

MULTISCALE NEURAL OPERATOR: LEARNING FAST AND GRID-INDEPENDENT PDE SOLVERS

Anonymous authors

Paper under double-blind review

ABSTRACT

Numerical simulations in climate, chemistry, or astrophysics are computationally too expensive for uncertainty quantification or parameter-exploration at high-resolution. Reduced-order or surrogate models are multiple orders of magnitude faster, but traditional surrogates are inflexible or inaccurate and pure machine learning (ML)-based surrogates too data-hungry. We propose a hybrid, flexible surrogate model that exploits known physics for simulating large-scale dynamics and limits learning to the hard-to-model term, which is called parametrization or closure and captures the effect of fine- onto large-scale dynamics. Leveraging neural operators, we are the first to learn grid-independent, non-local, and flexible parametrizations. Our *multiscale neural operator* is motivated by a rich literature in multiscale modeling, has quasilinear runtime complexity, is more accurate or flexible than state-of-the-art parametrizations and demonstrated on the chaotic equation multiscale Lorenz96.

1 INTRODUCTION

Climate change increases the likelihood of storms, floods, wildfires, heat waves, biodiversity loss and air pollution (IPCC, 2018). Decision-makers rely on climate models to understand and plan for changes in climate, but current climate models are computationally too expensive: as a result, they are hard to access, cannot predict local changes ($< 10km$), fail to resolve local extremes (e.g., rainfall), and do not reliably quantify uncertainties (Palmer et al., 2019). For example, running a global climate model at $1km$ resolution can take ten days on a $4888\times GPU$ node supercomputer, consuming the same electricity as a coal power plants generates in one hour (Fuhrer et al., 2018). Similarly, in molecular dynamics (Batzner et al., 2022), chemistry (Behler, 2011), biology (Yazdani et al., 2020), energy (Zhang et al., 2019), astrophysics or fluids (Duraismy et al., 2019), scientific progress is hindered by the computational cost of solving partial differential equations (PDEs) at high-resolution (Karniadakis et al., 2021). We are proposing the first PDE surrogate that quickly computes approximate solutions via correcting known large-scale simulations with learned, grid-independent, non-local parametrizations.




Figure 1: **Multiscale neural operator (MNO)**. Explicitly modeling all scales of Earth’s weather is too expensive for traditional and learning-based solvers (Palmer et al., 2019). MNO dramatically reduces the computational cost by modeling the large-scale explicitly and learning the effect of fine- onto large-scale dynamics; such as turbulence slowing down a river stream. We embed a grid-independent neural operator in the large-scale physical simulations as a “parametrization“, conceptually similar to stacking dolls (Snagglebit, 2022).

Surrogate models are fast, reduced-order, and lightweight copies of numerical simulations (Quarteroni & Rozza, 2014) and of significant interest in physics-informed machine learning (Kashinath et al., 2021; Reichstein et al., 2019; Karpatne et al., 2019; Ganguly et al., 2014). Machine learning (ML)-based surrogates have simulated PDEs up to 1 – 3 order of magnitude faster than traditional numerical solvers and are more flexible and accurate than traditional surrogate models (Karniadakis et al., 2021). However, pure ML-based surrogates are too data-hungry (Rasp et al., 2020); so, hybrid ML-physics models are created, for example, via incorporating known symmetries (Bronstein et al., 2021; Batzner et al., 2022) or equations (Willard et al., 2022). Most hybrid models represent the solution at the highest possible resolution which becomes computationally infeasible in multiscale or very high-resolution physics; even at optimal runtime (Pavliotis & Stuart, 2008; Peng et al., 2021).

As depicted in Figs. 1 and 2, we simulate multiscale physics by running easy-to-access large-scale models and focusing learning on the challenging task: *How can we model the influence of fine-onto large-scale dynamics, i.e., what is the subgrid parametrization term?* The lack of accuracy in current subgrid parametrizations, also called closure or residual terms, is one of the major sources of uncertainty in multiscale systems, such as turbulence or climate (Palmer et al., 2019; Gentile et al., 2018). Learning subgrid parametrizations can be combined with incorporating equations as soft (Raissi et al., 2019) or hard (Beucler et al., 2021a) constraints. Various works learn subgrid parametrizations, but are either inaccurate, hard to share or inflexible because they are local (Gentile et al., 2018), grid-dependent (Lapeyre et al., 2019), or domain-specific (Behler J, 2007), respectively as detailed in Section 2. We are the first to formulate the parametrization problem as learning neural operators (Anandkumar et al., 2020) to represent non-local, flexible, and grid-independent parametrizations.

We propose, *multiscale neural operator* (MNO), a novel learning-based PDE surrogate for multiscale physics with the key contributions:

- A learning-based multiscale PDE surrogate that has quasilinear runtime complexity, leverages known large-scale physics, is grid-independent, flexible, and does not require autodifferentiable solvers.
- The first surrogate to approximate grid-independent, non-local parametrizations via neural operators.
- Demonstration of the surrogate on the chaotic, coupled, multiscale PDE: multiscale Lorenz96.

2 RELATED WORKS

We embed our work in the broader field of physics-informed machine learning and surrogate modeling. We propose the first surrogate that corrects a coarse-grained simulation via learned, grid-independent, non-local parameterizations.

Direct numerical simulation. Despite significant progress in simulating physics numerically it remains prohibitively expensive to repeatedly solve high-dimensional partial differential equations (PDEs) (Karniadakis et al., 2021). For example, finite difference, element, volume, and (pseudo-) spectral methods have to be re-run for every choice of initial or boundary condition, grid, or parameters (Farlow, 1993; Boyd, 2013). The issue arises if the chosen method does not have optimal runtime, i.e., does not scale linearly with the number of grid points, which renders it infeasibly expensive for calculating ensembles (Boyd, 2013). Select methods have optimal or close-to-optimal runtime, e.g., quasi-linear $O(N \log N)$, and outperform machine learning-based methods in runtime and accuracy, but their implementation often requires significant problem-specific adaptations; for example multigrid (Briggs et al., 2000) or spectral methods (Boyd, 2013). We acknowledge the existence of impressive research directions towards optimal and flexible non-ML solvers, such as the spectral solver called Dedalus (Burns et al., 2020), but advocate to simultaneously explore easy-to-adapt ML methods to create fast, accurate, and flexible surrogate models.

Surrogate modeling. Surrogate models are approximations, lightweight copies, or reduced-order models of PDE solutions, often fit to data, and used for parameter exploration or uncertainty quantification (Smith, 2013; Quarteroni & Rozza, 2014). Surrogate models via SVD/POD (Chatterjee, 2000), Eigendecompositions/KLE (Fukunaga & Koontz, 1970), Koopman

operators/DMD (Williams et al., 2015), take simplifying assumptions to the dynamics, e.g., linearizing the equations, which can break down in high-dimensional or nonlinear regimes (Quarteroni & Rozza, 2014). Instead, our work leverages the expressiveness of neural operators as universal approximations (Chen & Chen, 1995) to learn fast high-dimensional surrogates that are accurate in nonlinear regimes (Lütjens et al., 2021; Yuval et al., 2021; Costa Nogueira et al., 2020; Nogueira Jr. et al., 2021). **Pure ML-based** surrogate models have shown impressive success in approximating dynamical systems from ground-truth simulation data – for example with neural ODEs (Rackauckas et al., 2020; Chen et al., 2018; Hasani et al., 2021), GNNs (Brandstetter et al., 2022; Cachay et al., 2021a), CNNs (Stachenfeld et al., 2022), neural operators (Li et al., 2021a; Anandkumar et al., 2020; Pathak et al., 2022; Lu et al., 2021; Jiang et al., 2021), RNNs (Kani & Elsheikh, 2017; Rasp et al., 2020), GPs (Chakraborty et al., 2021), reservoir computing (Pathak et al., 2018; Nogueira Jr. et al., 2021), or transformers (Chattopadhyay et al., 2020a) – but, without incorporating physical knowledge become data-hungry and poor at generalization (Karniadakis et al., 2021; Beucler et al., 2021b).

Physics-informed machine learning. Two main approaches of incorporating physical knowledge into ML systems is via known symmetries (Bronstein et al., 2021) or equations (Karniadakis et al., 2021). Our approach leverages known equations for computing a coarse-grid prior, which is complementary to using known equations as soft (Raissi et al., 2019; Lee & Carlberg, 2020; Zeng et al., 2021; Wu et al., 2020; Zhang et al., 2018; Yazdani et al., 2020) or hard constraints (Greydanus et al., 2019; Lutter et al., 2019; Beucler et al., 2021a; Donti et al., 2021; Beucler et al., 2019; Jin et al., 2020) as these methods can still be used to constrain the learned parametrization. In terms of symmetry, our approach exploits translational equivariance via Fourier transformations (Li et al., 2021a), but can be extended to other frameworks that exploit in- or equivariance of PDEs (Olver, 1986) to rotational (Fuchs et al., 2020; Thomas et al., 2018), Galilean (Wu et al., 2018; Prakash et al., 2021), scale (Beucler et al., 2021b), translational (Subel et al., 2021), reflectional (Cohen & Welling, 2017) or permutational (Zhou et al., 2020) transformations.

The field of physics-informed machine learning is very broad, as reviewed most recently in (Willard et al., 2022) and (Karniadakis et al., 2021; Carleo et al., 2019; Karpatne et al., 2017). We focus on the task of learning fast and accurate surrogate models of fine-scale models when a fast and approximate coarse-grained simulation is available. This task differs from other interesting research areas in equation discovery or symbolic regression (Brunton et al., 2016; Long et al., 2018b; 2019; Liu et al., 2021; Qian et al., 2022), downscaling or superresolution (Xie et al., 2018; Bode et al., 2021; Kurinchi-Vendhan et al., 2021; Stengel et al., 2020; Vandal et al., 2017; Groenke et al., 2020), design space exploration or data synthesis (Chen & Ahmed; Chan & Elsheikh, 2019), controls (Bieker et al., 2020) or interpretability (Toms et al., 2020; McGraw & Barnes, 2018). Our work is complementary to data assimilation or parameter calibration (Jia et al., 2019; 2021; Karpatne et al., 2017; Zhang et al., 2019; Bonavita & Laloyaux, 2020) which fit to observational data instead of models and differs from inverse modeling and parameter estimation (Parish & Duraisamy, 2016; Hamilton et al., 2017; Yin et al., 2021; Long et al., 2018a) which usually fit parametrizations that are independent of the previous state.

