

A TAILORED CONVOLUTIONAL NEURAL NETWORK FOR NONLINEAR MANIFOLD LEARNING OF COMPUTATIONAL PHYSICS DATA USING UNSTRUCTURED SPATIAL DISCRETIZATIONS*

JOHN TENCER[†] AND KEVIN POTTER[†]

Abstract. We propose a nonlinear manifold learning technique based on deep convolutional autoencoders that is appropriate for model order reduction of physical systems in complex geometries. Convolutional neural networks have proven to be highly advantageous for compressing data arising from systems demonstrating a slow-decaying Kolmogorov n -width. However, these networks are restricted to data on structured meshes. Unstructured meshes are often required for performing analyses of real systems with complex geometry. Our custom graph convolution operators based on the available differential operators for a given spatial discretization effectively extend the application space of deep convolutional autoencoders to systems with arbitrarily complex geometry that are typically discretized using unstructured meshes. We propose sets of convolution operators based on the spatial derivative operators for the underlying spatial discretization, making the method particularly well suited to data arising from the solution of partial differential equations. We demonstrate the method using examples from heat transfer and fluid mechanics and show better than an order of magnitude improvement in accuracy over linear methods.

Key words. manifold learning, model order reduction, convolutional neural network, autoencoder

AMS subject classifications. 35Q35, 35Q68, 35Q74, 35Q79,

1. Introduction. High-fidelity computational simulations have become an integral part of the engineering design process. For many applications, such as turbine design, these simulations require solving a set of partial differential equations (PDEs), whose accurate solution often requires a fine spatial resolution that creates a large state-space dimension for the resulting system of ordinary differential equations (ODEs). As a result, numerical integration of the resulting system of ODEs is computationally costly and, in some circumstances, (such as many-query or time-critical problems) this is infeasible.

For this reason, it is desirable to generate low-dimensional surrogate models which may be solved quickly and accurately. These surrogates are usually either data-driven surrogates (e.g. recurrent neural networks [142, 140, 55, 82, 95], dynamic mode decomposition [124, 131, 111, 5]) or projection-based reduced-order models (pROMs) [60, 9]. Data driven surrogates are typically much faster and easier to implement than pROMs while pROMs may provide improved accuracy and stability. Regardless of type, these surrogate models involve two stages: (1) a computationally expensive offline stage that involves computing state snapshots at different times for samples of the parameter space, using those snapshots to define a low-dimensional manifold on which the solution evolves, and in the case of data-driven surrogates, learning the dynamics, and (2) a computationally inexpensive online stage, computing approximate trajectories restricted to lie on the prescribed manifold. In this paper we focus exclusively on the first (offline) stage and particularly upon the generation of low-dimensional manifolds which approximate solution variability.

The vast majority of research in this area has focused on the use of linear solution

*Submitted to the editors June 8, 2020.

Funding: This work was funded by the Laboratory Directed Research and Development program at Sandia National Laboratories

[†]Sandia National Laboratories, Albuquerque, NM (jtencer@sandia.gov, kmpotte@sandia.gov).

manifolds selected via balanced truncation [40, 85, 46, 10], rational interpolation [6, 31], or proper orthogonal decomposition (POD) [17, 130, 9, 16, 23, 39]. These types of projection methods can be expected to be very accurate with a very low-dimensional manifold for problems that exhibit fast decaying Kolmogorov n -width. However, many problems (particularly advection-dominated problems) exhibit slowly decaying Kolmogorov n -width requiring a high-dimensional linear manifold to achieve accurate results. High-dimensional manifolds are undesirable as the dimension of the manifold is directly related to the online computational cost of the surrogate model.

For problems with slow-decaying Kolmogorov n -widths, nonlinear manifolds can reduce computational costs while maintaining or improving accuracy. This has been demonstrated for data-driven surrogates [5, 38, 82, 18, 96, 2, 97, 101, 139, 79, 114, 99, 109, 45] as well as pROMs [119, 21, 132, 22, 102, 69, 70, 62]. The use of a nonlinear manifold is distinct from learning nonlinear dynamics within a linear manifold [125, 83, 87, 126, 36, 120, 50]. Nonlinear manifolds suitable for surrogate modeling may be generated by learning approaches such as Local Linear Embedding [110], Diffusion Maps [24, 93], Laplacian Eigenmaps [7, 8], or deep convolutional autoencoders [81, 20, 51]. In this work, we will focus on deep convolutional autoencoders, a particular architecture of convolutional neural network (CNN) that is applicable to manifold learning tasks.

CNNs are a highly advantageous architecture for this type of nonlinear manifold learning because the number of learnable weights in the network is independent of the input feature size (full-order model state space dimension), significantly reducing the computational cost [67]. CNNs are a powerful tool that has seen success in performing varied machine learning tasks, such as image classification [123, 66, 65], image super resolution [107, 61, 98], natural language processing [57, 29, 53], and even physics [25, 121, 33, 122, 147, 143, 41, 18, 69, 70, 96, 38, 91, 1, 129, 74, 90]. Unfortunately, traditional CNNs are not currently widely applicable for surrogate modeling of real-world physical systems, because they are restricted to structured meshes (such as a sequence, 2D array of pixels, or 3D array of voxels). The requirement of a structured mesh makes it impossible to apply traditional CNNs to problems with arbitrarily complex geometry, such as any real engineered system, without resorting to either resampling [91] or employing overset [115, 94, 84] or immersed boundary [100, 86] methods.

More recently, graph convolutional networks (GCNs) [134, 146] have been developed to deal with data on unstructured graphs such as those arising from wide-ranging application domains, including social networks [133], knowledge graphs [47], and computer vision [13]. The unstructured meshes commonly used for spatial discretization of complex physical systems are a subset of these general unstructured graphs. The analysis of data arising from computer simulation of physics problems is distinct from most types of learning tasks typically associated with GCNs. GCNs are often designed to capture combinatorial structures like node degrees (the number of graph edges connected to a given node) that are irrelevant to the solution of PDEs [128]. Reliance on this connectivity information has caused issues in applications like computer vision where this structural information is irrelevant to the learning task [146]. Unlike the computer vision application, which is almost exclusively focused on triangle surface meshes (for which the most important attribute is the shape), nodal coordinates are also of little interest when attempting to reconstruct PDE solutions. In order to reproduce the solution manifold of a set of PDEs, the spatial distribution of the field discretized by the mesh (temperature, velocity, stress, etc.) is of primary concern. Ideally, the representation of the solution manifold encoded into the net-

work weights should be discretization-independent which enables remeshing without requiring retraining.

Applications of GCNs in the physical sciences have thus far been primarily focused on particle physics [37, 48, 105, 80, 103, 116], chemistry [30, 59, 72, 127, 137], biology [35, 148, 34], and material science [135, 19, 58, 136, 141, 145]. In all cases, either the type of data considered or the learning task were distinct from what we propose here. To our knowledge, this work represents the first application of GCNs to data arising from PDE solutions using unstructured spatial discretizations.

In order to generate low-dimensional surrogate models for real engineered systems whose governing equations (advection-diffusion, Navier-Stokes, Euler, etc.) preclude the existence of a sufficiently low-dimensional linear solution manifold, a manifold learning method must be employed that is (1) practical for high-dimensional state-spaces, (2) able to learn complex nonlinear manifolds, and (3) not restricted to structured meshes. Linear methods satisfy 1 and 3 but not 2. Fully-connected autoencoder networks satisfy 2 and 3 but not 1. Traditional convolutional autoencoder networks satisfy 1 and 2 but not 3. In this work, we propose using a unique GCN that is compatible with data on unstructured meshes with arbitrary connectivity allowing the technique to be used for any discretized system.

Our proposed manifold learning method is capable of satisfying all 3 requirements and is thus well-suited to this challenging learning task. The solution manifolds generated by this technique are compatible with any of the previously mentioned surrogate modeling techniques. We draw inspiration from prior work in the fields of graph neural networks and geometric learning and formulate a solution for our learning task that is customized to the learning of PDE solutions while fitting within the more general established framework of GCNs.

Additionally, we seek implementations that are applicable to generic spatial discretizations, including mesh-free methods. As such, we use operators that involve *discretization-independent* learned weights. We propose sets of convolution operators based on the spatial derivative operators for the underlying spatial discretization, making the method particularly well-suited to data arising from the solution of PDEs. In the 3 exemplar problems chosen, our approach consistently produces more than an order of magnitude reduction in projection error relative to an optimal linear method for low-dimensional manifolds. In the 1 exemplar problem with a structured mesh, our approach performs comparably to a traditional CNN autoencoder network.

The paper is organized as follows. The convolutional autoencoder network architecture is briefly described in [section 2](#). Our new algorithm is described in [section 3](#) in the context of GCNs. The implementation details are given in [section 4](#), [section 5](#), and [section 6](#). The performance of our method is summarized in [section 7](#). The numerical experiments and corresponding results are described in [section 8](#). The conclusions follow in [section 9](#).

2. Autoencoders for Manifold Learning. Autoencoders are a well-established unsupervised manifold learning technique [64, 113]. The autoencoder seeks to learn a mapping between a high-dimensional vector $x \in \mathbb{R}^n$ and a low-dimensional latent representation $\tilde{x} \in \mathbb{R}^{d_{latent}}$ with $n \gg d_{latent}$. The autoencoder consists of two components, the encoder $h_{enc} : \mathbb{R}^n \rightarrow \mathbb{R}^{d_{latent}}$ and the decoder $h_{dec} : \mathbb{R}^{d_{latent}} \rightarrow \mathbb{R}^n$. A neural network is used to approximate both h_{enc} and h_{dec} . The two are combined to define the autoencoder

$$(2.1) \quad h : x \mapsto h_{dec}(h_{enc}(x))$$

which is desired to be an approximation of the identity map. The autoencoder network is typically trained by minimizing $\|x - h(x)\|$ in some norm. Once trained, the encoder and decoder networks may be separated to provide a description of a d_{latent} -dimensional manifold in \mathbb{R}^n approximating the training data.

Fully-connected networks are commonly used to approximate h_{enc} and h_{dec} . Unfortunately, the number of learnable weights in a fully-connected network scales with n^2 . Additionally, fully-connected networks are tied to a particular input dimensionality and structure. As a result, practical considerations restrict the use of fully-connected networks to relatively low-dimensional applications with fixed geometries and spatial discretizations. Recent large turbulence simulations have $\mathcal{O}(10^{12})$ grid points [108]. Even less ambitious undertakings commonly use state vectors with $\mathcal{O}(10^6)$ or $\mathcal{O}(10^7)$ elements rendering fully-connected networks impractical.

CNNs take advantage of local support and weight sharing to make the number of learnable weights independent of the input dimension. As a result, CNNs have enjoyed wide success in the processing of high-dimensional data, including PDE solution data [41, 75, 27, 138]. Unfortunately, weight sharing in traditional CNNs is based on an assumed structure to the dataset. Each element i in the input vector is assumed to have a well defined neighborhood $\mathcal{N}(i)$ with a fixed size and a consistent and meaningful ordering of nodes $j \in \mathcal{N}(i)$ within that neighborhood). This assumption is valid for images or data on structured meshes, but is invalid for general unstructured spatial discretizations. GCNs provide the same weight sharing benefits of CNNs without making the same assumptions on the structure of the data.

3. Graph Convolutions. There are two main types of GCNs, spectral and spatial. Spectral GCNs interpret the graph convolution operation as the application of a filter from the perspective of graph signal processing [117]. In contrast, spatial GCNs interpret the convolution operation through the lens of information propagation. The spatial approach can enable greater parallelism due to the inherently local nature of the operation. Although our contribution may be viewed through either lense, we derive our contribution as an extension to spectral GCNs and a special case of MoNet [89] due to the convenience of matrix notation afforded to spectral GCNs.

In the literature on GCNs, a graph is represented by $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ where \mathcal{V} denotes the nodeset and $|\mathcal{V}| = n$ and \mathcal{E} denotes the edgeset. We use $v_i \in \mathcal{V}$ to denote a node and $e_{ij} = (v_i, v_j) \in \mathcal{E}$ denote an edge connecting v_i and v_j . We also have a node feature matrix $X \in \mathbb{R}^{n \times c}$ defining the c attributes present for each node.

We define the adjacency matrix $A \in \mathbb{R}^{n \times n}$, such that $A_{ij} = 1$ if $e_{ij} \in \mathcal{E}$ and $A_{ij} = 0$ if $e_{ij} \notin \mathcal{E}$. For stability reasons, it is customary to explicitly include self-loops in A , i.e. $A_{ii} = 1$. The adjacency matrix is normalized using the degree matrix $D \in \mathbb{R}^{n \times n}$, which is a diagonal matrix with $D_{ii} = \sum_j A_{ij}$. For GCNs, the graph Laplacian is defined as

$$(3.1) \quad L = D^{-1/2} A D^{-1/2},$$

and the graph convolution layer is then

$$(3.2) \quad X^{(l+1)} = \epsilon(LX^{(l)}W),$$

where ϵ is a nonlinear activation function, $W \in \mathbb{R}^{c_l \times c_{l+1}}$ is a trainable weight matrix, and c_l and c_{l+1} are the number of input and output features for the layer respectively.

We observe that a network whose convolutional layers are exclusively of the form (3.2) is insufficiently expressive owing to the comparatively small number of attributes

per node for most common PDEs relative to what is typical in, for example social science [43, 63] or knowledge graphs [42]. Additionally, the graph Laplacian has the undesirable property of directly depending upon the local mesh connectivity (through the adjacency matrix) in a way that prevents the learned weights from being discretization-independent.

Equation (3.2) can be split up into 3 parts, a filter application or mixing step, multiplication by a learned weight matrix, and the application of the nonlinear activation function. The mixing step, which “mixes” information from a node’s neighborhood, is the multiplication of L by the input feature matrix $X^{(l)} \in \mathbb{R}^{n \times c_l}$. We generalize this mixing step by allowing for an arbitrary number m of mixing operations $M_k \in \mathbb{R}^{n \times n}$, the results of which are concatenated along the feature/channel dimension:

$$(3.3) \quad F = \left[M_0 X^{(l)}, M_1 X^{(l)}, \dots, M_{m-1} X^{(l)} \right] \in \mathbb{R}^{n \times c_l m}.$$

The layer is completed by the subsequent multiplication by a weight matrix and application of the nonlinear activation function

$$(3.4) \quad X^{(l+1)} = \epsilon(FW),$$

with the generalized mixing step remain unchanged except that the number of weights $W \in \mathbb{R}^{c_l m \times c_{l+1}}$ is increased by a factor of m . The use of multiple mixing operators in this way has been previously shown to be beneficial for shape classification [12] and image segmentation [56] tasks. The standard GCN using the graph Laplacian can then be viewed as a special case, where $m = 1$ and $M_0 = L$.

