

Symmetry-Aware Autoencoders: s-PCA and s-nlPCA

Simon Kneer, Taraneh Sayadi, Denis Sipp, Peter Schmid, Georgios Rigas

▶ To cite this version:

Simon Kneer, Taraneh Sayadi, Denis Sipp, Peter Schmid, Georgios Rigas. Symmetry-Aware Autoencoders: s-PCA and s-nlPCA. 2021. hal-03420320

HAL Id: hal-03420320 https://hal.science/hal-03420320

Preprint submitted on 9 Nov 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers. L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Symmetry-Aware Autoencoders: s-PCA and s-nlPCA

Simon Kneer^{a,*}, Taraneh Sayadi^b, Denis Sipp^c, Peter Schmid^d, Georgios Rigas^a

^aDepartment of Aeronautics, Imperial College London, London SW7 2AZ, UK ^bJean le Rond d'Alembert Institute, CNRS/Sorbonne University, Paris, France ^cDAAA, ONERA, Université Paris Saclay, 8 rue des Vertugadins, 92190 Meudon, France ^dDepartment of Mathematics, Imperial College London, London SW7 2AZ, UK

Abstract

Nonlinear principal component analysis (nIPCA) via autoencoders has attracted attention in the dynamical systems community due to its larger compression rate when compared to linear principal component analysis (PCA). These model reduction methods experience an increase in the dimensionality of the latent space when applied to datasets that exhibit globally invariant samples due to the presence of symmetries. In this study, we introduce a novel machine learning embedding in the autoencoder, which uses spatial transformer networks and Siamese networks to account for continuous and discrete symmetries, respectively. The spatial transformer network discovers the optimal shift for the continuous translation or rotation so that invariant samples are aligned in the periodic directions. Similarly, the Siamese networks collapse samples that are invariant under discrete shifts and reflections. Thus, the proposed symmetry-aware autoencoder is invariant to predetermined input transformations dictating the dynamics of the underlying physical system. This embedding can be employed with both linear and nonlinear reduction methods, which we term symmetry-aware PCA (s-PCA) and symmetry-aware nIPCA (s-nIPCA). We apply the proposed framework to 3 fluid flow problems: Burgers' equation, the simulation of the flow through a step diffuser and the Kolmogorov flow to showcase the capabilities for cases exhibiting only continuous symmetries, only discrete symmetries or a combination of both.

1. Introduction

Modal decomposition methods have been widely used to further the understanding of complex physical processes for decades. In the field of fluid dynamics, ever since principal component analysis (PCA), also known as proper orthogonal decomposition (POD) was introduced by Lumley [14], it has been a staple tool for decomposing flows into relevant structures, since it gives an optimal linear representation of the flow in terms of the L_2 -norm.

The connection between neural networks and PCA was first drawn by Baldi & Hornik [1]. In their study, an autoencoder (AE) that is constructed as a linear multilayer perceptron (MLP) was employed to show its essential equivalence to PCA. An AE is a compilation of two networks, an encoder and a decoder, in series. Thus, an input can be reduced or encoded to a latent space, which is comparable to the space of amplitude coefficients for PCA, and afterwards reconstructed or decoded into the original input space. Plaut [25] was able to show that the reduced subspace that is spanned by PCA and a linear AE are identical. The modes for the latter method are not orthogonal by default, unlike for PCA, but can be approximately orthogonalized by employing a singular value decomposition on the weights after training [25]. However, PCA is only equivalent to linear AEs without biases, which are only a small subset of such networks. In a more general

^{*}Corresponding author

Email address: kneer@imperial.ac.uk (Simon Kneer)

setting, an AE is deployed in combination with activation functions that allow the network to learn nonlinear relations between the input and the output. This nonlinear PCA (nlPCA) was first employed by Milano & Koumoutsakos [20] and showed great promise in reducing the L_2 -reconstruction error for the randomly forced Burgers equation and turbulent channel flow when comparing to PCA. However, as of now, there is no unified way to visualize nlPCA modes, since the temporal and spatial behaviour are now nonlinearly coupled. Different promising approaches have been suggested. Fukami *et al.* [8] force the network to be additive in the final output space, which allows for a comprehensive way to investigate the spatial structures of each output channel. Page *et al.* [23] performed a latent Fourier analysis, that is able to extract patterns based on a filtering of the latent wavenumbers.

