

Vandermonde Neural Operators

L. Lingsch and M. Michelis and S. Perera and R. Katzschmann and

S. Mishra

Research Report No. 2023-24
June 2023

Seminar für Angewandte Mathematik
Eidgenössische Technische Hochschule
CH-8092 Zürich
Switzerland

VANDERMONDE NEURAL OPERATORS

LEVI LINGSCH, MIKE MICHELIS, SIRANI M. PERERA, ROBERT K. KATZSCHMANN,
AND SIDDARTHA MISHRA

ABSTRACT. Fourier Neural Operators (FNOs) have emerged as very popular machine learning architectures for learning operators, particularly those arising in PDEs. However, as FNOs rely on the fast Fourier transform for computational efficiency, the architecture can be limited to input data on equispaced Cartesian grids. Here, we generalize FNOs to handle input data on non-equispaced point distributions. Our proposed model, termed as Vandermonde Neural Operator (VNO), utilizes Vandermonde-structured matrices to efficiently compute forward and inverse Fourier transforms, even on arbitrarily distributed points. We present numerical experiments to demonstrate that VNOs can be significantly faster than FNOs, while retaining comparable accuracy, and improve upon accuracy of comparable non-equispaced methods such as the Geo-FNO.

1. INTRODUCTION

Partial Differential Equations (PDEs) are extensively used to mathematically model interesting phenomena in science and engineering [1]. As explicit solution formulas for PDEs are not available, traditional numerical methods such as finite difference, finite element, and spectral methods [2] are extensively used to simulate PDEs. Despite their tremendous success, the prohibitively high computational cost of these methods makes them infeasible for a variety of contexts in PDEs ranging from high-dimensional problems to the so-called *many query* scenarios [3]. This high computational cost also provides the rationale for the development of alternative *data driven* methods for the fast and accurate simulation of PDEs. Hence, a wide variety of machine learning algorithms have been proposed recently in this context. These include physics-informed neural networks (PINNs) [4], MLPs, and CNNs for simulating parametric PDEs [5, 6, 7, 8] as well as graph based algorithms [9, 10, 11, 12], to name a few.

However, as solutions of PDEs are expressed in terms of the so-called *solution operators*, which map input functions (initial and boundary data, coefficients, source terms) to the PDE solution, *Operator learning*, i.e., learning the underlying operators from data, has emerged as a dominant framework for applying machine learning to PDEs. Existing operator learning algorithms include, but are not limited to, operator networks [13], DeepOnets [14, 15, 16], attention based methods such as [17, 18, 19], and neural operators [20, 21, 22, 23].

Within this large class of operator learning algorithms, Fourier Neural Operators (FNO) [24] have gained much traction and are widely applied [25, 26]. Apart from favorable theoretical approximation properties [27, 28], FNOs are attractive due to their expressivity, simplicity and computational efficiency. A key element in the computational efficiency of FNO lies in the fact that its underlying convolution operation is efficiently carried out in Fourier space with the *fast Fourier transform*

(FFT) algorithm. It is well-known that FFT is only (log-)linear in computational complexity with respect to the number of points at which the underlying input functions are sampled. However, this computational efficiency comes at a cost as the recursive structure of FFT limits its applications to inputs sampled on the so-called *Regular* or *equispaced Cartesian (Rectangular) grids*, see Figure 2 left for an illustration. This is a major limitation in practice. In real-world applications, where information on the input and output signals is measured by sensors, it is not always possible to place sensors only on an equispaced grid. Similarly, when data is obtained through numerical simulations, often it is essential to discretize PDEs on irregular grids, such as those adapted to be refined to capture relevant spatial features of the underlying PDE solution or on unstructured grids that fit the complex geometry of the underlying domain. See Figure 2 for examples of non-equispaced distributions of sample points.

Several methods have recently been proposed in the literature to address this limitation of FNOs and modify/enhance it to handle data on non-equispaced points. For instance, geometry-aware FNO (Geo-FNO) [29] appends a neural network to the FNO to learn a deformation from the physical space to a regular grid. Then, the standard FFT can be applied to the latent space of equispaced grid points. This learned diffeomorphism corresponds to an adaptive moving mesh [30]. Factorized-FNO (F-FNO) builds upon the Geo-FNO, introducing an additional bias term in the Fourier layer and performing the Fourier transform over each dimension separately [31]. The non-equispaced Fourier PDE solver (NFS) uses a vision mixer [32] to interpolate from a non-equispaced signal onto a regular grid, again applying the standard FNO subsequently [33]. All these methods share the same design principle, i.e., given inputs on non-equispaced points, *interpolate* or transform this data into a regular grid and then apply FNO. This observation leads to a natural question: *why not consider the complimentary approach and modify the Fourier Neural Operator itself to enable it to be directly applied to input data on non-equispaced points?*

Addressing this question is the main goal of this paper where we propose a *novel modification* of FNO to enable its application on non-equispaced input data. More concretely,

- We propose a new operator learning algorithm that we term as *Vandermonde Neural Operator* or VNO which extends FNO to be applied to input data on non-equispaced points. To do this, we design an algorithm to efficiently compute discrete (inverse) Fourier transforms via Vandermonde structured matrices.
- We present a novel yet simple construction of Vandermonde-structured matrices to compute the forward and backward (inverse) Fourier transformations within the VNO algorithm that allows it to handle inputs, sampled on arbitrary non-equispaced point distributions.
- We present a suite of numerical experiments to demonstrate that VNO can train significantly faster than FNO for input data on non-equispaced points, particularly on lattices (tensor products of arbitrarily non-equispaced points in one-dimension, see Figure 2), while either retaining or improving on the test accuracy.

Thus, we present a novel algorithm to efficiently and accurately learn operators, such as those arising in PDEs, with input functions being possibly sampled on arbitrary non-equispaced points. Consequently VNO expands the FNO architecture

to learn operators on real-world problems where sensors may not be equispaced such as robotics, atmospheric sciences, and aerodynamics.

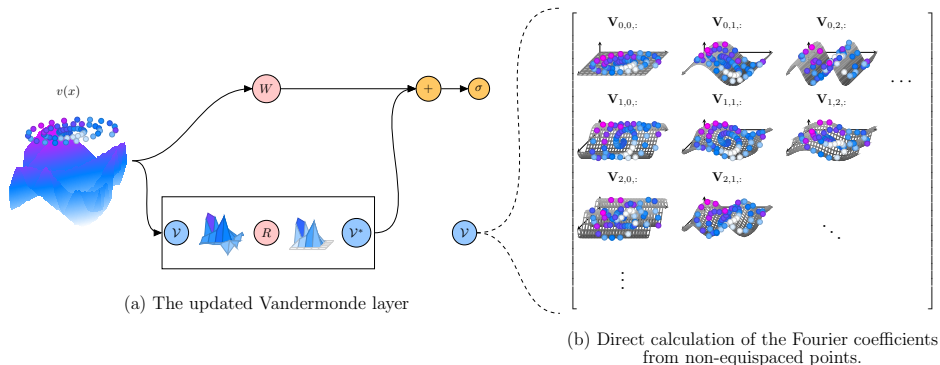


FIGURE 1. The proposed VNO operates directly on non-equispaced sample points. The input function $v(x)$ may be sampled at arbitrary points to construct the Fourier representation directly via the Vandermonde Matrix \mathbf{V} . Each element \mathbf{V}_{j_0, j_1} computes a convolution with the data for a given pair of harmonics j_0, j_1 over all points. Each element of $v(x)$ is multiplied by the magnitude of the harmonic pair at that location, and these values are summed, yielding the signal in Fourier space. W represents the bias term, and R is a matrix of parameters which modify the Fourier coefficients, implemented as in the standard FNO.