It is desirable to minimize the number of mixing operators m for a given problem because the number of learnable parameters for a given layer is $c_l \times m \times c_{l+1}$. However, too few mixing operators will be shown to result in reduced accuracy. For this reason, we propose several different predefined sets of mixing operators, such that the network may be tailored to the requirements of the application.

4. Mixing Layers. We propose deriving mixing operators from the spatial gradient operators defined for the underlying spatial discretization. The existence of these gradient operators is a prerequisite to any PDE solution procedure and, thus, should be readily available for any application. Unlike L , these operators should not introduce undesirable dependencies of the learned weights on the local element connectivity.

At this time, it is instructive to draw some comparisons to the traditional convolution operator from which we draw inspiration. Traditional convolution operators may be expressed in a way that is comparable to the description used for graph convolution layers in section 3 which is helpful for understanding the relationships between the two. Traditional convolution operators are usually described as the application of a kernel to an input of the same dimension. While the discussion to follow is applicable to any convolution operator, we will focus on the common case of the application of a 3×3 kernel to a 2D array (e.g., an image) for illustrative purposes.

The application of any kernel may be written as multiplication by an appropriately chosen sparse matrix. A 3×3 kernel has 9 parameters, k_{ij} , $(i, j) \in \{-1, 0, 1\} \times \{-1, 0, 1\}$. For a given input image $X^{(l)} \in \mathbb{R}^{M \times N}$, the output image from the convolution is

$$(4.1) \quad X_{m,n}^{(l+1)} = \sum_{i,j} k_{ij} X_{m+i,n+j}^{(l)}.$$




Fig. 1: Differential operators on a regular square mesh are closely related to common filters used for image processing. Note that the gradient operators (b and c) are Sobel filters and that the Laplacian (d) is an edge detection filter.

If the boundaries are treated with zero-padding, $X_{0,n} = X_{M+1,n} = X_{m,0} = X_{m,N+1} = 0$ for all m, n . If we flatten X into $\text{vec}(X) \in \mathbb{R}^{MN}$ we may write (4.1) as a sum of matrix-vector products. We define the linear flattening operation as

$$(4.2) \quad \begin{aligned} \text{vec}(X) &= \sum_{i=1}^N e_i \otimes X e_i \\ &= [X_{1,1}, \dots, X_{M,1}, X_{1,2}, \dots, X_{M,2}, \dots, X_{1,N}, \dots, X_{M,N}]^T, \end{aligned}$$

with e_i as the i -th canonical basis vector for the N -dimensional space and $\text{vec}_{M,N}^{-1}$ as the corresponding inverse operation. We see that

$$(4.3) \quad \text{vec}(X^{(l+1)}) = \sum_{i,j} k_{ij} P_{ij} \text{vec}(X^{(l)}),$$

where $P_{ij} \in \mathbb{R}^{MN \times MN}$ is a suitably chosen permutation matrix such that (4.1) and (4.3) are equivalent, i.e. $\text{vec}_{M,N}^{-1}(P_{ij} \text{vec}(X))_{m,n} = X_{m+i,n+j}$.

From (4.3) we see that the traditional 3×3 convolution operator is equivalent to a set of 9 mixing operators with each mixing operator being a permutation matrix corresponding to a shift of the input image. The kernel parameters k_{ij} play the role of the learned weights. There is a one-to-one correspondence between traditional convolutional kernels and these mixing operators. This connection has been previously exploited in the context of data-driven PDE identification for structured spatial discretizations [77, 76].

If we view an $M \times N$ image as being defined on a regular rectangular mesh with nodes at the pixel centers, we can compute matrix representations of various differential operators using a linear finite element basis. These differential operators are trivially expressed as sparse $MN \times MN$ matrices. For interior nodes, these differential operators are equivalent to convolutions with the (unnormalized) kernels shown in Figure 1. There are some differences between how these operators weight contributions at the boundary nodes compared to the normal approach of zero-padding for traditional CNNs, but these differences are not generally significant. Zero-padding can be viewed as applying a (potentially unphysical) boundary condition to the filtered outputs which has been observed to result in degraded performance near the boundary [54]. The results of using finite element differential operators at the boundary is more similar to partial convolution based padding [73] although still not equivalent.

Incorporating known boundary condition information into the mixing operators is a potential area of future research. Embedding this type of constraint has been shown to be beneficial for traditional CNNs [88].

Looking at the differential operators in Figure 1, it quickly becomes apparent that a given set of mixing operators defines a linear subspace of the 9-dimensional space spanned by the normal 3×3 convolution kernel. It has been previously observed that learned filters may be represented using a lower dimensional basis as a form of network compression [104]. Visualizing these operators by looking at their stencils is useful for gaining intuition regarding how expressive the resulting network is likely to be. Note, for instance, that for the particular case of a regular square grid that the Laplacian $\Delta = \nabla^2$ is a common edge-detection kernel while ∇_x and ∇_y are closely related to Sobel filters. The L operator from 3.1 is a box blur (if diagonal connections are included as graph edges) and may be recovered as a linear combination of the differential operators $L = I - \frac{1}{9}\Delta$.

Parameterized differential operators have been previously proposed for various learning tasks including shape classification of meshed surfaces [12] and image classification and segmentation on the unit sphere [56]. While these previous investigations were focused on more restricted application spaces, the operators suggested therein may be more generally applicable. In anisotropic CNNs [12], a patch operator is constructed using anisotropic heat kernels. In our framework, that roughly corresponds to mixing operators of the form

$$(4.4) \quad \nabla \cdot \left(\begin{bmatrix} \delta_{0i} & 0 & 0 \\ 0 & \delta_{1i} & 0 \\ 0 & 0 & \delta_{2i} \end{bmatrix} \nabla \right), i \in \{0, 1, 2\},$$

where δ_{ij} is the Kronecker delta. Parameterized differential operators have also been previously proposed for image classification and segmentation tasks on a unit sphere [56]. In that work, the set of mixing layers was chosen to be $\{I, \nabla_\theta, \nabla_\phi, \nabla^2 = \Delta\}$ with ∇_θ and ∇_ϕ corresponding to the polar and azimuthal partial derivative operators respectively. The derivative components were chosen to align with the global spherical coordinate system. We will present our proposed mixing operators using 2D cartesian coordinates but application to other coordinate systems and higher dimensions is trivial. For our test problems, we have observed that using the derivative operators directly is inadvisable, due to potentially large scale differences between a function and its derivatives. We propose normalizing the derivative operators, such that their absolute value row sums are unity:

$$(4.5) \quad \begin{aligned} \widetilde{\nabla}_x &= D_x^{-1} \nabla_x & (D_x)_{ii} &= \sum_j |(\nabla_x)_{ij}| \\ \widetilde{\nabla}_y &= D_y^{-1} \nabla_y & (D_y)_{ii} &= \sum_j |(\nabla_y)_{ij}| \\ \widetilde{\Delta} &= D_\Delta^{-1} \Delta & (D_\Delta)_{ii} &= \sum_j |\Delta_{ij}|. \end{aligned}$$

This type of normalization is common for GCNs. We have also observed performance benefits in some cases by further enriching the basis. Specifically, separating the positive and negative contributions of the gradient operators into separate operators

(which is analogous to adding a learned upwinding parameter), e.g.

$$(4.6) \quad \begin{aligned} (\nabla_x^+)_{ij} &= \begin{cases} 2(D_x^{-1}\nabla_x)_{ij}, & \text{for } (\nabla_x)_{ij} > 0 \\ 0, & \text{otherwise} \end{cases} \\ (\nabla_x^-)_{ij} &= \begin{cases} -2(D_x^{-1}\nabla_x)_{ij}, & \text{for } (\nabla_x)_{ij} < 0 \\ 0, & \text{otherwise} \end{cases} \\ (\nabla_y^+)_{ij} &= \begin{cases} 2(D_y^{-1}\nabla_y)_{ij}, & \text{for } (\nabla_y)_{ij} > 0 \\ 0, & \text{otherwise} \end{cases} \\ (\nabla_y^-)_{ij} &= \begin{cases} -2(D_y^{-1}\nabla_y)_{ij}, & \text{for } (\nabla_y)_{ij} < 0 \\ 0, & \text{otherwise} \end{cases} \end{aligned}$$

has proven to be valuable in some circumstances.

We therefore have a large set of proposed mixing operators, a subset of which may be used for any given mixing layer. In the examples to follow, we will refer to the following master list (4.7) when specifying the set of mixing operators used in a given layer. For compactness of notation, we omit the explicit scaling of the derivative operators and state that all mixing operators are normalized, such that their absolute value row sums are unity. The combinations of mixing operators examined in the experiments to follow will be:

$$(4.7) \quad \begin{aligned} \mathcal{M}_a &= \{I, \Delta_x, \Delta_y\} \\ \mathcal{M}_b &= \{I, \nabla_x, \nabla_y, \Delta\} \\ \mathcal{M}_c &= \{I, \nabla_x, \nabla_y, \Delta_x, \Delta_y\} \\ \mathcal{M}_d &= \{I, \nabla_x^+, \nabla_x^-, \nabla_y^+, \nabla_y^-, \Delta\} \\ \mathcal{M}_e &= \{I, \nabla_x^+, \nabla_x^-, \nabla_y^+, \nabla_y^-, \Delta_x, \Delta_y\} \\ \mathcal{M}_f &= \{I\} \\ \mathcal{M}_g &= \{L\} \\ \mathcal{M}_h &= \{\Delta\} \\ \mathcal{M}_i &= \{I, \Delta\}. \end{aligned}$$

Note that not all possible combinations of these operators are linearly independent (ie. $\nabla_x^+ - \nabla_x^- \propto \nabla_x$), so some care should be taken when selecting operators for a mixing layer to avoid duplication. Additionally, note that L is the graph Laplacian discussed in section 3 such that \mathcal{M}_g represents a traditional GCN.

5. Pooling and Unpooling Layers. The other crucial component to CNN architectures is the pooling operation. Pooling reduces the spatial resolution of the information as it propagates deeper into the network. The most common types of pooling operations in deep learning are max pooling and average pooling which take the maximum or average value of a feature in a given neighborhood. There have been numerous pooling operations proposed for GCNs [106, 28, 68, 15, 26, 118, 11]. In this work, we will use a sort of weighted average pooling derived from bilinear interpolation.

First, we define the unpooling operation. Given two meshes representing a coarse and a fine discretization of the same spatial domain, an $n_{fine} \times n_{coarse}$ linear interpolation operator may be defined that maps from the coarse mesh to the fine mesh.

Each row in this matrix corresponds to a node in the fine mesh. The values in that row correspond to the weights of each node in the coarse mesh when interpolating a function defined on the coarse mesh to the spatial location of the node on the fine mesh. This mapping from coarse mesh to fine mesh is the unpooling operator that will be used in the numerical experiments to follow. The corresponding pooling operator is simply the transpose of the unpooling operator defined in this way. Both of these operators are sparse matrices.

6. Network Architectures. Very similar architectures are used for all 3 exemplar problems. The only difference was the depth of the network. For each example, we precompute a sequence of coarsened meshes along with corresponding differential and interpolation operators. The network components are then applied at the different mesh resolutions, using the corresponding operators. The fundamental components of the autoencoder networks are shown in Figure 2. A diagram of the full autoencoder architecture is included in Appendix B. We have observed that the encoder task is significantly less complex than the decoder task for our example applications, which is intuitive, as the decoder task is analogous to solving the system of PDEs whereas the encoder task is analogous to performing a regression to find an unknown parameter from data. As such, we have chosen an asymmetric architecture that devotes a larger number of learned parameters to the decoder task.

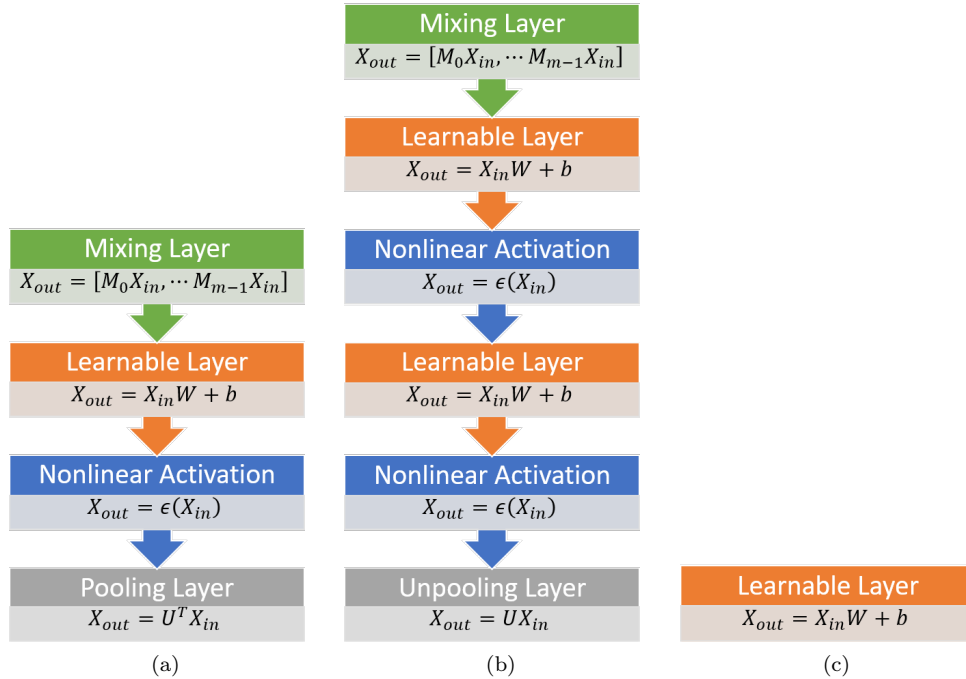


Fig. 2: The convolutional autoencoder networks are comprised of 3 fundamental components: the (a) encoder convolutional module, (b) decoder convolutional module, and (c) decoder output module.