Physical equations, such as the Navier-Stokes equations, are equivariant under symmetry transformations, leading to invariant solutions that are identical up to a prescribed transformation function. The presence of invariant solutions poses an issue for model reduction techniques, since conventional methods cannot find the connection between them and, thus, encode them into different regions of the reduced space. This leads to an artificial bloating of the latent space. Additionally, for linear methods, where the focus usually lies in the extraction of the modal shapes to gain an insight into the flow structures, these invariances can morph the modes away from physical structures. In fact, given a large enough dataset of an ergodic system, a necessary condition is that the PCA modes themselves are invariant with respect to the same symmetry operations the dynamical system is equivariant under. Observed instantaneous events are not bound by this condition, however, creating a split between observed states and the modal basis, see Holmes *et al.* [10].

In the case of PCA, this issue has long been observed for periodic boundaries, which create a group of translationally invariant solutions, $T(s): u(x) \to u(x+s)$, a subgroup of all continuous symmetry groups, F(s). For datasets of ergodic turbulent fluid flows, that are both statistically stationary and have reached statistical convergence with respect to the translational symmetries, the modal representation will become the Fourier basis [28]. These modes are merely an artifact of the method, i.e. PCA, and dynamically insignificant, since all wavenumbers are dynamically coupled through the nonlinearities of the underlying equations. In the case of nlPCA these invariances show themselves in the larger dispersion of the variables in the reduced space, see e.g. Page et al. [23]. Different methods to circumvent these issues have been proposed in the past. Mendible et al. [18] fit prescribed template functions to their data in order to extract an exact definition of the temporal evolution of the wave speed of travelling wave problems. Another method is called template matching, see Rowley et al. [26] and Holmes et al. [10], and aims to remove translational invariances by matching a template, chosen a priori, with the flow field and translating each snapshot by an offset from the template. Another popular method, known as the method of slicing using the first Fourier mode slice [4], aims to remove translational invariances by shifting the samples to set the phase of the first Fourier mode to 0. We shall henceforth refer to this method as the method of slicing, even if it is a specific case of the latter only. In combination with nlPCA, this method has successfully been applied by e.g. Linot & Graham [13], who also included temporal advancement in the reduced space. The method is only valid when the amplitude of the first Fourier mode is non-zero, since the shift of a sample with zero amplitude would become infinite. This problem is known to arise in turbulent channel flow [16] and can be circumvented by introducing a handcrafted modulation of the Fourier basis in the inhomogeneous direction. A different way to approach advection-dominated problems is to split the problem into advection modes, determined through optimal mass transfer. Additional features can then be captured by performing POD on the residual in the advected frame of reference, see Iollo & Lombardi [11].

Continuous invariances are not the only invariances included in many dynamical systems. Discrete symmetry groups, G, such as the reflectional symmetry in a cylinder wake, $\mathbf{R} : (u, v)(x, y) \rightarrow (u, -v)(x, -y)$, can play a role as well and are more difficult to remove from the system due to their discontinuous nature. Invariant polynomial bases can be constructed, but these have to be tailored to each problem, and in many cases no such bases have been found as of now [3].

In this work we propose a novel machine learning method, called symmetry-aware principal component analysis (s-PCA), in order to efficiently reduce the complexity of dynamical systems, exhibiting globally discrete or continuous invariant solutions. A conceptual visualization of this procedure is illustrated in fig. 1. Our network differs from a conventional AE, in that we apply two transformations and their inverse



Figure 1: Visualization of a symmetry-aware autoencoder architecture developed in this work. A sample \mathbf{X} is passed through a first subnetwork, called an *STN*, that transforms each sample to remove continuous invariances from the solution set. A second subnetwork removes the discrete invariances through Siamese branches. The sample is then passed through an autoencoder to generate a low-dimensional latent space, \mathbf{A} , and the associated reconstruction in the symmetry-reduced space. We then apply the inverse of the previous symmetry-reducing networks to finally recover the sample in the input space, $\hat{\mathbf{X}}$. At each training step the L_2 -loss, i.e. $|\hat{\mathbf{X}} - \mathbf{X}|_2$, is calculated and backpropagated until convergence is reached in all subnetworks.

before encoding and after decoding a sample, respectively.

