times 10^{-4}$) and parameter counts.
379

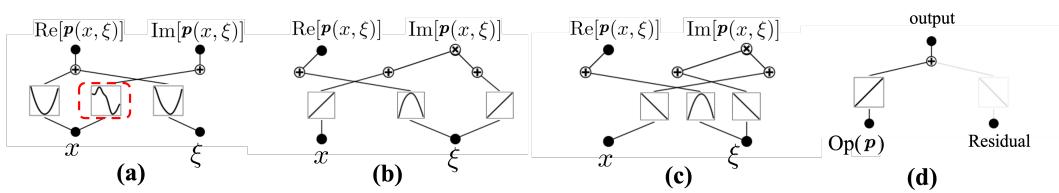
380 Model (params)	\mathcal{G}_1			\mathcal{G}_2			\mathcal{G}_3		
	381 A	382 B	383 B/A	384 A	385 B	386 B/A	387 A	388 B	389 B/A
FNO (566k)	6.36	98.8	15.53	10.6	87.0	8.21	11.4	81.4	7.14
U-FNO (579k)	2.79	22.9	8.21	8.14	339	41.65	92.4	292	3.16
AM-FNO (548k)	1.08	20.9	19.35	1.20	16.5	13.75	1.16	29.8	25.69
PDNO (538k)	1.41	6.31	4.5	1.92	12.1	6.3	4.03	27.2	6.7
KANO (152)	1.04	1.44	1.38	0.629	0.749	1.19	0.716	0.737	1.03
KANO_MLP (2k)	3.37	6.59	1.96	4.49	8.07	1.80	3.59	6.87	1.91
KANO_SYMBOLIC	0.512	0.526	1.03	0.498	0.500	1.00	0.520	0.536	1.03

388
389 **Results.** As shown in Table 1 and Figure 2, KANO shows consistent losses over Group A and
390 Group B, validating its robust generalization ability for position-dependent operators. **KANO_MLP**
391 also shows comparable out-of-distribution performance, which suggests that the generalization ability
392 of KANO stems from its dual-domain architecture apart from KAN. In contrast, FNO shows fragile
393 out-of-distribution behavior on Group B dataset with the significant loss increases. **U-FNO** (Wen et al.,
394 2022) and **AM-FNO** (Xiao et al., 2024) show even worse results. On the other hand, PDNO (Shin
395 et al., 2022) shows the most stable generalization among FNO families, although not as robust as
396 KANO and KANO_MLP. Along with the ablation study, this confirms that the pseudo-differential
397 operator framework is judicious for robust generalization on position-dependent operators, while
398 primarily relying on only spectral kernels makes the model fragile out of train distribution even with
399 localized enhancements.

400 Figure 3: Interpolation test results. (a) \mathcal{G}_1 (b) \mathcal{G}_2 (c) \mathcal{G}_3 .
401
402
403
404
405
406
407
408
409
410
411

412 Interpolation test results in Figure 3 further empirically validate our theories. The FNO curves (red)
413 of Figure 3 show slow increases on early and mid-interpolation, suggesting that the FNO’s learned
414 in-sample mappings are yet close to the ground-truth operators. However, the FNO curves abruptly
415 soar up in the latter ratio, suggesting that the interpolated functions are now far outside the train
416 distribution. These results, together with KANO’s one-order-of-magnitude lower loss at just 0.03%
417 of FNO’s size, are consistent with our claims in Theorem 1 and Theorem 2.

418 After convergence, we visualized the embedded KANs (Figure 4). We then froze these learned
419 symbols and continued training, referring to this variant as KANO_symbolic. KANO_symbolic
420 recovered the exact symbolic coefficients of the ground-truth operator to within the fourth decimal
421 place (Table 2). KANO’s loss matches KANO_symbolic’s loss in Table 1, confirming that KANO
422 converged close to the ground-truth operator.

423 Figure 4: (a) $p(x, \xi)$ of \mathcal{G}_1 . The middle edge does not contribute to the output. (b) $p(x, \xi)$ of \mathcal{G}_2 . (c)
424 $p(x, \xi)$ of \mathcal{G}_3 . (d) Φ of \mathcal{G}_3 . Edge of the residual in (d) looks linear, so we compared two scenarios,
425 linear and cubic, which the latter achieved lower loss and better generalization.
426
427
428
429
430
431

432 Table 2: Ground-truth vs. learned operators (coefficients rounded to 4th decimal place).

433 Ground-truth operator	434 Learned KANO operator
435 $\mathcal{G}_1 f = x^2 \cdot f - \partial_{xx} f$	$\tilde{\mathcal{G}}_1 f = (x^2 + 0.0003) \cdot f - \partial_{xx} f$
436 $\mathcal{G}_2 f = x \cdot \partial_x f + \partial_{xx} f$	$\tilde{\mathcal{G}}_2 f = 0.9996 x \cdot \partial_x f + \partial_{xx} f - 0.0003 f$
437 $\mathcal{G}_3 f = f^3 + x \cdot \partial_x f + \partial_{xx} f$	$\tilde{\mathcal{G}}_3 f = 1.0001 f^3 + 0.99997 x \cdot \partial_x f + 0.99997 \partial_{xx} f - 0.0002 f^2 - 0.0003 f - 0.0001$

440 **5.2 LONG-HORIZON QUANTUM DYNAMICS BENCHMARK**

441 We provide this benchmark on two position-dependent quantum dynamics: the quartic double-well
 442 Hamiltonian (DW) and the nonlinear Schrödinger equation with cubic nonlinearity (NLSE):

$$444 i\partial_t \psi = -\frac{1}{2}\partial_{xx} \psi + w(x) \cdot \psi \text{ (DW)}, \quad i\partial_t \psi = -\frac{1}{2}\partial_{xx} \psi + w(x) \cdot \psi + |\psi|^2 \cdot \psi \text{ (NLSE)},$$

$$446 \text{ where } w(x) = x^4 - (x - \frac{1}{32})^2 + 0.295.$$

447 We generate 200 initial states and yield the state trajectories by the Hamiltonians, sampling momen-
 448 tumb/position probability mass functions (PMFs) every 0.1ms for 100 time steps. The first 10 time
 449 steps are used for training, and the rest are used to evaluate the long-horizon prediction.

450 We modify KANO to capture the quantum state evolution: **Q-KANO**. Symbol p_θ is parametrized
 451 as $\exp[-i\Delta T \phi_\theta(\mathbf{x}, \xi)]$, where $\Delta T = 0.1$ ms. The adaptive activation is also defined as a complex
 452 exponential with learned phase $\vartheta = \Phi_\theta(|\mathbf{Op}_m(p_\theta)\psi|, \angle \mathbf{Op}_m(p_\theta)\psi)$ for input wave function $\psi(\mathbf{x})$:

$$454 \mathcal{G}_\theta^{\text{Q-KANO}}[\psi] = \mathbf{Op}_m(\exp[-i\Delta T \phi_\theta(\mathbf{x}, \xi)])\psi \cdot e^{-i\Delta T \vartheta}. \quad (19)$$

455 We investigate three supervision scenarios: **Full**-type training with full wave function, idealistic yet
 456 physically unattainable, **Pos**-type training with only position PMF, physically realistic yet the least
 457 informative, and **pos&mom**-type training with both position *and* momentum PMFs, which remains
 458 physically attainable while providing richer information although not full. We use Adam optimizer
 459 for all trainings. See Appendix B.2 for experiment details.

461 **Results.** We evaluate
 462 state infidelity⁷ between
 463 ground-truth evolution
 464 and model prediction at
 465 each time step (Table 3,
 466 Figure 6). In case of
 467 KANO, the **pos & mom**-
 468 type training achieves
 469 indistinguishable infi-
 470 delity from the ideal
 471 **full**-type training base-
 472 line. The ablation study
 473 with MLP variant of
 474 Q-KANO achieved
 475 comparably low state
 476 infidelity by **pos & mom**-type training as well. Meanwhile, the **pos**-type training displays a clear
 477 infidelity increase, especially on the NLSE.

461 Table 3: State infidelity after 90 additional time-evolution steps.

462 Model & Train Type	463 State Infidelity	
	464 Double-Well	465 NLSE
FNO (full)	1.5×10^{-2}	1.6×10^{-2}
Q-KANO (full)	6.3×10^{-6}	6.8×10^{-6}
Q-KANO (pos & mom)	6.3×10^{-6}	5.6×10^{-6}
Q-KANO (pos)	4.7×10^{-3}	6.1×10^{-2}
Q-KANO_MLP (pos & mom)	7.7×10^{-6}	8.5×10^{-6}
Q-KANO_SYMBOLIC (full)	2.0×10^{-8}	2.0×10^{-8}
Q-KANO_SYMBOLIC (pos & mom)	2.0×10^{-8}	3.0×10^{-8}
Q-KANO_SYMBOLIC (pos)	5.3×10^{-2}	6.1×10^{-2}

477 In contrast, even with **full** type training, FNO fails to maintain low state infidelity after the long-
 478 horizon propagation as expected. Iterative time evolution pushes the wave function far outside the
 479 train convex hull, and FNO’s learned in-sample mapping deviates from the ground-truth evolution
 480 rapidly, leading to four orders of infidelity increase compared to KANO.

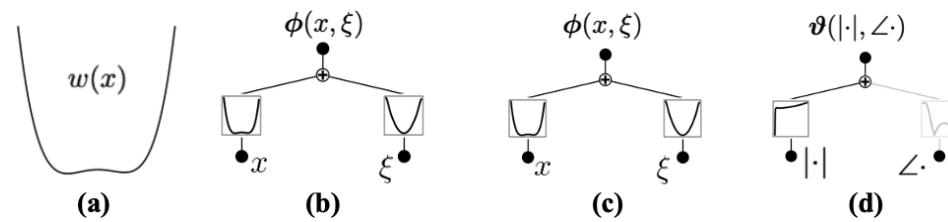
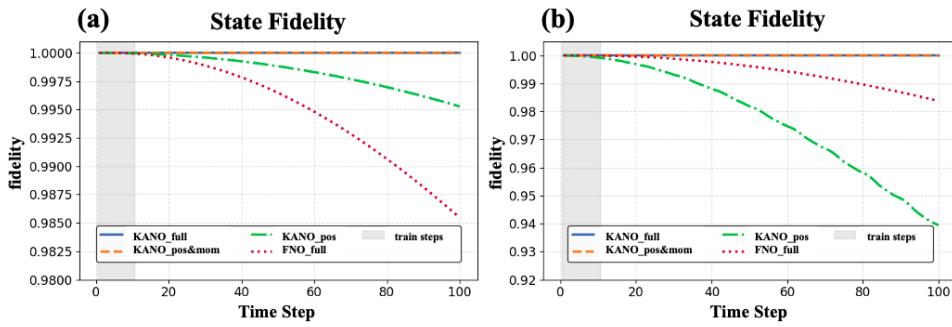
481 Table 4 juxtaposes the learned symbols with that of the ground-truth Hamiltonians and Figure 5
 482 shows the KAN visualizations from **pos & mom**-type training. With **full**-type training coefficients are
 483 recovered to the fourth decimal place, vindicating the ideal capacity of KANO when the information

484
 485 ⁷For predicted state $\tilde{\varphi}$ and ground-truth state φ , the state fidelity F is defined as the inner product between
 them ($F := \langle \tilde{\varphi}, \varphi \rangle$), and the state infidelity is defined as $(1 - F)$, hence shows how distant two states are.

486 is fully provided. Under the realistic **pos & mom**-type training, the reconstruction remains accurate
 487 except for two terms: the constant (global phase) and the NLSE’s cubic coefficient. Both discrepancies
 488 are predicted by quantum observability: global phases cancel in all PMFs, and the Kerr coefficient
 489 enters only through higher-order correlations that become harder to estimate from finite-shot statistics.
 490 Q-KANO faithfully reveals what the data support and nothing more.
 491

492
493 Table 4: Ground truth vs. learned symbols. Coefficients rounded to 4th decimal place.