The input is fed into the network as a 3D tensor of size $[B, C, n_0]$, where B is the batch size (snapshots per batch), C is the number of input channels (number of

nodal fields), and n_0 is the number of nodes in the input mesh. The ordering of the nodes is unimportant as long as it’s consistent across snapshots and with the row-ordering of the mixing and interpolation operators. Scaling may be optionally applied independently for each input channel as a preprocessing step but was not required for the examples to follow. For each network, a sequence of $d + 1$ progressively coarser meshes is generated with n_0, n_1, \dots, n_d nodes. The encoder network consists of one encoder convolutional module (Figure 2a) for each mesh resolution, a fully connected layer to reduce the size to 64, a nonlinear activation, and another fully connected layer to reduce to the desired latent space dimension d_{latent} , which is treated as a hyperparameter. All nonlinear activation layers use the ELU activation function

$$(6.1) \quad \epsilon(x) = \max(0, x) + \min(0, \alpha(\exp(x) - 1)),$$

with $\alpha = 1$, consistent with previously published results [69]. During experimentation, the choice of nonlinear activation function was not observed to significantly impact the quality of results. For the pooling operation, we use the renormalized transpose of the interpolation operator that maps from the coarse mesh to the fine mesh as described in section 5. We will use this same interpolation operator for the unpooling layers in the decoder.

The decoder network begins with 4 fully connected layers. The number of neurons in each layer is 64, $4n_d$, $16n_d$, and $64n_d$, where n_d is the number of nodes in the coarsest mesh. Each fully connected layer is followed by an ELU nonlinear activation layer. These are followed by one decoder convolutional module (Figure 2b) at each mesh resolution save the final (finest) mesh. An output layer (Figure 2c) is used rather than a convolutional layer for the finest mesh. Since no nonlinear activation functions are used in the output layer, the features may be optionally extracted at this point and treated as a nonlinear mode decomposition [92], which is useful for visualization and interpretability. While, in theory, each mixing layer could be comprised of a unique set of mixing operators; in practice, this is cumbersome to manage; so, for each experiment, a common set of mixing operators will be used for all mixing layers.

7. Computational Cost. The computational cost may be split into two parts, the offline (training) cost and the online (inference) cost. In terms of offline cost, deep autoencoder methods are orders of magnitude more expensive than linear methods. In terms of online cost, the linear methods consist of a single dense matrix-vector product which requires $2d_{latent}n_0$ flops. There have been several methods created to reduce or eliminate the n_0 scaling of the cost via hyper-reduction [112, 16, 17, 49]. For CNNs, each convolutional layer may be expressed as a series of sparse matrix vector products. Neglecting activation functions, the computational cost of a convolution at layer l is approximately

$$(7.1) \quad 2mn_l c_l c_{l+1}.$$

Recall from section 3 that m is the number of mixing operators (9 for a traditional 3×3 convolution, n_l is the number of nodes present at layer l , and c_l and c_{l+1} are the number of input and output channels respectively. There has been a great deal of research in recent years into reducing the online cost of evaluating deep neural networks, much of it focused on network compression and pruning [44, 71, 4, 144]. The pooling and unpooling operations are simple sparse matrix-vector products at a cost of 2 flops per nonzero element of the interpolation matrix.

Note that the online cost only requires evaluating the decoder half of the autoencoder network and that the overall cost will depend on the network architecture

(the depth, number of channels at each layer, etc.) and that (7.1) is valid for both traditional CNNs and the novel convolutional layer developed here. One could easily assume therefore that because $m < 9$ for all of the mixing operator sets included in (4.7) that the online computational costs of our networks should be lower than those of analogous traditional CNNs. However, the implementations of traditional convolutional layers in popular deep learning frameworks such as PyTorch and TensorFlow are highly optimized. In contrast, support for general sparse linear algebra in these frameworks is relatively recent. As a result, despite the reduction in flops, our implementation is noticeably slower (approximately 7x for a simple benchmark) in terms of wall clock time. We expect that as native support in deep learning frameworks improves that the speed of GCNs will eventually match and exceed that of traditional CNNs.

8. Numerical Experiments. We will be considering three different datasets from a range of computational physics applications. We will be constructing autoencoders and assessing performance in terms of the degree to which the solution lies within the generated manifold. Specifically, we will use the mean-squared error (MSE) and relative L^2 error between the solution snapshots and their reconstructions as our primary metrics. For the experiments to follow, we will vary the sets of mixing operators \mathcal{M} used within the network as well as d_{latent} . The three datasets are:

- Transient advection-diffusion,
- Unsteady laminar fluid flow past a cylinder, and
- Inviscid supersonic flow over a wedge.

The MSE loss function is used for both training and evaluation. This loss function weights the error at each node equally. For regular meshes, this will correspond closely with the L^2 error in the reconstruction. For highly irregular meshes, it may be advisable to use the L^2 or L^1 error directly to avoid over-emphasizing densely meshed regions. In practice, densely meshed regions often correspond to regions of particular interest; therefore, weighting those regions more heavily in the loss calculation may be appropriate in some cases. For the experiments to follow, MSE and L^2 yield similar trends and the same conclusions. For ease of interpretation, the L^2 error is used for the experiments with irregular meshes.

The basis of comparison for the nonlinear manifolds will be linear manifolds generated via POD, which is optimal in the MSE metric. In all cases the error cited for POD is the minimum projection error of the test data onto the POD basis. For the autoencoder results, the reconstruction error used is $\|x - h_{dec}(h_{enc}(x))\|$ which is bounded from below by the optimal projection error $\min_{\hat{x}} \|x - h_{dec}(\hat{x})\|$. In this way, the reported results are conservative.

All models were trained using the AdamW optimizer [78] implemented in PyTorch with an initial learning rate of 5×10^{-3} . The learning rate was reduced by a factor of 0.8 if no improvement in the training loss after 10 epochs until a minimum learning rate of 10^{-5} . Training was terminated after 20 consecutive epochs without improvement in the training loss once the learning rate attained this minimum value. There was no limit on the maximum number of epochs. This termination criteria was chosen to avoid penalizing slow-converging networks. The network weights were initialized using PyTorch defaults. Mini-batching was used for all experiments with batch sizes of 125, 465, and 41 respectively. The batch sizes for each experiment were chosen based on hardware limitations rather than optimized for accuracy or performance.

In the following subsections, we will summarize the results of each experiment. Referring back to the master list of proposed mixing operators (4.7), we use all of the

proposed mixing operator sets for all experiments. Detailed results are tabulated in [Appendix A](#). The data for each example was generated using a finite element spatial discretization. The datasets used for training and testing for each experiment are included as supplemental material.

8.1. Transient Advection-Diffusion. For our first example, consider the transient, 2D advection-diffusion equation:

$$(8.1) \quad \begin{aligned} \frac{\partial \alpha}{\partial t}(\vec{r}, t) &= \nabla \cdot (D \nabla \alpha(\vec{r}, t)) - \vec{u} \cdot \nabla \alpha(\vec{r}, t) \\ \vec{r} &\in \Omega \subset \mathbb{R}^2, t \in [0, 1] \end{aligned}$$

$$\alpha(\vec{r}, 0) = 0 \text{ in } \Omega$$

$$\alpha(\vec{r}, t) = 1 \text{ for } \vec{r} \in \partial\Omega.$$

The diffusivity D is chosen to be 0.1 and the advection velocity is chosen to be a unit vector, $\|\vec{u}\| = 1$. The equation is solved on the unit square $\Omega = (0, 1) \times (0, 1)$. The direction θ of the advection velocity is varied so $\vec{u} = (\cos(\theta), \sin(\theta))$. Training data is generated every 30° beginning with the positive x -axis ($\theta = 0$). Separate test data is generated for angles of $\theta \in \{45^\circ, 135^\circ, 225^\circ, 315^\circ\}$. For each value of θ , data is collected for $t \in (0, 1]$. Adaptive timestepping is used resulting in $\mathcal{O}(10^2)$ snapshots per value of θ . The 5 meshes used are regular square meshes of 160×160 , 80×80 , 40×40 , 20×20 , and 10×10 elements apiece.

The full results for all combinations of mixing operators, $\mathcal{M} \in \{\mathcal{M}_a, \dots, \mathcal{M}_i\}$ with $d_{latent} \in \{4, 16, 64\}$ are included in [Table 2](#) in [Appendix A](#). There are also selected results for $d_{latent} = 2$. The results of varying d_{latent} are uneven. In general, increasing d_{latent} should result in a more accurate model. However, since the actual solution manifold for this case is fully contained in \mathbb{R}^2 , it is perhaps expected that increasing d_{latent} beyond 4 fails to continue improving the solution quality. While it is possible to obtain accurate results for $d_{latent} = 2$ for some test problems, the networks suffer in terms of generalizability; a low training loss does not imply low test loss across the board. For POD, the solution continues to improve as expected as d_{latent} is increased. POD accuracy is comparable to the autoencoder results at $d_{latent} = 16$ and nearing machine precision for $d_{latent} = 64$.

When constructing surrogate models, it is desirable to minimize d_{latent} . For both pROMs and data-driven surrogates, reducing d_{latent} results in a reduction in online computational cost. For pROMs, the online cost is closely related to the cost of solving a least-squares problem (projecting the dynamics onto the solution manifold). The solution is usually generated via an iterative procedure, the cost of which will be problem dependent. If the quadratic Newton-Raphson method is used, the cost of each iteration will scale as $\mathcal{O}(d_{latent}^2)$ [32]. For data-driven surrogates, the computational cost reduction will depend on the surrogate used.

Because the mesh used for this example is a regular grid, a traditional CNN is also used. The architecture is as similar as possible to the autoencoder architecture described in [section 6](#). The same number of convolutional layers with the same number of filters are used and the fully connected layers are identical. The accuracy for both the traditional CNN and proposed extension are comparable indicating that our proposed technique effectively replicates the accuracy of the traditional CNN without relying on any underlying structure in the mesh. Notably, the traditional CNN reconstruction suffers from some high frequency artifacts arising from the upsampling




Fig. 3: Our graph-based autoencoder reproduces the spatial distribution of α much more accurately than the optimal linear (POD) manifold approach and produces qualitatively better reconstructions than even a traditional CNN autoencoder. All 4 subfigures show the distribution of α at the final simulated time ($t = 1$) for $\theta = 45$ deg. All reconstructions are performed with $d_{latent} = 4$.

procedure, as shown in [Figure 3](#). It should be noted that these same high-frequency artifacts have been observed in other works using deep convolutional autoencoders [[3](#), [69](#), [52](#)].

The results are more consistently impacted by the choice of mixing operators. For some sets of mixing operators (especially those with $m < 3$), the accuracy of the resulting model appears to be highly dependent on the random initialization of the network. This indicates an insufficiently expressive architecture. This is undesirable and a strong argument for preferring larger sets of mixing operators. Mixing operator

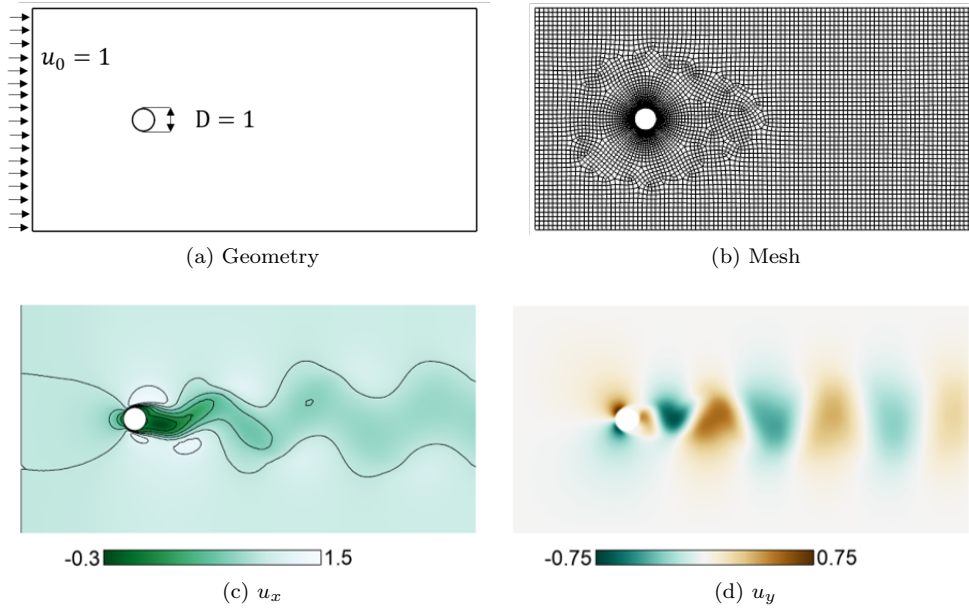


Fig. 4: Unsteady laminar flow past a cylinder (a) has solutions computed on an unstructured mesh (b) which exhibit the familiar vortex shedding patterns visible in the (c) horizontal and (d) vertical velocities.

set \mathcal{M}_e with $m = 7$ is a consistently well-performing option. Mixing operator sets with $m < 3$ consistently perform worse than those with $m \geq 3$. Operator sets \mathcal{M}_c and \mathcal{M}_e are consistently the top performers, which indicates that value is added by including the anisotropic Laplace operators in the mixing operator set for this problem. However, differences between operator sets with $m \geq 3$ were less significant than the differences between those sets and sets with $m < 3$, including the traditional GCN, \mathcal{M}_g .

Figure 3 compares the reconstructions achieved via convolutional autoencoder (with mixing operator set \mathcal{M}_e), a linear model (POD), and a traditional CNN with a similar autoencoder architecture. The autoencoder-based approaches consistently outperform POD.