Correcting coarse-grid simulations via parametrizations. Problems with large domains are often solved via multiscale methods (Pavliotis & Stuart, 2008). Multiscale methods simulate the dynamics on a coarse-grid and capture the effects of small-scale dynamics that occur within a grid cell via additive terms, called subgrid parametrizations, closures, or residuals (Pavliotis & Stuart, 2008; McGuffie & Henderson-Sellers, 2005). Existing subgrid parametrizations for many equations are still inaccurate (Webb et al., 2015) and ML outperformed them by learning parametrizations directly from high-resolution simulations; for example in turbulence (Duraisamy et al., 2019), climate (Gentine et al., 2018), chemistry (Hansen et al., 2013), biology (Peng et al., 2021), materials (Liu et al., 2022), or hydrology (Bennett & Nijssen, 2020). The majority of ML-based parametrizations, however, is local (Gentine et al., 2018; O’Gorman & Dwyer, 2018; Brenowitz & Bretherton, 2018; Brenowitz et al., 2020; Bretherton et al., 2022; Yuval et al., 2021; Cachay et al., 2021b; Bennett & Nijssen, 2020; Hansen et al., 2013; Liu et al., 2022; Prakash et al., 2021; Ling et al., 2016; Parish & Duraisamy, 2016; Wu et al., 2018; Rasp, 2020), i.e., the in- and output are variables of single grid points, which assumes perfect scale separation, for example, in isotropic homogeneous turbulent flows (P., 2006). However, local parametrizations are inaccurate; for example in the case of anisotropic nonhomogeneous dynamics (P., 2006; Wang et al., 2022), for correcting global error of coarse spectral discretizations (Boyd, 2013), or in large-scale climate models (Dueben & Bauer,




Figure 2: Left: **Model Architecture.** A physics-based model, \mathcal{N} , can quickly propagate the state, \bar{u}_t , at a large-scale, but will accumulate the error, $h = \mathcal{N}(u) - \mathcal{N}\bar{u}$. A neural operator, \mathcal{K}_θ , wraps the computational and implementation complexities of unmodeled fine-scale dynamics into a non-local and grid-independent term, \hat{h} , that iteratively corrects the large-scale model. Right: **Multiscale Lorenz96.** We demonstrate multiscale neural operator (MNO) on the multiscale Lorenz96 equation, a model for chaotic atmospheric dynamics. Image: (Rasp, 2020)

2018; Pathak et al., 2018). More recent works propose non-local parametrizations, but their formulations either rely on a fixed-resolution grid (Wang et al., 2022; Blakseth et al., 2022; Lapeyre et al., 2019; Chattopadhyay et al., 2020b), an autodifferentiable solver (Um et al., 2020; Sirignano et al., 2020; Frezat et al., 2022), or are formulated for a specific domain (Behler J, 2007). A single work proposes non-local and grid-independent parametrizations (Pathak et al., 2020), but requires the explicit representation of a high-resolution state which is computationally infeasible for large domains, such as in climate modeling. We are the first to propose grid-independent and non-local parametrizations via neural operators to create fast and accurate surrogate models of fine-scale simulations.

Neural operators for grid-independent, non-local parametrizations. Most current learning-based non-local parametrizations rely on FCNNs, CNNs (Lapeyre et al., 2019), or RNNs (Chattopadhyay et al., 2020b), which are mappings between finite-dimensional spaces and thus grid-dependent. In comparison, neural operators learn mappings in between infinite-dimensional function spaces (Kovachki et al., 2021) such as the Laplacian, Hessian, gradient, or Jacobian. Typically, neural operators lift the input into a grid-independent state such as Fourier (Li et al., 2021a), Eigen- (Bhattacharya et al., 2020), graph kernel (Li et al., 2020; Anandkumar et al., 2020) or other latent (Lu et al., 2021) modes and learn weights in the lifted domain. We are the first to formulate neural operators for learning parametrizations.

3 APPROACH

We propose *multiscale neural operator* (MNO): a surrogate model with quasilinear runtime complexity that exploits known coarse-grained simulations and learns a grid-independent, non-local parametrization. As detailed in the following MNO propagates the dynamics according to:

$$\underbrace{\frac{\partial \bar{u}}{\partial t}}_{\text{Corrected Large-scale Dyn.}} = \underbrace{\mathcal{N}(\bar{u})}_{\text{Large-scale Dyn.}} + \underbrace{\mathcal{K}_\theta(\bar{u})}_{\text{Parametrization}} \quad (1)$$

3.1 MULTISCALE NEURAL OPERATOR

Partial differential equations. We focus on partial differential equations (PDEs) that can be written as initial value problem (IVP) via the method of lines (William, 1991). The PDEs in focus have one temporal dimension, $t \in [0, T] =: D_t$, and (multiple) spatial dimensions, $x = [x_1, \dots, x_d]^T \in D_x$, and can be written in the iterative, explicit, symbolic form (Farlow, 1993):

$$\begin{aligned} \frac{\partial u}{\partial t} - \mathcal{N}(u) &= 0 \text{ with } t, x \in [0, T] \times D_x \\ u(0, x) &= u^0(x), \mathcal{B}[u](t, x) = 0 \text{ with } x \in D_x, \\ &\quad (t, x) \in [0, T] \times \partial D_x \end{aligned} \quad (2)$$

In our case, the (non-)linear operator, \mathcal{N} , encodes the **known** physical equations; for example a combination of Laplacian, integral, differential, etc. operators. Further, $u : D_t \times D_x \rightarrow D_u$ is the

solution to the initial values, $u^0 : D_x \rightarrow D_u$, and Dirichlet, $\mathcal{B}_D[u] = u - b_D$, or Neumann boundary conditions, $\mathcal{B}_N[u] = n^T \partial_x u - b_N$, with outward facing normal on the boundary, $n \perp \partial B$.

Scale separation. We transfer a concept from the rich and mathematical literature in multiscale modeling (Pavliotis & Stuart, 2008) to consider a filter kernel operator, \mathcal{G}^* , that creates the large-scale solution, $\bar{u}(x) = u(x) + u'(x)$, where u' are the small-scale deviations and $\bar{\cdot}$ denotes the filtered variable, $\bar{\phi}(x) = \mathcal{G}^* \phi = \int_{D_x} G(x, x') \phi(x') dx'$. Assuming the kernel, G , 1) preserves constant fields, $\bar{a} = a$, 2) commutes with differentiation, $[\mathcal{G}^*, \frac{\partial}{\partial s}]$ with $s=x, t$, and 3) is linear, $\overline{\phi + \psi} = \bar{\phi} + \bar{\psi}$ (P., 2006), we can rewrite (2) to:

$$\begin{aligned} \mathcal{G}^* \frac{\partial u}{\partial t} &= \frac{\partial \bar{u}}{\partial t} = \mathcal{G}^* \mathcal{N}(u) \\ &= \mathcal{N}(\bar{u}) + [\mathcal{G}^*, \mathcal{N}](u) \end{aligned} \quad (3)$$

where $[\mathcal{G}^*, \mathcal{N}](u) = \mathcal{G}^* \mathcal{N}(u) - \mathcal{N}(\mathcal{G}^* u)$ is the filter subgrid parametrization, closure term, or commutation error, i.e., the error introduced through propagating the coarse-grained solution.

Approximations of the subgrid parametrization as an operator that acts on \bar{u} require significant domain expertise and are derived on a problem-specific basis. In the case of isotropic homogeneous turbulence, for example, the subgrid parametrization can be approximated as the spatial derivative of the subgrid stress tensor, $[\mathcal{G}^*, \mathcal{N}](\bar{u})_{\text{turbulence}} \approx \frac{\partial \tau_{ij}}{\partial x_j} = \frac{\partial u'_i u'_j}{\partial x_j}$ (P., 2006). Many works approximate the subgrid stress tensor with physics-informed ML (Prakash et al., 2021; Ling et al., 2016; Parish & Duraisamy, 2016; Wu et al., 2018), but are domain-specific, local, or require a differentiable solver or fixed-grid. We propose a general purpose method to approximating the subgrid parametrization, independent of the grid, domain, isotropy, and underlying solver.

Multiscale neural operator. We aim to approximate the parametrization / filter commutation error, $[\mathcal{G}^*, \mathcal{N}] \approx: h$, via learning a neural operator on high-resolution training data. Let \mathcal{K}_θ be a neural operator with the mapping:

$$[\mathcal{G}^*, \mathcal{N}] \approx \mathcal{K}_\theta : \bar{U}(D_x; \mathbb{R}^{d_u}) \rightarrow H(D_x; \mathbb{R}^{d_u}) \quad (4)$$

where θ are the learned parameters and \bar{U}, H are separable Banach spaces of all continuous functions taking values, \mathbb{R}^{d_u} , defined on the bounded, open set, $D_x \subset \mathbb{R}^{d_x}$, with norm $\|f\|_{\bar{U}} = \|f\|_H = \max_{x \in D_x} |f(x)|$. We embed the neural operator as an autoregressive model with fixed time-discretization, Δt , such that the final *multiscale neural operator* (MNO) model is:

$$\bar{u}(t + \Delta t) = f(t, \bar{u}, \frac{\partial \bar{u}}{\partial x}, \frac{\partial^2 \bar{u}}{\partial x^2}, \dots) + \mathcal{K}_\theta(\bar{u}) \quad (5)$$

where $f(t, \bar{u}, \frac{\partial \bar{u}}{\partial x}, \frac{\partial^2 \bar{u}}{\partial x^2}) = \int_t^{t+\Delta t} \mathcal{N}(\bar{u}) d\tau$ is the known large-scale tendency, i.e. one-step solution. MNO is fit using MSE with the loss function:

$$L = \mathbb{E}_t \mathbb{E}_{\bar{u}|u(t) \sim p(t)} (\mathcal{L}(\mathcal{K}_\theta(\bar{u}(t)), [\mathcal{G}^*, \mathcal{N}](u(t)))) \quad (6)$$

where the ground-truth data, $u(t) \sim p(t)$, is generated by integrating a high-resolution simulation with varying parameters, initial or boundary conditions and uniformly sampling time snippets according to the distribution $p(t)$. Similar to problems in superresolution, there exist multiple realizations of the learned commutation error, $[\mathcal{G}^*, \mathcal{N}](\bar{u})$, for a given ground-truth, $[\mathcal{G}^*, \mathcal{N}](u)$; using MSE will learn a smooth average and future work will explore adversarial losses (Goodfellow et al., 2014) or an intersection between neural operators and normalizing flows (Rezende & Mohamed, 2015) or diffusion-based models (Sohl-Dickstein et al., 2015) to account for the stochasticity (Wilks, 2005). During training, the model input is generated via $\bar{u}(t) = \mathcal{G}^*(u(t))$ and the target via

$$h_{\text{target}} = \overline{\mathcal{N}(u)} - \mathcal{N}(\bar{u}). \quad (7)$$

During inference MNO is initialized with a large-scale state and integrates the dynamics in time via coupling the neural operator and a large-scale simulation.

Our approach does not need access to the high-resolution simulator or equations; it only requires a precomputed high-resolution dataset, which are increasingly available (Hersbach et al., 2020; Burns et al., 2022), and allows the user to incorporate existing easy-to-access solvers of large-scale equations. There is no requirement for the large-scale solver to be autodifferentiable which significantly simplifies the implementation for large-scale models, such as in climate. If desired, our loss function can easily be augmented with a physics-informed loss (Raissi et al., 2019) on the large-scale dynamics or parametrization term.

Choice of neural operator. Our formulation is general enough to allow the use of many operators, such as Fourier (Li et al., 2021a), PCA-based (Bhattacharya et al., 2020), low-rank (Khoo & Ying, 2019), Graph (Li et al., 2020) operators, or DeepOnet (Wang et al., 2021; Lu et al., 2021). Because DeepONet (Lu et al., 2021) focuses on interpolation and assumes fixed-grid sensor data, we decided to modify Fourier Neural Operator (FNO) (Li et al., 2021a) for our purpose. FNO is a universal approximator of nonlinear operators (Kovachki et al., 2021; Chen & Chen, 1995), grid-independent and can be formulated as autoregressive model (Li et al., 2021a). As there exists significant knowledge on symmetries and conservation properties of the commutation error (P., 2006), MNO’s explicit formulation increases interpretability and ease of incorporating symmetries and constraints. With FNO, we exploit approximate translational symmetries in the data and leave novel opportunities for neural operators that exploit the full range of known equi- and invariances of the subgrid parametrization term, such as Galilean invariance (Prakash et al., 2021), for future work.