The first transformation, F(s), dealing with the continuous invariances in our dataset, is connected to a machine learning method widely used in image recognition tasks, termed spatial transformer networks, STNs, [12]. In this approach, a secondary encoder finds a set of transformation parameters that are then applied to an input via affine transformations. Our method deviates from a classical STN in that we do not employ affine transformations, since our continuous symmetry groups appear in combination with periodic boundary conditions. Hence, we need to translate the samples, i.e. snapshots of our solution set, in a periodic way, which can be done by performing a Fourier transformation, shifting the phase of all coefficients and performing an inverse Fourier transform. Note that this is merely an efficient and accurate way to shift a sample and interpolate it onto the grid and not related to e.g. the method of slicing. As of now, our method is restricted to global translations associated with the periodic boundary conditions.

The method for removing discrete invariances from our dataset is closely related to the method developed by Mehr *et al.* [17]. In their work the candidate discrete transformations are applied to parallel encoders, creating a Siamese network, and the resulting latent space is then combined by choosing the maximum of each latent variable over all branches. Our approach differs in that the candidate branches consist of complete AEs, and the latent space is built by selecting the branch with the lowest reconstruction loss.

Hence, an invariant subspace is created at the same time the model reduction is applied, drastically decreasing the reconstruction error and the size of our latent space. We apply this symmetry reduction to both the linear (s-PCA) and nonlinear (s-nIPCA) decomposition methods, creating a higher degree of compression as well as improving the interpretability of the resulting modes.

After an in-depth analysis of our proposed method is conducted, investigations into the effectiveness of the above methodology are performed. A number of examples have been chosen, namely the one-dimensional Burgers' equation, the sudden expansion of the two-dimensional flow through a diffuser with stochastic forcing and the two-dimensional Kolmogorov flow. The first of these cases only exhibits the continuous group of translations, T(s), while the flow through a diffuser is equivariant under the discrete group of reflections, \mathbf{R} . The Kolmogorov flow is equivariant under the group of continuous translations, T(s), a shift-and-reflect symmetry, $\mathbf{S}: \, \omega(x, y) \to \omega(-x, y + \frac{\pi}{4})$, and a rotation through π , $\mathbf{P}: \, \omega(x, y) \to \omega(-x, -y)$.

2. Background

In this paper we employ multiple model reduction methods to address difficulties arising from invariant solutions within the dataset, and propose remedies to circumvent these problems. On the one hand, the input can be reduced with a linear parametric model, which can be constructed with the analytical method of PCA, and its neural network equivalent using linear autoencoders. As has been shown before [1], both methods span the same subspace independent of the construction method, however, using the latter orthogonality of the modes is not given without further modification. On the other hand, there is the nonlinear generalization of PCA, nlPCA, which can be constructed using nonlinear autoencoders. Both linear and nonlinear methods can be embedded into the invariance-reducing framework we have developed, leading to linear (s-PCA) and nonlinear symmetry-aware PCA (s-nlPCA), respectively. First, an overview of different invariant solutions commonly found in physical problems is provided.

2.1. Invariant solutions

We call a solution to a governing system of equations, $X(\boldsymbol{\xi}, t)$, dependent on the spatial coordinates $\boldsymbol{\xi}$ and time t, invariant under a given symmetry operation, \boldsymbol{S} , if $\boldsymbol{SX}(\boldsymbol{\xi}, t)$ is also a solution to the governing equations. We can split \boldsymbol{S} into two groups: continuous and discrete symmetry groups.