Hamiltonian	Train Type	Learned symbolic structure
DW	ground truth	$x^4 - x^2 + 0.0625 x + 0.295 + 0.5 \xi^2$
	full	$1.0004 x^4 + 0.0001 x^3 - 1.0013 x^2 + 0.0625 x + 0.2955 + 0.5 \xi^2$
	pos & mom	$1.0003 x^4 + 0.0001 x^3 - 1.0008 x^2 + 0.0623 x + 0.0001 + 0.5 \xi^2$
NLSE	ground truth	$x^4 - x^2 + 0.0625 x + 0.295 + 0.5 \xi^2 + \psi ^2$
	full	$1.0005 x^4 - 0.0001 x^3 - 1.0014 x^2 + 0.0626 x + 0.2942 + 0.5 \xi^2 + 0.9815 \psi ^2 + 0.0110 \psi $
	pos & mom	$0.9999 x^4 - 0.0003 x^3 - 1.0001 x^2 + 0.0630 x + 0.1141 + 0.5 \xi^2 + 0.9514 \psi ^2 - 0.5504 \psi $

505
506
507 Figure 5: **pos&mom** type training results. (a) Structure of the potential $w(x)$ (b) $p(x, \xi)$ of DW. (c)
508 $p(x, \xi)$ of NLSE. (d) $\varphi(|\cdot|, \angle \cdot)$ of NLSE. Potential $w(x)$ structure is clearly reconstructed.
509
510
511
512
513514
515 Figure 6: State fidelity over 100 time steps. (a) DW (b) NLSE
516517
518

6 CONCLUSION

519 We have presented the Kolmogorov–Arnold Neural Operator, an interpretable neural operator expres-
 520 sive on a generic position-dependent dynamics. KANO cherry-picks sparse representations of each
 521 term via jointly parametrizing on both spectral and spatial bases, and achieves robust generalization
 522 outside train distribution while exposing clear tractable representation via its KAN sub-networks. In
 523 all our benchmarks, KANO have successfully recovered the ground-truth operators accurately to the
 524 fourth decimal place in coefficients. In addition to the superior out-of-distribution generalization,
 525 KANO has also achieved orders of magnitude lower losses with less than 0.03% of the model
 526 size compared to the FNO baseline. KANO shifts operator learning from an opaque, surrogate-
 527 based paradigm towards interpretable data-driven scientific modeling, and provides robust empirical
 528 evidence supporting its enhanced dual-domain expressivity and interpretability.
 529

540

7 REPRODUCIBILITY STATEMENT

541
542 All codes and dataset of the experiments in this work are submitted as a zip file via the Supplementary
543 Material. Details in the data generation and implementation are provided in Appendix B, and the
544 full proofs of Theorems and Lemma are provided in Appendix C along with their mathematical
545 restatements.546
547

8 ETHICS STATEMENT

548
549 All authors of this work sincerely adhered to the ICLR Code of Ethics. We do not expect any potential
550 violation to the best of our knowledge.551
552

REFERENCES

553
554 Diab W Abueidda, Panos Pantidis, and Mostafa E Mobasher. Deepokan: Deep operator network
555 based on kolmogorov arnold networks for mechanics problems. *Computer Methods in Applied*
556 *Mechanics and Engineering*, 436:117699, 2025.557
558 Alireza Afzal Aghaei. rkan: Rational kolmogorov-arnold networks. *arXiv preprint arXiv:2406.14495*,
559 2024.560 Giuseppe Carleo, Ignacio Cirac, Kyle Cranmer, Laurent Daudet, Maria Schuld, Naftali Tishby, Leslie
561 Vogt-Maranto, and Lenka Zdeborová. Machine learning and the physical sciences. *Reviews of*
562 *Modern Physics*, 91(4):045002, 2019.563 Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary
564 differential equations. *Advances in neural information processing systems*, 31, 2018.565
566 Tianping Chen and Hong Chen. Universal approximation to nonlinear operators by neural networks
567 with arbitrary activation functions and its application to dynamical systems. *IEEE transactions on*
568 *neural networks*, 6(4):911–917, 1995.569
570 Lily Chiparova and Vasil Popov. Kolmogorov–arnold networks for system identification of first-and
571 second-order dynamic systems. *Engineering Proceedings*, 100(1):100059, 2025.572 Richard Courant, David Hilbert, and Joseph David Harris. Methods of mathematical physics, volume
573 ii: Partial differential equations. *American Journal of Physics*, 31(3):221–221, 1963.574
575 George Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of Control,*
576 *Signals and Systems*, 2(4):303–314, 1989.577 Alexander D’Amour, Katherine Heller, Dan Moldovan, Ben Adlam, Babak Alipanahi, Alex Beutel,
578 Christina Chen, Jonathan Deaton, Jacob Eisenstein, Matthew D Hoffman, et al. Underspecification
579 presents challenges for credibility in modern machine learning. *Journal of Machine Learning*
580 *Research*, 23(226):1–61, 2022.581
582 Tim De Ryck, Samuel Lanthaler, and Siddhartha Mishra. On the approximation of functions by tanh
583 neural networks. *Neural Networks*, 143:732–750, 2021.584 Laurent Demanet and Lexing Ying. Discrete symbol calculus. *SIAM review*, 53(1):71–104, 2011.585
586 Waleed Diab and Mohammed Al Kobaisi. Temporal neural operator for modeling time-dependent
587 physical phenomena. *Scientific Reports*, 15(1):32791, 2025.

588 Onur Eker. Scalable sampling of deep neural operators. 2024.

589
590 Lawrence C Evans. *Partial differential equations*, volume 19. American mathematical society, 2022.591 Francesco Faa di Bruno. Sullo sviluppo delle funzioni. *Annali di scienze matematiche e fisiche*, 6
592 (1855):479–80, 1855.

593 Maurice Fréchet. Sur quelques points du calcul fonctionnel. 1906.

594 Nart Gashi, Panagiotis Kakosimos, and George Papafotiou. System identification using kolmogorov-
 595 arnold networks: A case study on buck converters. In *2025 IEEE International Conference on*
 596 *Prognostics and Health Management (ICPHM)*, pp. 1–7. IEEE, 2025.

597

598 David Gilbarg, Neil S Trudinger, David Gilbarg, and NS Trudinger. *Elliptic partial differential*
 599 *equations of second order*, volume 224. Springer, 1977.

600 David Gottlieb and Steven A Orszag. *Numerical analysis of spectral methods: theory and applications*.
 601 SIAM, 1977.

602

603 Anirudh Goyal and Yoshua Bengio. Inductive biases for deep learning of higher-level cognition.
 604 *Proceedings of the Royal Society A*, 478(2266):20210068, 2022.

605

606 Lars Hörmander. *The analysis of linear partial differential operators III: Pseudo-differential operators*.
 607 Springer Science & Business Media, 2007.

608

609 Kurt Hornik, Maxwell Stinchcombe, and Halbert White. Multilayer feedforward networks are
 610 universal approximators. *Neural Networks*, 2(5):359–366, 1989.

611

612 George Em Karniadakis, Ioannis G Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang.
 613 Physics-informed machine learning. *Nature Reviews Physics*, 3(6):422–440, 2021.

614

615 Benjamin C Koenig, Suyong Kim, and Sili Deng. Kan-odes: Kolmogorov–arnold network ordinary
 616 differential equations for learning dynamical systems and hidden physics. *Computer Methods in*
 617 *Applied Mechanics and Engineering*, 432:117397, 2024.

618

619 Joseph J Kohn and Louis Nirenberg. An algebra of pseudo-differential operators. *Communications*
 620 *on Pure and Applied Mathematics*, 18(1-2):269–305, 1965.

621

622 Takeshi Koshizuka, Masahiro Fujisawa, Yusuke Tanaka, and Issei Sato. Understanding the expres-
 623 sivity and trainability of fourier neural operator: A mean-field perspective. *Advances in Neural*
 624 *Information Processing Systems*, 37:11021–11060, 2024.

625

626 Nikola Kovachki, Samuel Lanthaler, and Siddhartha Mishra. On universal approximation and error
 627 bounds for fourier neural operators. *Journal of Machine Learning Research*, 22(290):1–76, 2021.

628

629 Nikola Kovachki, Zongyi Li, Burigede Liu, Kamyar Azizzadenesheli, Kaushik Bhattacharya, Andrew
 630 Stuart, and Anima Anandkumar. Neural operator: Learning maps between function spaces with
 631 applications to pdes. *Journal of Machine Learning Research*, 24(89):1–97, 2023.

632

633 Nikola B Kovachki, Samuel Lanthaler, and Hrushikesh Mhaskar. Data complexity estimates for
 634 operator learning. *arXiv preprint arXiv:2405.15992*, 2024a.

635

636 Nikola B Kovachki, Samuel Lanthaler, and Andrew M Stuart. Operator learning: Algorithms and
 637 analysis. *arXiv preprint arXiv:2402.15715*, 2024b.

638

639 Pijush K Kundu, Ira M Cohen, David R Dowling, and Jesse Capecelatro. *Fluid mechanics*. Elsevier,
 640 2024.

641

642 Samuel Lanthaler. Operator learning of lipschitz operators: An information-theoretic perspective.
 643 *arXiv preprint arXiv:2406.18794*, 2024.

644

645 Samuel Lanthaler, Siddhartha Mishra, and George E Karniadakis. Error estimates for deeponets: A
 646 deep learning framework in infinite dimensions. *Transactions of Mathematics and Its Applications*,
 647 6(1):tnac001, 2022.

648

649 Samuel Lanthaler, Zongyi Li, and Andrew M Stuart. Nonlocality and nonlinearity implies universality
 650 in operator learning. *Constructive Approximation*, pp. 1–43, 2025.

651

652 Guoqiang Lei, D Exposito, and Xuerui Mao. Discontinuity-aware kan-based physics-informed neural
 653 networks. *arXiv preprint arXiv:2507.08338*, 2025.

654

655 Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew
 656 Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations.
 657 *arXiv preprint arXiv:2010.08895*, 2020.

648 Chaoyu Liu, Davide Murari, Chris Budd, Lihao Liu, and Carola-Bibiane Schönlieb. Enhancing
 649 fourier neural operators with local spatial features. *arXiv preprint arXiv:2503.17797*, 2025.
 650

651 Ning Liu, Yiming Fan, Xianyi Zeng, Milan Klöwer, Lu Zhang, and Yue Yu. Harnessing the power of
 652 neural operators with automatically encoded conservation laws. *arXiv preprint arXiv:2312.11176*,
 653 2023.

654 Ziming Liu, Pingchuan Ma, Yixuan Wang, Wojciech Matusik, and Max Tegmark. Kan 2.0:
 655 Kolmogorov-arnold networks meet science. *arXiv preprint arXiv:2408.10205*, 2024a.
 656

657 Ziming Liu, Yixuan Wang, Sachin Vaidya, Fabian Ruehle, James Halverson, Marin Soljačić,
 658 Thomas Y Hou, and Max Tegmark. Kan: Kolmogorov-arnold networks. *arXiv preprint
 659 arXiv:2404.19756*, 2024b.

660 Miguel Liu-Schiaffini, Julius Berner, Boris Bonev, Thorsten Kurth, Kamyar Azizzadenesheli, and
 661 Anima Anandkumar. Neural operators with localized integral and differential kernels. *arXiv
 662 preprint arXiv:2402.16845*, 2024.
 663

664 Lu Lu, Pengzhan Jin, and George Em Karniadakis. Deeponet: Learning nonlinear operators for
 665 identifying differential equations based on the universal approximation theorem of operators. *arXiv
 666 preprint arXiv:1910.03193*, 2019.

667 Lu Lu, Pengzhan Jin, Guofei Pang, Zhongqiang Zhang, and George Em Karniadakis. Learning
 668 nonlinear operators via deeponet based on the universal approximation theorem of operators.
 669 *Nature machine intelligence*, 3(3):218–229, 2021.

670 Kent E Morrison. Spectral approximation of multiplication operators. *New York J. Math*, 1(75):96,
 671 1995.

672 Amir Noorizadegan, Sifan Wang, and Leevan Ling. A practitioner’s guide to kolmogorov-arnold
 673 networks. *arXiv preprint arXiv:2510.25781*, 2025.

674 Shaoxiang Qin, Fuyuan Lyu, Wenhui Peng, Dingyang Geng, Ju Wang, Xing Tang, Sylvie Leroyer,
 675 Naiping Gao, Xue Liu, and Liangzhu Leon Wang. Toward a better understanding of fourier neural
 676 operators from a spectral perspective. *arXiv preprint arXiv:2404.07200*, 2024.

677 Jun John Sakurai and Jim Napolitano. *Modern quantum mechanics*. Cambridge University Press,
 678 2020.

679 Jin Young Shin, Jae Yong Lee, and Hyung Ju Hwang. Pseudo-differential neural operator: Generalized
 680 fourier neural operator for learning solution operators of partial differential equations. *arXiv
 681 preprint arXiv:2201.11967*, 2022.

682 Hiroki Shiraishi, Hisao Ishibuchi, and Masaya Nakata. X-kan: Optimizing local kolmogorov-arnold
 683 networks via evolutionary rule-based machine learning. *arXiv preprint arXiv:2505.14273*, 2025.

684 Khemraj Shukla, Juan Diego Toscano, Zhicheng Wang, Zongren Zou, and George Em Karniadakis. A
 685 comprehensive and fair comparison between mlp and kan representations for differential equations
 686 and operator networks. *Computer Methods in Applied Mechanics and Engineering*, 431:117290,
 687 2024.