2. METHODS

Our goal in this section is to present VNO as a new neural operator in which the FFT algorithm within the Fourier layer in a FNO [24] is replaced with a Vandermonde structured matrix. To this end, we start with a short description of this matrix construction below.

2.1. On Vandermonde Structured Matrices. A Vandermonde structured matrix can be defined via nodes $x_j \in \mathbb{R}$ or $\mathbb{C}, j \in \{0, 1, \dots, m\}$, in a geometric progression along each column (or row), defined by

$$(1) \quad \mathbf{V}_{j,k} = [x_j^k]_{j,k=0}^{m,n}.$$

In the case that the nodes represent the primitive n^{th} roots of unity, this geometric progression of powers along the columns (or rows) yields the discrete Fourier transform (DFT) matrix. This matrix is symmetric, unitary, and periodic, hence it can be factorized as a product of sparse matrices. The resulting factorization can be used to obtain a radix-2 algorithm that can efficiently compute the 1D DFT and its inverse with the FFT algorithm of $O(n \log n)$ complexity [34, 35, 36].

To motivate an alternative computational realization, we recall that the Fourier transform is an integral operator. Approximating it with a quadrature rule for discretization, we multiply each point value of the underlying function by the sinusoidal basis function for a given mode, and sum these terms to compute the




FIGURE 2. Distributions discussed in this paper. The FNO is restricted to the regular grid. The VNO may be applied to the lattice distribution via (3), or the random distribution via (6)

magnitude of this sinusoidal term. This process is then repeated for each mode. In contrast to the recursive, butterfly FFT algorithm, this interpretation can be generalized to the Vandermonde-structured matrices of the form in (1), which in turn, can be efficiently implemented through *batched matrix multiplications*.

Forward Transformations. The forward transformation computes the Fourier representation of a given function by using Vandermonde structured matrices. In one-dimension, the corresponding Vandermonde-structured matrix is

$$(2) \quad \mathbf{V}_{j,k} = \frac{1}{\sqrt{n}} \left[e^{-i(jp_k)} \right]_{j,k=0}^{m-1, n-1}.$$

Here, $\mathbf{p} = [p_0, p_1, \dots, p_{n-1}]^T$ is the vector of the positions of data points at which the underlying function is sampled. n is the number of data points, m the number of modes, and $i = \sqrt{-1}$. As in the FFT, note that it is necessary to normalize the positions of data points to a range between 0 and 2π , as the Fourier basis functions are assumed to be periodic on the unit circle.

The 2-dimensional Fourier transform is equivalent to one-dimensional Fourier transforms along each axis. Therefore, the transformation on any 2-dimensional lattice, i.e. the tensor product of one-dimensional point distributions along each axis (see Figure 2), can be performed by constructing two Vandermonde matrices, \mathbf{V}_1 and \mathbf{V}_2 , corresponding to the positions of data points along each axis. Given $X \in \mathbb{R}^{n \times n}$ or $\mathbb{C}^{n \times n}$ as the data matrix containing values of the underlying function, sampled on the non-equispaced lattice, the Vandermonde-structured matrix can be used as a basis transformation of this data to the Fourier space $\mathcal{X} \in \mathbb{C}^{m \times m}$ given via

$$(3) \quad \mathcal{X} = \mathbf{V}_1 X \mathbf{V}_2^T.$$

The transforms for non-equispaced one-dimensional and non-equispaced two-dimensional rectangular lattices are already present in the literature [37], yet generalizations to other distributions are not present. This related work is discussed in greater detail in Section 4.

Extension to arbitrary non-equispaced point distributions. Next, we would like to extend this construction beyond lattices to arbitrarily non-equispaced point distributions in two dimensions. To do this, we store again the positions of the sampling points as $P = [\mathbf{p}_0, \mathbf{p}_1] \in \mathbb{R}^{n \times 2}$ and normalize the points' values to a range between $[0, 2\pi]$. The corresponding Vandermonde matrix is extended into a 3^{rd} dimension as a tensor, i.e., it will have three indices. The first two indices of $\mathbf{V} \in \mathbb{C}^{m \times m \times n}$

correspond to the sinusoidal components along the first and second dimensions, respectively, while the third index corresponds to a specific point within the distribution. This results in the following *tensor*,

$$(4) \quad \mathbf{V}_{j_0, j_1, k} = \sqrt{\frac{2}{n}} \left[e^{-i(j_0 P_{k,0} + j_1 P_{k,1})} \right]_{j_0, j_1, k=0}^{m-1, m-1, n-1},$$

Given $X \in \mathbb{R}^{n \times n}$ or $\mathbb{C}^{n \times n}$ as the data matrix containing values of the underlying function, sampled on any arbitrary distribution of points, the Vandermonde-structured matrix can be used as a basis transformation of this data to the Fourier space $\mathcal{X} \in \mathbb{C}^{n \times n}$ given via,

$$(5) \quad \mathcal{X} = \mathbf{V}X,$$

with matrix \mathbf{V} below,

$$(6) \quad \mathbf{V} = \sqrt{\frac{2}{n}} \left[\begin{array}{c} e^{-i(0\mathbf{p}_0^T + 0\mathbf{p}_1^T)} \\ e^{-i(1\mathbf{p}_0^T + 0\mathbf{p}_1^T)} \\ \vdots \\ e^{-i((m-1)\mathbf{p}_0^T + 0\mathbf{p}_1^T)} \end{array} \right] \left[\begin{array}{c} e^{-i(0\mathbf{p}_0^T + 1\mathbf{p}_1^T)} \\ e^{-i(1\mathbf{p}_0^T + 1\mathbf{p}_1^T)} \\ \vdots \\ e^{-i((m-1)\mathbf{p}_0^T + 1\mathbf{p}_1^T)} \end{array} \right] \cdots \left[\begin{array}{c} e^{-i(0\mathbf{p}_0^T + (m-1)\mathbf{p}_1^T)} \\ e^{-i(1\mathbf{p}_0^T + (m-1)\mathbf{p}_1^T)} \\ \vdots \\ e^{-i((m-1)\mathbf{p}_0^T + (m-1)\mathbf{p}_1^T)} \end{array} \right]$$

This construction can be readily generalized to a fully non-equispaced domain of $N \in \mathbb{N}$ dimensions by the construction of an tensor with $N + 1$ indices, by equation 7. Here, $P = [\mathbf{p}_0, \mathbf{p}_1, \dots, \mathbf{p}_{N-1}] \in \mathbb{R}^{n \times N}$ is a matrix whose columns are vectors of the positions of the data along each dimension.

$$(7) \quad \mathbf{V}_{j_0, \dots, j_{N-1}, k} = \sqrt{\frac{N}{n}} \left[e^{-i \left(\sum_{l=0}^{N-1} j_l P_{k,l} \right)} \right]_{j_0, \dots, j_{N-1}=0, k=0}^{m-1, \dots, m-1, n-1}$$

Backward Transformations. The backward transformation computes the spatial representation of the function given its Fourier coefficients. This is realized in a straightforward manner via the adjoint of the Vandermonde-structured matrix in (6). We define the adjoint as the conjugate-transpose of the Vandermonde-structured matrix, denoted $\mathcal{V}_{j,k}^*$, where $\mathcal{V}_{j,k}$ is the stacked version of (6) as defined in **Supplementary Material (SM)** (15).