Continuous symmetry groups, F(s), are based on continuous operators that can be applied to a solution. These operators depend on continuous transformation parameters, denoted by s. In the case of an incompressible two-dimensional channel flow, (u, v, p)(x, y), with periodic boundary conditions in x, we observe such a symmetry group: the group of continuous translations, T(s). This means that any solution (u, v, p)(x + s, y) is a solution to the Navier-Stokes equations if (u, v, p)(x, y) is a solution, see e.g. Marensi *et al.* [16].

Discrete symmetry groups, G, are a finite set of operators, that transform the solution in a distinct way for each member of the set. An example is given by the group of reflections, R, where (u, -v)(x, -y) is a solution to the underlying equations, if (u, v)(x, y) is a solution. This symmetry is present in e.g. the two dimensional flow around a cylinder. Note that in the presence of multiple discrete symmetry groups, these can be unified into one by creating a new group G', that is composed of all unique combinations of the original symmetry groups.

Both these types of symmetries pose issues when performing model reduction on a discretized system, since approximately identical states, mainly differing in an application of a symmetry operation, are being encoded into vastly different regions of the latent space, see e.g. Page *et al.* [23]. Additionally, the PCA basis for these types of problems is less interpretable due to the morphing of the modes caused by the invariant solutions. In the following, a brief description of linear and nonlinear model reduction methods shall be given, followed by an excursion into the interpretability problems arising from invariances.

2.2. PCA and nlPCA

An autoencoder (AE) is a parametric model that encodes a given input dataset $\mathbf{X} \in \mathbb{R}^{N_{\xi} \times N_{t}}$, dependent on spatial coordinates $\boldsymbol{\xi}$ and time t of size N_{ξ} and N_{t} , into a latent space \mathbf{A} , and decodes this latter description back to the input space, hence creating a reconstruction $\hat{\mathbf{X}}(\boldsymbol{\xi}, t)$. The latent space is designed to have a lower dimension than the input and output space, thus, achieving a compression of the data. This parametric model, where E and D denote the encoder and decoder, respectively, can be expressed as

$$\mathbf{A} = E(\mathbf{X}) \tag{1}$$

$$\hat{\mathbf{X}} = D(\mathbf{A}). \tag{2}$$

The parameters of this model, i.e. the weights of the network, can be trained by backpropagating the L_2 error for each sample, $|\hat{\mathbf{X}} - \mathbf{X}|_2$. If this parametric model is linear, i.e. E and D are matrices that are being
multiplied with the input and the latent space, an analytical method to train this model is given by PCA.

2.2.1. PCA and linear autoencoders

Given a dataset **X**, PCA generates an orthonormal basis, $\Phi \in \mathbb{R}^{N_{\xi} \times N}$, with a corresponding latent space, $\mathbf{A} \in \mathbb{R}^{N \times N_{t}}$, that gives a modal approximation for the input as

$$\hat{\mathbf{X}}(\boldsymbol{\xi}, t) = \sum_{i=1}^{N} a_i(t) \boldsymbol{\Phi}_i(\boldsymbol{\xi}), \tag{3}$$

where $N < N_{\xi}$ is the series truncation, i.e. the number of modes used for the reconstruction. The orthogonal modes are defined via the eigenvalue problem

$$\frac{1}{N_t} \mathbf{X} \mathbf{X}^T \, \tilde{\mathbf{\Phi}} = \tilde{\mathbf{\Phi}} \mathbf{\Lambda},\tag{4}$$

where $\frac{1}{N_t} \mathbf{X} \mathbf{X}^T$ is identified as the correlation matrix. The modes $\tilde{\mathbf{\Phi}}$ are the set of eigenvectors of the correlation matrix with corresponding eigenvalues $\mathbf{\Lambda}$, quantifying the energetic content of each mode. The truncated basis can then be defined as

$$\boldsymbol{\Phi} = [\tilde{\boldsymbol{\Phi}}_1 \ \tilde{\boldsymbol{\Phi}}_2 \ \dots \ \tilde{\boldsymbol{\Phi}}_N]. \tag{5}$$

The amplitude coefficients spanning the latent space are subsequently obtained by projecting \mathbf{X} onto the desired set of PCA modes, i.e.

$$\mathbf{A} = \mathbf{\Phi}^T \mathbf{X}.\tag{6}$$

PCA guarantees that the reconstruction using N modes is the optimal linear reconstruction in terms of the L_2 -norm, i.e. $|\hat{\mathbf{X}} - \mathbf{X}|_2$ is minimal. If \mathbf{X} is comprised of velocity components, this is equivalent to a maximization of the kinetic energy.