688 Alasdair Tran, Alexander Mathews, Lexing Xie, and Cheng Soon Ong. Factorized fourier neural
 689 operators. *arXiv preprint arXiv:2111.13802*, 2021.

690 Andrew Trask, Felix Hill, Scott E Reed, Jack Rae, Chris Dyer, and Phil Blunsom. Neural arithmetic
 691 logic units. *Advances in neural information processing systems*, 31, 2018.

692 Lloyd N Trefethen. *Spectral methods in MATLAB*. SIAM, 2000.

693 Feng Wang, Hong Qiu, Yingying Huang, Xiaozhe Gu, Renfang Wang, and Bo Yang. Efkan: A
 694 kan-integrated neural operator for efficient magnetotelluric forward modeling. *arXiv preprint
 695 arXiv:2502.02195*, 2025.

702 Hanchen Wang, Tianfan Fu, Yuanqi Du, Wenhao Gao, Kexin Huang, Ziming Liu, Payal Chandak,
 703 Shengchao Liu, Peter Van Katwyk, Andreea Deac, et al. Scientific discovery in the age of artificial
 704 intelligence. *Nature*, 620(7972):47–60, 2023.

705

706 Yixuan Wang, Jonathan W Siegel, Ziming Liu, and Thomas Y Hou. On the expressiveness and
 707 spectral bias of kans. *arXiv preprint arXiv:2410.01803*, 2024.

708 Gege Wen, Zongyi Li, Kamyar Azizzadenesheli, Anima Anandkumar, and Sally M Benson. U-
 709 fno—an enhanced fourier neural operator-based deep-learning model for multiphase flow. *Advances*
 710 *in Water Resources*, 163:104180, 2022.

711

712 Zipeng Xiao, Siqi Kou, Hao Zhongkai, Bokai Lin, and Zhijie Deng. Amortized fourier neural
 713 operators. *Advances in Neural Information Processing Systems*, 37:115001–115020, 2024.

714

715 Siyao Yang, Kun Lin, and Annan Zhou. The kan-mha model: A novel physical knowledge based
 716 multi-source data-driven adaptive method for airfoil flow field prediction. *Journal of Computational*
 717 *Physics*, 528:113846, 2025.

718

719 Dmitry Yarotsky. Error bounds for approximations with deep relu networks. *Neural networks*, 94:
 103–114, 2017.

720

721 Zhilin You, Zhenli Xu, and Wei Cai. Mscalefno: Multi-scale fourier neural operator learning for
 722 oscillatory function spaces. *arXiv preprint arXiv:2412.20183*, 2024.

723

724 Tianchi Yu, Jingwei Qiu, Jiang Yang, and Ivan Oseledets. Sinc kolmogorov-arnold network and its
 725 applications on physics-informed neural networks. *arXiv preprint arXiv:2410.04096*, 2024.

726

727

728

729

730

731

732

733

734

735

736

737

738

739

740

741

742

743

744

745

746

747

748

749

750

751

752

753

754

755

APPENDIX

A TABLE OF NOTATION

Table 5: Main symbols and notation used in the paper.

Symbol	Meaning / Definition
<i>Domains, spaces, and operators</i>	
$D := (-L/2, L/2)^d$	Periodic d -dimensional spatial box of side length L
\mathcal{G}	Ground-truth solution operator to be learned
<i>Spectral & spatial sampling</i>	
$\Xi = \{\xi_1, \dots, \xi_m\}$	Retained Fourier modes (truncated spectrum); $m = \Xi $
\mathcal{F}	Fourier transform
$\mathcal{F}_m, \mathcal{F}_m^{-1}$	Truncated Fourier transform by Ξ and its band-limited inverse
$\mathcal{Y} = \{y_1, \dots, y_m\} \subset D$	Uniform spatial grid
h	Grid spacing of \mathcal{Y}
<i>Fourier Neural Operator (FNO)</i>	
\mathcal{L}_{FNO}	Single FNO layer
$\mathcal{G}_{\theta}^{\text{FNO}}$	FNO network
$\mathbf{R}_{\theta}(\xi)$	Learnable block-diagonal spectral multiplier
\mathbf{W}_{θ}	Point-wise learnable linear map
$\sigma(\cdot)$	Point-wise non-linear activation
<i>KANO layer, symbol calculus, and projection</i>	
$\mathcal{L}_{\text{KANO}}$	Single KANO layer
$\mathbf{p}(\mathbf{x}, \xi)$	Learnable pseudo-differential symbol
$\text{Op}_m(\mathbf{p})$	Kohn-Nirenberg quantization of width m defined by \mathbf{p}
$\Pi_{\text{KANO}}(\mathcal{G})$	KANO projection of \mathcal{G}
$\mathcal{G}_{\theta}^{\text{KANO}}$	KANO network
Φ_{θ}	Learnable activation
<i>Kolmogorov–Arnold Network (KAN) primitives</i>	
$\phi_{q,p}^{(\ell)}(\cdot)$	1D edge function on layer ℓ , connecting p^{th} node of layer ℓ to q^{th} node of layer $(\ell + 1)$
$b(t), \{B_i(t)\}$	Base function and B-spline basis used to parametrize $\phi_{q,p}^{(\ell)}(t)$
<i>Q-KANO (quantum dynamics) notation</i>	
$\psi(x)$	Input wave function
$w(x)$	Quartic double-well potential
ΔT	Time step of propagation
$\phi_{\theta}(\mathbf{x}, \xi)$	Parametrized phase for symbol $\mathbf{p}(\mathbf{x}, \xi)$ of Q-KANO
$\vartheta_{\theta}(\cdot , \angle \cdot)$	Parametrized phase for non-linear activation of Q-KANO
$\mathcal{G}_{\theta}^{\text{Q-KANO}}$	Q-KANO network
<i>Function spaces</i>	
$L^2(D)$	Square-integrable function space on domain D
$H^s(D)$	Sobolev function space of order $s \geq 0$ on domain D

B EXPERIMENT DETAILS

B.1 SYNTHETIC OPERATOR BENCHMARK

All experiments are carried out on periodic functions $f : \mathbb{T} \rightarrow \mathbb{R}$ with $\mathbb{T} = (-\pi, \pi]$ and a uniform trigonometric grid

$$x_j = x_{\min} + j \Delta x, \quad \Delta x = \frac{2\pi}{N}, \quad j = 0, \dots, N - 1,$$

810 with $N = 128$. Unless noted otherwise every random quantity is drawn *independently for every*
 811 *sample*.

813 **Outer envelope.** To avoid the Gibbs phenomenon all basis functions are multiplied by a smooth
 814 taper that decays to zero in a $\pi/6$ -wide buffer near the periodic boundary:

$$816 \quad 817 \quad 818 \quad 819 \quad 820 \quad A(x) = \begin{cases} 1, & |x| \leq 5\pi/6, \\ \cos^4\left[\frac{|x| - 5\pi/6}{\pi/6} \frac{\pi}{2}\right], & 5\pi/6 < |x| < \pi, \\ 0, & |x| \geq \pi. \end{cases}$$

821 The full “base” function is always $f_{\text{base}}(x) = A(x) g(x)$.

823 **Spectral derivatives and ground-truth operator.** Derivatives are computed with an exact Fourier
 824 stencil:

$$825 \quad f'(x) = \mathcal{F}^{-1}[i\xi \hat{f}(\xi)], \quad f''(x) = \mathcal{F}^{-1}[-\xi^2 \hat{f}(\xi)].$$

827 $U[a, b]$ denotes random digit drawn from range $[a, b]$.

829 Training families (Group A)

831 A1. sine_beats:

$$832 \quad 833 \quad g(x) = \sin(\omega_1 x + \phi_1) \sin(\omega_2 x + \phi_2), \quad \omega_i = 8 U[0.5, 3], \phi_i = U[0, 2\pi].$$

834 A2. chirped_cosine:

$$835 \quad 836 \quad g(x) = \cos(\alpha x^2), \quad \alpha = 12 U[0.5, 2].$$

837 A3. periodic (random harmonic series):

$$838 \quad g(x) = \sin(\omega x + \phi_1) + \cos(\omega x + \phi_2), \quad \omega = 8 U[0.5, 3], \phi_{1,2} = U[0, 2\pi].$$

840 Unseen families (Group B)

842 B1. wave_packet:

$$844 \quad 845 \quad g(x) = \exp\left[-\frac{(x-\mu)^2}{2\sigma^2}\right] \sin(\omega x + \phi), \mu = U[-2, 2], \sigma = \frac{1}{12} U[0.5, 2], \omega = 12 U[2, 6], \phi = U[0, 2\pi].$$

846 B2. sinc_pulse:

$$848 \quad 849 \quad 850 \quad g(x) = \begin{cases} \frac{\sin(\alpha x)}{\alpha x}, & |x| > 10^{-12}, \\ 1, & |x| \leq 10^{-12}, \end{cases} \quad \alpha = 12 U[0.5, 3].$$

851 B3. gaussian_hermite:

$$853 \quad 854 \quad g(x) = H_n\left(\frac{x-\mu}{\sigma}\right) \exp\left[-\frac{(x-\mu)^2}{2\sigma^2}\right], n \in \{1, 2, 3\} \text{ uniform}, \mu = U[-2, 2], \sigma = \frac{1}{8} U[0.5, 2],$$

855 where H_n is the degree- n Hermite polynomial.

857 NORMALIZATION

859 Each realization is divided by its maximum absolute value, $\|f\|_\infty$, to obtain $\|f\|_\infty = 1$. The envelope
 860 guarantees periodicity and keeps the numerical spectrum sharply band-limited.

862 **Sample counts.** $\#train = 2000$ samples from the three Group A families for train data and
 863 $\#test = 400$ samples each from the Group A and Group B families for generalization tests..

864
865

B.2 QUANTUM DYNAMICS BENCHMARK

866
867
868
869
870
871
872
873

We model a quantum apparatus with 200 state-preparation protocols each with perfect reproducibility, capable of generating an identical initial state whose wave function is drawn from one of the three families: *Periodic*, *Gaussian wave-packet*, and *Gaussian×Hermite*. The prepared initial states evolve under one of two unknown, time-independent Hamiltonians, and two arrays of 128 detectors measure position and momentum on uniform grids, yielding probability mass functions (PMFs) every 0.1ms for 100 time steps. PMFs collected from the first 10 time steps are used for training the models, and the rest of the PMFs collected from the remaining 90 time steps are used to evaluate the long-horizon fidelity drop beyond the train steps.

874
875

B.2.1 QUANTUM APPARATUS ASSUMPTIONS

876
877
878
879
880

1. **State preparation.** A collection of calibrated protocols can each prepare a designated initial wave-function each of one of three real-valued families *Periodic*, *Gaussian wave-packet*, or *Gaussian–Hermite*: $\psi_0^{(m)}(x) \in L^2(\mathbb{T})^8$, $m = 1, \dots, 200$. Repeated shots under the same protocol start from exactly the same $\psi_0^{(m)}$, enabling trajectory-level reproducibility for every member of the ensemble.
2. **Hamiltonian stability.** The (unknown) Hamiltonian is *time-independent*, so trajectories are perfectly repeatable once ψ_0 is fixed.
3. **Dual-basis detection.** Two 128-grid projective detectors measure the position basis $\{|x_i\rangle\}$ and the momentum basis $\{|\xi_j\rangle\}$, yielding empirical probability mass functions (PMFs) $\hat{p}_x(i) = |\psi(x_i)|^2$ and $\hat{p}_\xi(j) = |\hat{\psi}(\xi_j)|^2$ on a common torus grid \mathbb{T}_L , $L = 4$.

887

B.2.2 DATA GENERATION DETAILS

888

For each of 200 distinct sample trajectories we

890
891
892
893
894

1. draw the initial wave function and propagate on the Hamiltonian with a high-resolution Strang split: $\delta t = 1 \mu\text{s}$ for 10 000 micro-steps, producing coarse snapshots every 100 steps ($\Delta T = 0.1 \text{ ms}$, $T = 1, \dots, 100$);
2. store $(\psi_T, \mathbf{p}_x^T, \mathbf{p}_\xi^T)$ where $p_x^T(i) = |\psi_T(x_i)|^2$ and $p_\xi^T(j) = |\hat{\psi}_T(\xi_j)|^2$.

895
896
897

Only the first 10 coarse steps are used for training; the remaining 90 steps test fidelity drop on long-horizon. All simulations employ an $n = 128$ -point FFT grid to match the detectors.