2.2. The Vandermonde Neural Operator. Next, we propose a new neural operator for extending FNO to handle inputs on arbitrary point distributions. For clarity, we maintain consistency with the notation used by Li et al. [24] when presenting this procedure.

Our neural operator is expressed as an iterative map $v_t \mapsto v_{t+1} \forall t \in \{0, 1, \dots, T-1\}$. The subsequent solution, v_{t+1} is expressed as,

$$(8) \quad v_{t+1}(x) = \sigma(Wv_t(x) + (\mathcal{K}(\phi)v_t)(x)),$$

with nonlinear activation function $\sigma : \mathbb{R} \mapsto \mathbb{R}$ applied elementwise, residual connection $W : \mathbb{R}^{d_v} \mapsto \mathbb{R}^{d_v}$, and a bounded linear *kernel* operator.

We recall that for FNO [24], this *kernel operator* realizes a convolution in Fourier space with kernel,

$$(9) \quad (\mathcal{K}(\phi)v_t)(x) = \mathcal{F}^{-1}(R_\phi \cdot \mathcal{F}(v_t))(x), \quad \forall x \in D$$

Here, $\mathcal{F}, \mathcal{F}^{-1}$ are the Fourier and Inverse Fourier transforms, respectively, $D \subset \mathbb{R}^d$, and $R_\phi \in \mathbb{C}^{d_v \times d_v}$ is a matrix representing the Fourier transformation of a learned periodic kernel function parameterized by ϕ . In the discrete case, these Fourier and inverse Fourier transforms are performed with the FFT. However, this choice imposes the assumption that the input data is sampled on a regular equispaced grid.

To deal with input data sampled at points with non-equispaced distributions, we replace the *Fourier layer* (9) with the following *Vandermonde layer*,

$$(10) \quad (\mathcal{K}(\phi)v_t)(x) = \mathcal{V}^*(R_\phi \cdot \mathcal{V}(v_t))(x), \quad \forall x \in D.$$

Here, \mathcal{V} denotes a transformation by the Vandermonde structured matrix, defined by

$$(11) \quad (\mathcal{V}(f))_\xi(\mathbf{k}) = \sum_{j=0}^{n-1} f_\xi(\mathbf{x}_j) e^{-2\pi i \langle \mathbf{x}_j, \mathbf{k} \rangle}, \quad (\mathcal{V}^*(f))_\xi(\mathbf{x}) = \sum_{\mathbf{k} \in [\mathbb{Z}_m]^d} f_\xi(\mathbf{k}) e^{2\pi i \langle \mathbf{x}, \mathbf{k} \rangle},$$

where $i = \sqrt{-1}$, $\xi = 0, \dots, d_v - 1$, $f : D \mapsto \mathbb{R}^{d_v}$, $\mathbf{k} = (k_0, \dots, k_{d-1}) \in [\mathbb{Z}_m]^d$, $\mathbf{x} = (x_0, \dots, x_{d-1}) \in D$, where D has been discretized by $\mathbf{x}_0, \dots, \mathbf{x}_{n-1}$, and $\langle \mathbf{x}, \mathbf{k} \rangle = x_0 k_0 + x_1 k_1 + \dots + x_{d-1} k_{d-1}$. This is tantamount to a multiplication by the matrix in (4) for two dimensional data on a lattice, and (7) for the general case. The resulting neural operator of the form (8) with kernel given by (10), is termed as the *Vandermonde Neural operator* or VNO. Figure 1 provides an illustration of this architecture.

This generalization allows VNO to be applied beyond the FNO's restriction to equispaced grids. Furthermore, as the structure of the Vandermonde matrix makes no assumptions about the point distributions, it is possible to learn the solution operator to PDEs from random distributions of sampling points, even if each realization (input sample) has a different distribution. Additionally, as no additional transformations are to be learned like in [29], the training complexity remains the same as in FNO. Instead, the VNO can be applied directly to a non-equispaced distribution to compute the most accurate Fourier coefficients from the provided data.

2.3. Computational Complexity of VNO. The most notable feature of FFT is its computational efficiency. Calculating the Fourier coefficients of a 1-dimensional signal, sampled at n points, by using the brute force DFT, costs $O(n^2)$. In contrast, the FFT algorithm computes these coefficients with $O(n \log n)$ complexity.

Hence, it is natural to wonder why one should reconsider matrix multiplication techniques in our setting. In this context, we observe that the maximum performance gain with FFT occurs when the FFT computes all the Fourier coefficients, or modes, of an underlying signal. Furthermore, peak efficiency is reached for points on a dyadic interval. While the number of modes to compute may be truncated, the interconnected nature of the self-recursive radix-2 FFT algorithm makes it difficult in practice to attain peak efficiency. We refer the reader to **SM** Figure 4 for a visual representation. Therefore, reported performance gains by new FFT algorithms are often optimistic. Thus, in the case of truncated modes, matrix multiplications techniques should not be ruled out.

Moreover, for neural operators such as FNO and VNO, only a small subset of nonzero modes are required to approximate the operator [24]. This implies

that for a one-dimensional problem, the Vandermonde matrix has a fixed number of rows, while the number of columns grows with the problem size. Therefore, the computational complexity of the proposed transformations by a Vandermonde matrix cost $O(n)$ as the Vandermonde-structure can be fully determined using $O(n)$ as opposed to $O(n^2)$ [38, 39, 40], and hence the number of points is independent of the number of modes.

3. EXPERIMENTAL RESULTS

In this section, our aim is to investigate the performance of the proposed VNO architecture on a challenging suite of diverse PDE learning tasks.

Implementation, Training Details and Baselines. A key contribution of this paper is a new implementation of the Vandermonde structured matrix multiplications in *PyTorch*, which enables us to efficiently compute Fourier and Inverse Fourier transforms. Within a neural network, an efficient $O(n)$ algorithm must also be parallelizable to handle batches, as this massively speeds up the training process. Batches of data with the same or different point distributions are easily handled by the *torch.matmul()* and *torch.bmm()* functions, respectively.