3.2 ILLUSTRATION OF MNO VIA MULTISCALE LORENZ96

We illustrate the idea of MNO on a canonical model of atmospheric dynamics, the multiscale Lorenz96 equation Lorenz (2006); Thornes et al. (2017). This PDE is multiscale, chaotic, time-continuous, space-discretized, 2D (space+time), nonlinear, displayed in Fig. 2-right and detailed in Appendix A.3. Most importantly, the large- and small-scale solutions, $X_k \in \mathbb{R}, Y_{j,k} \in \mathbb{R} \forall j \in \{0, \dots, J\}, k \in \{0, \dots, K\}$, demonstrate the *curse of dimensionality*: the number of the small-scale states grows exponentially with scale and explicit modeling becomes computationally expensive, for example, quadratic for two-scales: $O(N^2) = O(JK)$. The PDE writes:

$$\begin{aligned} \frac{\partial X_k}{\partial t} &= \underbrace{X_{k-1}(X_{k+1} - X_{k-2}) - X_k + F}_{\text{Large-scale Dyn.: } \frac{\partial \bar{X}_k}{\partial t}} - \underbrace{\frac{h_s c}{b} \sum_{j=0}^{J-1} Y_{j,k}(X_k)}_{\text{Parametrization: } h}, \\ \frac{\partial Y_{j,k}}{\partial t} &= -cbY_{j+1,k}(Y_{j+2,k} - Y_{j-1,k}) - cY_{j,k} + \frac{h_s c}{b} X_k. \end{aligned} \quad (8)$$

where F is the forcing, h_s the coupling strength, b the relative magnitude of scales, and c the evolution speed. With the multiscale framework from Section 3.1, we define:

$$\begin{aligned} u(x) &= [X_0, Y_{0,0}, Y_{1,0}, \dots, Y_{J,0}, X_1, Y_{0,1}, \dots \\ &\quad , X_K, \dots, Y_{J,K}]_x \forall x \in D_x = \{0, \dots, K(J+1)\} \\ \mathcal{N}(u)(x) &= \begin{cases} \frac{\partial X_k}{\partial t} & \text{if } x = k(J+1) \forall k \in \{0, \dots, K\} \\ \frac{\partial Y_{j,k}}{\partial t} & \text{otherwise,} \end{cases} \\ G(x, x') &= \begin{cases} 1 & \text{if } x' = k(J+1) \forall k \in \{0, \dots, K\} \\ 0 & \text{otherwise,} \end{cases} \end{aligned}$$

with the solution, u , operator, \mathcal{N} , and kernel, G .

MNO learns the parametrization term via a neural operator, $\mathcal{K}_\theta = \hat{h} \approx h$, and then models:

$$\frac{\partial \hat{X}_k}{\partial t} = \frac{\partial \bar{X}_k}{\partial t} + \mathcal{K}_\theta(\hat{X}_{0:K})(k) \quad (9)$$

where the known large-scale dynamics are approximated with $\frac{\partial \bar{X}_k}{\partial t} \approx \frac{\partial \bar{X}_k}{\partial t}$ and ground-truth parametrization is $h(x) = \{-\frac{h_s c}{b} \sum_{j=0}^{J-1} Y_{j,k}(X_k) \text{ if } x = k(J+1) \forall k \in \{0, \dots, K\} \text{ and } 0 \text{ otherwise}\}$. See Appendix A.4 for all terms.




Figure 3: **MNO is faster than direct numerical simulation.** Our proposed multiscale neural operator (orange) can propagate multiscale PDE dynamics in quasilinear complexity, $O(N \log N)$. For a grid with $K = 2^{15}$, MNO is ~ 1000 -times faster than direct numerical simulation (black) which has quadratic complexity, $O(N^2)$

The parametrization, \mathcal{K}_θ , accepts inputs that are sampled anywhere inside the spatial domain, which differs from previous local (Rasp, 2020) or grid-dependent (Chattopadhyay et al., 2020b) Lorenz96 parametrizations.

We create the ground-truth data via randomly sampled initial conditions, periodic boundary conditions, and integrating the coupled equation with a 4th-order Runge-Kutta solver. After a Lyapunov timescale the state is independent of initial conditions and we extract 4K snippets with $T/\Delta t = 400$ steps length, corresponding to 10 Earth days, for 1-step training. During test the model is run autoregressively on 1K samples from a different initial condition, as detailed in Appendix A.3.

4 RESULTS

Our results demonstrate that multiscale neural operator (MNO) is faster than direct numerical simulation, generates stable solutions, and is more accurate than current parametrizations. We now proceed to discussing each of these in more detail.


4.1 RUNTIME COMPLEXITY: MNO IS FASTER THAN TRADITIONAL PDE SOLVERS

MNO (orange in Fig. 3) has quasilinear, $O(N \log N)$, runtime complexity in the number of large-scale grid points, $N=K$, in the multiscale Lorenz96 equation. The runtime is dominated by a lifting operation, here a fast Fourier transform (FFT), which is necessary to learn spatial correlations in a grid-independent space. In comparison, the direct numerical simulation (black) has quadratic runtime complexity, $O(N^2)$, because of the explicit representation of $N^2=JK$ small-scale states. Both models are linear in time, $O(T)$. Local parametrizations can achieve optimal runtime, $O(N)$, but it is an open question if there exists a decomposition that replaces FFT to yield an optimal, non-local, grid-independent model.

We ran MNO up to a resolution of $K = 2^{24}$, which would equal $75cm/px$ in a global 1D (space) climate model and only took $\approx 2s$ on a single CPU. MNO is three orders of magnitude (1000-times) faster than DNS, at a resolution of $K = 2^{15}$ or $200m/px$. For 2D or 3D simulations the gains of using MNO vs. DNS are even higher with $O(N^2 \log N)$ vs. $O(N^4)$ and $O(N^3 \log N)$ vs. $O(N^6)$, respectively (Khairoutdinov et al., 2005).

The runtimes have been calculated by choosing the best of 1-100k runs depending on grid size on a single-threaded Intel Xeon Gold 6248 CPU@2.50GHz with 164Gb RAM. We time a one step update which, for DNS, is the calculation of (8) and for MNO the calculation of (9), i.e., the sum of a large-scale step and a pass through the neural operator.

In Fig. 3, the runtime of MNO and DNS plateaus at low-resolution ($K < 2^9$), because runtime measurement is dominated by grid-independent operations. DNS plateaus at a lower runtime, be-



Method	RMSE
Climatology	6.902
Traditional parametrizations	2.326
ML-based parametrization (Rasp et al., 2018)	2.053
MNO (ours)	0.5067

Figure 4: **Left: MNO is more accurate than traditional parametrizations.** A sample plot shows, that our proposed multiscale neural operator (yellow/orange-dotted) can accurately forecast the large-scale physics (black-solid), $X_{k=0}(t)$. In comparison, ML-based (blue-dotted) and traditional (red-dotted) parametrizations quickly start to diverge. Note that the system is chaotic and small deviations are rapidly amplified; even inserting the exact parametrizations in float32 instead of float64 quickly diverges. **Right: Accuracy.** MNO is more accurate than traditional parametrizations as measured by the root mean-square error (RMSE).

cause MNO contains several fixed-cost matrix transformations. The runtime of DNS has a slight discontinuity at $K \approx 2^9$ due to extending from cache to RAM memory. We focus on a runtime comparison, but MNO also has significant savings in memory: representing the state at $K = 2^{17}$ in double precision occupies 64GB RAM for DNS and 0.5MB for MNO.

4.2 MNO IS MORE ACCURATE THAN TRADITIONAL PARAMETRIZATIONS

Figure 4-left shows a forecasted trajectory of a sample at the left boundary, $k = 0$, where MNO (orange-dotted) accurately forecasts the large-scale dynamics, $X_0(t)$, (black-solid) while current ML-based (blue-dotted) (Gentine et al., 2018) and traditional parametrizations (red-dotted) quickly diverge. The quantitative comparison of RMSE and a mean/std plot Fig. 4 over $1K$ samples and 200steps or 10days ($\Delta t = 0.005 = 36\text{min}$) confirms that MNO is the most accurate in comparison to ML-based parametrizations, traditional parametrizations, and a mean forecast (climatology). Note, the difficulty of the task: when forecasting *chaotic* dynamics even numerical errors rapidly amplify (P., 2006).

ML-based parametrizations is a state-of-the-art (SoA) model in learning parametrizations and trains a ResNet to forecast a local, grid-independent parametrization, $h_k = \text{NN}(X_k)$, similar to (Gentine et al., 2018). The **traditional parametrizations** (trad. param.) are often used in practice and use linear regression to learn a local, grid-independent parametrization (McGuffie & Henderson-Sellers, 2005). It was suggested that multiscale Lorenz96 is too easy as a test-case for comparing offline models because traditional parametrizations already perform well (Rasp, 2019), but the significant difference between MNO and Trad. Params. shows that online evaluation is still interesting. The **climatology** forecasts the mean of the training dataset, $X_k(t) = 1/T \sum_{t=0}^T 1/N \sum_{i=0}^N X_{k,i}(t)$. The full list of hyperparameters and model parameters can be found in Appendix A.5.2. For fairness, we only compare against grid-independent methods that do not require an autodifferentiable solver. Models with soft or hard constraints, e.g., PINNs (Raissi et al., 2019) or DC3 (Donti et al., 2021), are complementary to MNO. Further, note that our implementation of MNO uses an a priori loss function and could likely be improved by implementing an a posteriori loss functions, i.e., a loss functions that propagates the loss over multiple time steps similar to (Frezat et al., 2022) which requires an autodifferentiable solver or (Brandstetter et al., 2022) which does not.

4.3 MNO IS STABLE

Figure 5 shows that predicting large-scale dynamics with MNO is stable. We first plot a randomly selected sample of the first large-scale state, $X_{k=0}(t)$ (left-black), to illustrate that the prediction is bounded. The MNO prediction (left-yellow) follows the ground-truth up to an approximate horizon of, $t = 1.8$ or 9 days, then diverges from the ground-truth solution, but stays within the bounds of the ground-truth prediction and does not diverge to infinity. The RMSE over time in Figure 5 shows that MNO (yellow) is approximately more accurate than current ML-based (blue) and traditional (red) parametrizations for $\approx 100\%$ -longer time, measuring the time to intersect with climatology. Despite the difficulty in predicting chaotic dynamics, the RMSE of MNO reaches a plateau, which is slightly above the optimal plateau given by the climatology (black).

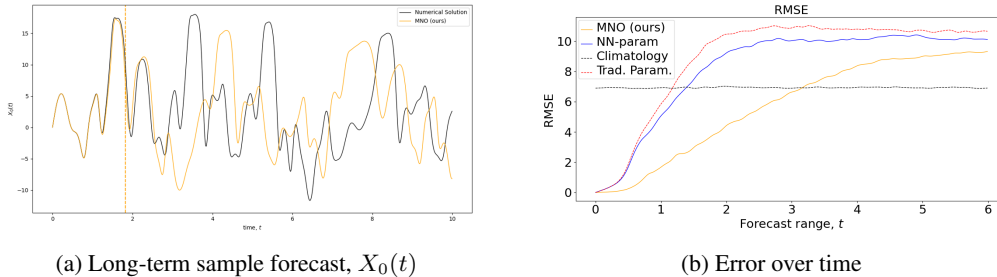


Figure 5: **MNO is stable.** MNO can propagate a sample state, $X_{k=0}(t)$, over a long time horizon without diverging to infinity (left). The right plot shows that the RMSE of MNO plateaus for long-term forecasts, further confirming stability. Further, MNO (yellow) maintains accuracy longer than ML-based parametrizations (blue) and a climatology (black).

The RMSE over time is calculated as:

$$\text{RMSE}(t) = \frac{1}{K} \sum_{k=0}^K \sqrt{\left(\frac{1}{N} \sum_{i=0}^N (\hat{X}_{k,i}(t) - X_{k,i}(t))^2 \right)}. \quad (10)$$

5 LIMITATIONS AND FUTURE WORK

We demonstrated the accuracy, speed, and stability of MNO on the chaotic multiscale Lorenz96 equation. Future work, can extend MNO towards higher-dimensional or time-irregular systems and further integrate symmetries or constraints:

The results show promise to extend MNO to higher-dimensional, chaotic, multiscale, multiphysics problems. We demonstrate the first steps towards quasi-geostrophic turbulence and Rayleigh-Benard convection in Appendix A.1 and aim towards integration in large-scale global atmospheric models, e.g., as approximation of cloud processes (Wang et al., 2022; Palmer et al., 2019). Reducing the cost of climate models could dramatically improve uncertainties (Lütjens et al., 2021) or decision-exploration (Rooney-Varga et al., 2020).