We can equivalently generate the modes by a neural network, or more specifically a multilayer perceptron (MLP). In this context our parametric model is termed a linear AE without biases and formed as

$$\mathbf{A} = \mathbf{W}_E \mathbf{X} \tag{7}$$

$$\hat{\mathbf{X}} = \mathbf{W}_D \mathbf{A},\tag{8}$$

where the matrices $\mathbf{W}_E \in \mathbb{R}^{N \times N_{\xi}}$ and $\mathbf{W}_D \in \mathbb{R}^{N_{\xi} \times N}$ are called the weights of the encoder and decoder, respectively. Their size is dependent on the spatial dimension of each sample, N_{ξ} , and the desired size of the latent space, N. Through backpropagation these weights are trained on samples of the dataset to minimize the L_2 -norm of the error, $|\hat{\mathbf{X}} - \mathbf{X}|_2$, termed the loss. Given the convexity of the problem, it is easy to see that, apart from the orthonormality property of $\boldsymbol{\Phi}$, the columns of \mathbf{W}_E are entirely equivalent to the PCA modes. The orthogonality of the modes can approximately be recovered after the training process by performing a singular value decomposition (SVD) on the weights, see Plaut [25].

The fully converged linear AE using a latent dimension of N spans the same subspace as PCA using N modes, see Baldi & Hornik [1] and Plaut [25]. Independently of the way this model is generated, performance is suboptimal when invariances are present in the dataset, since the dimensionality of the latent space is increased.

2.2.2. PCA and invariant solutions

For the following discussion an ergodic turbulent flow with statistical stationarity and a roughly equal appearance of the possible transformed states is assumed. While this is not a necessity for the arising problem of PCA modes being constricted in their shape due to the appearance of invariant solutions, it is required for the exact way the shapes are constricted. These constrictions are not due to each sample of the flow field exhibiting certain symmetries, but rather a byproduct of PCA itself.

If the continuous group of translations, T(s), applies to a specific coordinate direction, the PCA will produce a Fourier basis in this direction, see e.g. Holmes *et al.* [10]. This property is often exploited when periodic boundaries are present, since one can perform a discrete Fourier transform (DFT) in the periodic direction prior to performing PCA and hence reduce the computational complexity of the operation, see e.g. Freund & Colonius [7]. Nonetheless, these modes are arguably less informative with regard to flow structures.

Given discrete symmetries a similar issue emerges. For the discrete reflectional symmetry, \mathbf{R} , the PCA basis displays symmetric or antisymmetric shapes with respect to the axis of reflection, see Sirovich [28]. Turbulence initially breaks the symmetry of the laminar base flow. After a sufficiently long time, however, these symmetries are statistically recovered, i.e., each solution appears in its original and reflected state, approximately equidistributed in time. This property has again been exploited for reducing the computational requirements to calculate the PCA modes, see e.g. Bourgeois *et al.* [2]. In the case of reflectional symmetries, we can divide the flow into a symmetric and an antisymmetric part, each only half the size of the original domain. As a consequence, the complexity of the eigenvalue problem reduces from $O(N_{\xi}^2)$ to $O(1/2(N_{\xi})^2)$. However, since modes are either symmetric or antisymmetric, it might be tempting to conclude that each sample is either symmetric or antisymmetric, when in fact the majority of samples is asymmetric.