898
899
900
901
902
903

Spatial discretization. We place the problem on a periodic box of length $L = 4$ with $N = 128$ grid points $x_j = x_{\min} + j\Delta x$, $\Delta x = L/N$. Periods suppress wrap-around artifacts because every initial state is tapered by the smooth envelope $A(x)$ defined in Appendix B.1. Spatial derivatives are taken spectrally: let $\xi_m = 2\pi m/L$ for $m = -N/2, \dots, N/2 - 1$. Writing $\hat{\psi}_m = \mathcal{F}[\psi](\xi_m)$,

$$\partial_x \psi = \mathcal{F}^{-1}[i \xi_m \hat{\psi}_m], \quad \partial_{xx} \psi = \mathcal{F}^{-1}[-\xi_m^2 \hat{\psi}_m].$$

904
905
906
907
908

Strang-splitting time integrator. Let $\mathcal{K} := -\frac{1}{2}\partial_{xx}$ (kinetic), $\mathcal{V} : \psi \mapsto w(x) \cdot \psi$ (potential) and $\mathcal{N} : \psi \mapsto |\psi|^2 \cdot \psi$ (cubic nonlinearity). With time step Δt the second-order Strang factorization reads

$$e^{(\mathcal{K}+\mathcal{V}+\mathcal{N})\Delta t} = e^{\frac{\Delta t}{2}(\mathcal{K}+\mathcal{V})} e^{\Delta t \mathcal{N}} e^{\frac{\Delta t}{2}(\mathcal{K}+\mathcal{V})} + \mathcal{O}(\Delta t^3).$$

909

Because \mathcal{K} is diagonal in Fourier space and \mathcal{V} in real space we implement each half step explicitly:

910
911
912
913
914
915
916
917

$$\boxed{\begin{aligned} \psi &\xleftarrow{\mathcal{V}/2} e^{-\frac{i}{2}\Delta t w(x)} \psi, \\ \hat{\psi} &\xleftarrow{\mathcal{K}/2} \mathcal{F}[\psi], \quad \hat{\psi}_m \leftarrow e^{-\frac{i}{4}\Delta t \xi_m^2} \hat{\psi}_m, \\ \psi &\xleftarrow{\mathcal{N}} \mathcal{F}^{-1}[\hat{\psi}], \quad \psi \leftarrow e^{-i\Delta t |\psi|^2} \psi, \\ &\text{repeat } \mathcal{K}/2 \text{ and } \mathcal{V}/2. \end{aligned}}$$

⁸Square-integrable function space.

918 The L^2 norm is renormalized after every macro step to compensate numerical drift. For DW
 919 Hamiltonian, non-linearity time evolution is omitted.
 920

921 **Integrator parameters.** We use an inner step $\delta t = 10^{-6}$. A *macro* step of $\Delta T = 100 \delta t = 10^{-4}$
 922 is recorded and the sequence is propagated for $T_{\max} = 100 \Delta T = 10^{-2}$.
 923

924 **Initial-state families (real wave-functions).** Each sample draws a real profile $f(x)$ from
 925 $\{A_1, A_2, A_3\}$ below, multiplies it by the envelope A and normalizes it in L^2 :
 926

$$\psi_0(x) = \frac{A(x) f(x)}{\|A f\|_2}.$$

929 **A1. random harmonic series**

$$930 \quad f(x) = \sum_{m=0}^5 \frac{1}{2} \xi_m^{(c)} \cos(2\pi m \xi) + \frac{1}{2} \xi_m^{(s)} \sin(2\pi m \xi), \quad \xi = \frac{x+2}{4}, \quad \xi_m^{(\cdot)} \sim \mathcal{N}(0, 1).$$

934 **A2. Gaussian wave packet** $f(x) = \exp[-(x - x_0)^2/(2\sigma^2)]$ with $x_0 \sim U[-0.4\pi, 0.4\pi]$ and
 935 $\sigma \sim U[0.1, 0.3]$.
 936

A3. Gaussian–Hermite mode $f(x) = H_n\left(\frac{x-x_0}{\sigma}\right) \exp[-(x - x_0)^2/(2\sigma^2)]$ with $n \in \{0, 1, 2\}$
 937 uniform, x_0, σ as above, and H_n the Hermite polynomial.
 938

939 **Stored quantities.** For every sample index s and every snapshot $T \in \{0, 1, \dots, 100\}$ we save

$$(x, \psi(x, T), |\psi(x, T)|^2, |\hat{\psi}(\xi, T)|^2) \longrightarrow \text{wavefunc, pos_pdf, mom_pdf}.$$

942 All arrays are written in `float32` except the complex wave-function, stored as `complex64`. Alto-
 943 gether one call to `generate_dataset (num_samples=200)` produces $200 \times 101 \times 4 = 80,800$
 944 labeled records.
 945

946 **B.2.3 TRAIN TYPE DETAILS**

947 Let $\psi_{\Theta}^{(T)}$ be the T -step prediction of Q-KANO given ψ_0 .
 948

949 Train Type	950 Train Dataset	951 Loss function
952 full	953 complex $\psi^{(T)}$	$\mathcal{L}_{full} = \frac{\ \psi_{\Theta}^{(T)} - \psi^{(T)}\ _2}{\ \psi^{(T)}\ _2}$
954 pos	955 $\mathbf{p}_x^{(T)}$	$\mathcal{L}_{pos} = D_{KL}(\mathbf{p}_x^{(T)} \parallel \psi_{\Theta}^{(T)} ^2)$
956 pos & mom	957 $\mathbf{p}_x^{(T)}, \mathbf{p}_{\xi}^{(T)}$	$\mathcal{L}_{pos\&mom} = D_{KL}(\mathbf{p}_x^{(T)} \parallel \psi_{\Theta}^{(T)} ^2) + D_{KL}(\mathbf{p}_{\xi}^{(T)} \parallel \hat{\psi}_{\Theta}^{(T)} ^2)$

958 **C PROOFS**

959 **C.1 NOTATION AND PRELIMINARIES**

961 Throughout, $\mathbb{T}^d := [-\pi, \pi]^d$ denotes the flat d -torus and \mathbb{Z}^d the lattice of Fourier indices. For $\xi \in \mathbb{Z}^d$
 962 let $e_{\xi}(\mathbf{x}) := e^{i\xi \cdot \mathbf{x}}$. The Fourier coefficient of a square integrable function \mathbf{f} is
 963

$$\hat{\mathbf{f}}(\xi) := (2\pi)^{-d} \int_{\mathbb{T}^d} \mathbf{f}(\mathbf{x}) e^{-i\xi \cdot \mathbf{x}} d\mathbf{x}.$$

966 In a Sobolev space H^s with an order of smoothness $s \in \mathbb{R}$, the Sobolev norm of function \mathbf{f} is
 967

$$968 \quad \|\mathbf{f}\|_{H^s}^2 := \frac{(2\pi)^d}{2} \sum_{\xi \in \mathbb{Z}^d} (1 + |\xi|^{2s}) |\hat{\mathbf{f}}(\xi)|^2.$$

971 **Asymptotics.** Write $A \lesssim B$ if $A \leq C B$ for a constant C depending only on fixed parameters
 (dimension, regularity exponents, etc.).
 970

972 **Vector Notation** Fix a spatial dimension $d \geq 1$ and an index $j \in \{1, \dots, d\}$. For $\xi =$
 973 $(\xi_1, \dots, \xi_d) \in \mathbb{Z}^d$ we write
 974

$$\xi_{-j} := (\xi_1, \dots, \xi_{j-1}, \xi_{j+1}, \dots, \xi_d) \in \mathbb{Z}^{d-1}$$

975 for the vector obtained by *removing* the j -th coordinate of ξ . Conversely, for $\alpha = (\alpha_1, \dots, \alpha_{d-1}) \in$
 976 \mathbb{Z}^{d-1} and $n \in \mathbb{Z}$ we define the *insertion* map
 977

$$(\alpha, n)_j := (\alpha_1, \dots, \alpha_{j-1}, n, \alpha_j, \dots, \alpha_{d-1}) \in \mathbb{Z}^d.$$

980 We use $|\cdot|_\infty$ for the max–norm on \mathbb{Z}^d , i.e. $|\xi|_\infty = \max_{1 \leq i \leq d} |\xi_i|$. When we write $\hat{\mathbf{u}}(\alpha, r)$, this is
 981 shorthand for the d –dimensional coefficient $\hat{\mathbf{u}}((\alpha, r)_j)$.
 982

983 C.2 PROOF OF LEMMA 1

985 **Restatement of Lemma 1 (with explicit notation).** Let $\mathbf{u} \in H^s(\mathbb{T}^d)$ with $s > 0$, and assume that
 986 its Fourier coefficients are compactly supported:
 987

$$\hat{\mathbf{u}}(\xi) = \mathbf{0} \quad \text{for all } |\xi|_\infty > N_0, \quad \text{and} \quad \hat{\mathbf{u}} \not\equiv \mathbf{0}.$$

988 Fix $j \in \{1, \dots, d\}$ and set $\mathbf{v}(\mathbf{x}) := x_j \mathbf{u}(\mathbf{x})$. Then there exist
 989

$$\alpha \in \mathbb{Z}^{d-1}, \quad m \in \{1, 2, \dots, 2N_0 + 1\}, \quad c > 0, \quad R \in \mathbb{N},$$

990 and an infinite set of the fiber at j^{th} coordinate,
 991

$$\Xi_{\alpha, R} \subset \{\xi \in \mathbb{Z}^d : \xi_{-j} = \alpha, |\xi_j| \geq R\}$$

992 such that
 993

$$|\hat{\mathbf{v}}(\xi)| \geq \frac{c}{(1 + |\xi_j|)^m} \quad \forall \xi \in \Xi_{\alpha, R}. \quad (20)$$

994 In particular, if for some $\alpha \in \mathbb{Z}^{d-1}$ one has
 995

$$\sum_{r=-N_0}^{N_0} \hat{\mathbf{u}}((\alpha, r)_j) \neq \mathbf{0} \quad (\text{equivalently, } \sum_{r=-N_0}^{N_0} \hat{\mathbf{u}}(\alpha, r) \neq \mathbf{0} \text{ in the shorthand above}),$$

1000 then the bound equation 20 holds with the sharper exponent $m = 1$.
 1001

1002 *Proof.* We work with square-integrable function $\mathbf{v} \in L^2(\mathbb{T}^d)$ where $\mathbb{T}^d = [0, 2\pi]^d$.
 1003

1004 **Step 1 — The periodic “coordinate” and its Fourier coefficients.** The function $\mathbf{x} \mapsto x_j$ is not
 1005 periodic. Introduce the zero-mean, periodic 1D sawtooth
 1006

$$\psi_j(\mathbf{x}) := x_j - \pi, \quad \mathbf{x} \in [0, 2\pi]^d,$$

1007 extended periodically to \mathbb{T}^d . A direct computation (factorization of the integral and one-dimensional
 1008 integration by parts) shows that its Fourier coefficients are supported on the j^{th} -coordinate: for
 1009 $\xi \in \mathbb{Z}^d$,

$$\widehat{\psi}_j(\xi) = \begin{cases} -\frac{1}{i \xi_j}, & \text{if } \xi_{-j} = \mathbf{0} \text{ and } \xi_j \neq 0, \\ 0, & \text{if } \xi = \mathbf{0} \text{ or } \xi_{-j} \neq \mathbf{0}. \end{cases} \quad (21)$$

1010 Moreover $x_j = \psi_j + \pi$, hence
 1011

$$\widehat{\mathbf{v}}(\xi) = \mathcal{F}[\psi_j \mathbf{u}](\xi) + \pi \widehat{\mathbf{u}}(\xi). \quad (22)$$

1012 **Step 2 — Exact coefficient formula outside the support of $\widehat{\mathbf{u}}$.** Since $\widehat{\mathbf{u}}(k) = \mathbf{0}$ for $|\xi|_\infty > N_0$, the
 1013 second term in equation 22 vanishes whenever $|\xi|_\infty > N_0$. Using equation 21 and the convolution
 1014 theorem, we obtain for any $\xi \in \mathbb{Z}^d$ with $|\xi|_\infty > N_0$:

$$\widehat{\mathbf{v}}(\xi) = \sum_{\ell \in \mathbb{Z}^d} \widehat{\psi}_j(\xi - \ell) \widehat{\mathbf{u}}(\ell) = -\frac{1}{i} \sum_{\substack{\ell_{-j} = \xi_{-j} \\ |\ell_j| \leq N_0}} \frac{\widehat{\mathbf{u}}(\ell)}{\xi_j - \ell_j}. \quad (23)$$

1026 Thus, along any fixed transverse index $\alpha := \xi_{-j} \in \mathbb{Z}^{d-1}$, the tail $\hat{\mathbf{v}}(\alpha, n)$ for large $|n|$ is a finite
 1027 sum of simple fractions in the single variable n .
 1028

1029 **Step 3 — Moments on a nontrivial fiber and the first non-vanishing moment.** Because $\hat{\mathbf{u}} \neq 0$,
 1030 there exists at least one $\alpha \in \mathbb{Z}^{d-1}$ for which the fiber

$$1031 \quad \mathcal{F}_\alpha := \{r \in \mathbb{Z} : \hat{\mathbf{u}}(\alpha, r) \neq 0\}$$

1032 is nonempty. Define the (finite) coefficients $\mathbf{c}_r := \hat{\mathbf{u}}(\alpha, r)$ for $r \in [-N_0, N_0]$, and their moments
 1033

$$1034 \quad \mu_p := \sum_{r=-N_0}^{N_0} r^p \mathbf{c}_r \quad (p \geq 0).$$

1037 Let $m \in \{1, \dots, 2N_0 + 1\}$ be the smallest index for which $\mu_{m-1} \neq 0$. Such an m exists since not
 1038 all \mathbf{c}_r vanish.