8.2. Unsteady Laminar Fluid Flow Past a Cylinder. The second example involves unsteady laminar flow past a cylinder that generates the familiar vortex shedding pattern known as a von Kármán vortex street. The governing equations for this system are the incompressible Navier-Stokes equations:

$$(8.2) \quad \begin{aligned} \frac{\partial \vec{u}}{\partial t} + (\vec{u} \cdot \nabla) \vec{u} - \nu \nabla^2 \vec{u} &= -\frac{1}{\rho} \nabla p \\ \nabla \cdot \vec{u} &= 0. \end{aligned}$$

The geometry is that of a cylinder in cross-flow, as shown in Figure 4. The cylinder has a diameter of 1 and the free stream velocity is 1. The kinematic viscosity ν is varied such that the Reynolds number is between 100 and 400. Symmetry boundary conditions are applied at the top and bottom edges of the domain and an open

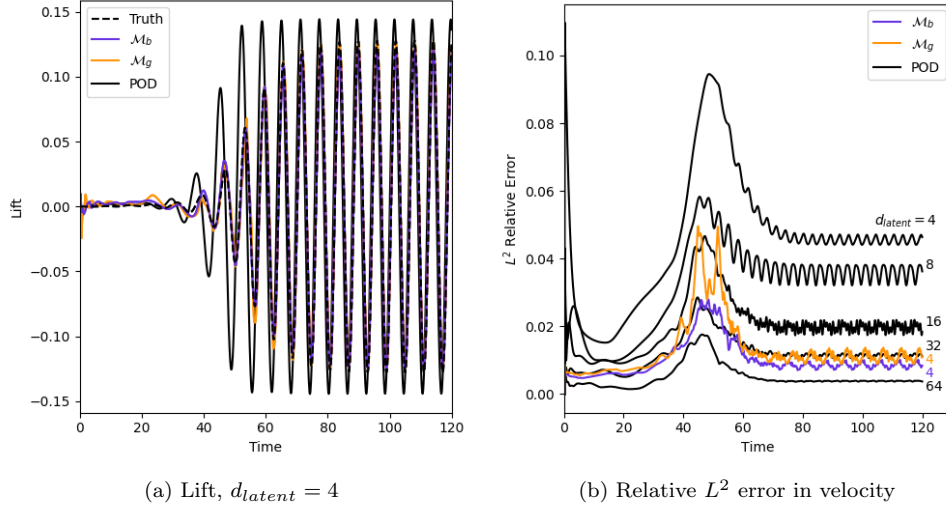


Fig. 5: (a) Autoencoder reconstructions with $d_{latent} = 4$ are significantly more accurate than POD in terms of (a) transient lift generated from the reconstructed pressure and (b) transient L^2 error in the velocity distribution for $Re = 150$. The autoencoder accuracy with $d_{latent} = 4$ is comparable to POD with $d_{latent} = 32$. Additionally, the differential operator approach (\mathcal{M}_b) clearly outperforms the traditional GCN (\mathcal{M}_g).

pressure boundary condition is applied at the outlet. Solutions are generated on the unstructured mesh of 6384 quad elements shown in Figure 4b. Simulations are run with adaptive time stepping with a termination time of 120 resulting in $\mathcal{O}(10^3)$ solution snapshots generated per Reynolds number value. Training data is generated for Reynolds numbers of 100, 200, 300, and 400. Separate test data is generated at Reynolds numbers of 150, 250, and 350. The fields considered for reconstruction are the horizontal velocity $u_x \in (-0.4, 1.4)$, vertical velocity $u_y \in (-0.75, 0.75)$, and pressure $p \in (-0.75, 0.6)$.

A selection of the results are tabulated in Table 1. Of particular note is that the autoencoder-based reconstructions are consistently more accurate than the optimal linear (POD) reconstruction for small d_{latent} . The complete results are included in Table 3 in Appendix A.

Table 1: Performance for laminar flow experiment with $d_{latent} = 4$. Reported values correspond to relative L^2 error of the velocity reconstruction averaged over snapshots.

Mixing Set	Training Snapshots	Test Snapshots		
		$Re = 150$	$Re = 250$	$Re = 350$
(b)	6.021e-03	1.010e-02	9.335e-03	8.804e-03
(e)	7.580e-03	1.082e-02	1.078e-02	1.039e-02
(g)	7.985e-03	1.315e-02	1.259e-02	1.114e-02
POD	4.833e-02	4.633e-02	4.080e-02	4.751e-02

Relative to the first experiment, the accuracy of the trained autoencoders was




Fig. 6: Absolute error in the horizontal velocity reconstruction at the final timestep for the 3 test datasets using POD (left) and our autoencoder (right) with $d_{latent} = 4$.

insensitive to both of the hyperparameters we varied. There was not a clear winner amongst the mixing operator sets. Among operator sets with $m > 3$, the loss values are all so close together that slightly different training procedures may lead to different relative rankings, indicating that many of the combinations of mixing operators chosen may be suitable for this transient laminar flow application. We suspect that this is due to the relative lack of variability between the training and test sets owed to the periodic nature of the flow and the consistent flow direction.

However, operator sets \mathcal{M}_b and \mathcal{M}_d are consistently among the top performers, which indicates that splitting the Laplace operator $\Delta = \Delta_x + \Delta_y$ is likely not worthwhile for this application, possibly due to the invariance of the inlet flow direction across the dataset. Once again, POD achieves comparable accuracy to the autoencoder with $d_{latent} = 16$ although it is still consistently worse than even the worst performing autoencoder at that size (see Figure 5b). Figure 5 and Figure 6 show comparisons between the autoencoders with $d_{latent} = 4$ and POD. Figure 5 shows the temporal distributions of the lift force on the cylinder calculated from the reconstructed pressure and the L^2 error in the velocity reconstruction, while Figure 6 shows

the spatial distribution of the reconstruction error in the horizontal velocity at the final timestep. For small values of d_{latent} , our autoencoders are more than an order of magnitude more accurate than POD. Across all values of d_{latent} value is observed in using differential operators over a traditional GCN (\mathcal{M}_g).

The fact that increasing d_{latent} did not result in increased accuracy indicates that the system's dynamic response is adequately captured by the lower dimensional space, which is not surprising, given that the state evolves on a 2D manifold parameterized by Reynolds number and time. However, as seen in the advection-diffusion example, the existence of a low-dimensional solution manifold is no guarantee that our autoencoder will be able to consistently learn its structure. Although our autoencoders are able to learn a reasonable 2D manifold, there is a considerable accuracy benefit in providing the network with extra flexibility by setting $d_{latent} = 4$.




Fig. 7: Steady inviscid supersonic flow over a wedge results in an oblique shock originating at the corner. The flow state is piecewise constant with the inlet and outlet conditions shown.

8.3. Inviscid Supersonic Flow Over a Wedge. The third and final experiment involves steady, 2D supersonic inviscid flow. The geometry includes a surface with a sharp corner turning into the flow to induce an oblique shock, as shown in Figure 7. The angle of the wall is $\theta = 15^\circ$. The governing equations

$$(8.3) \quad \begin{aligned} \frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \vec{u}) &= 0, \\ \frac{\partial (\rho \vec{u})}{\partial t} + \nabla \cdot (\rho \vec{u} \vec{u}^T) + \nabla P &= 0, \text{ and} \\ \frac{\partial (\rho E)}{\partial t} + \nabla \cdot (\rho H \vec{u}) &= 0 \end{aligned}$$

are the inviscid Euler equations (enforcing conservation of mass, momentum, and energy), with H as the total enthalpy and E as the total internal energy. The equations are closed using the ideal gas law $P = \rho RT/M$.

The inflow boundary conditions are fixed as constant values. A symmetry boundary condition is prescribed along the bottom surface and open flow boundary conditions are prescribed along the top and right (outflow) surfaces. The problem is initialized with the inflow conditions and pseudo-timestepping is used to achieve steady-state. The inflow pressure and temperature are chosen to be $p_1 = 101352.93$ and $T_1 = 288.89$, respectively. The free stream Mach number M_1 is varied across training and testing scenarios. A shock forms at angle β to the flow where the wall turns sharply into the flow at the corner. θ and β are related to M_1 by the $\theta - \beta - M$ equation:

$$(8.4) \quad \tan \theta = 2 \cot \beta \frac{M_1^2 \sin^2 \beta - 1}{M_1^2 (\gamma + \cos 2\beta) + 2}.$$

Once β is computed from (8.4), the downstream conditions can be computed from the following equations:

$$(8.5) \quad \begin{aligned} \frac{p_2}{p_1} &= 1 + \frac{2\gamma}{\gamma + 1} (M_1^2 \sin^2 \beta - 1), \\ \frac{\rho_2}{\rho_1} &= \frac{(\gamma + 1) M_1^2 \sin^2 \beta}{(\gamma - 1) M_1^2 \sin^2 \beta + 2}, \\ \frac{T_2}{T_1} &= \frac{p_2 \rho_1}{p_1 \rho_2}, \text{ and} \\ M_2 &= \frac{1}{\sin(\beta - \theta)} \left(\frac{1 + \frac{\gamma-1}{2} M_1^2 \sin^2 \beta}{\gamma M_1^2 \sin^2 \beta - \frac{\gamma-1}{2}} \right)^{1/2}. \end{aligned}$$

The conditions ahead of and behind the shock are then given by $\{M_1, \rho_1, p_1, T_1\}$ and $\{M_2, \rho_2, p_2, T_2\}$ respectively. The local sound speed is given by $c = \sqrt{\gamma p/\rho}$. From the sound speed, the air speed $V = Mc = M\sqrt{\gamma p/\rho}$ can be computed. Therefore, the velocity field is $\mathbf{v} = (u_1, 0)$ ahead of the shock and $\mathbf{v} = (u_2 \cos(\theta), u_2 \sin(\theta))$ behind the shock, where $u_i = M_i \sqrt{\gamma p_i/\rho_i}$ and $i = 1, 2$.

The domain is discretized using 12800 quad elements and training data is generated for $M_1 \in \{2, 2.1, 2.2, \dots, 5.9, 6.0\}$. Separate test data is generated for $M_1 \in \{2.25, 3.25, 4.25, 5.25\}$.

Full results are tabulated in [Table 4](#) in [Appendix A](#). The 41 datapoints used for training in this case are considerably fewer than those used for the other experiments, owing to the fact that significantly less information is harvested from a steady-state solution than from a transient solution. The data could theoretically be augmented by inclusion of pseudo-timestepping data generated during the solution process, but doing so proved unnecessary in this case. However, only having 41 training points restricts the maximum latent dimension of the POD reconstruction to no more than 41. For this problem, only the density field is reconstructed since the spatial distribution is the same for all solution variables.

Most mixing operator sets performed comparably for this experiment, although we once again observe significant performance degradation for $m < 3$. The autoencoder outperformed the POD reconstruction over the full range of Mach numbers considered for $d_{latent} < 8$. The POD reconstruction includes large oscillations that pollute the solution both upwind and downwind of the shock. These oscillations persist (especially for lower Mach numbers) until a relatively large number of modes are retained ($d_{latent} > 16$). In contrast, the autoencoder reconstruction error is localized

around the shock front. For the density reconstruction pictured in [Figure 8](#), the amplitude of these errors increases with increasing M_1 , but only at a rate commensurate with the increase of ρ_2 . The autoencoder is able to reconstruct the field data with a very small latent space. Little performance degradation is seen, even for $d_{latent} = 1$, in which case the autoencoder has precisely learned the solution manifold for this problem.

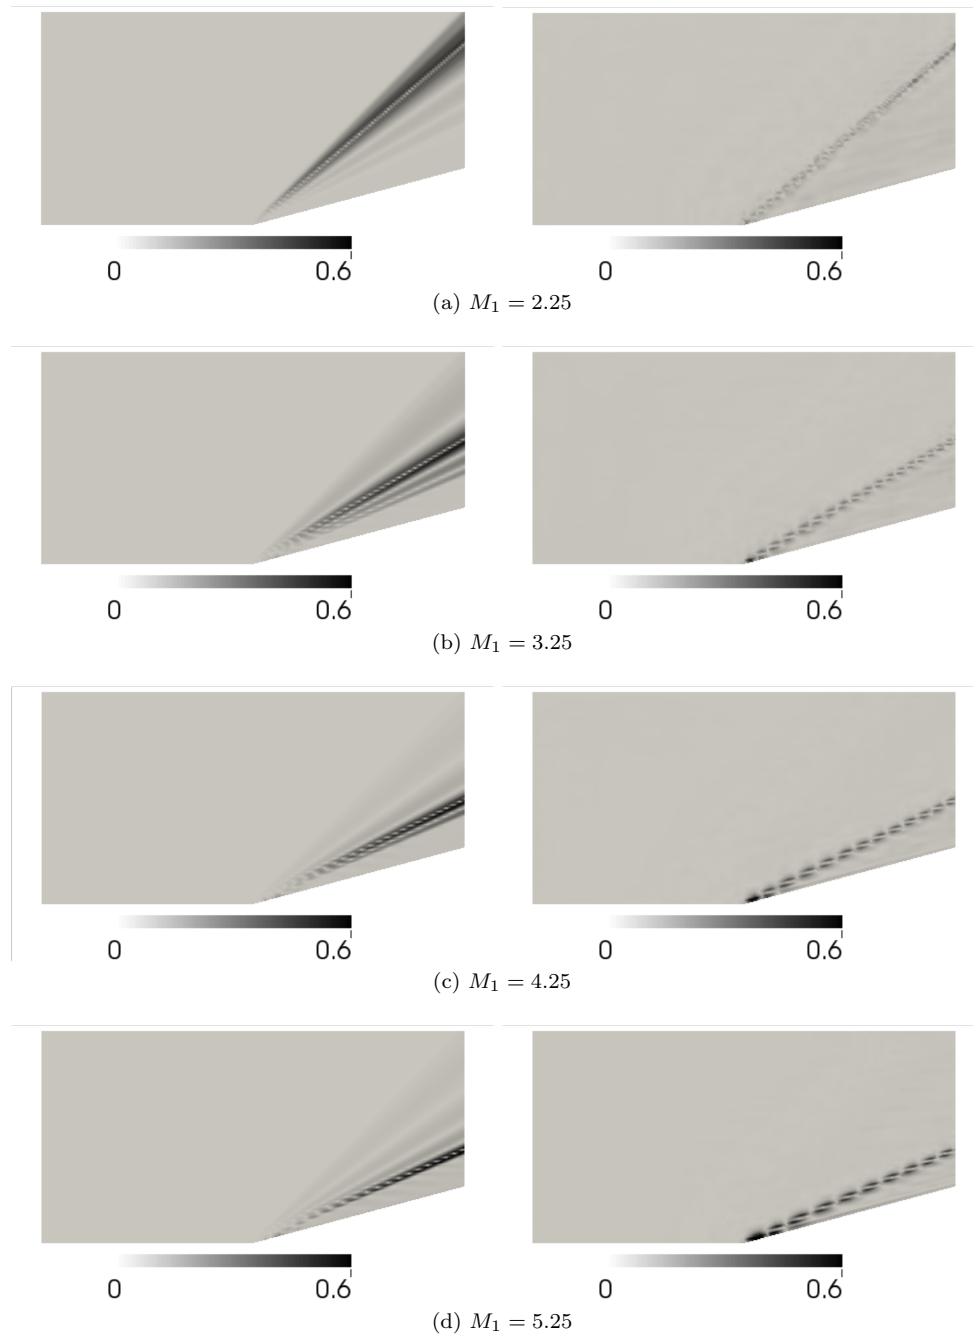


Fig. 8: Absolute error in the density reconstruction for the 4 test points using POD (left) and our autoencoder (right) with equivalent latent space dimension ($d_{latent} = 4$).