MNO is grid-independent in space but not in time which could be alleviated via integrations with Neural ODEs (Chen et al., 2018). MNO is a myopic model which might suffice for chaotic dynamics (Li et al., 2021b), but could be combined with LSTMs (Mohan et al., 2019) or reservoir computing (Pathak et al., 2018) to contain a memory. Further, we leveraged global Fourier decompositions to exploit grid-independent periodic spatial correlations, but future work could also capture local discontinuities, e.g., along coastlines (Jiang et al., 2021) with multiwavelets (Gupta et al., 2021), or incorporate non-periodic boundaries via Chebyshev polynomials.

Lastly, MNO can be combined with Geometric deep learning, PINNs, or hard constraint models. This avenue of research is particularly exciting with MNO as there exist many known symmetries for various parametrization terms (Prakash et al., 2021).

6 CONCLUSION

We proposed a hybrid physics-ML surrogate of multiscale PDEs that is quasilinear, accurate, and stable. The surrogate limits learning to the influence of fine- onto large-scale dynamics and is the first to use neural operators for a grid-independent, non-local corrective term of large-scale simulations. We demonstrated that multiscale neural operator (MNO) is faster than direct numerical simulation ($O(N \log N)$ vs. $O(N^2)$), and more accurate ($\approx 100\%$ longer prediction horizon) than state-of-the-art parametrizations on the chaotic, multiscale equations multiscale Lorenz96. With the dramatic reduction in runtime MNO could enable rapid parameter exploration and robust uncertainty quantification in complex climate models.

7 ETHICS STATEMENT

Climate change is a defining challenge of our time and environmental disasters will become more frequent: from storms, floods, wildfires and heat waves to biodiversity loss and air pollution (IPCC, 2018). The impacts of these environmental disasters will likely be unjustly distributed: island states, minority populations, and the Global South are already facing the most severe consequences of climate change, while the Global North is responsible for the most emissions since the industrial revolution (Althor et al., 2016). Decision-makers require more accurate, accessible, and local tools to understand and limit the economic and human impact of a changing climate (Palmer et al., 2019). We propose multiscale neural operator (MNO) to improve the parametrizations in climate models, thus leading to more accurate predictions. Related techniques to MNO, specifically neural operator-based surrogate models, could help reduce computational complexity of large-scale weather and climate models. The reduced computational complexity would make them more accessible to low-resource countries or allow for higher resolution predictions. Unfortunately, discoveries for faster differential equations solvers can and likely will be leveraged in ethically questionable fields, such as missile development or oil discovery. We acknowledge the possible negative impacts and hope that our targeted discussion and application to equations from climate modeling can steer the our work towards a positive impact.

8 REPRODUCIBILITY STATEMENT

The code will be submitted in a comment directed to the reviewers and area chairs with a link to an anonymous repository once the discussion forum is opened. The code will be made publicly available upon acceptance along with an open-source license, data, and instructions to reproduce the main results. A list of hyperparameters per model is detailed in Appendix A.6.2, data splits are explained in results; all simulation details in Appendix A.4; and background of the neural operator model in Appendix A.3. Conducting the study from ideation to publication used a total of approximately 10K CPU hours on an internal cluster.

REFERENCES

- Glenn Althor, James E. M. Watson, and Richard A. Fuller. Global mismatch between greenhouse gas emissions and the burden of climate change. *Scientific Reports*, 6, 2016.
- Anima Anandkumar, Kamyar Azizzadenesheli, Kaushik Bhattacharya, Nikola Kovachki, Zongyi Li, Burigede Liu, and Andrew Stuart. Neural operator: Graph kernel network for partial differential equations. In *ICLR 2020 Workshop on Integration of Deep Neural Models and Differential Equations*, 2020.
- Simon Batzner, Albert Musaelian, Lixin Sun, Mario Geiger, Jonathan P. Mailoa, Mordechai Kornbluth, Nicola Molinari, Tess E. Smidt, and Boris Kozinsky. E(3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nature Communications*, 13, 2022.
- Jörg Behler. Neural network potential-energy surfaces in chemistry: a tool for large-scale simulations. *Phys. Chem. Chem. Phys.*, 13, 2011.
- Parrinello M. Behler J. Generalized neural-network representation of high-dimensional potential-energy surfaces. *Phys Rev Lett.*, 14, 2007.
- Andrew Bennett and Bart Nijssen. Deep learned process parameterizations provide better representations of turbulent heat fluxes in hydrologic models. *Earth and Space Science Open Archive*, pp. 20, 2020.
- Tom Beucler, Stephan Rasp, Michael Pritchard, and Pierre Gentine. Achieving Conservation of Energy in Neural Network Emulators for Climate Modeling. jun 2019.
- Tom Beucler, Michael Pritchard, Stephan Rasp, Jordan Ott, Pierre Baldi, and Pierre Gentine. Enforcing analytic constraints in neural networks emulating physical systems. *Phys. Rev. Lett.*, 126: 098302, Mar 2021a.

- Tom Beucler, Michael S. Pritchard, Janni Yuval, Ankitesh Gupta, Liran Peng, Stephan Rasp, Fiaz Ahmed, Paul A. O’Gorman, J. David Neelin, Nicholas J. Lutsko, and Pierre Gentine. Climate-invariant machine learning. *CoRR*, 2021b.
- Kaushik Bhattacharya, Bamdad Hosseini, Nikola B. Kovachki, and Andrew M. Stuart. Model reduction and neural networks for parametric pdes, 2020.
- Katharina Bieker, Sebastian Peitz, Steven L. Brunton, J. Nathan Kutz, and Michael Dellnitz. Deep model predictive flow control with limited sensor data and online learning, 2020.
- Sindre Stenen Blakseth, Adil Rasheed, Trond Kvamsdal, and Omer San. Deep neural network enabled corrective source term approach to hybrid analysis and modeling. *Neural Networks*, 146: 181–199, 2022.
- Mathis Bode, Michael Gauding, Zeyu Lian, Dominik Denker, Marco Davidovic, Konstantin Kleinheinz, Jenia Jitsev, and Heinz Pitsch. Using physics-informed enhanced super-resolution generative adversarial networks for subfilter modeling in turbulent reactive flows. *Proceedings of the Combustion Institute*, 38(2):2617–2625, 2021.
- Massimo Bonavita and Patrick Laloyaux. Machine learning for model error inference and correction. *Journal of Advances in Modeling Earth Systems*, 12(12), 2020.
- J.P. Boyd. *Chebyshev and Fourier Spectral Methods: Second Revised Edition*. Dover Books on Mathematics. Dover Publications, 2013.
- Johannes Brandstetter, Daniel E. Worrall, and Max Welling. Message passing neural PDE solvers. In *International Conference on Learning Representations (ICLR)*, 2022.
- N. D. Brenowitz and C. S. Bretherton. Prognostic validation of a neural network unified physics parameterization. *Geophysical Research Letters*, 45(12):6289–6298, 2018.
- Noah D. Brenowitz, Tom Beucler, Michael Pritchard, and Christopher S. Bretherton. Interpreting and stabilizing machine-learning parametrizations of convection. *Journal of the Atmospheric Sciences*, 77(12):4357 – 4375, 2020.
- Christopher S. Bretherton, Brian Henn, Anna Kwa, Noah D. Brenowitz, Oliver Watt-Meyer, Jeremy McGibbon, W. Andre Perkins, Spencer K. Clark, and Lucas Harris. Correcting coarse-grid weather and climate models by machine learning from global storm-resolving simulations. *Journal of Advances in Modeling Earth Systems*, 14(2), 2022.
- William L. Briggs, Van Emden Henson, and Steve F. McCormick. *A Multigrid Tutorial (2nd Ed.)*. Society for Industrial and Applied Mathematics, USA, 2000. ISBN 0898714621.
- Michael M. Bronstein, Joan Bruna, Taco Cohen, and Petar Velickovic. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *CoRR*, 2021.
- Steven L. Brunton, Joshua L. Proctor, and J. Nathan Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 113(15):3932–3937, 2016.
- Keaton J. Burns, Geoffrey M. Vasil, Jeffrey S. Oishi, Daniel Lecoanet, and Benjamin P. Brown. Dedalus: A flexible framework for numerical simulations with spectral methods. *Physical Review Research*, 2(2), April 2020.
- Randal Burns, Gregory Eyink, Charles Meneveau, Alex Szalay, Tamer Zaki, Ethan Vishniac, Akshat Gupta, Mengze Wang, Yue Hao, Zhao Wu, and Gerard Lemson. Johns hopkins turbulence database, 2022. last accessed May, 2022.
- Salva Rühling Cachay, Emma Erickson, Arthur Fender C. Bucker, Ernest Pokropek, Willa Potosnak, Suyash Bire, Salomey Osei, and Björn Lütjens. The world as a graph: Improving el niño forecasts with graph neural networks, 2021a.

- Salva Rühling Cachay, Venkatesh Ramesh, Jason N. S. Cole, Howard Barker, and David Rolnick. ClimART: A benchmark dataset for emulating atmospheric radiative transfer in weather and climate models. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021b.
- J. Campin, C. Hill, H. Jones, and J. Marshall. Super-parameterization in ocean modeling: Application to deep convection. *Ocean Modelling*, 36:90–101, 2011.
- Giuseppe Carleo, Ignacio Cirac, Kyle Cranmer, Laurent Daudet, Maria Schuld, Naftali Tishby, Leslie Vogt-Maranto, and Lenka Zdeborová. Machine learning and the physical sciences. *Rev. Mod. Phys.*, 91, Dec 2019.
- S. Chakraborty, S. Adhikari, and R. Ganguli. The role of surrogate models in the development of digital twins of dynamic systems. *Applied Mathematical Modelling*, 90:662–681, 2021.
- Shing Chan and Ahmed H. Elsheikh. Parametric generation of conditional geological realizations using generative neural networks, 2019.
- Anindya Chatterjee. An introduction to the proper orthogonal decomposition. *Current Science*, 78 (7):808–817, 2000.
- Ashesh Chattopadhyay, Mustafa Mustafa, Pedram Hassanzadeh, and Karthik Kashinath. Deep spatial transformers for autoregressive data-driven forecasting of geophysical turbulence. In *Proceedings of the 10th International Conference on Climate Informatics, CI2020*, pp. 106–112, New York, NY, USA, 2020a. Association for Computing Machinery.
- Ashesh Chattopadhyay, Adam Subel, and Pedram Hassanzadeh. Data-driven super-parameterization using deep learning: Experimentation with multiscale lorenz 96 systems and transfer learning. *Journal of Advances in Modeling Earth Systems*, 12(11), 2020b.
- Tian Qi Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. In *Advances in Neural Information Processing Systems 31*, pp. 6571–6583. Curran Associates, Inc., 2018.
- Tianping Chen and Hong Chen. Universal approximation to nonlinear operators by neural networks with arbitrary activation functions and its application to dynamical systems. *IEEE Transactions on Neural Networks*, 6(4):911–917, 1995.
- Wei Chen and Faez Ahmed. Padgan: Learning to generate high-quality novel designs. *Journal of Mechanical Design*, 143(3).
- Taco S. Cohen and Max Welling. Steerable cnns. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, 2017*.
- Alberto Costa Nogueira, João Lucas de Sousa Almeida, Guillaume Auger, and Campbell D. Watson. Reduced order modeling of dynamical systems using artificial neural networks applied to water circulation. In Heike Jagode, Hartwig Anzt, Guido Juckeland, and Hatem Ltaief (eds.), *High Performance Computing*, pp. 116–136, Cham, 2020. Springer International Publishing.
- Priya L. Donti, David Rolnick, and J Zico Kolter. DC3: A learning method for optimization with hard constraints. In *International Conference on Learning Representations (ICLR)*, 2021.
- P. D. Dueben and P. Bauer. Challenges and design choices for global weather and climate models based on machine learning. *Geoscientific Model Development*, 11(10):3999–4009, 2018.
- Karthik Duraisamy, Gianluca Iaccarino, and Heng Xiao. Turbulence modeling in the age of data. *Annual Review of Fluid Mechanics*, 51(1):357–377, 2019.
- A. Dutt and V. Rokhlin. Fast fourier transforms for nonequispaced data. *SIAM Journal on Scientific Computing*, 14(6):1368–1393, 1993.
- S.J. Farlow. *Partial Differential Equations for Scientists and Engineers*. Dover books on advanced mathematics. Dover Publications, 1993.