In all experiments, we use the same hyperparameters for training purposes; ADAM optimizer with a learning rate of 0.005, scheduler step 10, gamma decay of 0.97, and trained for 500 epochs. We also use the L1-loss function, which produced both a lower L1-error and L2-error than the L2-loss. The test error was measured as the relative L1 error.

As baselines, we use FNO in all the experiments and its variant, Geo-FNO [29] in experiments where the underlying domain has a complicated, non-equispaced geometry. For comparisons with the FNO or Geo-FNO, we choose the number of modes in each layer as well as the width of each layer to be the same across all architectures in an experiment. Given that the Fourier layer is fundamentally the same for VNO, FNO, and Geo-FNO, albeit using different methods to compute the transform, the relative performance differences between methods are consistent as the number of modes and width are varied. We see this as a fair comparison, as all model parameters and the model size remains consistent within an experiment. All experiments are performed on the Nvidia GeForce RTX 3090 with 24GB memory. **Benchmark 1: Burgers' Equation.** The one-dimensional viscous Burgers' equation is a widely considered model problem for fluid flow given by

$$(12) \quad \begin{aligned} \partial_t u(x, t) + \partial_x \left(\frac{1}{2} u^2(x, t) \right) &= \nu \partial_{xx} u(x, t) & x \in (0, 1) \quad t \in (0, 1] \\ u(x, 0) &= u_0(x) & x \in (0, 1) \end{aligned}$$

where u denotes the fluid velocity and ν the viscosity. We follow [24] in fixing $\nu = 0.1$ and considering the operator that maps the initial data u_0 to the solution $u(\cdot, T)$ at final time $T = 1$. The training and test data, presented in [24] for this problem, is used. Points are sub-sampled to create non-equispaced distributions.

We test the VNO and FNO on this benchmark on three different point distributions, shown in **SM** Figure 5, namely equispaced point distribution, contracting-expanding distribution, and fully non-equispaced randomly chosen set of points. The training time (per epoch) and test errors are shown in Table 1. We observe from this table that for equispaced points, both FNO and VNO have low test errors, with FNO being marginally more accurate than VNO. On the other hand, VNO

TABLE 1. Performance results for the experiments on shear flow, surface-level specific humidity, and flow past airfoils. The FNO is applied to a dense, equispaced, rectangular grid for the first three problems, while the VNO is applied to a lattice engineered to balance resolution and efficiency.

Model	model size	training time	testing error
Burgers' Equation			
<i>Equispaced Distribution:</i>			
VNO	549569	0.86s	0.095%
FNO	549569	0.97s	0.071%
<i>Contracting-Expanding Distribution:</i>			
VNO	549569	0.35s	0.86%
FNO–interpolation	549569	1.24s	1.00%
<i>Random Distribution:</i>			
VNO	549569	0.25s	1.67%
FNO–interpolation	549569	1.41s	1.11%
Shear Layer			
VNO	6571010	44s	5.89%
FNO	6571010	189s	6.16%
Surface-level Specific Humidity			
VNO	16786657	3.6s	4.37%
FNO	16786657	38s	5.25%
Flow past Airfoil			
VNO	2368225	7.38s	0.49%
Geo-FNO	3020963	7.42s	1.14%

has slightly lower training time per epoch showing that batched matrix multiplications, which form the basis of VNO, are as computationally efficient as the FFT algorithm that underpins the FNO.

However, FNO cannot be directly applied to the two non-equispaced point distributions that we consider here (contracting-expanding and random). Hence, we have to interpolate the input point values of the underlying function to an equispaced grid and subsequently apply FNO. To this end, we use a cubic-spline interpolation procedure and report the test results of this model in Table 1.

We observe that for the contracting-expanding distribution, VNO is more accurate than FNO–interpolation while being faster in training by a factor of almost 4. For the random point distribution, FNO–interpolation is clearly more accurate but has a significantly higher training time (almost a factor of 6) vis a vis VNO.

Thus, this experiment already indicates that not only is VNO able to handle input data on arbitrary non-equispaced grids, its accuracy is still comparable to FNO. When FNO is augmented with interpolation procedures to deal with non-equispaced sample points, VNO is considerably faster to train while being comparable in accuracy.

Benchmark 2: Shear Layer. We follow a recent work on convolutional neural operators [23] in considering the incompressible Navier-Stokes equations

$$(13) \quad \frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} + \nabla p = \nu \Delta \mathbf{u}, \quad \nabla \cdot \mathbf{u} = 0.$$

Here, $\mathbf{u} \in \mathbb{R}^2$ is the fluid velocity and p is the pressure. The underlying domain is the unit square with periodic boundary conditions and the viscosity $\nu = 4 \times 10^{-4}$, only applied to high-enough Fourier modes (those with amplitude ≥ 12) to model fluid flow at *high Reynolds-number*. The solution operator maps the initial conditions $\mathbf{u}(t = 0)$ to the solution at final time $T = 1$.

We consider initial conditions representing the well-known *thin shear layer* problem [41, 42] (See [23] for details), where the shear layer evolves via vortex shedding to a complex distribution of vortices (see Figure 3a for an example of the flow). The training and test samples are generated, with a spectral viscosity method [42] of a fine resolution of 1024^2 points, from an initial sinusoidal perturbation of the shear layer [42], with layer thickness of 0.1 and 10 perturbation modes, the amplitude of each sampled uniformly from $[-1, 1]$ as suggested in [23].

As seen from Figure 3a, the flow shows interesting behavior with sharp gradients in two mixing regions, which are in the vicinity of the initial interfaces. On the other hand, the flow is nearly constant further away from this mixing region. Hence, we will consider VNO with input functions being sampled on a lattice shown in **SM** Figure 6 which is adapted to resolve large gradients features of the flow. On the other hand, FNO is tested on the equispaced point distribution. From Table 1, we observe that VNO is marginally more accurate than FNO while being 4 times faster per training epoch, demonstrating a significant computational advantage over FNO on this benchmark.

Benchmark 3: Surface-Level Specific Humidity. Next, we focus on a *real world* data set and learning task where the objective is to predict the surface-level specific humidity over South America at a later time (6 hours into the future), given inputs such as wind speeds, precipitation, evaporation, and heat exchange at a particular time. The exact list of inputs is given in **SM** Table 2. The physics of this problem are intriguingly complex, necessitating a *data-driven approach* to learn the underlying operator. To this end, we use data provided by the Modern-Era Retrospective analysis for Research and Application v2 (MERRA-2) satellite data to forecast the surface-level specific humidity [43]. Moreover, we are interested in a more accurate regional prediction, namely over the Amazon rainforest. Hence, for the VNO model, we will sample data on points on a lattice that is more dense over this rainforest, while being sparse (with smooth transitions) over the rest of the globe, see **SM** Figure 7 for visualization of this lattice.


In contrast to VNO, FNO samples input on the equispaced global grid. Test error is calculated over the region, shown in Figure 3b. The results, presented in Table 1, show that VNO is not only more accurate than FNO, but also one order of magnitude faster to train. The greater accuracy of VNO over FNO, on a regional scale, is also clearly observed in Figure 3b, where we observe that VNO is able to capture elements such as the formation of vortices visible in the lower right hand corner and the mixing of airstreams over the Pacific Ocean. These elements are smoothed out by the FNO.