2.2.3. nlPCA

Nonlinear autoencoders are the generalization of section 2.2.1 and their usage can be referred to as nonlinear PCA (nlPCA). They employ nonlinear activation functions, f_{ac} , to accommodate nonlinear relations expressed by the network. Employing an MLP network, nlPCA takes on the form

$$\mathbf{A} = \mathbf{W}_{E,2} f_{ac} (\mathbf{W}_{E,1} \mathbf{X} + \mathbf{b}_{E,1}) + \mathbf{b}_{E,2}$$
(9)

$$\mathbf{X} = \mathbf{W}_{D,2} f_{ac} (\mathbf{W}_{D,1} \mathbf{A} + \mathbf{b}_{D,1}) + \mathbf{b}_{D,2}.$$
 (10)

Here $\mathbf{W}_{E,1} \in \mathbb{R}^{H \times N_{E}}$ and $\mathbf{W}_{E,2} \in \mathbb{R}^{N \times H}$ are the first and second weight matrices of the encoder with corresponding biases $\mathbf{b}_{E,1} \in \mathbb{R}^{H}$ and $\mathbf{b}_{E,2} \in \mathbb{R}^{N}$. The same principle applies to the decoder with transpose dimensions for the weights. Biases are an additive vector for each of the layers. The dimension of the hidden layer, H, is a user-specified parameter. Both the encoder and decoder now need at least one activated layer with a final linear layer, according to the universal approximation theorem [6]. The chosen activation function and hidden dimension H shall be of little concern, since these parameters have a minor influence on the proposed method in this paper. However, we note that H was chosen sufficiently large, so as to ensure a proper representation of the nonlinear transform. Similarly, the type of network we describe above is referred to as a multi-layer perceptron (MLP); convolutional neural networks (CNNs) can be embedded in our symmetry-reducing framework just as well.

2.3. Energy evaluation of arbitrary reconstructions

As discussed above, there is no unified way of examining the nlPCA modes, owing to their nonlinear interactions. To visualize the effectiveness of the method, beyond the reduction in L_2 -error, one can devise a method that reflects the nonlinear reconstruction capabilities within the linear framework, see e.g. Murata *et al.* [21]. Recall that the PCA amplitude coefficients are obtained by projecting the dataset $\mathbf{X}(\boldsymbol{\xi}, \mathbf{t})$ onto the set of PCA modes $\boldsymbol{\Phi}(\boldsymbol{\xi})$. Given these amplitude coefficients $\mathbf{A}(\mathbf{t})$, the relation

$$\mathbf{\Lambda} = \frac{1}{N_t} \sum_{k=1}^{N_t} \mathbf{A}(t_k) \circ \mathbf{A}(t_k), \tag{11}$$

where $\mathbf{\Lambda}$ are the eigenvalues of the covariance matrix, N_t is the total number of samples and \circ the Hadamard product, holds since the PCA modes are orthogonal. As stated above, the eigenvalue associated with each mode is a measurement of the energetic content of each mode. Projecting an arbitrary reconstruction $\hat{\mathbf{X}}(\boldsymbol{\xi}, t)$, e.g. nlPCA, onto the PCA modes

$$\hat{\mathbf{A}} = \boldsymbol{\Phi}^T \hat{\mathbf{X}},\tag{12}$$



Figure 2: Visualization of the symmetry-aware autoencoder architecture developed in this paper. A sample \mathbf{X} is passed to the network. A secondary network, called an STN [12], extracts a continuous transformation parameter s from the sample $(\mathbf{X} \to s)$. The original sample is then manipulated according to a predefined **continuous** transformation function, $\mathbf{F}(s)$, and s passed on $(\mathbf{X} \to \mathbf{X}^*)$. Each member of a set of predefined **discrete** transformations, \mathbf{G} , is then applied $(\mathbf{X}^* \to \mathbf{X}')$. The network breaks into branches where each branch is a candidate discrete transform of the modified input. The same autoencoder is then applied to all branches in parallel $(\mathbf{X}' \to \mathbf{\hat{X}}')$. The resulting reconstructions are then merged by selecting the branch, l, with the least L_2 -error to its respective input $(\mathbf{\hat{X}}' \to \mathbf{\hat{X}}')$. We then apply the inverse discrete transform of the member \mathbf{G}^l , \mathbf{G}^{-l} $(\mathbf{\hat{X}}'_l \to \mathbf{\hat{X}}^*)$. The inverse continuous symmetry group $\mathbf{F}(s)^{-1}$ is then used to retransform the reconstruction to the original input space $(\mathbf{\hat{X}}^* \to \mathbf{\hat{X}})$. The trainable parameters of the network, i.e. the STN and the autoencoder, are trained by backpropagating the L_2 -error, $|\mathbf{\hat{X}} - \mathbf{X}|_2$.