- Hugo Frezat, Julien Le Sommer, Ronan Fablet, Guillaume Balarac, and Redouane Lguensat. A posteriori learning for quasi-geostrophic turbulence parametrization, 2022.
- Fabian Fuchs, Daniel Worrall, Volker Fischer, and Max Welling. Se(3)-transformers: 3d rotation equivariant attention networks. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33. Curran Associates, Inc., 2020.
- O. Fuhrer, T. Chadha, T. Hoefler, G. Kwasniewski, X. Lapillonne, D. Leutwyler, D. Lüthi, C. Osuna, C. Schär, T. C. Schulthess, and H. Vogt. Near-global climate simulation at 1 km resolution: establishing a performance baseline on 4888 gpus with cosmo 5.0. *Geosci. Model Dev.*, 11:1665–1681, 2018.
- K. Fukunaga and W.L.G. Koontz. Application of the karhunen-loève expansion to feature selection and ordering. *IEEE Transactions on Computers*, C-19(4):311–318, 1970.
- A. R. Ganguly, E. A. Kodra, A. Agrawal, A. Banerjee, S. Boriah, Sn. Chatterjee, So. Chatterjee, A. Choudhary, D. Das, J. Faghmous, P. Ganguli, S. Ghosh, K. Hayhoe, C. Hays, W. Hendrix, Q. Fu, J. Kawale, D. Kumar, V. Kumar, W. Liao, S. Liess, R. Mawalagedara, V. Mithal, R. Oglesby, K. Salvi, P. K. Snyder, K. Steinhauser, D. Wang, and D. Wuebbles. Toward enhanced understanding and projections of climate extremes using physics-guided data mining techniques. *Nonlinear Processes in Geophysics*, 21(4):777–795, 2014.
- P. Gentine, M. Pritchard, S. Rasp, G. Reinaudi, and G. Yacalis. Could machine learning break the convection parameterization deadlock? *Geophysical Research Letters*, 45(11):5742–5751, 2018.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014.
- Samuel Greydanus, Misko Dzamba, and Jason Yosinski. Hamiltonian neural networks. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems 32*, pp. 15379–15389. Curran Associates, Inc., 2019.
- Brian Groenke, Luke Madaus, and Claire Monteleoni. Climalign: Unsupervised statistical downscaling of climate variables via normalizing flows. In *Proceedings of the 10th International Conference on Climate Informatics, CI2020*, pp. 60–66, New York, NY, USA, 2020. Association for Computing Machinery.
- Gaurav Gupta, Xiongye Xiao, and Paul Bogdan. Multiwavelet-based operator learning for differential equations. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- Franz Hamilton, Alun L. Lloyd, and Kevin B. Flores. Hybrid modeling and prediction of dynamical systems. *PLOS Computational Biology*, 13(7):1–20, 07 2017.
- Katja Hansen, Grégoire Montavon, Franziska Biegler, Siamac Fazli, Matthias Rupp, Matthias Schefler, O. Anatole von Lilienfeld, Alexandre Tkatchenko, and Klaus-Robert Müller. Assessment and validation of machine learning methods for predicting molecular atomization energies. *Journal of Chemical Theory and Computation*, 9(8):3404–3419, 2013.
- Ramin Hasani, Mathias Lechner, Alexander Amini, Daniela Rus, and Radu Grosu. Liquid time-constant networks. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(9):7657–7666, May 2021.
- Hans Hersbach, Bill Bell, Paul Berrisford, Shoji Hirahara, András Horányi, Joaquín Muñoz-Sabater, Julien Nicolas, Carole Peubey, Raluca Radu, Dinand Schepers, Adrian Simmons, Cornel Soci, Saleh Abdalla, Xavier Abellan, Gianpaolo Balsamo, Peter Bechtold, Gionata Biavati, Jean Bidlot, Massimo Bonavita, Giovanna De Chiara, Per Dahlgren, Dick Dee, Michail Diamantakis, Rossana Dragani, Johannes Flemming, Richard Forbes, Manuel Fuentes, Alan Geer, Leo Haimberger, Sean Healy, Robin J. Hogan, Elías Hólm, Marta Janisková, Sarah Keeley, Patrick Laloyaux, Philippe Lopez, Cristina Lupu, Gabor Radnoti, Patricia de Rosnay, Iryna Rozum, Freja Vamborg, Sebastien Villaume, and Jean-Noël Thépaut. The era5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, 146(730):1999–2049, 2020.

- IPCC. Global warming of 1.5c. an ipcc special report on the impacts of global warming of 1.5c above pre-industrial levels and related global greenhouse gas emission pathways, in the context of strengthening the global response to the threat of climate change, sustainable development, and efforts to eradicate poverty, 2018.
- Xiaowei Jia, Jared Willard, Anuj Karpatne, Jordan Read, Jacob Zwart, Michael S Steinbach, and Vipin Kumar. Physics guided rnns for modeling dynamical systems: A case study in simulating lake temperature profiles. In *SIAM International Conference on Data Mining, SDM 2019*, SIAM International Conference on Data Mining, SDM 2019, pp. 558–566. Society for Industrial and Applied Mathematics Publications, 2019.
- Xiaowei Jia, Jared Willard, Anuj Karpatne, Jordan S. Read, Jacob A. Zwart, Michael Steinbach, and Vipin Kumar. Physics-guided machine learning for scientific discovery: An application in simulating lake temperature profiles. *ACM/IMS Trans. Data Sci.*, 2(3), 2021.
- Peishi Jiang, Nis Meinert, Helga Jordão, Constantin Weisser, Simon Holgate, Alexander Lavin, Björn Lütjens, Dava Newman, Haruko Wainwright, Catherine Walker, and Patrick Barnard. Digital Twin Earth – Coasts: Developing a fast and physics-informed surrogate model for coastal floods via neural operators. *2021 NeurIPS Workshop on Machine Learning for the Physical Sciences (MLAPS)*, 2021.
- Pengzhan Jin, Zhen Zhang, Aiqing Zhu, Yifa Tang, and George Em Karniadakis. Sympnets: Intrinsic structure-preserving symplectic networks for identifying hamiltonian systems. *Neural Networks*, 132, 12 2020.
- J. Nagoor Kani and Ahmed H. Elsheikh. DR-RNN: A deep residual recurrent neural network for model reduction. *CoRR*, abs/1709.00939, 2017.
- George Em Karniadakis, Ioannis G. Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3:422–440, June 2021.
- A. Karpatne, I. Ebert-Uphoff, S. Ravela, H. A. Babaie, and V. Kumar. Machine learning for the geosciences: Challenges and opportunities. *IEEE Transactions on Knowledge and Data Engineering*, 31(8):1544–1554, Aug 2019.
- Anuj Karpatne, Gowtham Atluri, James H. Faghmous, Michael Steinbach, Arindam Banerjee, Auroop Ganguly, Shashi Shekhar, Nagiza Samatova, and Vipin Kumar. Theory-guided data science: A new paradigm for scientific discovery from data. *IEEE Transactions on Knowledge and Data Engineering*, 29(10):2318–2331, 2017.
- Anuj Karpatne, William Watkins, Jordan Read, and Vipin Kumar. Physics-guided Neural Networks (PGNN): An Application in Lake Temperature Modeling. *arXiv e-prints*, pp. arXiv:1710.11431, October 2017.
- K. Kashinath, M. Mustafa, A. Albert, J-L. Wu, C. Jiang, S. Esmailzadeh, K. Azizzadenesheli, R. Wang, A. Chattopadhyay, A. Singh, A. Manepalli, D. Chirila, R. Yu, R. Walters, B. White, H. Xiao, H. A. Tchelepi, P. Marcus, A. Anandkumar, P. Hassanzadeh, and null Prabhat. Physics-informed machine learning: case studies for weather and climate modelling. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 379(2194), 2021.
- Marat Khairoutdinov, David Randall, and Charlotte DeMott. Simulations of the atmospheric general circulation using a cloud-resolving model as a superparameterization of physical processes. *Journal of the Atmospheric Sciences*, 62(7 I):2136–2154, jul 2005.
- Yuehaw Khoo and Lexing Ying. Switchnet: A neural network model for forward and inverse scattering problems. *SIAM Journal on Scientific Computing*, 41(5), 2019.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun (eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.

- Nikola B. Kovachki, Zongyi Li, Burigede Liu, Kamyar Azizzadenesheli, Kaushik Bhattacharya, Andrew M. Stuart, and Anima Anandkumar. Neural operator learning maps between function spaces. *CoRR*, abs/2108.08481, 2021.
- Rupa Kurinchi-Vendhan, Björn Lütjens, Ritwik Gupta, Lucien Werner, and Dava Newman. Wisosuper: Benchmarking super-resolution methods on wind and solar data. *Conference on Neural Information Processing Systems (NeurIPS) Workshop on Tackling Climate Change with Machine Learning (CCML)*, 2021.
- Corentin J. Lapeyre, Antony Misdariis, Nicolas Cazard, Denis Veynante, and Thierry Poinso. Training convolutional neural networks to estimate turbulent sub-grid scale reaction rates. *Combustion and Flame*, 203:255–264, 2019.
- Kookjin Lee and Kevin T. Carlberg. Model reduction of dynamical systems on nonlinear manifolds using deep convolutional autoencoders. *Journal of Computational Physics*, 404:108973, 2020. ISSN 0021-9991.
- Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Andrew Stuart, Kaushik Bhattacharya, and Anima Anandkumar. Multipole graph neural operator for parametric partial differential equations. In *Advances in Neural Information Processing Systems*, volume 33, pp. 6755–6766. Curran Associates, Inc., 2020.
- Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. *ICML*, 2021a.
- Zongyi Li, Nikola B. Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew M. Stuart, and Anima Anandkumar. Markov neural operators for learning chaotic systems. *CoRR*, abs/2106.06898, 2021b.
- Julia Ling, Andrew Kurzawski, and Jeremy Templeton. Reynolds averaged turbulence modelling using deep neural networks with embedded invariance. *Journal of Fluid Mechanics*, 807:155–166, 2016.
- Burigede Liu, Nikola Kovachki, Zongyi Li, Kamyar Azizzadenesheli, Anima Anandkumar, Andrew M. Stuart, and Kaushik Bhattacharya. A learning-based multiscale method and its application to inelastic impact problems. *Journal of the Mechanics and Physics of Solids*, 158:104668, 2022.
- Ziming Liu, Bohan Wang, Qi Meng, Wei Chen, Max Tegmark, and Tie-Yan Liu. Machine-learning nonconservative dynamics for new-physics detection. *Phys. Rev. E*, 104:055302, Nov 2021.
- Yun Long, Xueyuan She, and Saibal Mukhopadhyay. Hybridnet: Integrating model-based and data-driven learning to predict evolution of dynamical systems. In *Proceedings of The 2nd Conference on Robot Learning*, volume 87 of *Proceedings of Machine Learning Research*, pp. 551–560. PMLR, 29–31 Oct 2018a.
- Zichao Long, Yiping Lu, Xianzhong Ma, and Bin Dong. PDE-net: Learning PDEs from data. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 3208–3216. PMLR, 10–15 Jul 2018b.
- Zichao Long, Yiping Lu, and Bin Dong. Pde-net 2.0: Learning pdes from data with a numeric-symbolic hybrid deep network. *Journal of Computational Physics*, 399:108925, 2019.
- Edward Lorenz. Predictability - a problem partly solved. In Tim Palmer and Renate Hagedorn (eds.), *Predictability of Weather and Climate*. Cambridge University Press, Cambridge, 2006.
- Edward N. Lorenz and Kerry A. Emanuel. Optimal sites for supplementary weather observations: Simulation with a small model. *Journal of the Atmospheric Sciences*, 55(3):399 – 414, 1998.
- Lu Lu, Pengzhan Jin, Guofei Pang, Zhongqiang Zhang, and George Em Karniadakis. Learning nonlinear operators via deeponet based on the universal approximation theorem of operators. *Nature Machine Intelligence*, 3:218–229, 2021.