Benchmark 4: Flow Past Airfoil. We also investigate transonic flow over an airfoil, as governed by the compressible Euler equations,


$$(14) \quad \frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{u}) = 0 \quad \frac{\partial \rho \mathbf{u}}{\partial t} + \nabla \cdot (\rho \mathbf{u} \otimes \mathbf{u} + p \mathbb{I}) = 0 \quad \frac{\partial E}{\partial t} + \nabla \cdot ((E + p) \mathbf{u}) = 0.$$

Here, ρ is the fluid density, t time, \mathbf{u} the velocity vector, p the pressure, and E the total energy, related to each other by an ideal gas equation of state. The data for this experiment has been taken from [29], where the authors have chosen farfield conditions $\rho_\infty = 1$ [kg/m³], $p_\infty = 1$ [atm], Mach number $M_\infty = 0.8$, and angle of attack $AoA = 0$.


The underlying operator maps the airfoil shape to the pressure field. In this case, the underlying distribution of sample points changes between each input (airfoil shape). VNO can readily handle this situation. As a baseline, we use Geo-FNO as proposed in [29]. To have a fair comparison with VNO, we allow Geo-FNO to learn the diffeomorphism online just as VNO does. The test errors, presented in Table 1 show that VNO is significantly more accurate than Geo-FNO for this problem, while being comparable in training time. In particular, as observed from Figure 3c, VNO captures the trailing shock much better than Geo-FNO.



(A) Horizontal velocity for the Shear Layer experiment



(B) Surface-level specific humidity over South America.



(C) Pressure distributions around an airfoil.

FIGURE 3. These figures display examples of the ground truth, the target which the VNO, FNO, or Geo-FNO attempt to match. Left: Ground Truth. Center: VNO Right: FNO for (a) and (b) and Geo-FNO for (c).

4. DISCUSSION

Summary. FNO has emerged as a widely used architecture for operator learning, particularly in the context of PDEs. However, as the FNO relies on the FFT to efficiently carry out convolutions in Fourier space, its application is restricted to input data sampled on equispaced Cartesian grids. In practice, input data is often sampled on non-equispaced point distributions. Hence, extending FNO to handle such scenarios is of great interest. In contrast to the common strategy of interpolating or transforming data from non-equispaced to equispaced point distributions, we aimed to propose an alternative framework that directly replaced FFT within FNO to allow for input functions to be sampled on more general point distributions.

To this end, we leveraged the Vandermonde-structured matrices that arise quite naturally in computing the Fourier and Inverse Fourier transforms. By coming up with an efficient implementation of the resulting *batched matrix multiplications*, we propose a new operator learning framework called the Vandermonde Neural operator (VNO) with the ability to handle inputs on arbitrary non-equispaced point distributions. We compare VNO with FNO and recently proposed extensions of FNO such as Geo-FNO on a suite of experiments that range from the simple one-dimensional viscous Burgers’ equation to more complicated incompressible and compressible flow equations and also a realistic Earth science data set.

We find that the VNO can be significantly faster at training, while being comparable in accuracy or significantly more accurate than FNO (and its extensions) on the considered problems. Hence, we demonstrate that the VNO can serve as an efficient and accurate neural operator which can be broadly applied in various scenarios.

Related Work. We start with a succinct summary of the extensive literature on Vandermonde-structured matrices and related constructions. In this context, the delay Vandermonde matrix (DVM) is a superclass of the DFT matrix. The DVM structure is utilized to analyze TTD wideband multibeam beamforming while solving the longstanding beam squint problem [44, 45, 46, 47, 48]. Although the Vandermonde matrices can be ill-conditioned [49, 50, 51, 52, 53, 39, 54], the proposed VNO does not encounter ill-conditioning, as nodes are placed on the unit disk [54, 55, 48], and we do not compute the explicit inverse of the Vandermonde-structured matrix because it is too expensive and numerically less accurate [51]. Nonequispaced FFT (NFFT) has already been developed [56, 57, 58]. These algorithms rely on a mixture of interpolation, windowing techniques, and strategic applications of FFT to maintain $O(n \log n)$ cost. However, the inverse NFFT cannot be calculated in a comparably direct manner [59, 60]. The nonequispaced discrete Fourier transform is also presented as a summation [56, 57, 58], which can be arranged into a Vandermonde-structured matrix, but this is not presented in the literature. Excluding fast transforms, the structure of the Vandermonde matrix has been employed to perform the nonequispaced, or nonuniform discrete Fourier transform (NUDFT) [37] in one and two dimensions; however, the two-dimensional distributions, in this case, are limited to the lattice and nonuniform parallel lines. The methods we propose in this paper extend the use of Vandermonde-structured matrices to irregular point distributions in two or more dimensions. The NUDFT is rarely used, as many applications of Fourier transforms require all Fourier coefficients, resulting in $O(n^2)$ cost [37]. This is not the case for the FNO, and thus

the proposed VNO avoids the rapidly growing computational costs associated with NUFFT.

Limitations and Future Work. The elements of the Vandermonde-structured matrix are directly related to the positions of the data points for a given problem. If all samples are using an identical point distribution, the Vandermonde-structured matrix may be constructed once, at the beginning of the run time. However, this cannot be done if point distributions vary among samples. In this case, we must either precompute Vandermonde-structured matrices corresponding to each point distribution and load them with the corresponding data, or we must construct the matrices at run-time. Precomputing the matrices can offer performance advantages, but for problems with many data points, a large number of samples, or a large number of modes, the size of the Vandermonde-structured matrices in memory can grow quite large, even exceeding the size of the original data set. Furthermore, the number of rows grows as the power of the number of spatial dimensions even though the operation complexity. This may limit this method for data in 3 or more spatial dimensions—even though the operation complexity is approximately equivalent to the FFT. Performance gains from precomputing the matrices are diminished as such large matrices can not be loaded from memory as quickly. Constructing the matrices at run time, *i.e.*, during training, also hinders performance. In the future, it would be worth investigating how the run-time matrix construction might be sped up by using compile time languages [61]. Another direction to investigate is the use of different basis functions. The Fourier basis functions assume periodicity along the torus, but it is possible to extend and modify the Vandermonde-structured approach to handle spherical basis functions as well, expanding the use of VNO techniques to new fields [62, 63, 64, 65].