1039 **Step 4 — Asymptotics along a line and a polynomial lower bound.** For $n \in \mathbb{Z}$ with $|n| > N_0$,
 1040 formula equation 23 along the line $\xi_{-j} = \alpha$ reads
 1041

$$1042 \quad \hat{\mathbf{v}}(\alpha, n) = -\frac{1}{i} \sum_{r=-N_0}^{N_0} \frac{\mathbf{c}_r}{n-r}.$$

1045 Expanding $\frac{1}{n-r} = \frac{1}{n} \sum_{q \geq 0} \left(\frac{r}{n}\right)^q$ for $|n| > 2N_0$ and collecting terms yields the asymptotic expansion
 1046

$$1047 \quad \hat{\mathbf{v}}(\alpha, n) = -\frac{1}{i} \left(\frac{\mu_0}{n} + \frac{\mu_1}{n^2} + \dots + \frac{\mu_{m-1}}{n^m} + O\left(\frac{1}{|n|^{m+1}}\right) \right), \quad |n| \rightarrow \infty.$$

1049 By the choice of m , the first nonzero term is μ_{m-1}/n^m . Consequently, there exist $R \in \mathbb{N}$ and $c > 0$
 1050 such that

$$1051 \quad |\hat{\mathbf{v}}(\alpha, n)| \geq \frac{c}{|n|^m} \quad \text{for all } |n| \geq R, n \in \mathbb{Z}.$$

1053 **Step 5 — Conclusion and the special case $m = 1$.** Then equation 20 holds for all $\xi \in \Xi_{\alpha, R}$ with the
 1054 exponent m determined in Step 3, by the definition of $\Xi_{\alpha, R}$ and that $|n| \geq R$. If $\mu_0 = \sum_r \mathbf{c}_r \neq 0$
 1055 for the chosen fiber (equivalently, $\sum_{r=-N_0}^{N_0} \hat{\mathbf{u}}(\alpha, r) \neq 0$), then $m = 1$ and we obtain the sharper
 1056 $|\hat{\mathbf{v}}(\xi)| \gtrsim (1 + |\xi_j|)^{-1}$ along the line $\xi_{-j} = \alpha$. \square
 1057

1058 **Remark 4.** The explicit one-dimensional tail equation 23 shows that multiplying a band-limited
 1059 field by the coordinate x_j produces a polynomial Fourier tail decay along lines parallel to the j^{th}
 1060 coordinate, with rate $(1 + |\xi_j|)^{-m}$ where m is the first non-vanishing moment of the finitely many
 1061 coefficients on the relevant fiber. In particular, when $m = 1$, the decay is exactly $(1 + |\xi_j|)^{-1}$. Such
 1062 algebraic tails are consistent with the pseudo-spectral projection error estimate quoted in Kovachki
 1063 et al. (2021, Thm. 40).

1064 **C.3 PROOF OF THEOREM 1**

1066 **Restate of Theorem 1.** Let $\alpha = (\alpha_1, \dots, \alpha_d) \in \mathbb{N}^d$ with total degree $M := |\alpha| \geq 1$ and define
 1067 the position-multiplier

$$1068 \quad \mathcal{M}(\mathbf{x}) := x_1^{\alpha_1} x_2^{\alpha_2} \cdots x_d^{\alpha_d}.$$

1069 For inputs band-limited to radius N_0 and lying in $H^s(\mathbb{T}^d)$ with $s > \frac{d}{2}$, any Fourier Neural Operator
 1070 $\mathcal{G}_\theta^{\text{FNO}}$ that achieves $\|\mathcal{M}(\mathbf{x}) - \mathcal{G}_\theta^{\text{FNO}}\|_{H^s \rightarrow H^{s-\delta}} \leq \varepsilon$ ($0 < \delta < 1$) must employ a spectral bandwidth
 1071 (FNO width) $m \gtrsim \varepsilon^{-M/s}$ and a parameter count $|\theta| \geq \exp(c \varepsilon^{-Md/s})$ for some $c > 0$ depending
 1072 only on (d, s, δ, N_0) .

1074 *Proof.* Set $s' := s - \delta$ with $0 < \delta < 1$.

1076 **Step 1 — Algebraic tail produced by $\mathcal{M}(\mathbf{x})$.** Applying Lemma 1 once per factor of x_j shows that
 1077 for some constant $C_0 > 0$ and an infinite set $\Xi_\infty \subset \mathbb{Z}^d$,

$$1079 \quad |\mathcal{F}[\mathcal{M}(\mathbf{x})\mathbf{u}](\xi)| \geq \frac{C_0}{(1 + |\xi|)^{M+1}} \quad \forall \xi \in \Xi_\infty. \quad (24)$$

1080
1081 **Step 2 – Pseudo-spectral projection lower bound.** For any $\mathbf{f} \in H^s$, pseudo-spectral projection
1082 error estimate gives (Gottlieb & Orszag, 1977)

1083
$$\|(\mathcal{I} - \Pi_{\text{FNO}})\mathbf{f}\|_{H^{s'}} \geq C_1 \left[\sum_{|\xi|_\infty > N} (1 + |\xi|^{2s'}) |\widehat{\mathbf{f}}(\xi)|^2 \right]^{1/2},$$

1084
1085

1086 where \mathcal{I} is an identity operator. Insert $\mathbf{f} = \mathcal{M}(\mathbf{x})\mathbf{u}$ and the tail bound equation 24; summing over
1087 $\Xi_\infty \cap \{|\xi|_\infty > m\}$ yields

1088
$$\|(\mathcal{I} - \Pi_{\text{FNO}})\mathcal{M}(\mathbf{x})\mathbf{u}\|_{H^{s'}} \gtrsim m^{-(M-\delta)}.$$

1089

1090 Imposing this residual $\leq \frac{1}{2}\varepsilon$ forces
1091

1092
$$m \geq C_2 \varepsilon^{-M/s}, \quad C_2 = C_2(d, s, \delta, M, N_0) > 0. \quad (25)$$

1093

1094 **Step 3 – Canonical neural scaling in the latent map.** An FNO with spectral radius (width) m
1095 manipulates a latent vector of dimension $(2m+1)^d \sim m^d$. Approximating a generic Lipschitz map
1096 $\mathbf{G} : \mathbb{C}^{m^d} \rightarrow \mathbb{C}^{m^d}$ to accuracy $\varepsilon/2$ with a fully connected network requires (Yarotsky, 2017; De Ryck
1097 et al., 2021) *at least*

1098
$$\text{parameters} \gtrsim \varepsilon^{-m^d}.$$

1099

1100 (Kovachki et al., 2021, Remark 22)

1101 **Step 4 – Substitute the bandwidth lower bound.** Using equation 25, $m^d \sim \varepsilon^{-Md/s}$. Hence the
1102 latent network must have at least

1103
$$|\theta| \gtrsim \varepsilon^{-\varepsilon^{-Md/s}},$$

1104

1105 a *super-exponential* curse of dimensionality in the target accuracy ε .
1106

□

1109 C.4 PROOF OF THEOREM 2

1110 **Restate of Theorem 2.** Let $s \geq s' \geq 0$ and $s_p > d/2$. Consider any finite composition

1111
$$\mathcal{G} = (\varsigma_\ell \circ \mathcal{L}_\ell) \circ \cdots \circ (\varsigma_1 \circ \mathcal{L}_1), \quad \mathcal{L}_i = \text{Op}_m(\mathbf{p}_i),$$

1112

1113 where each Kohn–Nirenberg symbol \mathbf{p}_i belongs to $W^{s_p, 2}(\mathbb{T}_x^d \times \mathbb{T}_\xi^d)$ ⁹, and each ς_i is a point-wise non-
1114 linearity of uniformly bounded degree (so the number of such nonlinearities is $O(\ell)$ and independent
1115 of ε). For every $\varepsilon > 0$ there exists a *single-layer* KANO $\mathcal{G}_\theta^{\text{KANO}}$ such that
1116

1117
$$\|\mathcal{G} - \mathcal{G}_\theta^{\text{KANO}}\|_{H^s \rightarrow H^{s'}} \leq \varepsilon, \quad |\theta| = O(\varepsilon^{-d/(2s_p)}).$$

1118

1119 *Proof.* **Step 1 – Kohn–Nirenberg quantization error.** By the quadrature bound from Demanet &
1120 Ying (2011), for each \mathcal{L}_i we have
1121

1122
$$\|\mathcal{L}_i - \text{Op}_m(\mathbf{p}_i)\|_{H^s \rightarrow H^{s'}} \leq C' B m^{-(s-s')},$$

1123

1124 where $\text{Op}_m(\mathbf{p}_i)$ keeps only frequencies $|\xi|_\infty \leq m$. Since each ς_i is a bounded-degree pointwise
1125 map, its Nemytskii operator is Lipschitz on bounded sets; write $L_i := \text{Lip}(\varsigma_i)$ ¹⁰ on the relevant
1126 range and set $L_{\max} := \max_i L_i$. Let $M := \max_j \{\|\mathcal{L}_j\|_{H^{s'} \rightarrow H^{s'}}, \|\text{Op}_m(\mathbf{p}_j)\|_{H^{s'} \rightarrow H^{s'}}\}$. Because
1127 $s_p > \frac{d}{2}$ implies $\|\mathbf{p}_j\|_{L^\infty} \lesssim \|\mathbf{p}_j\|_{W^{s_p, 2}}$, both \mathcal{L}_j and $\text{Op}_m(\mathbf{p}_j)$ are bounded on $H^{s'}$ with a bound
1128 independent of m . A telescoping estimate for the interleaved composition then yields
1129

1130
$$\|\mathcal{G} - \Pi_{\text{KANO}}(\mathcal{G})\|_{H^s \rightarrow H^{s'}} \leq C_* C' B m^{-(s-s')}, \quad C_* \leq \ell (L_{\max} M)^{\ell-1} L_{\max}.$$

1131

1132 ⁹Standard square integrable periodic Sobolev space on the product torus \mathbb{T}^{2d} .

1133 ¹⁰Lipschitz constant of the point-wise nonlinearity ς_i on the relevant value range, i.e. $L_i := \sup_{a \neq b} \frac{|\varsigma_i(a) - \varsigma_i(b)|}{|a-b|}$ with a, b restricted to the compact interval attained by the i -th preactivations.

1134 Pick $m := (2C_* C' B / \varepsilon)^{1/(s-s')}$. Then the full composition \mathcal{G} deviates from its projected counter-
 1135 part

$$\Pi_{\text{KANO}}(\mathcal{G}) := (\varsigma_\ell \circ \mathbf{Op}_m(\mathbf{p}_\ell)) \circ \cdots \circ (\varsigma_1 \circ \mathbf{Op}_m(\mathbf{p}_1))$$

1136 by at most $\varepsilon/2$ in $H^s \rightarrow H^{s'}$ operator norm.

1139

1140 **Step 2 – KAN approximation of both symbols and pointwise nonlinearities.** By the width-fixed
 1141 KAN scaling law from Wang et al. (2024), for any $\eta > 0$ there exists a Kolmogorov–Arnold Network
 1142 $\Phi_{i,\eta}$ with $|\Phi_{i,\eta}| = O(\eta^{-d/(2s_p)})$ such that $\|\mathbf{p}_i - \Phi_{i,\eta}\|_{L^\infty} \leq \eta$ on $\mathbb{T}_x^d \times \mathbb{T}_\xi^d$. Likewise, since each ς_i
 1143 is a fixed-degree pointwise map, its action over the compact value range visited by the projected flow
 1144 can be uniformly approximated by a width-fixed KAN $\Psi_{i,\eta}$ with size $|\Psi_{i,\eta}| = O(\eta^{-1/(2s_p)})$ and
 1145 $\|\varsigma_i - \Psi_{i,\eta}\|_{L^\infty} \leq \eta$. Choose

$$\eta_{\text{sym}} := \frac{\varepsilon}{4\ell}, \quad \eta_{\text{nl}} := \frac{\varepsilon}{4\ell}.$$

1146 Define the single-layer KANO
 1147

$$\mathcal{G}_\theta^{\text{KANO}} := (\Psi_{\ell,\eta_{\text{nl}}} \circ \mathbf{Op}_m(\Phi_{\ell,\eta_{\text{sym}}})) \circ \cdots \circ (\Psi_{1,\eta_{\text{nl}}} \circ \mathbf{Op}_m(\Phi_{1,\eta_{\text{sym}}})),$$

1148 where the same spectral radius m from Step 1 is used in every $\mathbf{Op}_m(\cdot)$.