9. Conclusions. We combined the traditional CNN’s ability to learn physically relevant filters with the GCN’s applicability to handle data defined on unstructured meshes, allowing us to accurately reconstruct the nonlinear solution manifolds of PDEs, which had previously proven challenging for linear manifold methods. Our approach achieves an order of magnitude improvement in accuracy relative to linear methods for the 3 selected exemplar problems. The accuracy is comparable to that achieved by a traditional CNN autoencoder for the transient advection-diffusion problem for which both are applicable. For the other 2 example problems, the traditional CNN is unable to deal with the lack of structure in the spatial discretization, but our proposed approach continues to achieve similar gains in accuracy relative to POD enabling low-cost surrogate models of systems that previously required a much higher dimensional latent space. Additionally, our networks consistently outperform the traditional GCN which is expressed as a special case of our proposed architecture. In terms of computational cost, our method is comparable to traditional convolutional autoencoders in terms of FLOPS although further implementation improvements are required to be truly competitive in terms of wall clock time.

The proposed networks represent an extension of CNNs to the processing of data represented using unstructured spatial discretizations. Doing so significantly increases the applicability of CNNs, due to the resulting support for arbitrarily complex geometries. The only prerequisite is the ability to compute the required differential and interpolation operators, which should be readily available for any discretization used for the solution of PDEs (including meshless methods). In our implementation, these operators are precomputed, but they could also be computed at run time, which would be useful if training/test data include a variety of spatial discretizations, but would come with additional computational overhead. The extension of the powerful CNN technique to these types of data is important, not only in that it enables low-dimensional nonlinear surrogate models on arbitrarily complex geometries, but also that it enables other tasks requiring low-dimensional solution manifolds of computational data (e.g. synthesizing computational and experimental data, data compression, feature extraction/visualization, anomaly detection) for complex systems.

The reported accuracy improvements are restricted to the datasets considered. The effects of the various hyperparameters may differ depending on the physics, geometry, solution parameter values, or the training data collected. The choice of mixing operators is shown to have limited impact on the ultimate accuracy of the method, although highly restrictive sets ($m < 3$, which includes the traditional GCN) exhibited markedly worse accuracy. For the 3 different experiments, different mixing operator sets were shown to be advantageous. For the advection-diffusion experiment, operator set \mathcal{M}_e was generally best; for the laminar flow experiment, operator set \mathcal{M}_d was generally best; and for the inviscid Euler experiment, operator set \mathcal{M}_b was generally best. This difference indicates a problem-dependent nature to the optimal choice of mixing operators. Additionally, more expressive operator sets (larger m) are observed to usually (although not always) result in easier training. This is consistent with other results in the machine learning literature [14]. Different sets of mixing operators (and even different operators entirely) may prove advantageous for other applications. When approaching a new application, we suggest attempting to use operator set \mathcal{M}_b first and then adding in additional mixing operators if needed.

For the examples considered, the solutions evolve along low-dimensional manifolds, which are able to be successfully represented with $d_{latent} \leq 4$. Further increasing the latent space dimension did not generally result in increased accuracy for any of the autoencoder approaches (including the traditional CNN). Larger latent spaces

may be required for different datasets, but it is expected that d_{latent} should not be required to exceed the dimensionality of the underlying solution manifold by much, if at all, which is consistent with prior findings on the use of deep autoencoders for physics applications [69, 70]. A similar network architecture was used for all 3 experiments and minimal effort was made to optimize it. It is expected that further accuracy improvements are possible if the architecture and training procedure are tailored to a particular dataset.

All of the datasets examined here were 2D. However, as noted in [section 4](#), the extension to 3D is trivial. Additionally, the transition from 2D to 3D when using our proposed mixing operator sets is significantly less costly than for traditional CNNs, corresponding to an increase in the number of learnable parameters per convolutional layer of between 25% and 43% depending on the mixing operator set relative to the 200% required for traditional CNNs.

Appendix A. Tables of Loss Values. Here, we include detailed summaries of the accuracy achieved in all three experiments, for every value of each hyperparameter considered. For the transient-advection diffusion experiment, the solution variable $\alpha \in [0, 1]$. For the laminar flow experiment, the variables used for training and included in the computation of the loss function are the horizontal velocity $u_x \in (-0.4, 1.4)$, the vertical velocity $u_y \in (-0.75, 0.75)$, and the pressure $p \in (-0.75, 0.6)$. For the inviscid Euler experiment, only the density $\rho \in (1.2, 6.2)$ is included since all of the variables share the same spatial distribution.

Table 2: Loss performance for transient advection-diffusion experiment. The ‘loss’ refers to the MSE loss of the solution averaged across solution variables and summed across snapshots. For convenience, the most accurate set of mixing operators for each test and for each value of d_{latent} is highlighted. In those cases where POD was more accurate than the autoencoder, those results are also highlighted.

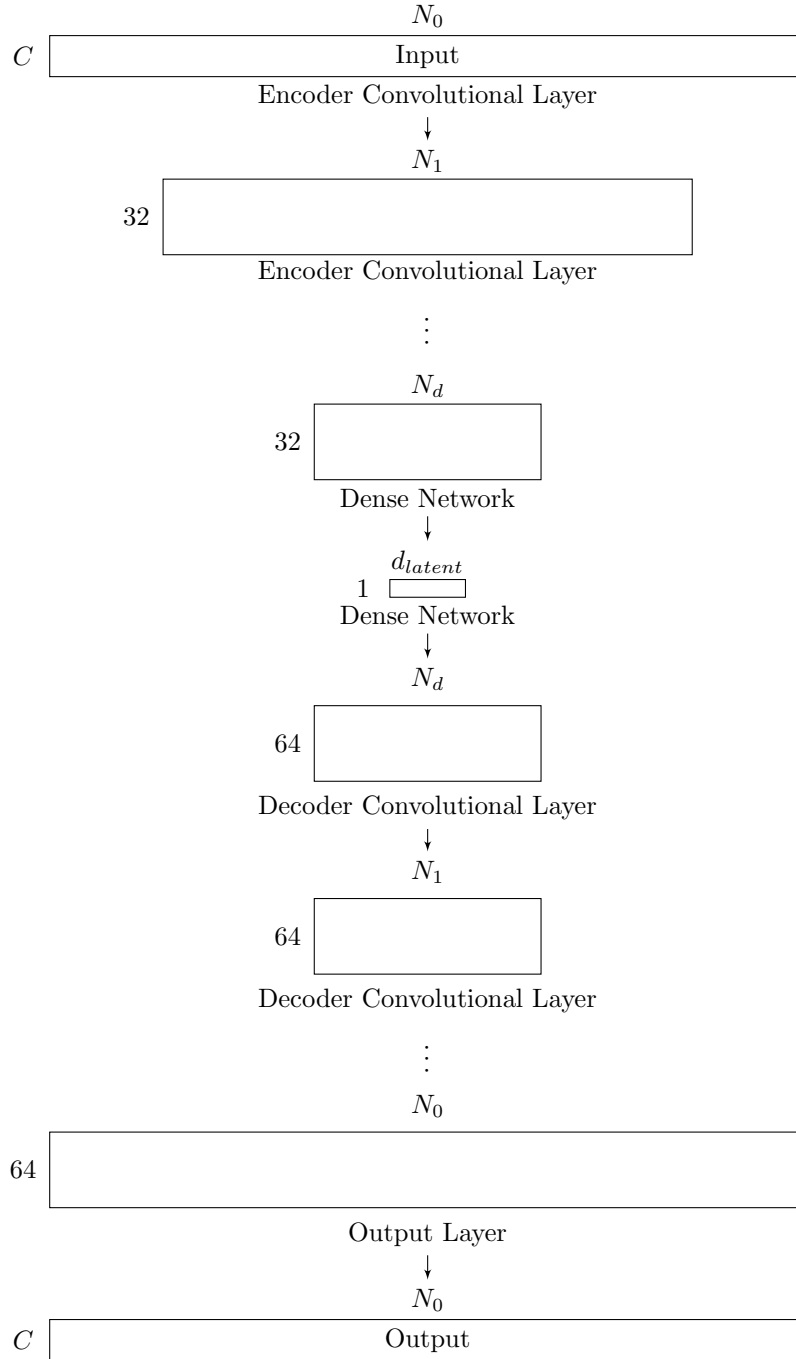
	Mixing Set	Training Loss	Test Loss			
			$\theta = 45^\circ$	$\theta = 135^\circ$	$\theta = 225^\circ$	$\theta = 315^\circ$
$d_{latent} = 2$	(a)	1.246e-5	7.828e-6	3.759e-6	1.868e-5	1.036e-3
	(b)	9.524e-5	4.199e-2	1.208e-2	1.765e-4	1.711e-2
	(c)	1.596e-5	5.159e-6	1.056e-2	6.170e-4	2.219e-5
	(d)	1.261e-5	3.217e-2	8.677e-6	4.880e-6	1.623e-4
	(e)	1.335e-5	7.714e-6	8.158e-6	2.840e-3	1.048e-2
$d_{latent} = 4$	(a)	4.116e-5	7.339e-6	8.486e-6	1.102e-5	7.708e-6
	(b)	5.994e-5	1.286e-5	1.027e-5	9.766e-6	1.067e-5
	(c)	7.141e-5	1.906e-5	2.149e-5	1.269e-5	1.577e-5
	(d)	4.083e-5	1.425e-5	1.046e-5	7.233e-6	2.288e-5
	(e)	2.273e-5	7.048e-6	5.419e-6	4.553e-6	4.967e-6
	(f)	2.221e-4	3.289e-5	3.338e-5	3.274e-5	3.458e-5
	(g)	4.592e-4	6.978e-5	6.587e-5	6.787e-5	6.627e-5
	(h)	9.388e-1	1.345e-1	1.345e-1	1.345e-1	1.345e-1
	(i)	1.645e-3	2.442e-4	2.352e-4	3.084e-4	2.285e-4
	POD	1.198e-2	1.456e-3	1.456e-3	1.456e-3	1.456e-3
CNN	2.770e-5	3.523e-5	3.456e-5	3.599e-5	3.608e-5	
$d_{latent} = 16$	(a)	2.305e-4	3.982e-5	4.294e-5	3.205e-5	4.031e-5
	(b)	7.178e-5	1.164e-5	1.169e-5	1.144e-5	1.181e-5
	(c)	9.102e-6	2.828e-6	9.565e-6	4.527e-6	3.649e-6
	(d)	4.203e-5	1.194e-5	1.057e-5	1.095e-5	1.050e-5
	(e)	9.437e-6	1.005e-5	7.028e-6	5.801e-6	9.888e-6
	(f)	9.388e-1	1.345e-1	1.345e-1	1.345e-1	1.345e-1
	(g)	5.791e-1	8.338e-2	8.239e-2	8.353e-2	8.181e-2
	(h)	9.388e-1	1.345e-1	1.345e-1	1.345e-1	1.345e-1
	(i)	8.142e-5	1.378e-5	1.781e-5	2.017e-5	1.397e-5
	POD	7.096e-5	6.973e-6	6.973e-6	6.973e-6	6.973e-6
CNN	3.817e-5	4.797e-5	4.589e-5	5.064e-5	4.728e-5	
$d_{latent} = 64$	(a)	1.203e-5	4.078e-6	7.811e-6	5.454e-6	6.364e-6
	(b)	7.620e-5	1.188e-5	1.505e-5	1.392e-5	1.410e-5
	(c)	2.376e-5	3.756e-6	1.024e-5	9.465e-6	3.630e-5
	(d)	2.064e-4	3.063e-5	3.011e-5	3.036e-5	3.010e-5
	(e)	1.995e-5	3.676e-6	4.691e-6	4.461e-6	5.268e-6
	(f)	7.820e-2	1.120e-2	1.122e-2	1.120e-2	1.119e-2
	(g)	1.663e-4	2.549e-5	2.357e-5	3.691e-5	2.415e-5
	(h)	9.388e-2	1.344e-1	1.345e-1	1.345e-1	1.345e-1
	(i)	4.366e-2	1.824e-3	1.003e-2	1.623e-3	1.122e-2
	POD	1.002e-9	1.277e-10	1.274e-10	1.274e-10	1.275e-10
CNN	2.872e-5	3.978e-5	3.923e-5	3.937e-5	3.843e-5	

Table 3: L^2 norm performance for laminar flow experiment. Reported value are the relative L^2 error norm of the reconstructed velocity. All results converged in fewer than 2,000 epochs.

	Mixing Set	Training Loss	Test Loss		
			$Re = 150$	$Re = 250$	$Re = 350$
$d_{latent} = 2$	(a)	1.471e-02	4.783e-02	2.486e-02	1.994e-02
	(b)	1.304e-02	5.602e-02	2.511e-02	1.825e-02
	(c)	1.509e-02	5.565e-02	2.571e-02	1.936e-02
	(d)	2.439e-02	3.677e-02	3.452e-02	2.583e-02
	(e)	2.100e-02	3.764e-02	2.634e-02	2.240e-02
	(f)	2.111e-02	3.879e-02	2.756e-02	2.317e-02
	(g)	2.163e-02	3.674e-02	2.883e-02	2.483e-02
	(h)	1.144e-01	1.339e-01	1.444e-01	1.502e-01
	(i)	1.568e-02	4.924e-02	2.423e-02	2.341e-02
	(POD)	8.933e-02	7.426e-02	6.928e-02	7.513e-02
$d_{latent} = 4$	(a)	6.644e-03	1.237e-02	1.202e-02	1.008e-02
	(b)	6.021e-03	1.010e-02	9.335e-03	8.804e-03
	(c)	6.476e-03	1.337e-02	1.109e-02	1.021e-02
	(d)	6.481e-03	1.165e-02	1.025e-02	9.434e-03
	(e)	7.580e-03	1.082e-02	1.078e-02	1.039e-02
	(f)	8.170e-03	1.322e-02	1.214e-02	1.120e-02
	(g)	7.985e-03	1.315e-02	1.259e-02	1.114e-02
	(h)	2.339e-02	2.913e-02	2.633e-02	2.551e-02
	(i)	8.643e-03	1.322e-02	1.223e-02	1.111e-02
	(POD)	4.833e-02	4.633e-02	4.080e-02	4.751e-02
$d_{latent} = 8$	(a)	5.964e-03	1.039e-02	9.639e-03	9.110e-03
	(b)	5.805e-03	9.268e-03	8.879e-03	8.399e-03
	(c)	5.536e-03	9.206e-03	9.070e-03	8.606e-03
	(d)	8.607e-03	1.223e-02	1.187e-02	1.167e-02
	(e)	9.566e-03	1.236e-02	1.233e-02	1.242e-02
	(f)	7.375e-03	1.192e-02	1.126e-02	1.036e-02
	(g)	7.186e-03	1.205e-02	1.170e-02	1.043e-02
	(h)	2.084e-02	2.551e-02	2.430e-02	2.402e-02
	(i)	7.356e-03	1.162e-02	1.096e-02	1.032e-02
	(POD)	2.742e-02	3.224e-02	3.446e-02	3.688e-02
$d_{latent} = 16$	(a)	5.871e-03	1.061e-02	9.553e-03	8.961e-03
	(b)	6.314e-03	1.063e-02	9.855e-03	9.084e-03
	(c)	5.315e-03	8.677e-03	9.069e-03	8.498e-03
	(d)	8.730e-03	1.232e-02	1.154e-02	1.130e-02
	(e)	8.388e-03	1.167e-02	1.137e-02	1.123e-02
	(f)	7.362e-03	1.199e-02	1.143e-02	1.071e-02
	(g)	7.561e-03	1.190e-02	1.125e-02	1.087e-02
	(h)	1.142e-01	1.338e-01	1.444e-01	1.502e-01
	(i)	7.361e-03	1.097e-02	1.035e-02	1.011e-02
	(POD)	1.889e-02	2.021e-02	1.926e-02	2.075e-02

Table 4: L^2 performance for Euler experiment. Reported values are the relative L^2 error norm of the reconstructed density. All results converged in fewer than 10,000 epochs.