REFERENCES

- [1] Lawrence C Evans. *Partial differential equations*, volume 19. American Mathematical Soc., 2010.
- [2] A. Quarteroni and A. Valli. *Numerical approximation of Partial differential equations*, volume 23. Springer, 1994.
- [3] George Em Karniadakis, Ioannis G. Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. Physics informed machine learning. *Nature Reviews Physics*, pages 1–19, may 2021.
- [4] 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.
- [5] Y. Zhu and N. Zabaras. Bayesian deep convolutional encoder–decoder networks for surrogate modeling and uncertainty quantification. *Journal of Computational Physics*, 336:415–447, 2018.
- [6] Kjetil O. Lye, Siddhartha Mishra, and Deep Ray. Deep learning observables in computational fluid dynamics. *Journal of Computational Physics*, 410:109339, 2020.
- [7] Kjetil O. Lye, Siddhartha Mishra, Deep Ray, and Praveen Chandrashekar. Iterative surrogate model optimization (ISMO): An active learning algorithm for PDE constrained optimization with deep neural networks. *Computer Methods in Applied Mechanics and Engineering*, 374:113575, 2021.
- [8] Nils Wandel, Michael Weinmann, and Reinhard Klein. Unsupervised deep learning of incompressible fluid dynamics. *CoRR*, abs/2006.08762, 2020.
- [9] Alvaro Sanchez-Gonzalez, Jonathan Godwin, Tobias Pfaff, Rex Ying, Jure Leskovec, and Peter W. Battaglia. Learning to simulate complex physics with graph networks. *CoRR*, abs/2002.09405, 2020.
- [10] Tobias Pfaff, Meire Fortunato, Alvaro Sanchez-Gonzalez, and Peter W. Battaglia. Learning mesh-based simulation with graph networks. *CoRR*, abs/2010.03409, 2020.
- [11] J. Brandstetter, D. E. Worrall, and M. Welling. Message Passing Neural PDE solvers. *arXiv preprint arXiv:2202.03376*, 2022.
- [12] L. Equer, T.K. Rusch, and S. Mishra. Multi-scale message passing neural pde solvers. *arXiv preprint arXiv:2302.03580*, 2023.
- [13] 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.
- [14] Lu Lu, Pengzhan Jin, and George Em Karniadakis. DeepONet: Learning nonlinear operators for identifying differential equations based on the universal approximation theorem of operators. *CoRR*, abs/1910.03193, 2019.
- [15] Z. Mao, L. Lu, O. Marxen, T. Zaki, and G. E. Karniadakis. DeepMandMnet for hypersonics: Predicting the coupled flow and finite-rate chemistry behind a normal shock using neural-network approximation of operators. Preprint, available from arXiv:2011.03349v1, 2020.
- [16] Shengze Cai, Zhicheng Wang, Lu Lu, Tamer A Zaki, and George Em Karniadakis. DeepM&Mnet: Inferring the electroconvection multiphysics fields based on operator approximation by neural networks. *Journal of Computational Physics*, 436:110296, 2021.
- [17] Georgios Kissas, Jacob H Seidman, Leonardo Ferreira Guilhoto, Victor M Preciado, George J Pappas, and Paris Perdikaris. Learning operators with coupled attention. *Journal of Machine Learning Research*, 23(215):1–63, 2022.
- [18] Shuhao Cao. Choose a transformer: Fourier or galerkin. In *35th conference on neural information processing systems*, 2021.
- [19] M. Prasthofer, T. De Ryck, and S. Mishra. Variable input deep operator networks. *arXiv preprint arXiv:2205.11404*, 2022.
- [20] N. Kovachki, Z. Li, B. Liu, K. Azizzadensheli, K. Bhattacharya, A. Stuart, and A. Anandkumar. Neural operator: Learning maps between function spaces. *arXiv preprint arXiv:2108.08481v3*, 2021.
- [21] Zongyi Li, Nikola B Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Andrew M Stuart, Kaushik Bhattacharya, and Anima Anandkumar. Multipole graph neural operator for parametric partial differential equations. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems (NeurIPS)*, volume 33, pages 6755–6766. Curran Associates, Inc., 2020.

- [22] Zongyi Li, Nikola B Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Andrew M Stuart, Kaushik Bhattacharya, and Anima Anandkumar. Multipole graph neural operator for parametric partial differential equations. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems (NeurIPS)*, volume 33, pages 6755–6766. Curran Associates, Inc., 2020.
- [23] Bogdan Raonić, Roberto Molinaro, Tobias Rohner, Siddhartha Mishra, and Emmanuel de Bezenac. Convolutional neural operators. *arXiv preprint arXiv:2302.01178*, 2023.
- [24] Zongyi Li, Nikola B. Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew M. Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. *CoRR*, abs/2010.08895, 2020.
- [25] Zongyi Li, Hongkai Zheng, Nikola Kovachki, David Jin, Haoxuan Chen, Burigede Liu, Kamyar Azizzadenesheli, and Anima Anandkumar. Physics-informed neural operator for learning partial differential equations. *arXiv preprint arXiv:2111.03794*, 2021.
- [26] J. Pathak, S. Subramanian, P. Harrington, S. Raja, A. Chattopadhyay, M. Mardani, T. Kurth, D. Hall, Z. Li, K. Azizzadenesheli, p. Hassanzadeh, K. Kashinath, and A. Anandkumar. Fourcastnet: A global data-driven high-resolution weather model using adaptive fourier neural operators. *arXiv preprint arXiv:2202.11214*, 2022.
- [27] Nikola Kovachki, Samuel Lanthaler, and Siddhartha Mishra. On universal approximation and error bounds for fourier neural operators. 07 2021.
- [28] S. Lanthaler, R. Molinaro, P. Hadorn, and S. Mishra. Nonlinear reconstruction for operator learning of pdes with discontinuities. In *International Conference on Learning Representations*, 2023.
- [29] Zongyi Li, Daniel Zhengyu Huang, Burigede Liu, and Anima Anandkumar. Fourier neural operator with learned deformations for pdes on general geometries, 2022.
- [30] Weizhang Huang and Robert D Russell. *Adaptive Moving Mesh Methods*, volume 174. Springer Science and Business Media, 2010.
- [31] Alasdair Tran, Alexander Mathews, Lexing Xie, and Cheng Soon Ong. Factorized fourier neural operators. In *The Eleventh International Conference on Learning Representations*, 2023.
- [32] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *CoRR*, abs/2010.11929, 2020.
- [33] Haitao Lin, Lirong Wu, Yufei Huang, Siyaun Li, Guojiang Zhao, and Stan Z Li. Non-equispaced fourier neural solvers for pdes, 2022.
- [34] J. W. Cooley and J. W. Tukey. An algorithm for the machine calculation of complex fourier series. *Math. Comp.*, 19:297–301, 1965.
- [35] R. Yavne. An economical method for calculating the discrete fourier transform. In *in Proc. AFIPS Fall Joint Computer Conf. 33*, pages 115–125, 1968.
- [36] S. G. Johnson and Frigo M. A modified split-radix fft with fewer arithmetic operations. *IEEE Trans.*, 55(1):111–119, 2007.
- [37] Sonali Bagchi and Sanjit K. Mitra. *The nonuniform discrete Fourier transform and its applications in signal processing*. Springer Science and Business Media, 1999.
- [38] I. Gohberg and V. Olshevsky. Complexity of multiplication with vectors for structured matrices. *Linear Algebra and its Applications*, 202:163–192, 1994.
- [39] V. Y. Pan. *Structured Matrices and Polynomials: Unified Superfast Algorithms*. Birkhauser/Springer, Boston/New York, 2001.
- [40] I. Gohberg and V. Olshevsky. Fast algorithms with preprocessing for matrix-vector multiplication problems. *Journal of Complexity*, 10(4):411–427, 1994.
- [41] J.B. Bell, P. Collela, and H. M. Glaz. A second-order projection method for the incompressible Navier-Stokes equations. *J. Comput. Phys.*, 85:257–283, 1989.
- [42] S. Lanthaler, S. Mishra, and C. Parés-Pulido. Statistical solutions of the incompressible euler equations. *Mathematical Models and Methods in Applied Sciences*, 31(02):223–292, Feb 2021.
- [43] NASA EarthData. Modern-era retrospective analysis for research and applications (merra-2) 2-dimensional data collection., 2021 - 2023.
- [44] S. M. Perera, L. Lingsch, A. Madanayake, S. Mandal, and N. Mastronardi. Fast dvm algorithm for wideband time-delay multi-beam beamformers. *the IEEE Transactions on Signal Processing*, 70(5913-5925), 2022.