1149

1150 **Step 3 – Error accumulation beyond projection.** Using linearity of the symbol-to-operator map
 1151 and stability of Nemytskii (pointwise) maps under uniform approximation, the post-projection error
 1152 splits into a sum of the symbol parts and the nonlinearity parts:

$$\begin{aligned} \|\Pi_{\text{KANO}}(\mathcal{G}) - \mathcal{G}_\theta^{\text{KANO}}\|_{H^s \rightarrow H^{s'}} &\leq \sum_{i=1}^{\ell} \left\| \mathbf{Op}_m(\mathbf{p}_i) - \mathbf{Op}_m(\Phi_{i,\eta_{\text{sym}}}) \right\|_{H^s \rightarrow H^{s'}} + \sum_{i=1}^{\ell} \|\varsigma_i - \Psi_{i,\eta_{\text{nl}}}\|_{L^\infty} \\ &\leq \ell \eta_{\text{sym}} + \ell \eta_{\text{nl}} = \varepsilon/2. \end{aligned}$$

1153 Combining with Step 1 yields $\|\mathcal{G} - \mathcal{G}_\theta^{\text{KANO}}\|_{H^s \rightarrow H^{s'}} \leq \varepsilon$.

1154

1155 **Step 4 – Parameter complexity.** Summing the sizes of all KAN blocks gives

$$|\theta| = \sum_{i=1}^{\ell} O(\eta_{\text{sym}}^{-d/(2s_p)}) + \sum_{i=1}^{\ell} O(\eta_{\text{nl}}^{-1/(2s_p)}) + O(1) = O((\varepsilon/\ell)^{-d/(2s_p)}) + O((\varepsilon/\ell)^{-1/(2s_p)}).$$

1156 Since ℓ and the number/degree of the ς_i are fixed (do not scale with ε), the dominating term is
 1157 $O(\varepsilon^{-d/(2s_p)})$, establishing the claimed complexity. \square

1158

D ADDITIONAL DETAILS ON THE PURE-SPECTRAL BOTTLENECK OF FNO

1159

1160 In Section 3.1 we performed a linear analysis of a single FNO layer via its Jacobian to illustrate the
 1161 pure-spectral bottleneck. We showed that, in the first-order approximation, all spectral off-diagonals
 1162 arise from the spectrum of the input-dependent gate $\sigma'(\mathbf{z}(\mathbf{u}, \cdot))$ (12). Hence, although a single
 1163 FNO layer is capable of generating spectral off-diagonals, they are tied to the input distribution and
 1164 leads to the structural fragility in out-of-distribution performance (generalization on the unseen input
 1165 distribution). Experimental results in Section 5 are aligned with our concerns.

1166

1167 In this Appendix, we show that the same phenomenon persists at *every* order of the Fréchet expansion
 1168 of a single FNO layer, and composing deep layers does not remove the input distribution dependence
 1169 of the spectral off-diagonals generated by the model as well.

1170 Throughout, we work on the flat torus $\mathbb{T}^d = [-\pi, \pi]^d$ and Sobolev spaces H^s as in Appendix C.

1171

D.1 HIGHER-ORDER SPECTRAL OFF-DIAGONALS OF FNO

1172

1173 Recall the FNO layer \mathcal{L}_{FNO} (4):

1174

$$\mathcal{L}_{\text{FNO}}(\mathbf{u})(\mathbf{x}) = \sigma \left(\mathcal{F}_m^{-1} \left(\mathbf{R}_{\theta_1}(\xi) \cdot \mathcal{F}_m(\mathbf{u})(\xi) \right) (\mathbf{x}) + \mathbf{W}_{\theta_2} \cdot \mathbf{u}(\mathbf{x}) \right).$$

1188 For the analysis it is convenient to collect the linear terms into a linear operator
 1189

$$1190 \quad \mathcal{A} := \mathcal{F}_m^{-1} \circ \mathbf{R}_\theta(\boldsymbol{\xi}) \circ \mathcal{F}_m + \mathbf{W}_\theta, \quad (26)$$

1191 so that

$$1192 \quad \mathcal{L}_{\text{FNO}}(\mathbf{u})(\mathbf{x}) = \sigma(\mathbf{z}(\mathbf{u}, \mathbf{x})), \quad \mathbf{z}(\mathbf{u}, \mathbf{x}) := (\mathcal{A}\mathbf{u})(\mathbf{x}), \quad (27)$$

1193 where σ is the non-lineaer activation acting point-wise.
 1194

1195 We first compute the k -th Fréchet derivative (Fréchet, 1906) of $\mathcal{L}_{\text{FNO}}(\mathbf{u})(\mathbf{x})$.
 1196

1196 **Proposition 1 (Structure of higher-order derivatives of a single Fourier layer).** *Let $\mathbf{u} \in H^s$ and k be the order of the Fréchet derivative. Also let $\mathbf{h}_1, \dots, \mathbf{h}_k \in H^s$ be the arbitrary direction functions of each order. Then for every integer $k \geq 1$, the Fréchet derivative of $\mathcal{L}_{\text{FNO}}(\mathbf{u})(\mathbf{x})$ is:*
 1197

$$1199 \quad 1200 \quad \mathcal{D}^k \mathcal{L}_{\text{FNO}}(\mathbf{u})[\mathbf{h}_1, \dots, \mathbf{h}_k](\mathbf{x}) := \sigma^{(k)}(\mathbf{z}(\mathbf{u}, \mathbf{x})) \prod_{j=1}^k (\mathcal{A}\mathbf{h}_j)(\mathbf{x}), \quad (28)$$

1202 where $\sigma^{(k)}$ is the usual scalar k -th derivative. More explicitly, for each channel q ,
 1203

$$1204 \quad 1205 \quad [\mathcal{D}^k \mathcal{L}_{\text{FNO}}(\mathbf{u})[\mathbf{h}_1, \dots, \mathbf{h}_k]]_q(\mathbf{x}) := \sigma^{(k)}(\mathbf{z}_q(\mathbf{u}, \mathbf{x})) \prod_{j=1}^k [\mathcal{A}\mathbf{h}_j]_q(\mathbf{x}). \quad (29)$$

1207 *Proof.* Since \mathcal{A} (26) is linear, $\mathcal{D}\mathbf{z}(\mathbf{u}) \equiv \mathcal{A}$ and $\mathcal{D}^r \mathbf{z}(\mathbf{u}) \equiv 0$ for all $r \geq 2$. By the chain rule,
 1208

$$1209 \quad \mathcal{D}\mathcal{L}_{\text{FNO}}(\mathbf{u})[\mathbf{h}_i](\mathbf{x}) = \sigma'(\mathbf{z}(\mathbf{u}, \mathbf{x})) (\mathcal{A}\mathbf{h}_i)(\mathbf{x}),$$

1211 where σ' acts point-wise on each channel. For higher order derivatives, we apply the Faà di Bruno
 1212 formula (Faa di Bruno, 1855). Because all higher derivatives of \mathbf{z} vanish, every term involving
 1213 $\mathcal{D}^r \mathbf{z}(\mathbf{u})$ with $r \geq 2$ drops out, leaving

$$1214 \quad \mathcal{D}^k \mathcal{L}_{\text{FNO}}(\mathbf{u})[\mathbf{h}_1, \dots, \mathbf{h}_k] = \sigma^{(k)}(\mathbf{z}(\mathbf{u})) [\mathcal{A}\mathbf{h}_1, \dots, \mathcal{A}\mathbf{h}_k]. \quad (30)$$

1216 Since σ acts point-wise, this reduces to Eq. (28), and to Eq. (equation 29) on each channel. \square
 1217

1218 Thus, as apparent in Eq. (29), *every* Fréchet derivative is a point-wise product of
 1219

- 1220 • an input-dependent gate $\sigma^{(k)}(\mathbf{z}(\mathbf{u}, \mathbf{x}))$ depending on the current input \mathbf{u} , and
- 1221 • the product of k fixed linear responses $(\mathcal{A}\mathbf{h}_i)(\mathbf{x})$ incapable of generating spectral off-diagonals.

1223 Therefore, the structure of the high-order derivatives of a single layer FNO is just as tied to the input
 1224 \mathbf{u} distribution as the first-order Jacobian shown in Section 3.1.
 1225

1226 D.2 SPECTRAL REPRESENTATION OF HIGH-ORDER DERIVATIVES AND MULTI-LAYER FNO

1228 Fix $k \geq 1$ and define

$$1229 \quad \mathbf{s}_k(\mathbf{u}, \mathbf{x}) := \sigma^{(k)}(\mathbf{z}(\mathbf{u}, \mathbf{x})), \quad \mathbf{b}_i(\mathbf{x}) := (\mathcal{A}\mathbf{h}_i)(\mathbf{x}), \quad i = 1, \dots, k. \quad (31)$$

1231 Then Eq. (28) reads
 1232

$$1233 \quad 1234 \quad \mathcal{D}^k \mathcal{L}_{\text{FNO}}(\mathbf{u})[\mathbf{h}_1, \dots, \mathbf{h}_k](\mathbf{x}) = \mathbf{s}_k(\mathbf{u}, \mathbf{x}) \cdot \prod_{i=1}^k \mathbf{b}_i(\mathbf{x}). \quad (32)$$

1236 The Fourier transform of a product of $(k+1)$ functions is a $(k+1)$ -fold convolution, so for each
 1237 output frequency $\boldsymbol{\xi} \in \mathbb{Z}^d$ we have

$$1238 \quad 1239 \quad \widehat{\mathcal{D}^k \mathcal{L}_{\text{FNO}}(\mathbf{u})[\mathbf{h}_1, \dots, \mathbf{h}_k]}(\boldsymbol{\xi}) = (\widehat{\mathbf{s}_k(\mathbf{u})} * \widehat{\mathbf{b}_1} * \dots * \widehat{\mathbf{b}_k})(\boldsymbol{\xi}). \quad (33)$$

1240 Spectrally, the linear operator \mathcal{A} (26) acts diagonally:
 1241

$$\widehat{\mathcal{A}\mathbf{h}_i}(\boldsymbol{\xi}) = \mathbf{A}(\boldsymbol{\xi}) \widehat{\mathbf{h}_i}(\boldsymbol{\xi}), \quad (34)$$

1242 where $\mathbf{A}(\xi) = \mathbf{R}(\xi) + \mathbf{W}$ as in Eq. (12). Expanding the convolution in Eq. (33) yields
 1243

$$1244 \mathcal{D}^k \mathcal{L}_{\text{FNO}}(\widehat{\mathbf{u}})[\mathbf{h}_1, \dots, \mathbf{h}_k](\xi) = \sum_{\xi_1, \dots, \xi_k \in \mathbb{Z}^d} \underbrace{\widehat{\mathbf{s}_k(\mathbf{u})}(\xi - \xi_1 - \dots - \xi_k)}_{\text{depends on input } \mathbf{u}} \cdot \prod_{i=1}^k \mathbf{A}(\xi_i) \widehat{\mathbf{h}_i}(\xi_i), \quad (35)$$

1247 where $\{\xi_i\}$ are the frequency variable for each direction $\{\mathbf{h}_i\}$. Eq. (35) shows that the k -th derivative
 1248 could be read as a multi-linear operator whose *spectral kernel*

$$1249 \mathbf{K}_k(\mathbf{u}; \xi; \xi_1, \dots, \xi_k) := \widehat{\mathbf{s}_k(\mathbf{u})}(\xi - \xi_1 - \dots - \xi_k) \cdot \prod_{i=1}^k \mathbf{A}(\xi_i) \quad (36)$$

1252 connects (ξ_1, \dots, ξ_k) to the output function frequency ξ .
 1253

1254 Two structural facts follow immediately from Eq. (36):

1255 1. Input function \mathbf{u} decides the value $\widehat{\mathbf{s}_k(\mathbf{u})} = \mathcal{F}[\sigma^{(k)}(\mathbf{z}(\mathbf{u}, \cdot))]$,
 1256 2. The spectral multipliers $\mathbf{A}(\xi_i)$ depend only on the learned weights and remain diagonal in Fourier:
 1257 they are independent to input \mathbf{u} but does not contribute to spectral off-diagonals.