	Mixing Set	Training Loss	Test Loss			
			$M_1 = 2.25$	$M_1 = 3.25$	$M_1 = 4.25$	$M_1 = 5.25$
$d_{latent} = 1$	(a)	1.593e-02	4.002e-03	5.916e-03	9.158e-03	1.281e-02
	(b)	1.551e-02	3.483e-03	5.739e-03	9.087e-03	1.279e-02
	(c)	1.560e-02	3.579e-03	5.762e-03	9.107e-03	1.278e-02
	(d)	1.587e-02	3.810e-03	5.896e-03	9.171e-03	1.284e-02
	(e)	1.841e-02	5.029e-03	6.970e-03	1.060e-02	1.441e-02
	(f)	1.961e-02	6.103e-03	7.166e-03	1.059e-02	1.489e-02
	(g)	1.231e-01	8.232e-02	5.029e-02	4.711e-02	7.018e-02
	(h)	1.230e-01	8.227e-02	5.027e-02	4.704e-02	7.021e-02
	(i)	1.588e-02	3.654e-03	5.975e-03	9.300e-03	1.299e-02
	POD	7.352e-02	3.872e-02	4.478e-02	3.395e-02	2.817e-02
$d_{latent} = 2$	(a)	1.644e-02	4.344e-03	6.165e-03	9.344e-03	1.291e-02
	(b)	1.549e-02	3.443e-03	5.746e-03	9.077e-03	1.277e-02
	(c)	1.585e-02	3.846e-03	5.876e-03	9.185e-03	1.283e-02
	(d)	1.686e-02	4.917e-03	6.255e-03	9.402e-03	1.299e-02
	(e)	1.566e-02	3.566e-03	5.777e-03	9.136e-03	1.281e-02
	(f)	1.912e-02	6.923e-03	7.482e-03	1.023e-02	1.369e-02
	(g)	1.230e-01	8.229e-02	5.033e-02	4.700e-02	7.015e-02
	(h)	2.190e-02	1.160e-02	9.602e-03	1.069e-02	1.353e-02
	(i)	1.599e-02	3.875e-03	5.972e-03	9.255e-03	1.297e-02
	POD	5.281e-02	3.093e-02	2.685e-02	2.867e-02	1.395e-02
$d_{latent} = 4$	(a)	1.594e-02	4.030e-03	5.850e-03	9.160e-03	1.282e-02
	(b)	1.551e-02	3.461e-03	5.754e-03	9.081e-03	1.278e-02
	(c)	1.562e-02	3.567e-03	5.783e-03	9.120e-03	1.282e-02
	(d)	1.607e-02	4.085e-03	5.953e-03	9.213e-03	1.288e-02
	(e)	1.603e-02	3.792e-03	5.935e-03	9.202e-03	1.292e-02
	(f)	1.578e-02	3.503e-03	5.856e-03	9.253e-03	1.305e-02
	(g)	1.613e-02	3.757e-03	5.891e-03	9.358e-03	1.327e-02
	(h)	1.230e-01	8.236e-02	5.025e-02	4.699e-02	7.014e-02
	(i)	1.561e-02	3.502e-03	5.811e-03	9.123e-03	1.283e-02
	POD	3.297e-02	2.593e-02	1.698e-02	1.257e-02	8.830e-03
$d_{latent} = 8$	(a)	1.562e-02	3.605e-03	5.773e-03	9.097e-03	1.279e-02
	(b)	1.546e-02	3.369e-03	5.730e-03	9.087e-03	1.279e-02
	(c)	1.561e-02	3.589e-03	5.765e-03	9.110e-03	1.280e-02
	(d)	1.641e-02	4.688e-03	6.022e-03	9.201e-03	1.287e-02
	(e)	1.565e-02	3.577e-03	5.809e-03	9.106e-03	1.285e-02
	(f)	1.567e-02	3.555e-03	5.864e-03	9.153e-03	1.289e-02
	(g)	1.647e-02	3.970e-03	6.119e-03	9.511e-03	1.345e-02
	(h)	1.230e-01	8.229e-02	5.022e-02	4.704e-02	7.021e-02
	(i)	1.568e-02	3.604e-03	5.844e-03	9.152e-03	1.284e-02
	POD	1.690e-02	1.389e-02	8.644e-03	5.928e-03	6.802e-03

Appendix B. Network Architecture.

Acknowledgments. Supported by the Laboratory Directed Research and Development program at Sandia National Laboratories, a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International Inc. for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

REFERENCES

- [1] J. ABRAS AND N. S. HARIHARAN, *Application of Machine Learning to Automate Flow-Physics Identification in Computed Solutions: Hover Rotor Wake Vortex Identification and Breakdown Analysis*, <https://doi.org/10.2514/6.2021-0474>, <https://arc.aiaa.org/doi/abs/10.2514/6.2021-0474>, <https://arxiv.org/abs/https://arc.aiaa.org/doi/pdf/10.2514/6.2021-0474>.
- [2] S. E. AHMED, S. M. RAHMAN, O. SAN, A. RASHEED, AND I. M. NAVON, *Memory embedded non-intrusive reduced order modeling of non-ergodic flows*, *Physics of Fluids*, 31 (2019), p. 126602.
- [3] A. AITKEN, C. LEDIG, L. THEIS, J. CABALLERO, Z. WANG, AND W. SHI, *Checkerboard artifact free sub-pixel convolution: A note on sub-pixel convolution, resize convolution and convolution resize*, arXiv preprint arXiv:1707.02937, (2017).
- [4] S. ANWAR, K. HWANG, AND W. SUNG, *Structured pruning of deep convolutional neural networks*, *ACM Journal on Emerging Technologies in Computing Systems (JETC)*, 13 (2017), pp. 1–18.
- [5] O. AZENCOT, N. B. ERICHSON, V. LIN, AND M. W. MAHONEY, *Forecasting sequential data using consistent koopman autoencoders*, 2020, <https://arxiv.org/abs/2003.02236>.
- [6] C. BEATTIE AND S. GUGERCIN, *Model Reduction by Rational Interpolation*, ch. 7, pp. 297–334, <https://doi.org/10.1137/1.9781611974829.ch7>, <https://epubs.siam.org/doi/abs/10.1137/1.9781611974829.ch7>, <https://arxiv.org/abs/https://epubs.siam.org/doi/pdf/10.1137/1.9781611974829.ch7>.
- [7] M. BELKIN AND P. NIYOGI, *Laplacian eigenmaps and spectral techniques for embedding and clustering*, in *Advances in neural information processing systems*, 2002, pp. 585–591.
- [8] M. BELKIN AND P. NIYOGI, *Laplacian eigenmaps for dimensionality reduction and data representation*, *Neural computation*, 15 (2003), pp. 1373–1396.
- [9] P. BENNER, S. GUGERCIN, AND K. WILLCOX, *A survey of projection-based model reduction methods for parametric dynamical systems*, 2015.
- [10] P. BENNER, P. KÜRSCHNER, AND J. SAAK, *Frequency-limited balanced truncation with low-rank approximations*, *SIAM Journal on Scientific Computing*, 38 (2016), pp. A471–A499, <https://doi.org/10.1137/15M1030911>, <https://doi.org/10.1137/15M1030911>, <https://arxiv.org/abs/https://doi.org/10.1137/15M1030911>.
- [11] F. M. BIANCHI, D. GRATAROLA, AND C. ALIPPI, *Spectral clustering with graph neural networks for graph pooling*, 2020, <https://arxiv.org/abs/1907.00481>.
- [12] D. BOSCAINI, J. MASCI, E. RODOLÀ, AND M. M. BRONSTEIN, *Learning shape correspondence with anisotropic convolutional neural networks*, 2016, <https://arxiv.org/abs/1605.06437>.
- [13] J. BRUNA, W. ZAREMBA, A. SZLAM, AND Y. LECUN, *Spectral networks and locally connected networks on graphs*, (2013), <https://arxiv.org/abs/1312.6203>.
- [14] R.-D. BUHAI, Y. HALPERN, Y. KIM, A. RISTESKI, AND D. SONTAG, *Empirical study of the benefits of overparameterization in learning latent variable models*, arXiv preprint arXiv:1907.00030, (2019).
- [15] C. CANGEA, P. VELIČKOVIĆ, N. JOVANOVIĆ, T. KIPF, AND P. LIÒ, *Towards sparse hierarchical graph classifiers*, 2018, <https://arxiv.org/abs/1811.01287>.
- [16] K. CARLBERG, C. BOU-MOSLEH, AND C. FARHAT, *Efficient non-linear model reduction via a least-squares petrov–galerkin projection and compressive tensor approximations*, *International Journal for Numerical Methods in Engineering*, 86, pp. 155–181, <https://doi.org/10.1002/nme.3050>, <https://onlinelibrary.wiley.com/doi/abs/10.1002/nme.3050>, <https://arxiv.org/abs/https://onlinelibrary.wiley.com/doi/pdf/10.1002/nme.3050>.
- [17] K. CARLBERG, C. FARHAT, J. CORTIAL, AND D. AMSALLEM, *The gnat method for non-linear model reduction: Effective implementation and application to computational fluid dynamics and turbulent flows*, *Journal of Computational Physics*, 242 (2013), pp. 623 – 647, <https://doi.org/https://doi.org/10.1016/j.jcp.2013.02.028>, <http://www.sciencedirect.com/science/article/pii/S0021999113001472>.
- [18] K. T. CARLBERG, A. JAMESON, M. J. KOCHENDERFER, J. MORTON, L. PENG, AND F. D. WITHERDEN, *Recovering missing cfd data for high-order discretizations using deep neural networks and dynamics learning*, *Journal of Computational Physics*, 395 (2019), pp. 105 – 124, <https://doi.org/https://doi.org/10.1016/j.jcp.2019.05.041>, <http://www.sciencedirect.com/science/article/pii/S0021999119303857>.
- [19] C. CHEN, W. YE, Y. ZUO, C. ZHENG, AND S. P. ONG, *Graph networks as a universal machine learning framework for molecules and crystals*, *Chemistry of Materials*, 31 (2019), pp. 3564–3572.
- [20] M. CHEN, X. SHI, Y. ZHANG, D. WU, AND M. GUIZANI, *Deep features learning for medical*

- image analysis with convolutional autoencoder neural network*, IEEE Transactions on Big Data, (2017).
- [21] E. CHIAVAZZO, C. W. GEAR, C. J. DSILVA, N. RABIN, AND I. G. KEVREKIDIS, *Reduced models in chemical kinetics via nonlinear data-mining*, 2013, <https://arxiv.org/abs/1307.6849>.
- [22] E. CHIAVAZZO, C. W. GEAR, B. E. SONDAY, AND I. G. KEVREKIDIS, *A manifold learning approach to model reduction in combustion*, 2018.
- [23] F. CHINESTA, P. LADEVÈZE, AND E. CUETO, *A short review on model order reduction based on proper generalized decomposition*, Archives of Computational Methods in Engineering, 18 (2011), pp. 395–404.
- [24] R. R. COIFMAN AND S. LAFON, *Diffusion maps*, Applied and computational harmonic analysis, 21 (2006), pp. 5–30.
- [25] V. DAS, A. POLLACK, U. WOLLNER, AND T. MUKERJI, *Convolutional neural network for seismic impedance inversion*, 08 2018, pp. 2071–2075, <https://doi.org/10.1190/segam2018-2994378.1>.
- [26] I. S. DHILLON, Y. GUAN, AND B. KULIS, *Weighted graph cuts without eigenvectors a multilevel approach*, IEEE transactions on pattern analysis and machine intelligence, 29 (2007), pp. 1944–1957.
- [27] M. DIAS RIBEIRO, A. REHMAN, S. AHMED, AND A. DENGEL, *Deepecfd: Efficient steady-state laminar flow approximation with deep convolutional neural networks*, arXiv, (2020), pp. arXiv–2004.
- [28] F. DIEHL, *Edge contraction pooling for graph neural networks*, 2019, <https://arxiv.org/abs/1905.10990>.
- [29] C. DOS SANTOS AND M. GATTI, *Deep convolutional neural networks for sentiment analysis of short texts*, in Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers, Dublin, Ireland, Aug. 2014, Dublin City University and Association for Computational Linguistics, pp. 69–78, <https://www.aclweb.org/anthology/C14-1008>.
- [30] D. K. DUVENAUD, D. MACLAURIN, J. IPARRAGUIRRE, R. BOMBARELL, T. HIRZEL, A. ASPURUGUZIK, AND R. P. ADAMS, *Convolutional networks on graphs for learning molecular fingerprints*, in Advances in neural information processing systems, 2015, pp. 2224–2232.
- [31] F. EBERT AND T. STYKEL, *Rational interpolation, minimal realization and model reduction*. 2007.
- [32] M. ERDOGDU, M. BAYATI, AND L. DICKER, *Scaled least squares estimator for glms in large-scale problems*, Advances in Neural Information Processing Systems, (2016), pp. 3332–3340. 30th Annual Conference on Neural Information Processing Systems, NIPS 2016 ; Conference date: 05-12-2016 Through 10-12-2016.
- [33] A. B. FARIMANI, J. GOMES, AND V. S. PANDE, *Deep learning the physics of transport phenomena*, CoRR, abs/1709.02432 (2017), <http://arxiv.org/abs/1709.02432>, <https://arxiv.org/abs/1709.02432>.
- [34] E. N. FEINBERG, D. SUR, Z. WU, B. E. HUSIC, H. MAI, Y. LI, S. SUN, J. YANG, B. RAMSUNDAR, AND V. S. PANDE, *Potentialnet for molecular property prediction*, ACS central science, 4 (2018), pp. 1520–1530.
- [35] A. FOUT, J. BYRD, B. SHARIAT, AND A. BEN-HUR, *Protein interface prediction using graph convolutional networks*, in Advances in neural information processing systems, 2017, pp. 6530–6539.
- [36] S. FRESCA, L. DEDE, AND A. MANZONI, *A comprehensive deep learning-based approach to reduced order modeling of nonlinear time-dependent parametrized pdes*, 2020, <https://arxiv.org/abs/2001.04001>.
- [37] J. GILMER, S. S. SCHOENHOLZ, P. F. RILEY, O. VINYALS, AND G. E. DAHL, *Neural message passing for quantum chemistry*, vol. 70 of Proceedings of Machine Learning Research, International Convention Centre, Sydney, Australia, 06–11 Aug 2017, PMLR, pp. 1263–1272, <http://proceedings.mlr.press/v70/gilmer17a.html>.
- [38] F. J. GONZALEZ AND M. BALAJEWICZ, *Deep convolutional recurrent autoencoders for learning low-dimensional feature dynamics of fluid systems*, arXiv preprint arXiv:1808.01346, (2018).
- [39] C. GRÄSSLE, M. HINZE, AND S. VOLKWEIN, *Model order reduction by proper orthogonal decomposition*, 2019, <https://arxiv.org/abs/1906.05188>.
- [40] S. GUGERCIN AND A. C. ANTOULAS, *A survey of model reduction by balanced truncation and some new results*, International Journal of Control, 77 (2004), pp. 748–766, <https://doi.org/10.1080/00207170410001713448>, <https://doi.org/10.1080/00207170410001713448>, <https://arxiv.org/abs/https://doi.org/10.1080/00207170410001713448>.
- [41] X. GUO, W. LI, AND F. IORIO, *Convolutional neural networks for steady flow approxima-*