- [45] S.M. Perera, V. Ariyaratna, N. Udayanga, A. Madanayake, G. Wu, L. Belostotski, R. Cintra, and T. Rappaport. Wideband n-beam arrays with low-complexity algorithms and mixed-signal integrated circuits. *IEEE Journal of Selected Topics in Signal Processing*, 12(2):368–382, 2018.
- [46] V. Ariyaratna, N. Udayanga, A. Madanayake, S. M. Perera, L. Belostotski, and R.J. Cintra. Design methodology of an analog 9-beam squint-free wideband if multi-beamformer for mmw applications. In: *Proceedings of IEEE 2017 Moratuwa Engineering Research Conference (MERCOn)*, pages 236–241, 2017.
- [47] S. M. Perera, A. Madanayake, and R. Cintra. Efficient and self-recursive delay vandermonde algorithm for multi-beam antenna arrays. *IEEE Open Journal of Signal Processing*, 1(1):64–76, 2020.
- [48] S. M. Perera, A. Madanayake, and R. Cintra. Radix-2 self-recursive algorithms for vandermonde-type matrices and true-time-delay multi-beam antenna arrays. *IEEE Access*, 8:25498–25508, 2020.
- [49] G. Walter. Norm estimates for inverses of vandermonde matrices. *Numerische Mathematik*, 23(4):337–347, 1975.
- [50] N. J. Higham. *Accuracy and Stability of Numerical Algorithms, Second edition*. SIAM, Philadelphia, USA, 2002.
- [51] N. J. Higham. *Accuracy and Stability of Numerical Algorithms*. SIAM, Philadelphia, USA, 1996.
- [52] N. J. Higham. Error analysis of the björck-pereyra algorithms for solving vandermonde systems. *Numer. Math.*, 50(5):613–632, 1987.
- [53] J. Demmel and P. Koev. The accurate and efficient solution of a totally positive generalized vandermonde linear system. *SIAM Journal on Matrix Analysis and Applications*, 27(1):142–152, 2005.
- [54] V. Y. Pan. How bad are vandermonde matrices? *SIAM Journal on Matrix Analysis and Applications*, 37(2):676–694, 2016.
- [55] S. M. Perera, A. Madanayake, A. Ogle, D. Silverio, and J. Qi. A fast algorithm to solve delay vandermonde systems in phased-array digital receivers. *the IEEE Transaction on Aerospace and Electronic Systems*, 57(4):2288–2297, 2021.
- [56] J.A. Fessler and B.P. Sutton. Nonuniform fast fourier transforms using min-max interpolation. *IEEE Transactions on Signal Processing*, 51(2):560–574, 2003.
- [57] Q.H. Liu and N. Nguyen. An accurate algorithm for nonuniform fast fourier transforms (NUFFT). *IEEE Microwave and Guided Wave Letters*, 8(1):18–20, 1998.
- [58] Melanie Kircheis and Daniel Potts. Fast and direct inversion methods for the multivariate nonequispaced fast fourier transform. 2 2023.
- [59] Melanie Kircheis and Daniel Potts. Direct inversion of the nonequispaced fast fourier transform. 11 2018.
- [60] Stefan Kunis. Nonequispaced fft generalisation and inversion. *PhD. dissertation, Inst. Math., Univ. Lübeck, Lübeck*, 2006.
- [61] Ensar Ajkunic, Hana Fatkic, Emina Omerovic, Kristina Talic, and Novica Nosovic. A comparison of five parallel programming models for c++. In *2012 Proceedings of the 35th International Convention MIPRO*, pages 1780–1784, 2012.
- [62] Pasi Ylä-Oijala and M. Taskinen. Electromagnetic scattering by large and complex structures with surface. *Waves in Random and Complex Media - WAVES RANDOM COMPLEX MEDIA*, 19:105–125, 02 2009.
- [63] Nail Gumerov and Ramani Duraiswami. Fast multipole methods for the helmholtz equation in three dimensions. *Fast Multipole Methods for the Helmholtz Equation in Three Dimensions*, 01 2004.
- [64] Dmitri G. Fedorov, Takeshi Nagata, and Kazuo Kitaura. Exploring chemistry with the fragment molecular orbital method. *Phys. Chem. Chem. Phys.*, 14:7562–7577, 2012.
- [65] Jens Keiner and Daniel Potts. Fast evaluation of quadrature formulae on the sphere. *Math. Comput.*, 77:397–419, 01 2008.

Supplementary Material for:
Vandermonde Neural Operators.

APPENDIX A. TECHNICAL DETAILS

A.1. The FFT Signal Flow Graph. The signal flow graph provides a graphical representation of the FFT algorithm.




FIGURE 4. The 8-point fast Fourier transform signal flow graph. x and y represent the signal in the physical and Fourier domain, respectively. Dashed lines represent a multiplication by -1 , red elements denote a multiplication by that factor, and converging arrows represent a sum.

A.2. Practical Construction and Implementation Details. In practical applications, the matrices within this tensor may be stacked, such that $\mathbf{V} \in \mathbb{C}^{(m^2) \times n}$ constructed via the method proposed in Equation 15, with structure illustrated by Equation 16.

$$(15) \quad \mathbf{V}_{j,k} = \sqrt{\frac{2}{n}} \left[e^{-i \left(\sum_{l=0}^N (\lfloor \frac{j}{m^l} \rfloor \bmod m) P_{k,l} \right)} \right]_{j,k=0}^{(m-1)^N, n-1}$$

$$(16) \quad \mathbf{V} = \sqrt{\frac{2}{n}} \begin{bmatrix} \begin{bmatrix} e^{-i(0\mathbf{p}_1^T + 0\mathbf{p}_2^T)} \\ \vdots \\ e^{-i((m-1)\mathbf{p}_1^T + 0\mathbf{p}_2^T)} \end{bmatrix} \\ \begin{bmatrix} e^{-i(0\mathbf{p}_1^T + 1\mathbf{p}_2^T)} \\ \vdots \\ e^{-i((m-1)\mathbf{p}_1^T + 1\mathbf{p}_2^T)} \end{bmatrix} \\ \vdots \\ \begin{bmatrix} e^{-i(0\mathbf{p}_1^T + (m-1)\mathbf{p}_2^T)} \\ \vdots \\ e^{-i((m-1)\mathbf{p}_1^T + (m-1)\mathbf{p}_2^T)} \end{bmatrix} \end{bmatrix}$$

A.3. Point Distributions. We investigate several point distributions for the one dimensional case. For the two dimensional experiments, we investigate a lattice with a nonuniform distribution along one axis (*Shear Layer*), a lattice with a nonuniform distribution along both axes (*Surface-Level Specific Humidity*), and a fully nonequispaced point distribution (*Airfoil*). Visualizations of these are provided in this subsection.