- Björn Lütjens, Catherine H. Crawford, Mark Veillette, and Dava Newman. Pce-pinns: Physics-informed neural networks for uncertainty propagation in ocean modeling. *International Conference on Learning Representations (ICLR) Workshop on AI for Modeling Oceans and Climate Change*, May 2021.
- Björn Lütjens, Catherine H Crawford, Mark Veillette, and Dava Newman. Spectral PINNs: Fast uncertainty propagation with physics-informed neural networks. In *Advances in Neural Information Processing Systems (NeurIPS) Workshop on The Symbiosis of Deep Learning and Differential Equations (DLDE)*, 2021.
- Michael Lutter, Christian Ritter, and Jan Peters. Deep lagrangian networks: Using physics as model prior for deep learning. In *International Conference on Learning Representations (ICLR)*, 2019.
- Andrew Majda and Xiaoming Wang. *Nonlinear Dynamics and Statistical Theories for Basic Geophysical Flows*. Cambridge University Press, 2006.
- Marie C. McGraw and Elizabeth A. Barnes. Memory matters: A case for granger causality in climate variability studies. *Journal of Climate*, 31(8):3289 – 3300, 2018.
- Kendal McGuffie and Ann Henderson-Sellers. *A Climate Modeling Primer, Third Edition*. John Wiley and Sons, Ltd., Jan. 2005.
- Arvind Mohan, Don Daniel, Michael Chertkov, and Daniel Livescu. Compressed convolutional lstm: An efficient deep learning framework to model high fidelity 3d turbulence, 2019.
- Alberto C. Nogueira Jr., Felipe C. T. Carvalho, João Lucas S. Almeida, Andres Codas, Eloisa Ben-tivegna, and Campbell D. Watson. Reservoir computing in reduced order modeling for chaotic dynamical systems. In Heike Jagode, Hartwig Anzt, Hatem Ltaief, and Piotr Luszczek (eds.), *High Performance Computing*, pp. 56–72, Cham, 2021. Springer International Publishing.
- Paul A. O’Gorman and John G. Dwyer. Using machine learning to parameterize moist convection: Potential for modeling of climate, climate change, and extreme events. *Journal of Advances in Modeling Earth Systems*, 10(10):2548–2563, 2018.
- Peter J. Olver. *Symmetry Groups of Differential Equations*, pp. 77–185. Springer New York, New York, NY, 1986.
- Sagaut P. *Large Eddy Simulation for Incompressible Flows: An Introduction*. Scientific Computation. Springer-Verlag Berlin Heidelberg, 3 edition, 2006.
- Tim Palmer, Bjorn Stevens, and Peter Bauer. We need an international center for climate modeling, 2019. URL <https://blogs.scientificamerican.com/observations/we-need-an-international-center-for-climate-modeling/>. last accessed 04/13/20.
- Ambrish Pandey, Janet D. Scheel, and Jörg Schumacher. Turbulent superstructures in rayleigh-bénard convection. *Nature Communications*, 9, 2018.
- Eric J. Parish and Karthik Duraisamy. A paradigm for data-driven predictive modeling using field inversion and machine learning. *Journal of Computational Physics*, 305:758–774, 2016. ISSN 0021-9991.
- Jaideep Pathak, Brian Hunt, Michelle Girvan, Zhixin Lu, and Edward Ott. Model-free prediction of large spatiotemporally chaotic systems from data: A reservoir computing approach. *Phys. Rev. Lett.*, 120, Jan 2018.
- Jaideep Pathak, Mustafa Mustafa, Karthik Kashinath, Emmanuel Motheau, Thorsten Kurth, and Marcus Day. Using machine learning to augment coarse-grid computational fluid dynamics simulations, 2020.
- Jaideep Pathak, Shashank Subramanian, Peter Harrington, Sanjeev Raja, Ashesh Chattopadhyay, Morteza Mardani, Thorsten Kurth, David Hall, Zongyi Li, Kamyar Azizzadenesheli, Pedram Hassanzadeh, Karthik Kashinath, and Animashree Anandkumar. FourCastNet: A Global Data-driven High-resolution Weather Model using Adaptive Fourier Neural Operators. February 2022.


- Greg Pavliotis and Andrew Stuart. *Multiscale Methods Averaging and Homogenization*, volume 53 of *Texts in Applied Mathematics*. Springer-Verlag New York, 1 edition, 2008. ISBN 978-0-387-73829-1.
- Grace C. Y. Peng, Mark Alber, Adrian Buganza Tepole, William R. Cannon, Suvranu De, Savador Dura-Bernal, Krishna Garikipati, George Karniadakis, William W. Lytton, Paris Perdikaris, Linda Petzold, and Ellen Kuhl. Multiscale modeling meets machine learning: What can we learn? *Archives of Computational Methods in Engineering*, 28:1017–1037, 2021.
- Aviral Prakash, Kenneth E. Jansen, and John A. Evans. Invariant Data-Driven Subgrid Stress Modeling in the Strain-Rate Eigenframe for Large Eddy Simulation. *arXiv e-prints*, June 2021.
- Zhaozhi Qian, Krzysztof Kacprzyk, and Mihaela van der Schaar. D-CODE: Discovering closed-form ODEs from observed trajectories. In *International Conference on Learning Representations*, 2022.
- Alfio Quarteroni and Gianluigi Rozza. Reduced order methods for modeling and computational reduction, 2014.
- Christopher Rackauckas, Yingbo Ma, Julius Martensen, Collin Warner, Kirill Zubov, Rohit Supekar, Dominic Skinner, and Ali Ramadhan. Universal differential equations for scientific machine learning. *ArXiv*, abs/2001.04385, 2020.
- M. Raissi, P. Perdikaris, and G.E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686 – 707, 2019.
- S. Rasp. Coupled online learning as a way to tackle instabilities and biases in neural network parameterizations: general algorithms and lorenz 96 case study (v1.0). *Geoscientific Model Development*, 13(5):2185–2196, 2020.
- Stephan Rasp. Lorenz 96 is too easy! machine learning research needs a more realistic toy model, July 2019. URL <https://raspstephan.github.io/blog/lorenz-96-is-too-easy/>. last accessed May 2022.
- Stephan Rasp, Michael S. Pritchard, and Pierre Gentine. Deep learning to represent subgrid processes in climate models. *Proceedings of the National Academy of Sciences*, 115(39):9684–9689, 2018.
- Stephan Rasp, Peter D. Dueben, Sebastian Scher, Jonathan A. Weyn, Soukayna Mouatadid, and Nils Thuerey. Weatherbench: A benchmark data set for data-driven weather forecasting. *Journal of Advances in Modeling Earth Systems*, 12(11), 2020.
- Markus Reichstein, Gustau Camps-Valls, Bjorn Stevens, Martin Jung, Joachim Denzler, Nuno Carvalhais, and Prabhat. Deep learning and process understanding for data-driven earth system science. *Nature*, 566:195 – 204, 2019.
- Danilo Jimenez Rezende and Shakir Mohamed. Variational inference with normalizing flows. In *Proceedings of the 32nd International Conference on International Conference on Machine Learning - Volume 37, ICML’15*, pp. 1530–1538. JMLR.org, 2015.
- Juliette N. Rooney-Varga, Florian Kapmeier, John D. Sterman, Andrew P. Jones, Michele Putko, and Kenneth Rath. The climate action simulation. *Simulation & Gaming*, 51(2):114–140, 2020. URL <https://en-roads.climateinteractive.org/>.
- Justin Sirignano, Jonathan F. MacArt, and Jonathan B. Freund. Dpm: A deep learning pde augmentation method with application to large-eddy simulation. *Journal of Computational Physics*, 423, 2020.
- Ralph C. Smith. Uncertainty quantification: Theory, implementation, and applications. In *Computational science and engineering*, pp. 382. SIAM, 2013.

- Snagglebit. Sloth nesting dolls - hand painted modern russian matryoshka doll set, 2022. URL <https://www.etsy.com/ie/listing/690540535/sloth-nesting-dolls-hand-painted-modern>. last accessed 01/22.
- Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In Francis Bach and David Blei (eds.), *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pp. 2256–2265, Lille, France, 07–09 Jul 2015. PMLR.
- Kim Stachenfeld, Drummond Buschman Fielding, Dmitrii Kochkov, Miles Cranmer, Tobias Pfaff, Jonathan Godwin, Can Cui, Shirley Ho, Peter Battaglia, and Alvaro Sanchez-Gonzalez. Learned simulators for turbulence. In *International Conference on Learning Representations*, 2022.
- Karen Stengel, Andrew Glaws, Dylan Hettinger, and Ryan N. King. Adversarial super-resolution of climatological wind and solar data. *Proceedings of the National Academy of Sciences*, 117(29):16805–16815, 2020.
- Adam Subel, Ashesh Chattopadhyay, Yifei Guan, and Pedram Hassanzadeh. Data-driven subgrid-scale modeling of forced burgers turbulence using deep learning with generalization to higher reynolds numbers via transfer learning. *Physics of Fluids*, 33(3):031702, 2021.
- Nathaniel Thomas, Tess E. Smidt, Steven Kearnes, Lusann Yang, Li Li, Kai Kohlhoff, and Patrick Riley. Tensor field networks: Rotation- and translation-equivariant neural networks for 3d point clouds. *CoRR*, abs/1802.08219, 2018.
- Tobias Thornes, Peter Duben, and Tim Palmer. On the use of scale-dependent precision in earth system modelling. *Quarterly Journal of the Royal Meteorological Society*, 143(703):897–908, 2017.
- Benjamin A. Toms, Elizabeth A. Barnes, and Imme Ebert-Uphoff. Physically interpretable neural networks for the geosciences: Applications to earth system variability. *Journal of Advances in Modeling Earth Systems*, 12(9), 2020.
- Kiwon Um, Robert Brand, Yun Fei, Philipp Holl, and Nils Thuerey. Solver-in-the-Loop: Learning from Differentiable Physics to Interact with Iterative PDE-Solvers. *Advances in Neural Information Processing Systems*, 2020.
- Thomas Vandal, Evan Kodra, Sangram Ganguly, Andrew Michaelis, Ramakrishna Nemani, and Auroop R. Ganguly. DeepSD: Generating high resolution climate change projections through single image super-resolution. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD ’17, pp. 1663–1672, New York, NY, USA, 2017. Association for Computing Machinery.
- Peidong Wang, Janni Yuval, and Paul A. O’Gorman. Non-local parameterization of atmospheric subgrid processes with neural networks, 2022.
- Sifan Wang, Hanwen Wang, and Paris Perdikaris. Learning the solution operator of parametric partial differential equations with physics-informed deepnets. *Science Advances*, 7(40), 2021.
- Mark J. Webb, Adrian P. Lock, Christopher S. Bretherton, Sandrine Bony, Jason N.S. Cole, Abderrahmane Idelkadi, Sarah M. Kang, Tsuyoshi Koshiro, Hideaki Kawai, Tomoo Ogura, Romain Roehrig, Yechul Shin, Thorsten Mauritsen, Steven C. Sherwood, Jessica Vial, Masahiro Watanabe, Matthew D. Woelfle, and Ming Zhao. The impact of parametrized convection on cloud feedback. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 373(2054), nov 2015.
- Daniel S. Wilks. Effects of stochastic parametrizations in the lorenz ’96 system. *Quarterly Journal of the Royal Meteorological Society*, 131(606):389–407, 2005.
- Jared Willard, Xiaowei Jia, Shaoming Xu, Michael Steinbach, and Vipin Kumar. Integrating scientific knowledge with machine learning for engineering and environmental systems. *ACM Comput. Surv.*, jan 2022.