- tion, in Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '16, New York, NY, USA, 2016, Association for Computing Machinery, p. 481–490, <https://doi.org/10.1145/2939672.2939738>.
- [42] T. HAMAGUCHI, H. OIWA, M. SHIMBO, AND Y. MATSUMOTO, *Knowledge base completion with out-of-knowledge-base entities: A graph neural network approach*, Transactions of the Japanese Society for Artificial Intelligence, 33 (2018), p. 1–10, <https://doi.org/10.1527/tjsai.f-h72>, <http://dx.doi.org/10.1527/tjsai.F-H72>.
- [43] W. L. HAMILTON, R. YING, AND J. LESKOVEC, *Inductive representation learning on large graphs*, 2017, <https://arxiv.org/abs/1706.02216>.
- [44] S. HAN, H. MAO, AND W. J. DALLY, *Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding*, arXiv preprint arXiv:1510.00149, (2015).
- [45] K. HASEGAWA, K. FUKAMI, T. MURATA, AND K. FUKAGATA, *Machine-learning-based reduced-order modeling for unsteady flows around bluff bodies of various shapes*, Theoretical and Computational Fluid Dynamics, 34 (2020), pp. 367–383.
- [46] M. HEINKENSCHLOSS, T. REIS, AND A. ANTOULAS, *Balanced truncation model reduction for systems with inhomogeneous initial conditions*, Automatica, 47 (2011), pp. 559 – 564, <https://doi.org/https://doi.org/10.1016/j.automatica.2010.12.002>, <http://www.sciencedirect.com/science/article/pii/S0005109810004905>.
- [47] M. HENAFF, J. BRUNA, AND Y. LECUN, *Deep convolutional networks on graph-structured data*, (2015), <https://arxiv.org/abs/1506.05163>.
- [48] I. HENRION, J. BREHMER, J. BRUNA, K. CHO, K. CRANMER, G. LOUPPE, AND G. ROCHETTE, *Neural message passing for jet physics*, (2017).
- [49] J. HERNÁNDEZ, M. CAICEDO, AND A. FERRER, *Dimensional hyper-reduction of nonlinear finite element models via empirical cubature*, Computer Methods in Applied Mechanics and Engineering, 313 (2017), pp. 687 – 722, <https://doi.org/https://doi.org/10.1016/j.cma.2016.10.022>, <http://www.sciencedirect.com/science/article/pii/S004578251631355X>.
- [50] C. HOANG, K. CHOWDHARY, K. LEE, AND J. RAY, *Projection-based model reduction of dynamical systems using space-time subspace and machine learning*, arXiv preprint arXiv:2102.03505, (2021).
- [51] D. HOLDEN, J. SAITO, T. KOMURA, AND T. JOYCE, *Learning motion manifolds with convolutional autoencoders*, in SIGGRAPH Asia 2015 Technical Briefs, 2015, pp. 1–4.
- [52] S. HONG AND S.-K. SONG, *Kick: Shift-n-overlap cascades of transposed convolutional layer for better autoencoding reconstruction on remote sensing imagery*, IEEE Access, 8 (2020), pp. 107244–107259.
- [53] B. HU, Z. LU, H. LI, AND Q. CHEN, *Convolutional neural network architectures for matching natural language sentences*, in Advances in Neural Information Processing Systems 27, Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, eds., Curran Associates, Inc., 2014, pp. 2042–2050, <http://papers.nips.cc/paper/5550-convolutional-neural-network-architectures-for-matching-natural-language-sentences.pdf>.
- [54] C. INNAMORATI, T. RITSCHER, T. WEYRICH, AND N. J. MITRA, *Learning on the edge: Explicit boundary handling in cnns*, arXiv preprint arXiv:1805.03106, (2018).
- [55] X. JIA, J. WILLARD, A. KARPATNE, J. READ, J. ZWART, M. STEINBACH, AND V. KUMAR, *Physics guided rnns for modeling dynamical systems: A case study in simulating lake temperature profiles*, in Proceedings of the 2019 SIAM International Conference on Data Mining, SIAM, 2019, pp. 558–566.
- [56] C. M. JIANG, J. HUANG, K. KASHINATH, PRABHAT, P. MARCUS, AND M. NIESSNER, *Spherical CNNs on unstructured grids*, in International Conference on Learning Representations, 2019, <https://openreview.net/forum?id=Bkl-43C9FQ>.
- [57] N. KALCHBRENNER, E. GREFFENSTETTE, AND P. BLUNSOM, *A convolutional neural network for modelling sentences*, CoRR, abs/1404.2188 (2014), <http://arxiv.org/abs/1404.2188>, <https://arxiv.org/abs/1404.2188>.
- [58] M. KARAMAD, R. MAGAR, Y. SHI, S. STAHROSTAMI, I. D. GATES, AND A. BARATI FARIMANI, *Orbital graph convolutional neural network for material property prediction*, Physical Review Materials, 4 (2020), <https://doi.org/10.1103/physrevmaterials.4.093801>, <http://dx.doi.org/10.1103/PhysRevMaterials.4.093801>.
- [59] S. KEARNES, K. MCCLOSKEY, M. BERNDL, V. PANDE, AND P. RILEY, *Molecular graph convolutions: moving beyond fingerprints*, Journal of computer-aided molecular design, 30 (2016), pp. 595–608.
- [60] G. KERSCHEN, J. CLAUDE GOLINVAL, A. F. VAKAKIS, AND L. A. BERGMAN, *The method of*

- proper orthogonal decomposition for dynamical characterization and order reduction of mechanical systems: an overview*, *Nonlinear Dynamics*, (2005), pp. 147–169.
- [61] J. KIM, J. KWON LEE, AND K. MU LEE, *Accurate image super-resolution using very deep convolutional networks*, in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [62] Y. KIM, Y. CHOI, D. WIDEMANN, AND T. ZOHDI, *A fast and accurate physics-informed neural network reduced order model with shallow masked autoencoder*, 2020, <https://arxiv.org/abs/2009.11990>.
- [63] T. N. KIPF AND M. WELLING, *Semi-supervised classification with graph convolutional networks*, 2016, <https://arxiv.org/abs/1609.02907>.
- [64] M. A. KRAMER, *Nonlinear principal component analysis using autoassociative neural networks*, *AIChE journal*, 37 (1991), pp. 233–243.
- [65] A. KRIZHEVSKY, I. SUTSKEVER, AND G. E. HINTON, *Imagenet classification with deep convolutional neural networks*, in *Advances in Neural Information Processing Systems 25*, F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, eds., Curran Associates, Inc., 2012, pp. 1097–1105, <http://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks.pdf>.
- [66] S. LAWRENCE, C. L. GILES, AH CHUNG TSOI, AND A. D. BACK, *Face recognition: a convolutional neural-network approach*, *IEEE Transactions on Neural Networks*, 8 (1997), pp. 98–113, <https://doi.org/10.1109/72.554195>.
- [67] Y. LECUN AND Y. BENGIO, *Convolutional networks for images, speech, and time-series*, MIT Press, 1995.
- [68] J. LEE, I. LEE, AND J. KANG, *Self-attention graph pooling*, 2019, <https://arxiv.org/abs/1904.08082>.
- [69] K. LEE AND K. CARLBERG, *Model reduction of dynamical systems on nonlinear manifolds using deep convolutional autoencoders*, 2018, <https://arxiv.org/abs/1812.08373>.
- [70] K. LEE AND K. CARLBERG, *Deep conservation: A latent dynamics model for exact satisfaction of physical conservation laws*, 2019, <https://arxiv.org/abs/1909.09754>.
- [71] H. LI, A. KADAV, I. DURDANOVIC, H. SAMET, AND H. P. GRAF, *Pruning filters for efficient convnets*, arXiv preprint arXiv:1608.08710, (2016).
- [72] X. LI, X. YAN, Q. GU, H. ZHOU, D. WU, AND J. XU, *Deepchemstable: chemical stability prediction with an attention-based graph convolution network*, *Journal of chemical information and modeling*, 59 (2019), pp. 1044–1049.
- [73] G. LIU, K. J. SHIH, T.-C. WANG, F. A. REDA, K. SAPRA, Z. YU, A. TAO, AND B. CATANZARO, *Partial convolution based padding*, arXiv preprint arXiv:1811.11718, (2018).
- [74] T. LIU, Y. LI, Q. JING, Y. XIE, AND D. ZHANG, *Supervised learning method for the physical field reconstruction in a nanofluid heat transfer problem*, *International Journal of Heat and Mass Transfer*, 165 (2021), p. 120684.
- [75] Y. LIU, Y. WANG, L. DENG, F. WANG, F. LIU, Y. LU, AND S. LI, *A novel in situ compression method for cfd data based on generative adversarial network*, *Journal of Visualization*, 22 (2019), pp. 95–108.
- [76] Z. LONG, Y. LU, AND B. DONG, *Pde-net 2.0: Learning pdes from data with a numeric-symbolic hybrid deep network*, *Journal of Computational Physics*, 399 (2019), p. 108925.
- [77] Z. LONG, Y. LU, X. MA, AND B. DONG, *Pde-net: Learning pdes from data*, in *International Conference on Machine Learning*, 2018, pp. 3208–3216.
- [78] I. LOSHCHELOV AND F. HUTTER, *Decoupled weight decay regularization*, 2019, <https://arxiv.org/abs/1711.05101>.
- [79] B. LUSCH, J. N. KUTZ, AND S. L. BRUNTON, *Deep learning for universal linear embeddings of nonlinear dynamics*, *Nature communications*, 9 (2018), pp. 1–10.
- [80] J. A. MARTÍNEZ, O. CERRI, M. SPIROPULU, J. VLMANT, AND M. PIERINI, *Pileup mitigation at the large hadron collider with graph neural networks*, *The European Physical Journal Plus*, 134 (2019), p. 333.
- [81] J. MASCI, U. MEIER, D. ÇİREŞAN, AND J. SCHMIDHUBER, *Stacked convolutional auto-encoders for hierarchical feature extraction*, in *International conference on artificial neural networks*, Springer, 2011, pp. 52–59.
- [82] R. MAULIK, B. LUSCH, AND P. BALAPRAKASH, *Reduced-order modeling of advection-dominated systems with recurrent neural networks and convolutional autoencoders*, arXiv preprint arXiv:2002.00470, (2020).
- [83] R. MAULIK, A. MOHAN, B. LUSCH, S. MADIREDDY, P. BALAPRAKASH, AND D. LIVESCU, *Time-series learning of latent-space dynamics for reduced-order model closure*, *Physica D: Nonlinear Phenomena*, 405 (2020), p. 132368.
- [84] R. MEAKIN, *Moving body overset grid methods for complete aircraft tiltrotor simulations*,