FIGURE 5. Point distributions used in the Burgers' equation experiments. Data is selected from the uniform distribution to construct the contracting-expanding distribution and random distribution. The space between points in the contracting-expanding distribution grows from a point in both directions according to a geometric distribution, tuned so that the resulting distribution has 64 points. The random distribution is generated by generating a random permutation of all 8192 indices and selecting the first 64 points.

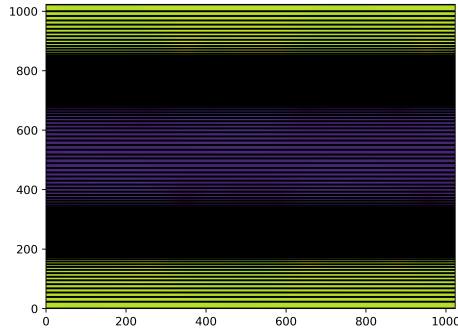
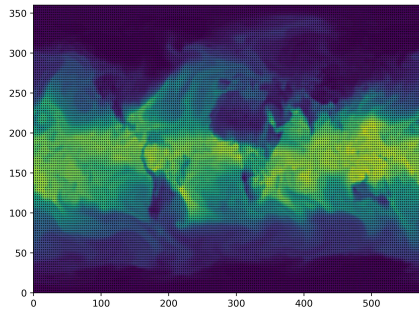



FIGURE 6. Nonequispaced lattice for the shear layer problem. Sampling is dense close the interface region, smoothly becoming sparse further from this region.



(A) FNO point distribution. The points displayed in this image have been sub-sampled from the original distribution to maintain clarity in the figure.



(B) VNO point distribution. A densely sampled region is located over South America and the lattice becomes more sparse further from this region.

FIGURE 7. Distributions used within the surface-level specific humidity experiment.




FIGURE 8. Several point distributions around airfoils and their associated pressure distributions.

A.4. Details on the Surface-Level Specific Humidity Data. We use 16 different parameters to make predictions for the *Surface-Level Specific Humidity* experiments. These are listed in this subsection.

TABLE 2. Inputs for the surface level specific humidity predictions.

Acronym in MERRA-2	Input	Units
CDH	heat exchange coefficient	$\frac{kg}{m^2s}$
CDQ	moisture exchange coefficient	$\frac{kg}{m^2s}$
EFLUX	total latent energy flux	$\frac{W}{m^2}$
EVAP	evaporation from turbulence	$\frac{kg}{m^2s}$
FRCAN	areal fraction of anvil showers	1
FRCCN	areal fraction of convective showe	1
FRCLS	areal fraction of large scale show	1
HLML	surface level height	m
QLML	surface level specific humidity	1
QSTAR	surface moisture scale	$\frac{kg}{kg}$
SPEED	surface wind speed	$\frac{m}{s}$
TAUX	eastward surface stress	$\frac{N}{m^2}$
TAUY	northward surface stress	$\frac{N}{m^2}$
TLML	surface air temperature	K
ULML	surface eastward wind	$\frac{m}{s}$
VLML	surface northward wind	$\frac{m}{s}$

A.5. Histograms and Deviations of Experimental Results. This subsection provides information on the distributions of results for experiments as well as results for the median test error given different initial seeds.

TABLE 3. The mean and standard deviation of the median test errors over 10 different initializations of the VNO for the *Burgers'* experiments.

Data Distribution	Median
<i>Uniform</i>	$0.10 \pm 0.03\%$
<i>Contracting-Expanding</i>	$0.86 \pm 0.08\%$
<i>Random</i>	$1.67 \pm 0.05\%$

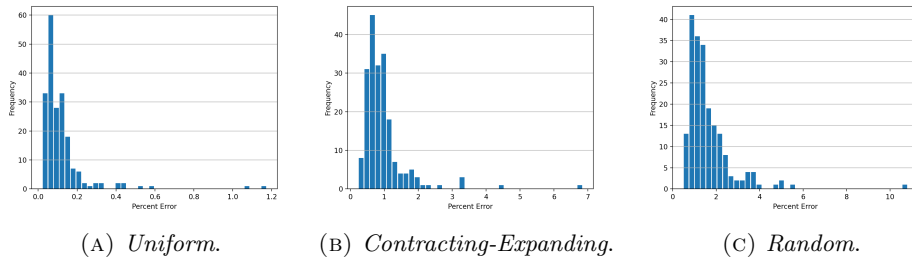


FIGURE 9. Histograms displaying the distribution of test errors for the different point distributions of the *Burgers'* numerical experiments.

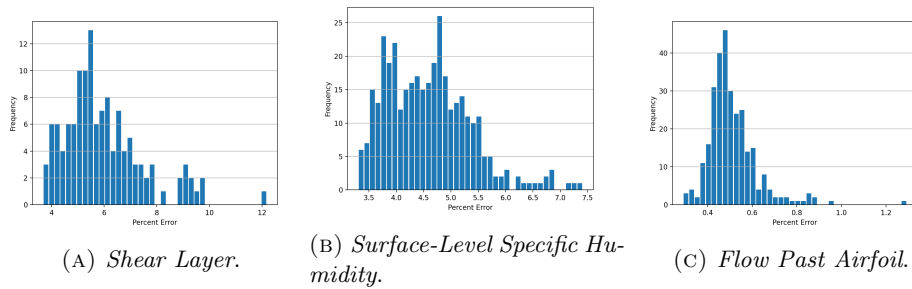


FIGURE 10. Histograms displaying the distribution of test errors for the two-dimensional numerical experiments.

DEPARTMENT OF MATHEMATICS, ETH ZURICH, ZURICH, SWITZERLAND
Email address: `llingsch@ethz.ch`

ETH AI CENTER, ETH ZURICH, ZURICH, SWITZERLAND
Email address: `mike.michelis@ai.ethz.ch`

DEPARTMENT OF MATHEMATICS, EMBRY-RIDDLE AERONAUTICAL UNIVERSITY, FLORIDA, USA
Email address: `sirani.perera@erau.edu`

SOFT ROBOTICS LAB, DEPARTMENT OF MECHANICAL AND PROCESS ENGINEERING, ETH ZURICH,
ZURICH, SWITZERLAND
Email address: `rkk@ethz.ch`

SEMINAR FOR APPLIED MATHEMATICS AND ETH AI CENTER, ETH ZURICH, ZURICH, SWITZER-
LAND
Email address: `siddhartha.mishra@sam.math.ethz.ch`