- Schiesser William. *The Numerical Method of Lines: Integration of Partial Differential Equations*. Elsevier, 1991.
- Matthew O. Williams, Ioannis G. Kevrekidis, and Clarence W. Rowley. A data-driven approximation of the koopman operator: Extending dynamic mode decomposition. *Journal of Nonlinear Science*, 25:1432–1467, 2015.
- Jin-Long Wu, Heng Xiao, and Eric Paterson. Physics-informed machine learning approach for augmenting turbulence models: A comprehensive framework. *Phys. Rev. Fluids*, 3:074602, Jul 2018.
- Jin-Long Wu, Karthik Kashinath, Adrian Albert, Dragos Chirila, Prabhat, and Heng Xiao. Enforcing statistical constraints in generative adversarial networks for modeling chaotic dynamical systems. *Journal of Computational Physics*, 406, 2020.
- You Xie, Erik Franz, Mengyu Chu, and Nils Thuerey. Tempogan: A temporally coherent, volumetric gan for super-resolution fluid flow. *ACM Trans. Graph.*, 37(4), jul 2018.
- Alireza Yazdani, Lu Lu, Maziar Raissi, and George Em Karniadakis. Systems biology informed deep learning for inferring parameters and hidden dynamics. *PLOS Computational Biology*, 16(11):1–19, 11 2020.
- Yuan Yin, Vincent Le Guen, Jérémie Dona, Emmanuel de Bézenac, Ibrahim Ayed, Nicolas Thome, and Patrick Gallinari. Augmenting physical models with deep networks for complex dynamics forecasting. *Journal of Statistical Mechanics: Theory and Experiment*, 2021(12):124012, dec 2021.
- Janni Yuval, Paul A. O’Gorman, and Chris N. Hill. Use of neural networks for stable, accurate and physically consistent parameterization of subgrid atmospheric processes with good performance at reduced precision. *Geophysical Research Letter*, 48:e2020GL091363, 2021.
- Yang Zeng, Jin-Long Wu, and Heng Xiao. Enforcing imprecise constraints on generative adversarial networks for emulating physical systems. *Communications in Computational Physics*, 30(3):635–665, 2021.
- Liang Zhang, Gang Wang, and Georgios B. Giannakis. Real-time power system state estimation and forecasting via deep unrolled neural networks. *IEEE Transactions on Signal Processing*, 67(15):4069–4077, 2019.
- Linfeng Zhang, Jiequn Han, Han Wang, Roberto Car, and Weinan E. Deep potential molecular dynamics: A scalable model with the accuracy of quantum mechanics. *Phys. Rev. Lett.*, 120, Apr 2018.
- Jie Zhou, Ganqu Cui, Shengding Hu, Zhengyan Zhang, Cheng Yang, Zhiyuan Liu, Lifeng Wang, Changcheng Li, and Maosong Sun. Graph neural networks: A review of methods and applications. *AI Open*, 1:57–81, 2020.

A APPENDIX

A.1 QUASI-GEOSTROPHIC TURBULENCE



(a) Ground-truth

Figure 6: **Quasi-Geostrophic Turbulence (QGT)**. We demonstrate the first steps to extend MNO to 2D QGT. This plot shows our ground-truth training data with the direct numerical simulation, ω , and subgrid parametrization, ω' .

We are planning to demonstrate multiscale neural operator a high-dimensional system, specifically the one-layer quasi-geostrophic (QG) turbulence depicted in Fig. 6. QG turbulence is a derivative of the Navier-Stokes equations and a good model for atmospheric turbulence at the equator. The equations are derived from the incompressible, i.e., $\nabla \cdot \mathbf{u} = 0$, Navier-Stokes equations by 1) taking the curl of velocity field, $w = \nabla \times \mathbf{u}$, and 2) assuming the beta-plane approximation, $f = f_0 + \beta y$, and hydrostatic, $\frac{\partial p}{\partial z} = -\rho g$, and geostrophic, $f v = \frac{1}{\rho} \frac{\partial p}{\partial x}$; $f u = -\frac{1}{\rho} \frac{\partial p}{\partial y}$, balances (Majda & Wang, 2006). The resulting equations, called quasi-geostrophic turbulence, are given by:

$$\begin{aligned} \partial_t \omega + J(\psi, \omega) &= \nu \nabla^2 \omega - \mu \omega - \beta \partial_x \psi + F \\ \omega &= \nabla^2 \psi \end{aligned} \quad (11)$$

where ω is the vorticity, $\mathbf{u} = [u, v]^T = [-\partial_y \psi, \partial_x \psi]^T$ is the velocity vector, ψ is the streamfunction, $J(\psi, \omega) = \partial_x \psi \partial_y \omega - \partial_y \psi \partial_x \omega$ is the nonlinear Jacobian operator. Further, the parameters are the turbulent viscosity, ν , linear drag coefficient, μ , Rossby parameter, β , and source term, F . Vorticity can be computed with $\omega = \hat{z} \cdot \nabla \times \mathbf{u} = \partial_x v - \partial_y u$

Filtering the equation with a kernel results in the parametrized large-scale equation given by:

$$\partial_t \bar{\omega} + J(\bar{\psi}, \bar{\omega}) = \nu \nabla^2 \bar{\omega} - \mu \bar{\omega} - \beta \partial_x \bar{\psi} + \bar{F} + \underbrace{J(\bar{\psi}, \bar{\omega}) - \overline{J(\psi, \omega)}}_{\text{Parametrization: } h(\psi, \omega)} \quad (12)$$

We then aim to approximate the subgrid-scale (SGS) parametrization, $\mathcal{K}_\theta \approx h$, with the neural operator, such that the final model is:


$$\begin{aligned} \frac{\partial \omega(x, y)}{\partial t} &= \frac{\partial \bar{\omega}(x, y)}{\partial t} + \mathcal{K}_\theta(\psi, \omega)(x, y)_\omega \\ \frac{\partial \psi(x, y)}{\partial t} &= \frac{\partial \bar{\psi}(x, y)}{\partial t} + \mathcal{K}_\theta(\psi, \omega)(x, y)_\psi \end{aligned} \quad (13)$$

The QG turbulence equation is solved with a pseudospectral solver in space and RK4 explicit time integration. We choose the parameters, $N_x = N_y = 512$, $\Delta t = 480\text{s}$, $\mu = 1.25 \times 10^{-8}\text{s}^{-1}$, $\nu = 352\text{m}^2/\text{s}$, $\beta = 0$, and Reynolds number, $\text{Re} = 22 \times 10^4$. The variables are non-dimensionalized with $T_d = 1.2 \times 10^6\text{s}$, i.e., $\Delta t_{\text{solver}} = \Delta t/T_d$ and $L_d = 504 \times 10^4/\pi m$, i.e., $\Delta x_{\text{solver}} = 2\pi/N_x$. The

reduced system is run with scale $\delta = 4$, such that $\bar{N}_x = \bar{N}_y = 128$. The forcing initiates turbulent mixing and simulates slowly varying wind stress according to the solution of, $F = C_f(t)[\cos(4y + \pi \sin(1.4t)) - \cos(4x + \pi \sin(1.5t))]$, and $0.5\|F\|_2 = 3$ with enstrophy injection rate, $C_F(t)$. To generate turbulent chaotic dynamics that are decoupled from the initial state, the simulation is initialized with some large-scale Fourier states and warmed up for 1300days. After warm-up we generate 18000 iterations ($\delta 18000$ at the fine-scale) from which we independently sample training and validation snippets.

A.2 RAYLEIGH-BÉNARD CONVECTION

The proposed multiscale neural operator can also be leveraged for systems without explicit multi-scale formulation. We demonstrate this by formulating the MNO equations for Rayleigh-Bénard Convection equations, as displayed in Fig. 7.



(a) Ground-truth

Figure 7: **Rayleigh-Bénard Convection.** We depicted a sample plot for ground-truth training data of the 2D RBC.

A.2.1 DETAILS AND INTERPRETATION

Rayleigh-Bénard Convection (RBC) is a challenging set of equations for turbulent, chaotic, and convection-dominated flows. The equation finds applications in fluid dynamics, atmospheric dynamics, radiation, phase changes, magnetic fields, and more (Pandey et al., 2018).

So far, we have generated a ground-truth dataset that we implemented with the 2D turbulent Rayleigh-Bénard Convection equations with Dedalus spectral solver (Burns et al., 2020) similar to (Pandey et al., 2018):

$$\begin{aligned}
 \frac{\partial u}{\partial t} + u \cdot \nabla u &= \sqrt{\frac{\text{Pr}}{\text{Ra}}} \nabla^2 u - \nabla p + b \\
 \frac{\partial T}{\partial t} + u \cdot \nabla T &= \frac{1}{\sqrt{\text{RaPr}}} \nabla^2 T \\
 \nabla \cdot u &= 0
 \end{aligned}
 \tag{14}$$

with temperature/buoyancy, T , Rayleigh number, $Ra = g\alpha\Delta TH^3/(\nu\kappa)$, thermal expansion coefficient, α , Prandtl number, $Pr = \nu/\kappa$, momentum diffusivity or kinematic viscosity, ν , thermal diffusivity, $\kappa = \frac{1}{\sqrt{RaPr}}$, acceleration due to gravity, g , temperature difference, ΔT , unit vector, e , pressure, p , Nusselt number, $Nu = \sqrt{\frac{Pr}{Ra}}$, Reynolds number, $Re = \sqrt{\langle \nabla^2 u \rangle_{V,t} \frac{Ra}{Pr}}$, and full volume-time average, $\langle \cdot \rangle_{V,t}$, cell length, L_x . The equations have been non-dimensionalized with the free-fall velocity, $U_f = \sqrt{g\alpha\Delta H}$, and cell height, H . In the horizontal direction, x , we use periodic boundary conditions and in the vertical direction, z , we use no-slip boundary conditions for the velocity, $u(z=0) = u(z=L_z) = 0$, and fixed temperatures, $T(z=0) = L_z$, $T(z=L_z) = 0$. The initial conditions are sampled randomly, $b(z, t=0) = L_z + z + z(L_z - z)\omega$, with $\omega \sim \mathcal{N}(0, 1 \times 10^{-3})$.

We chose: $Ra = 2 \times 10^6$, $Pr = 1$, $L_x = 4$, $H = 1$.

A.3 FOURIER NEURAL OPERATOR

Our neural operator for learning subgrid parametrizations is based on Fourier neural operators (Li et al., 2021a). Intuitively, the neural operator learns a parameter-to-solution mapping by learning a global convolution kernel. In detail, it learns the operator to transform the current large-scale state, $\underline{X}(x_{0:K}, t) \in \mathbb{R}^{K \times dx}$ to the subgrid parametrization, $\hat{f}_x(x_{0:K}, t) := \underline{X}_{0:K} \in \mathbb{R}^{K \times dx}$ with number of grid points, K , and input dimensionality, dx , according to the following equations:

$$\begin{aligned} \underline{v}_0 &= \underline{X}_{0:K} P^T + 1^{K \times 1} b_P \\ \underline{v}_{i+1} &= \sigma \left(\underline{v}_i W^T + \int_{D_x} \kappa_\phi(x, x') v_i(x') dx' \right) \\ &\approx \sigma \left(\underline{v}_i W^T + 1^{n_v \times 1} b_W + \mathcal{F}^{-1}(R_\phi \cdot \mathcal{F} \underline{v}_i) \right) \\ \hat{f}_{x,0:K} &= \underline{v}_{n_d} Q^T + 1^{K \times 1} b_Q \end{aligned} \tag{15}$$

First, MNO lifts the input via a linear transform with matrix, $P \in \mathbb{R}^{n_v \times dx}$, bias, $b_P \in \mathbb{R}^{1 \times n_v}$, vector of ones, $1^{K \times 1}$, and number of channels, n_v . The linear transform is local in space, i.e., the same transform is applied to each grid point.