- <https://doi.org/10.2514/6.1993-3350>, <https://arc.aiaa.org/doi/abs/10.2514/6.1993-3350>, <https://arxiv.org/abs/https://arc.aiaa.org/doi/pdf/10.2514/6.1993-3350>.
- [85] V. MEHRMANN AND T. STYKEL, *Balanced truncation model reduction for large-scale systems in descriptor form*, in Dimension Reduction of Large-Scale Systems, P. Benner, D. C. Sorensen, and V. Mehrmann, eds., Berlin, Heidelberg, 2005, Springer Berlin Heidelberg, pp. 83–115.
- [86] R. MITTAL, H. DONG, M. BOZKURTAS, F. NAJJAR, A. VARGAS, AND A. [VON LOEBBECKE], *A versatile sharp interface immersed boundary method for incompressible flows with complex boundaries*, Journal of Computational Physics, 227 (2008), pp. 4825 – 4852, <https://doi.org/https://doi.org/10.1016/j.jcp.2008.01.028>, <http://www.sciencedirect.com/science/article/pii/S0021999108000235>.
- [87] A. T. MOHAN AND D. V. GAITONDE, *A deep learning based approach to reduced order modeling for turbulent flow control using lstm neural networks*, arXiv preprint arXiv:1804.09269, (2018).
- [88] A. T. MOHAN, N. LUBBERS, D. LIVESCU, AND M. CHERTKOV, *Embedding hard physical constraints in neural network coarse-graining of 3d turbulence*, arXiv preprint arXiv:2002.00021, (2020).
- [89] F. MONTI, D. BOSCAINI, J. MASCI, E. RODOLA, J. SVOBODA, AND M. M. BRONSTEIN, *Geometric deep learning on graphs and manifolds using mixture model cnns*, in The IEEE Conference on Computer Vision and Pattern Recognition (CVPR), July 2017.
- [90] M. MORIMOTO, K. FUKAMI, K. ZHANG, A. G. NAIR, AND K. FUKAGATA, *Convolutional neural networks for fluid flow analysis: toward effective metamodeling and low-dimensionalization*, 2021, <https://arxiv.org/abs/2101.02535>.
- [91] J. MORTON, F. D. WITHERDEN, AND M. J. KOCHENDERFER, *Parameter-conditioned sequential generative modeling of fluid flows*, 2019, <https://arxiv.org/abs/1912.06752>.
- [92] T. MURATA, K. FUKAMI, AND K. FUKAGATA, *Nonlinear mode decomposition with convolutional neural networks for fluid dynamics*, (2019).
- [93] B. NADLER, S. LAFON, R. R. COIFMAN, AND I. G. KEVREKIDIS, *Diffusion maps, spectral clustering and reaction coordinates of dynamical systems*, Applied and Computational Harmonic Analysis, 21 (2006), pp. 113–127.
- [94] K. NAKAHASHI, F. TOGASHI, AND D. SHAROV, *Intergrid-boundary definition method for overset unstructured grid approach*, AIAA Journal, 38 (2000), pp. 2077–2084, <https://doi.org/10.2514/2.869>, <https://doi.org/10.2514/2.869>, <https://arxiv.org/abs/https://doi.org/10.2514/2.869>.
- [95] T. NAKAMURA, K. FUKAMI, K. HASEGAWA, Y. NABAE, AND K. FUKAGATA, *Convolutional neural network and long short-term memory based reduced order surrogate for minimal turbulent channel flow*, Physics of Fluids, 33 (2021), p. 025116.
- [96] N. OMATA AND S. SHIRAYAMA, *A novel method of low-dimensional representation for temporal behavior of flow fields using deep autoencoder*, AIP Advances, 9 (2019), p. 015006, <https://doi.org/10.1063/1.5067313>.
- [97] S. E. OTTO AND C. W. ROWLEY, *Linearly recurrent autoencoder networks for learning dynamics*, SIAM Journal on Applied Dynamical Systems, 18 (2019), pp. 558–593.
- [98] J. PARK, D. HWANG, K. KIM, S. KANG, Y. KIM, AND J. LEE, *Computed tomography super-resolution using deep convolutional neural network*, Physics in medicine and biology, 63 (2018), <https://doi.org/10.1088/1361-6560/aacdd4>.
- [99] S. PAWAR, S. RAHMAN, H. VADDIREDDY, O. SAN, A. RASHEED, AND P. VEDULA, *A deep learning enabler for nonintrusive reduced order modeling of fluid flows*, Physics of Fluids, 31 (2019), p. 085101.
- [100] C. S. PESKIN, *The immersed boundary method*, Acta Numerica, 11 (2002), p. 479–517, <https://doi.org/10.1017/S0962492902000077>.
- [101] S. C. PULIGILLA AND B. JAYARAMAN, *Deep multilayer convolution frameworks for data-driven learning of fluid flow dynamics*, in 2018 Fluid Dynamics Conference, 2018, p. 3091.
- [102] L. PYTA AND D. ABEL, *Nonlinear model reduction of the navier-stokes-equations*, in 2016 American Control Conference (ACC), 2016, pp. 5249–5254.
- [103] S. R. QASIM, J. KIESELER, Y. IYAMA, AND M. PIERINI, *Learning representations of irregular particle-detector geometry with distance-weighted graph networks*, The European Physical Journal C, 79 (2019), pp. 1–11.
- [104] Q. QIU, X. CHENG, R. CALDERBANK, AND G. SAPIRO, *Dcfnet: Deep neural network with decomposed convolutional filters*, (2018), <https://arxiv.org/abs/1802.04145>.
- [105] H. QU AND L. GOUSKOS, *Particlenet: jet tagging via particle clouds*, arXiv preprint arXiv:1902.08570, (2019).
- [106] E. RANJAN, S. SANYAL, AND P. P. TALUKDAR, *Asap: Adaptive structure aware pooling for*

- learning hierarchical graph representations*, 2020, <https://arxiv.org/abs/1911.07979>.
- [107] P. RASTI, T. UIBOUPIN, S. ESCALERA, AND G. ANBARJAFARI, *Convolutional neural network super resolution for face recognition in surveillance monitoring*, in *Articulated Motion and Deformable Objects*, F. J. Perales and J. Kittler, eds., Cham, 2016, Springer International Publishing, pp. 175–184.
- [108] K. RAVIKUMAR, D. APPELHANS, AND P. YEUNG, *Gpu acceleration of extreme scale pseudo-spectral simulations of turbulence using asynchronism*, in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, 2019, pp. 1–22.
- [109] S. A. RENGANATHAN, R. MAULIK, AND V. RAO, *Machine learning for nonintrusive model order reduction of the parametric inviscid transonic flow past an airfoil*, *Physics of Fluids*, 32 (2020), p. 047110.
- [110] S. T. ROWEIS AND L. K. SAUL, *Nonlinear dimensionality reduction by locally linear embedding*, *science*, 290 (2000), pp. 2323–2326.
- [111] C. W. ROWLEY, I. MEZI?, S. BAGHERI, P. SCHLATTER, D. HENNINGSON, ET AL., *Spectral analysis of nonlinear flows*, *Journal of fluid mechanics*, 641 (2009), pp. 115–127.
- [112] D. RYCKELYNCK, *A priori hyperreduction method: an adaptive approach*, *Journal of computational physics*, 202 (2005), pp. 346–366.
- [113] J. SCHMIDHUBER, *Deep learning in neural networks: An overview*, *Neural networks*, 61 (2015), pp. 85–117.
- [114] A. D. SCILLITOE, P. SESHADRI, AND C. Y. WONG, *Instantaneous flowfield estimation with gaussian ridges*, in *IAA Scitech 2021 Forum*, 2021, p. 1138.
- [115] S. E. SHERER AND J. N. SCOTT, *High-order compact finite-difference methods on general overset grids*, *Journal of Computational Physics*, 210 (2005), pp. 459 – 496, <https://doi.org/https://doi.org/10.1016/j.jcp.2005.04.017>, <http://www.sciencedirect.com/science/article/pii/S0021999105002366>.
- [116] J. SHLOMI, P. BATTAGLIA, AND J.-R. VLMANT, *Graph neural networks in particle physics*, *arXiv preprint arXiv:2007.13681*, (2020).
- [117] D. I. SHUMAN, S. K. NARANG, P. FROSSARD, A. ORTEGA, AND P. VANDERGHEYNST, *The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains*, *IEEE Signal Processing Magazine*, 30 (2013), p. 83–98, <https://doi.org/10.1109/msp.2012.2235192>, <http://dx.doi.org/10.1109/MSP.2012.2235192>.
- [118] M. SIMONOVSKY AND N. KOMODAKIS, *Dynamic edge-conditioned filters in convolutional neural networks on graphs*, 2017, <https://arxiv.org/abs/1704.02901>.
- [119] B. E. SONDAY, A. SINGER, C. W. GEAR, AND I. G. KEVREKIDIS, *Manifold learning techniques and model reduction applied to dissipative pdes*, 2010, <https://arxiv.org/abs/1011.5197>.
- [120] P. A. SRINIVASAN, L. GUASTONI, H. AZIZPOUR, P. SCHLATTER, AND R. VINUESA, *Predictions of turbulent shear flows using deep neural networks*, *Physical Review Fluids*, 4 (2019), p. 054603.
- [121] R. STEWART AND S. ERMON, *Label-free supervision of neural networks with physics and domain knowledge*, *CoRR*, abs/1609.05566 (2016), <http://arxiv.org/abs/1609.05566>, <https://arxiv.org/abs/1609.05566>.
- [122] C. M. STROFER, J.-L. WU, H. XIAO, AND E. PATERSON, *Data-driven, physics-based feature extraction from fluid flow fields using convolutional neural networks*, *Communications in Computational Physics*, 25 (2019), <https://doi.org/10.4208/cicp.oa-2018-0035>, <http://dx.doi.org/10.4208/cicp.OA-2018-0035>.
- [123] C. SZEGEDY, W. LIU, Y. JIA, P. SERMANET, S. REED, D. ANGUELOV, D. ERHAN, V. VANHOUCKE, AND A. RABINOVICH, *Going deeper with convolutions*, in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015.
- [124] J. H. TU, C. W. ROWLEY, D. M. LUCHTENBURG, S. L. BRUNTON, AND J. N. KUTZ, *On dynamic mode decomposition: Theory and applications*, *arXiv preprint arXiv:1312.0041*, (2013).
- [125] Z. Y. WAN, P. VLACHAS, P. KOUMOUTSAKOS, AND T. SAPSIS, *Data-assisted reduced-order modeling of extreme events in complex dynamical systems*, *PloS one*, 13 (2018), p. e0197704.
- [126] Q. WANG, J. S. HESTHAVEN, AND D. RAY, *Non-intrusive reduced order modeling of unsteady flows using artificial neural networks with application to a combustion problem*, *Journal of computational physics*, 384 (2019), pp. 289–307.
- [127] X. WANG, Z. LI, M. JIANG, S. WANG, S. ZHANG, AND Z. WEI, *Molecule property prediction based on spatial graph embedding*, *Journal of chemical information and modeling*, 59 (2019), pp. 3817–3828.

- [128] Y. WANG, V. KIM, M. BRONSTEIN, AND J. SOLOMON, *Learning geometric operators on meshes*, Representation Learning on Graphs and Manifolds 2019 (ICLR workshop), (2019), <http://par.nsf.gov/biblio/10124350>.
- [129] A. WAREY, J. GAO, AND R. GROVER, *Prediction of engine-out emissions using deep convolutional neural networks*, tech. report, SAE Technical Paper, 2021.
- [130] J. WILLCOX, K; PERAIRE, *Balanced model reduction via the proper orthogonal decomposition*, AIAA journal, (2002).
- [131] M. O. WILLIAMS, I. G. KEVREKIDIS, AND C. W. ROWLEY, *A data-driven approximation of the koopman operator: Extending dynamic mode decomposition*, Journal of Nonlinear Science, 25 (2015), pp. 1307–1346.
- [132] C. J. WINSTEAD, *Nonlinear model reduction based on manifold learning with application to the burgers' equation*, (2017).
- [133] L. WU, P. SUN, R. HONG, Y. FU, X. WANG, AND M. WANG, *Socialgcn: An efficient graph convolutional neural network based model for social recommendation*, (2018), <https://arxiv.org/abs/1811.02815>.
- [134] Z. WU, S. PAN, F. CHEN, G. LONG, C. ZHANG, AND P. S. YU, *A comprehensive survey on graph neural networks*, ArXiv, abs/1901.00596 (2019).
- [135] T. XIE AND J. C. GROSSMAN, *Crystal graph convolutional neural networks for an accurate and interpretable prediction of material properties*, Physical review letters, 120 (2018), p. 145301.
- [136] T. XIE AND J. C. GROSSMAN, *Hierarchical visualization of materials space with graph convolutional neural networks*, The Journal of chemical physics, 149 (2018), p. 174111.
- [137] S. YE, J. LIANG, R. LIU, AND X. ZHU, *Symmetrical graph neural network for quantum chemistry, with dual r/k space*, arXiv preprint arXiv:1912.07256, (2019).
- [138] S. YE, Z. ZHANG, X. SONG, Y. WANG, Y. CHEN, AND C. HUANG, *A flow feature detection method for modeling pressure distribution around a cylinder in non-uniform flows by using a convolutional neural network*, Scientific Reports, 10 (2020), pp. 1–10.
- [139] E. YEUNG, S. KUNDU, AND N. HODAS, *Learning deep neural network representations for koopman operators of nonlinear dynamical systems*, in 2019 American Control Conference (ACC), IEEE, 2019, pp. 4832–4839.
- [140] J. YU, C. YAN, AND M. GUO, *Non-intrusive reduced-order modeling for fluid problems: A brief review*, Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering, 233 (2019), pp. 5896–5912.
- [141] M. ZENG, J. N. KUMAR, Z. ZENG, R. SAVITHA, V. R. CHANDRASEKHAR, AND K. HIPPALGAONKAR, *Graph convolutional neural networks for polymers property prediction*, arXiv preprint arXiv:1811.06231, (2018).
- [142] R. ZHANG, Y. LIU, AND H. SUN, *Physics-informed multi-lstm networks for metamodeling of nonlinear structures*, arXiv preprint arXiv:2002.10253, (2020).
- [143] Y. ZHANG, W.-J. SUNG, AND D. MAVRIS, *Application of convolutional neural network to predict airfoil lift coefficient*, 2017, <https://arxiv.org/abs/1712.10082>.
- [144] C. ZHAO, B. NI, J. ZHANG, Q. ZHAO, W. ZHANG, AND Q. TIAN, *Variational convolutional neural network pruning*, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2019, pp. 2780–2789.
- [145] X. ZHENG, P. ZHENG, AND R.-Z. ZHANG, *Machine learning material properties from the periodic table using convolutional neural networks*, Chemical science, 9 (2018), pp. 8426–8432.
- [146] J. ZHOU, G. CUI, Z. ZHANG, C. YANG, Z. LIU, AND M. SUN, *Graph neural networks: A review of methods and applications*, ArXiv, abs/1812.08434 (2018).
- [147] Y. ZHU, N. ZABARAS, P.-S. KOUTSOURELAKIS, AND P. PERDIKARIS, *Physics-constrained deep learning for high-dimensional surrogate modeling and uncertainty quantification without labeled data*, Journal of Computational Physics, 394 (2019), pp. 56 – 81, <https://doi.org/https://doi.org/10.1016/j.jcp.2019.05.024>, <http://www.sciencedirect.com/science/article/pii/S0021999119303559>.
- [148] M. ZITNIK, M. AGRAWAL, AND J. LESKOVEC, *Modeling polypharmacy side effects with graph convolutional networks*, Bioinformatics, 34 (2018), pp. i457–i466.