1259 Consequently, for any $k \geq 1$:

1261 *Every non-diagonal spectral coupling is induced by Fourier coefficients of non-*
 1262 *linear activation derivatives $\sigma^{(k)}(\mathbf{z}(\mathbf{u}, \cdot))$ evaluated on the current input \mathbf{u} .*

1264 Higher orders $k \geq 2$ introduce convolutions of higher derivatives $\sigma^{(k)}$, but they do not mitigate the
 1265 input dependence argument of spectral off-diagonals provided in Section 3.1.

1266 **Deep FNO networks.** A full FNO is a composition of multiple FNO layers:

$$1268 \mathcal{G}_{\theta}^{\text{FNO}} = \mathcal{L}_{\text{FNO}}^{(\ell)} \circ \dots \circ \mathcal{L}_{\text{FNO}}^{(1)}. \quad (37)$$

1270 Applying the Faà di Bruno formula (Faà di Bruno, 1855) to this composition, every term in
 1271 $\mathcal{D}^k \mathcal{G}_{\theta}^{\text{FNO}}(\mathbf{u})$ becomes a product of:

1272 • spectral multipliers $\{\mathbf{A}_j(\xi)\}$ from the linear parts of the layers, and
 1273 • Fourier transforms of activation derivatives $\sigma^{(r)}(\mathbf{z}^{(j)}(\mathbf{u}, \cdot))$ from intermediate layers j and
 1274 derivative orders $r \geq 1$.
 1275

1276 Thus, in a *deep* FNO, all spectral off-diagonals at any order k still factors through Fourier transforms
 1277 of activation derivatives evaluated on intermediate pre-activations governed by the input \mathbf{u} . Adding
 1278 layers introduces more such gates but do not mitigate the input dependence on spectral off-diagonals
 1279 introduced in Section 3.1.

1280 **D.3 THE MEANING OF INPUT DEPENDENCE OF SPECTRAL OFF-DIAGONALS WHEN**
 1281 **LEARNING A SPECTRALLY DENSE GROUND-TRUTH OPERATORS**

1283 Let \mathcal{T} be a fixed linear operator such as the position multiplier $a(x) \mapsto x^2 a(x)$ in Section 3.1. In
 1284 Fourier basis, \mathcal{T} is represented by a dense Toeplitz matrix such as $\mathbf{T}_n[x^2]$ as in Eq. (11), and its
 1285 Fréchet derivatives are

$$1287 \mathcal{D}\mathcal{T}(\mathbf{u}) \equiv \mathcal{T} \quad \text{for all } \mathbf{u}, \quad \mathcal{D}^k \mathcal{T}(\mathbf{u}) \equiv 0 \quad \text{for all } k \geq 2,$$

1288 i.e., the spectral off-diagonals of \mathcal{T} are completely *independent* of the input \mathbf{u} .

1289 Suppose we wish to learn \mathcal{T} with robust out-of-distribution performance over a Sobolev ball
 1290

$$1291 \mathcal{B}_B := \{\mathbf{u} \in H^s : \|\mathbf{u}\|_{H^s} \leq B\}.$$

1292 A natural notion of generalization is a small error in operator norm for an unseen Sobolev ball \mathcal{B}_B :

$$1293 \sup_{\mathbf{u} \in \mathcal{B}_B} \frac{\|\mathcal{G}_{\theta}^{\text{FNO}}(\mathbf{u}) - \mathcal{T}\mathbf{u}\|_{H^s}}{\|\mathbf{u}\|_{H^s}} \leq \varepsilon. \quad (38)$$

1295 This requires that on \mathcal{B}_B

1296 • the Jacobian $\mathcal{D}\mathcal{G}_\theta^{\text{FNO}}(u)$ stays close to \mathcal{T} , and
 1297 • all higher-order derivatives $\mathcal{D}^k\mathcal{G}_\theta^{\text{FNO}}(\square)$, $k \geq 2$, remain uniformly small (negligible).

1299 Combining Eq. (36) with the composition structure above reveals a paradox:

1301 1. If the gates $\sigma^{(k)}(\mathbf{z}^{(j)}(\mathbf{u}, \cdot))$ and their Fourier transforms $\widehat{s_k^{(j)}(\mathbf{u})}$ vary significantly with
 1302 $\mathbf{u} \in \mathcal{B}_B$, then both the Jacobian and all higher-order kernels $\mathbf{K}_k(\mathbf{u}; \xi, \xi_1, \dots, \xi_k)$ vary with
 1303 the input. In that case the effective spectral off-diagonals of $\mathcal{G}_\theta^{\text{FNO}}$ cannot coincide with a
 1304 single, input-independent Toeplitz kernel across all $\mathbf{u} \in \mathcal{B}_B$. Any dense off-diagonal pattern
 1305 learned from the train distribution will be tied to the train subspace and becomes fragile
 1306 under distribution shift, as observed in our experiments in Section 5.

1307 2. If, on the other hand, we try to make these kernels effectively independent of \mathbf{u} on \mathcal{B}_B , then
 1308 Eq. (28) and Eq. (36) force all activation derivatives to be nearly constant (for $k = 1$) or
 1309 nearly zero (for $k \geq 2$) on the relevant pre-activation range. In this regime, the network is
 1310 forced into an almost linear operating region where

1311 • σ' is approximately constant, so $\widehat{s_1(\mathbf{u})}$ is concentrated near zero frequency and
 1312 $\mathcal{D}\mathcal{L}_{\text{FNO}}(\mathbf{u})$ becomes (block-)diagonal in spectral basis; and
 1313 • $\sigma^{(k)} \approx 0$ for $k \geq 2$, so higher-order terms vanish.

1314 The resulting layer effectively reduces to a spectral multiplier $\mathbf{A}(\xi)$ and cannot represent a
 1315 dense Toeplitz map such as $\mathbf{T}_n[x^2]$ whose off-diagonals are non-trivial and fixed independent
 1316 of inputs.

1317 In other words, for spectrally dense, position-dependent operators, FNO faces a fundamental trade-off:

1318 • it can use strongly input-dependent gates to synthesize spectral off-diagonals, but then those
 1319 off-diagonals are necessarily tied to the input distribution, or
 1320 • it can suppress the input dependence of the gates to emulate a fixed operator, but then the
 1321 spectral kernel collapses towards a diagonal (or nearly diagonal) multiplier.

1322 In neither case does a practical-size FNO realize a fixed, input-independent dense Toeplitz kernel with
 1323 robust out-of-distribution generalization for spectrally dense operators, even though a sufficiently
 1324 large FNO *can* parametrize the off-diagonals of the in-sample mapping on the training distribution.

1325 **What we do (and do not) claim about FNO.** The generalized analysis in this Appendix refines
 1326 the statement of Section 3.1:

1327 • We *do not* claim that FNO cannot generate spectral off-diagonals. Eq. (12) and Eq. (36)
 1328 show that at first and higher order, off-diagonals appear whenever the activation derivatives
 1329 have non-trivial Fourier coefficients.

1330 • We *do* claim that for spectrally dense operators, these off-diagonals are *always tied to*
 1331 *input-dependent gates*. As a result, a large FNO can fit the off-diagonals of the in-sample
 1332 mapping on the training subspace, but it cannot efficiently learn a fixed dense off-diagonals
 1333 with robust out-of-distribution generalization on unseen function spaces which is exactly
 1334 what we observe in our experiments in Section 5.

1335 Stacking more layers introduces more input-dependent gates but does not create an input-independent
 1336 spectral mixing mechanism, so depth does not mitigate this bottleneck.

1337 D.4 WHY KANO DOES NOT SUFFER THE SAME PROBLEM

1338 By contrast, KANO directly learns an input-independent pseudo-differential symbol $\mathbf{p}_\theta(\mathbf{x}, \xi)$ in the
 1339 dual bases via Eq. (15). KANO has a spectral kernel capable of mode mixing as it is governed by
 1340 spatial basis \mathbf{x} (which is convolution in spectral basis) as well as the frequency mode ξ via symbol
 1341 $\mathbf{p}_\theta(\mathbf{x}, \xi)$, not only by the non-linear activation gates tied to the input function \mathbf{u} . Once \mathbf{p}_θ is learned,
 1342 it is shared across all input functions in the ball \mathcal{B}_B , including even the unseen function subspace.
 1343 This dual-domain, symbol-based parameterization allows KANO to learn the fixed off-diagonals
 1344 of spectrally dense, position-dependent operators with robust out-of-distribution generalization, as
 1345 confirmed by our experiments in Section 5.

1350 **E COMPUTATION AND MEMORY COMPLEXITY OF KANO**
1351

1352 In this Appendix we compare the memory and computation complexity of KANO compared to FNO.
1353 We first quantify the per-layer costs of a single FNO layer (4) and a single KANO layer (15), and
1354 then combine them with the parameter complexity results of Section 3.2 and Section 4.2 to argue
1355 that, on the target class of variable-coefficient PDE and position-dependent dynamics, the higher
1356 per-layer cost of KANO is compensated by better model size scaling.

1357 Throughout this appendix we work on the discrete torus \mathbb{T}^d with a uniform spatial grid $\mathcal{Y} =$
1358 $\{\mathbf{y}_1, \dots, \mathbf{y}_m\}$ and a truncated Fourier set $\Xi = \{\boldsymbol{\xi}_1, \dots, \boldsymbol{\xi}_m\}$ as in Section 2.1 and Section 4.1.
1359

1360 **E.1 PER-LAYER FLOPS AND ACTIVATION MEMORY**
1361

1362 We measure complexity in floating-point operations (FLOPs) and activation memory per forward pass
1363 for a single layer of each model. Backward passes in standard automatic differentiation are assumed
1364 to be within a constant factor of the forward cost and do not change the asymptotic conclusions.

1365 **FNO layer.** Consider a single FNO layer (4)

1366
$$\mathcal{L}_{\text{FNO}}(\mathbf{a})(\mathbf{x}) = \sigma \left(\mathcal{F}_m^{-1} \left(\mathbf{R}_{\theta_1}(\boldsymbol{\xi}) \cdot \mathcal{F}_m(\mathbf{a})(\boldsymbol{\xi}) \right) (\mathbf{x}) + \mathbf{W}_{\theta_2} \cdot \mathbf{a}(\mathbf{x}) \right)$$

1367 with C_{in} input channels and C_{out} output channels. On a d -dimensional grid \mathcal{Y} with m^d coordinates, a
1368 single forward application of \mathcal{L}_{FNO} has the following costs:

1369

1370 - *FFT and inverse FFT:* \mathcal{F}_m and \mathcal{F}_m^{-1} are applied channel-wise and cost

1371
$$\text{FFT cost} \sim \mathcal{O}((C_{\text{in}} + C_{\text{out}}) m^d \log m^d).$$

1372

1373 - *Spectral multiplier:* for each retained mode $\boldsymbol{\xi} \in \Xi$, $\mathbf{R}_{\theta}(\boldsymbol{\xi}) \in \mathbb{C}^{C_{\text{out}} \times C_{\text{in}}}$ is a dense matrix;
1374 multiplying by $\hat{\mathbf{u}}(\boldsymbol{\xi}) \in \mathbb{C}^{C_{\text{in}}}$ costs $\mathcal{O}(C_{\text{in}} C_{\text{out}})$ per mode, hence

1375
$$\text{spectral block} \sim \mathcal{O}(m^d C_{\text{in}} C_{\text{out}}).$$

1376

1377 - *Point-wise linear map:* \mathbf{W}_{θ} is applied at each spatial point $\mathbf{y} \in \mathcal{Y}$ as a dense matrix in
1378 channel space, costing

1379
$$\text{spatial linear map} \sim \mathcal{O}(m^d C_{\text{in}} C_{\text{out}}).$$

1380

1381 - *Nonlinearity:* the point-wise nonlinearity σ is $\mathcal{O}(m^d C_{\text{out}})$.

1382 Collecting terms, for fixed channel counts we obtain the per-layer forward cost:

1383
$$\text{FLOPs}(\mathcal{L}_{\text{FNO}}) \sim \mathcal{O}(m^d \log m^d + m^d) \sim \mathcal{O}(m^d \log m^d). \quad (39)$$

1384 The activation memory footprint is dominated by storing \mathbf{u} , $\mathcal{F}_m \mathbf{u}$, the pre-activation $\mathbf{z}(\mathbf{u}, \cdot)$ and the
1385 post-activation:

1386
$$\text{memory}(\mathcal{L}_{\text{FNO}}) = \mathcal{O}(m^d C_{\text{in}} + m^d C_{\text{out}}) \sim \mathcal{O}(m^d). \quad (40)$$

1387 The parameter memory is $\mathcal{O}(m^d C_{\text{in}} C_{\text{out}})$ for the spectral multipliers plus $\mathcal{O}(C_{\text{in}} C_{\text{out}})$ for \mathbf{W}_{θ} .