Second, multiple nonlinear ‘‘Fourier layers’’ are applied to the encoded/lifted state. The encoded/lifted state’s, $\underline{v}_i \in \mathbb{R}^{K \times n_v}$, spatial dimension is transformed into the Fourier domain via a fast Fourier transform. We implement the FFT as a multiplication with the pre-built forward and inverse Type-I DST matrix, $\mathcal{F} \in \mathbb{C}^{k_{\max} \times K}$ and $\mathcal{F}^{-1} \in \mathbb{C}^{K \times k_{\max}}$, respectively, returning the vector, $\mathcal{F} \underline{v}_i \in \mathbb{C}^{k_{\max} \times n_v}$.

The dynamics are learned via convoluting the encoded state with a weight matrix. In Fourier space, convolution is a multiplication, hence each frequency is multiplied with a complex weight matrix across the channels, such that $R \in \mathbb{C}^{k_{\max} \times n_v \times n_v}$. In parallel to the convolution with R , the encoded state is multiplied with the linear transform, $W \in \mathbb{R}^{n_v \times n_v}$, and bias, $b_W \in \mathbb{R}^{1 \times n_v}$. From a representation learning-perspective, the Fourier decomposition as a fast and interpretable feature extraction method that extracts smooth, periodic, and global features. The linear transform can be interpreted as residual term concisely capturing nonlinear residuals.

So far, we have only applied linear transformations. To introduce nonlinearities, we apply a nonlinear activation function, σ , at the end of each Fourier layer. While the non-smoothness of the activation function ReLU, $\sigma(z) = \max(0, z)$, could introduce unwanted discontinuities in the solution, we choose it resulted in more accurate models than smoother activation functions such as tanh or sigmoid.

Finally, the transformed state, \underline{v}_{n_d} , is projected back onto solution space via another linear transform, $Q \in \mathbb{R}^{dx \times n_v}$, and bias, b_Q .

The values of all trainable parameters, P, R, W, Q, b_* , are found by using a nonlinear optimization algorithm, such as stochastic gradient descent or, here, Adam Kingma & Ba (2015). We have used MSE between the predicted, \hat{f}_x , and ground-truth, f_x , subgrid parametrizations as loss. The neural operator is implemented in pytorch, but does not require an autodifferentiable PDE solver to generate training data. During implementation, we used the DFT which assumes a uniformly spaced grids,

but can be exchanged with non-uniform DFTs (NUDFT) to transform non-uniform grids (Dutt & Rokhlin, 1993).

A.4 MULTISCALE LORENZ96

A.4.1 DETAILS AND INTERPRETATION

The equation contains K variables, $X_k \in \mathbb{R}$, and JK small-scale variables, $Y_{j,k} \in \mathbb{R}$ that represent large-scale or small-scale atmospheric dynamics such as the movement of storms or formation of clouds, respectively. At every time-step each large-scale variable, X_k , influences and is influenced by J small-scale variables, $Y_{0:J,k}$. The coupling could be interpreted as X_k causing static instability and $Y_{j,k}$ causing drag from turbulence or latent heat fluxes from cloud formation. The indices k, j are both interpreted as latitude, while $k \in \{0, \dots, K-1\}$ indexes boxes of latitude and $j \in \{0, \dots, J-1\}$ indexes elements inside the box. Illustrated on a 1D Earth with a circumference of 360° that is discretized with $K = 36, J = 10$, one a spatial step in k, j would equal $10^\circ, 1^\circ$, respectively Lorenz (2006); we choose $K = J = 4$. A time step with $\Delta t = 0.005$ would equal 36 minutes Lorenz (2006).

We choose a large forcing, $F > 10$, for which the equation becomes chaotic. The last terms in each equation capture the interaction between small- and large-scale, $f_{x,k} = -\frac{hc}{b} \sum_{j=0}^J Y_{j,k}(X_k), f_y$. The scale interaction is defined by the parameters where $h = 0.5$ is the coupling strength between spatial scales (with no coupling if h would be zero), $b = 10$ is the relative magnitude, and $c = 8$ the evolution speed of $X - Y$. The linear, $-X_k$, and quadratic terms, X_k^2 , model dissipative and advective (e.g., moving) dynamics, respectively.

The equation assumes perfect ‘‘scale separation’’ which means that small-scale variables of different grid boxes, k , are independent of each other at a given timestep, $Y_{j_1,k_2}(t) \perp Y_{j_2,k_1}(t) \forall t, j_1, j_2, k_1 \neq k_2$. The separation of small- and large-scale variables can be along the same or different domain and the discretized variables would then be $y \in [0, \Delta x]$ or $y \in [y_0, y_{\text{end}}]$, respectively. The equation wraps around the full large- or small-scale domain by using periodic boundaries, $X_{-k} := X_{K-k}, X_{K+k} := X_k, Y_{-j,k} := Y_{J-j,k}, Y_{J+j,k} := Y_{j,k}$. Note that having periodic boundary conditions in the small-scale domain allows for superparametrization, i.e., independent simulation of the small-scale dynamics Campin et al. (2011) and differs from the three-tier Lorenz96 where variables at the borders of the small-scale domain depend on small-scale variables of the neighbouring k Thornes et al. (2017).

A.4.2 SIMULATION

The initial conditions are sampled uniformly from a set of integers, $X(t_0) \sim U(-5, -4, \dots, 5, 6)$, as a mean-zero unit-variance Gaussian $Y(t_0) \sim \mathcal{N}(0, 1)$, and lower scale Gaussian $Z(t_0) \sim 0.05\mathcal{N}(0, 1)$. The train and test set contains 4k and 1k samples, respectively. Each sample is $T = 1$ model time unit (MTU) or 200 ($=T/\Delta t$) time-steps long, which corresponds to 5 Earth days ($=T/\Delta t * 36\text{min}$ with $\Delta t = 0.005$) Lorenz (2006). Hence, our results test the generalization towards different initial conditions, but not robustness to extrapolation or different choices of parameters, c, b, h, F . The sampling starts after $T = 10$. warmup time. The dataset uses double precision.

We solve the equation by fourth order Runge-Kutta in time with step size $\Delta t = 0.005$, similar to (Lorenz & Emanuel, 1998). For a PDE that is discretized with fixed time step, Δt , the ground-truth train and test data, $h_{x,0:K}(t)$, is constructed by integrating the coupled large- and small-scale dynamics.

Note, that the neural operator only takes in the current state of the large-scale dynamics. Hence, , i.e., it uses the full large-scale spatial domain as input, which exploits spatial correlations and learns parametrizations that are independent of the large-scale spatial discretization.

Our method can be queried for infinite time-steps into the future as it does not use time as input.

We are incorporating the prior knowledge from physics by calculating the large-scale dynamics, $dX_{LS,0:K}$. Note that the small-scale physics do not need to be known. Hence, MNO could be applied to any fixed time-step dataset for which an approximate model is known.




Figure 8: **Mean accuracy.** MNO (orange) forecasts the mean (solid) of the ground-truth DNS (blue) more accurately in comparison to ML-based parametrizations (green) and climatology (red). The standard deviations is plotted as dotted lines.

A.5 APPENDIX TO ILLUSTRATION OF MNO VIA MULTISCALE LORENZ96

The other large-scale (LS) and fine-scale (FS) terms are

$$\begin{aligned}
 \text{filtered FS dynamics, } \overline{\mathcal{N}(u)}(x) &= \begin{cases} \frac{\partial X_k}{\partial t} & \text{if } x = k(J+1) \forall k \in \{0, \dots, K\} \\ 0 & \text{otherwise} \end{cases} \\
 \text{LS dynamics, } \mathcal{N}(\bar{u})(x) &= \begin{cases} \frac{\partial \bar{X}_k}{\partial t} & \text{if } x = k(J+1) \forall k \in \{0, \dots, K\} \\ 0 & \text{otherwise} \end{cases} \\
 \text{with abbreviation, } \frac{\partial \bar{X}_k}{\partial t} &:= X_{k-1}(X_{k+1} - X_{k-2}) - X_k + F \\
 \text{LS state, } \bar{u}(x) = \mathcal{G} * u(x) &= [X_0, 0, \dots, 0, X_1, 0, \dots, X_K]
 \end{aligned} \tag{16}$$

A.6 APPENDIX TO RESULTS

A.6.1 ACCURACY

Figure 8 shows that the predicted mean and standard deviation of MNO (orange) closely follows the ground-truth (blue). The ML-based parametrization (green) follows the ground-truth only for a few time steps (until $\sim t = 0.125$). The climatology (red) depicts the average prediction in the training dataset.

A.6.2 MODEL CONFIGURATION

Multiscale Lorenz96: MNO As hyperparameters we chose the number of channels, $n_v = 64$, number of retained modes, $k_{\max} = 3$, number of Fourier layers, $n_d = 3$, and no batch norm layer. The time-series modeling task uses a history of only one time step to learn chaotic dynamics (Li et al., 2021b). We are using ADAM optimizer with learning rate, $\lambda = 0.001$, step size, 20, number of epochs, $n_e = 2$, and an exponential learning rate scheduler with gamma, $\gamma = 0.9$ (Kingma & Ba, 2015). Training took 1 : 50min on a single core Intel i7-7500U CPU@2.70GHz.

Multiscale Lorenz96: ML-based parametrization The ML-based parametrizations uses a ResNet with $n_{\text{layers}} = 2$ residual layers that contain a fully connected network with $n_{\text{units}} = 32$ units. The model is optimized with Adam (Kingma & Ba, 2015) with learning rate 0.01, $\beta = (0.9, 0.999)$, $\epsilon = 1 * 10^{-8}$, trained for $20n_{\text{epochs}} = 20$.

Multiscale Lorenz96: Traditional parametrization The traditional parametrization uses least-squares to find the best linear fit. The weight matrix is computed with $A = (X^T X)^{-1} X^T Y$, where X and Y are the concatenation of input large-scale features and target parametrizations, respectively. Inference is conducted with $\hat{y} = Ax$.

A.7 NEURAL NETWORKS VS. NEURAL OPERATORS

Most work in physics-informed machine learning relies on fully-connected neural networks (FCNNs) or convolutional neural networks (Karniadakis et al., 2021). FCNNs however are mappings between finite-dimensional spaces and learn mappings for single equation instances rather than learning the PDE solver. In our case FCNNs only learn mappings on fixed spatial grids. We leverage the recently formulated neural operators to extend the formulation to arbitrary grids. The key distinction is that the FCNN learns a parameter-dependent set of weights, Φ_{a_y} , that has to be re-trained for every new parameter setting. The neural operator is a learned function mapping with parameter-independent weights, Θ , that takes parameter settings as input and returns a function over the spatial domain, $G_{\Theta}(a_y)$. In comparison, the forcing term is approximated by an FCNN as $\hat{f}_{x,\Phi}(x_k; a_y) = g_{\Phi_{a_y}}(x_k)$ and by a neural operator as $\hat{f}_{x,\Theta}(x_k; a_y) = G_{\Theta}(a_y)(x_k)$. The mappings are given by:

$$\begin{aligned} \text{FCNN: } g_{\Phi_{a_y}} : D_x &\rightarrow \mathbb{R}^{d_x}, \\ \text{NO: } G_{\Theta} : H_{a_y}(D_x; \mathbb{R}^{d_{a_y}}) &\rightarrow H_X(D_x; \mathbb{R}^{d_x}). \end{aligned} \tag{17}$$

H_{a_y} is a function space (Banach) of PDE parameter functions, a_y , that map the spatial domain, D_y , onto d_{a_y} dimensional parameters, such as ICs, BCs, parameters, or forcing terms. H_X is the function space of residuals that map the spatial domain, D_x , onto the space of d_X -dimensional residuals, \mathbb{R}^{d_x} .