1388 **KANO layer.** Now consider a KANO layer (15)

1389
$$\mathcal{L}_{\text{KANO}}(\mathbf{a})(\mathbf{x}) = \Phi_{\theta_1} \left(\mathcal{F}_m^{-1} [\mathbf{p}_{\theta_2}(\mathbf{x}, \boldsymbol{\xi}) * \mathcal{F}_m(\mathbf{a})(\boldsymbol{\xi})](\mathbf{x}), \mathbf{a}(\mathbf{x}) \right),$$

1390 where $\mathbf{p}_{\theta_2}(\mathbf{x}, \boldsymbol{\xi})$ is implemented by a width-fixed KAN symbol network and Φ_{θ_1} is another
1391 width-fixed KAN activation network. Using Kohn–Nirenberg quantization (16):

1392
$$\mathcal{F}_m^{-1} [\mathbf{p}(\mathbf{x}, \boldsymbol{\xi}) * \mathcal{F}_m(\mathbf{a})(\boldsymbol{\xi})](\mathbf{x}) := \left(\frac{h}{L} \right)^d \sum_{\boldsymbol{\xi} \in \Xi} \sum_{\mathbf{y} \in \mathcal{Y}} e^{i(\mathbf{x}-\mathbf{y}) \cdot \boldsymbol{\xi}} \mathbf{p}(\mathbf{x}, \boldsymbol{\xi}) \mathbf{a}(\mathbf{y}),$$

1393 on a d -dimensional grid \mathcal{Y} with m^d points and a retained frequency set Ξ with m^d modes, a single
1394 forward application of $\mathcal{L}_{\text{KANO}}$ has the following costs:

1404
 1405 • *KN quantization (double sum)*: for each output location \mathbf{x} we evaluate the double sum over
 1406 $\xi \in \Xi$ and $\mathbf{y} \in \mathcal{Y}$, i.e. m^{2d} terms per \mathbf{x} . Each term involves a dense matrix–vector product
 1407 $\mathbf{p}(\mathbf{x}, \xi) \mathbf{a}(\mathbf{y})$ of cost $\mathcal{O}(C_{\text{in}} C_{\text{out}})$. Summed over all $\mathbf{x} \in \mathcal{Y}$ this yields

$$1408 \quad \text{KN operator} \sim \mathcal{O}(m^{3d} C_{\text{in}} C_{\text{out}}).$$

1409
 1410 • *Symbol network evaluation*: the symbol KAN \mathbf{p}_{θ_2} is evaluated once per pair $(\mathbf{x}, \xi) \in \mathcal{Y} \times \Xi$,
 1411 i.e. at m^{2d} points. Let P_p denote the cost of a single forward evaluation of \mathbf{p}_{θ_2} . Then

$$1412 \quad \text{symbol KAN} \sim \mathcal{O}(m^{2d} P_p).$$

1413
 1414 • *Activation KAN*: the activation network Φ_{θ_1} is applied point-wise at each spatial point
 1415 $\mathbf{x} \in \mathcal{Y}$. Let P_{Φ} denote the cost of one forward pass of Φ_{θ_1} . Then

$$1416 \quad \text{activation KAN} \sim \mathcal{O}(m^d P_{\Phi}).$$

1417
 1418 Collecting terms, for fixed channel counts and fixed KAN architectures we obtain the per-layer
 1419 forward cost:

$$1420 \quad \text{FLOPs}(\mathcal{L}_{\text{KANO}}) \sim \mathcal{O}(m^{3d} C_{\text{in}} C_{\text{out}} + m^{2d} P_p + m P_{\Phi}). \quad (41)$$

1421 The activation memory footprint is dominated by storing the feature maps and, if materialized, the
 1422 symbol grid $\mathbf{p}_{\theta_2}(\mathbf{x}, \xi)$:

$$1423 \quad \text{memory}(\mathcal{L}_{\text{KANO}}) \sim \mathcal{O}(m^d C_{\text{in}} + m^d C_{\text{out}} + m^{2d} C_{\text{in}} C_{\text{out}}) \sim \mathcal{O}(m^{2d} C_{\text{in}} C_{\text{out}}), \quad (42)$$

1424 in addition to the parameter size of the KAN subnetworks.

1425
 1426 **Inference cost of compact KAN subnetworks.** For completeness, we quantify the inference cost
 1427 P_p and P_{Φ} of the KAN subnetworks. Consider a fully connected KAN layer with input width d_{in} ,
 1428 output width d_{out} , and G basis functions per edge. Each edge $(j \rightarrow i)$ carries a learnable univariate
 1429 function $f_{ij} : \mathbb{R} \rightarrow \mathbb{R}$ represented as

$$1431 \quad f_{ij}(x) = \sum_{g=1}^G w_{i,j,g} \phi_g(x),$$

1432 where $\{\phi_g\}_{g=1}^G$ are fixed basis functions (e.g. B–splines or rational functions) and $w_{i,j,g}$ are learned
 1433 coefficients. Evaluating $f_{ij}(x_j)$ for a given scalar input x_j requires computing the active basis
 1434 functions $\phi_g(x_j)$ and a dot product over G elements.

1435 A single KAN layer thus computes, for each output coordinate i ,

$$1439 \quad y_i = \sum_{j=1}^{d_{\text{in}}} f_{ij}(x_j),$$

1440 and the total cost of one forward pass through this layer would be $\mathcal{O}(d_{\text{out}} d_{\text{in}} G)$, up to lower order
 1441 terms from basis evaluation. The parameter count of this layer is of the same order, $\mathcal{O}(d_{\text{out}} d_{\text{in}} G)$.
 1442 Under the assumption of compact KAN subnetworks, widths, depths, and the number of basis
 1443 functions G are all small and independent to the operator resolution m . If we denote by N an upper
 1444 bound on their layer widths and by L_{KAN} their depth, then their total inference costs satisfy

$$1445 \quad P_p, P_{\Phi} = \mathcal{O}(L_{\text{KAN}} N^2 G).$$

1446 In the per-layer KANO complexity (41) the KAN subnetworks contribute only a constant-factor
 1447 overhead that does not grow with the spatial or spectral resolution, hence for large resolution
 1448 $m \gg N, L_{\text{KAN}}$ the dominant term would be the KN quantization term $\mathcal{O}(m^{3d} C_{\text{in}} C_{\text{out}})$.

1449 E.2 COMPUTATION COMPLEXITY AS A FUNCTION OF ACCURACY FOR POSITION-DEPENDENT 1450 DYNAMICS

1451 As apparent in Eq. (39 and Eq. (41), KANO requires orders heavier FLOPs compared to FNO due
 1452 to the expensive double sum nature of its KN quantization, when two models are of similar size.
 1453 For small resolution m , the inference cost P_p and P_{Φ} still set hard lower bound on the computation

1458 complexity as well. However, one of the main arguments we make in this work is that KANO enjoys
 1459 incomparable parameter efficiency compared to FNO on the target class of spectrally dense operators
 1460 such as position-dependent dynamics: parameter complexity of KANO scales polynomially where
 1461 that of FNO can scale super-exponentially on a spectrally dense ground-truth operator. Therefore,
 1462 KANO’s expensive FLOPs requirement can be compensated by parameter efficiency where FNO
 1463 suffers the curse of dimensionality.

1464 For a position-dependent dynamics \mathcal{G} and accuracy $\varepsilon > 0$, we write $|\Theta_{\text{FNO}}(\varepsilon)|$ and $|\Theta_{\text{KANO}}(\varepsilon)|$ for
 1465 the smallest parameter counts required by FNO and KANO respectively to achieve $\|\mathcal{G} - \mathcal{G}_\theta\| \leq \varepsilon$,
 1466 following Theorem 1 and Theorem 2. Also, let $T_{\text{FNO}}(\varepsilon)$ and $T_{\text{KANO}}(\varepsilon)$ denote the cost of one forward
 1467 (or forward+backward) pass of an FNO or KANO architecture chosen to achieve accuracy ε on the
 1468 ground-truth operator. We do not attempt to model the number of optimization steps; instead, we
 1469 focus on how the cost of a *single* training or inference step scales with ε . Assuming a model must
 1470 perform at least a constant number of FLOPs per parameter in each step, lower bound for computation
 1471 complexity can be estimated as:

$$T_{\text{FNO}}(\varepsilon) \gtrsim c_0 |\Theta_{\text{FNO}}(\varepsilon)|, \quad T_{\text{KANO}}(\varepsilon) \gtrsim c'_0 |\Theta_{\text{KANO}}(\varepsilon)| \quad (43)$$

1472 for constants $c_0, c'_0 > 0$. Such lower bound assumption is not universal for any neural network, but in
 1473 FNO architecture all its parameters equally appear in the matrix computation via the dense parameter
 1474 block \mathbf{R} and \mathbf{W} ; especially for implementations that evaluate and update all parameters in every step.
 1475 Hence, computation complexity lower bound assumption by its model size is reasonable for FNO.
 1476

1477 **Computation complexity of FNO by accuracy for position operators from Theorem 1.** For
 1478 the position-multiplier $\mathcal{M}(\mathbf{x}) := x_1^{\alpha_1} \cdots x_d^{\alpha_d}$ of total degree $M := |\alpha| \geq 1$, the restatement of
 1479 Theorem 1 in Appendix C.3 shows that an FNO $\mathcal{G}_\theta^{\text{FNO}}$ that achieves
 1480

$$\|\mathcal{M} - \mathcal{G}_\theta^{\text{FNO}}\|_{H^s} \leq \varepsilon$$

1481 on band-limited inputs can suffer curse of dimensionality:
 1482

$$|\Theta_{\text{FNO}}(\varepsilon)| \sim \exp(c \varepsilon^{-Md/s}), \quad (44)$$

1483 for some constant $c > 0$ depending only on geometric constants of Sobolev space H^s . From Eq. (43),
 1484 the computation complexity of FNO by accuracy on \mathcal{M} can be written as:
 1485

$$T_{\text{FNO}}(\varepsilon) \gtrsim c_0 |\Theta_{\text{FNO}}(\varepsilon)| \sim \mathcal{O}(\exp(\varepsilon^{-Md/s})) \quad (45)$$

1486 **Computation complexity of KANO by accuracy for position operators from Theorem 2.** On
 1487 the other hand, Theorem 2 and Corollary 1 show that as the KANO projection of \mathcal{M} yields symbols
 1488 $\mathcal{P}_{\mathcal{M}}(\mathbf{x}, \xi)$ of sufficient smoothness, KANO can achieve $\|\mathcal{M} - \mathcal{G}_\theta^{\text{KANO}}\| \leq \varepsilon$ with
 1489

$$|\Theta_{\text{KANO}}(\varepsilon)| \sim \mathcal{O}(\varepsilon^{-\beta}), \quad (46)$$

1490 where $\beta = d/(2s_p)$ or $\beta = d/(2r)$ is a geometric exponent determined by the symbol regularity
 1491 (Wang et al., 2024). From Eq. (43), the computation complexity of KANO by accuracy on \mathcal{M}
 1492 can be written as:
 1493

$$T_{\text{KANO}}(\varepsilon) \gtrsim c'_0 |\Theta_{\text{KANO}}(\varepsilon)| \sim \mathcal{O}(\exp(\varepsilon^{-\beta})) \quad (47)$$

1494 Therefore, when the ground-truth operator is spectrally dense yet its KANO projection give smooth
 1495 enough symbol so that both Theorem 1 and Theorem 2 are effective (for instance, position-dependent
 1496 dynamics), lower bound of computation complexity of FNO can prevail over that of KANO because
 1497 of the curse of dimensionality discussed in Section 3.2.

1498 We emphasize that this is an asymptotic statement under the specific operator class where both
 1499 Theorem 1 and Theorem 2 hold. For operators that are nearly spectrally diagonal (e.g. standard FNO
 1500 benchmarks), FNO is theoretically much faster than KANO when the two models are in similar size.
 1501 However, when learning a spectrally dense, variable-coefficient operators to high precision, KANO’s
 1502 parameter efficiency can, in theory, compensate for its expensive double sum KN quantization while
 1503 FNO faces a super-exponential growth in model size hence per-step FLOPs to achieve the same
 1504 level of accuracy.

1512 **F USE OF LARGE LANGUAGE MODEL**

1513
1514 Large Language Model (LLM) is used to polish the writing in this paper, such as detecting grammar
1515 errors and typos. LLM is also used to search for appropriate references for correct citations, and all
1516 the proposed references are fully inspected and verified before citing in the paper.

1517

1518

1519

1520

1521

1522

1523

1524

1525

1526

1527

1528

1529

1530

1531

1532

1533

1534

1535

1536

1537

1538

1539

1540

1541

1542

1543

1544

1545

1546

1547

1548

1549

1550

1551

1552

1553

1554

1555

1556

1557

1558

1559

1560

1561

1562

1563

1564

1565