and taking the temporal mean of the square amplitudes

$$\hat{\mathbf{\Lambda}} = \frac{1}{N_t} \sum_{k=1}^{N_t} \hat{\mathbf{A}}(t_k) \circ \hat{\mathbf{A}}(t_k), \tag{13}$$

we arrive at a measurement of the energy for each linear mode contained within the chosen reconstruction. The closer the reconstruction $\hat{\mathbf{X}}$ tends to \mathbf{X} , the closer we approach $\hat{\mathbf{A}}$ to \mathbf{A} and $\hat{\mathbf{\Lambda}}$ to $\mathbf{\Lambda}$. We define the equivalent PCA energy for each reconstruction as

$$e_s = \frac{\hat{\lambda}_s}{\sum_{i=1} \lambda_i}.$$
(14)

3. Symmetry-aware model reduction: s-PCA and s-nlPCA

To remove invariances from the solution set in an automated manner, a neural network embedding has been designed and is shown in fig. 2. Please see https://github.com/simonkneer/Symmetry-Aware-Autoencoding for Keras implementations of this network for our considered test cases. The schematic shows how our proposed architecture deals with continuous and discrete symmetry groups simultaneously. When

the batch of samples **X** is first passed to the network, a subnetwork, called a Spatial Transformer Network (STN) [12], evaluates the samples first. This subnetwork is in itself an encoder that transforms the input sample to a set of transformation parameters **s**:

$$\mathbf{s} = \mathbf{W}_{STN,2} f_{ac} (\mathbf{W}_{STN,1} \mathbf{X} + \mathbf{b}_{STN,1}) + \mathbf{b}_{STN,2}.$$
(15)

Here $\mathbf{W}_{STN,1} \in \mathbb{R}^{H \times N_{\xi}}$ and $\mathbf{W}_{STN,2} \in \mathbb{R}^{N_{s} \times H}$ are the first and second weight matrices of the encoder with corresponding biases $\mathbf{b}_{STN,1} \in \mathbb{R}^{H}$ and $\mathbf{b}_{STN,2} \in \mathbb{R}^{N_{s}}$, where H is the hidden dimension and N_{s} is the number of transformation parameters. The samples \mathbf{X} are then modified according to the prescribed continuous transformation function $\mathbf{F}(\mathbf{s})$,

$$\mathbf{X}^* = \mathbf{F}(\mathbf{s})[\mathbf{X}]. \tag{16}$$

For the group of continuous translations, this amounts to the identification of a shift parameter s, i.e. determining how much each of the samples needs to be translated in space to minimize the error of the subsequent autoencoder. In this case an efficient and accurate way to perform the shift and the subsequent interpolation is to perform a discrete Fourier transform (DFT) of the sample **X**. For simplicity this method shall be described for a one-dimensional dataset $\mathbf{X}(\mathbf{x}, \mathbf{t})$:

$$\underline{\mathbf{X}}(\mathbf{k}, \mathbf{t}) = \mathcal{F}(\mathbf{X}(\mathbf{x}, \mathbf{t})), \tag{17}$$

where **k** are the wavenumbers of the frequency domain. A generalization is achieved by performing a Fourier transform for the translated direction at each grid point of the other dimensions. Each of the Fourier coefficients $\underline{X}(k, \mathbf{t})$ are then shifted by

$$\underline{X}^{*}(k,\mathbf{t}) = \underline{X}(k,\mathbf{t})e^{-ik2\pi\frac{s}{L_{x}}} \quad \forall k \in \mathbf{k},$$
(18)

where L_x is the extent in the spatial direction x. The samples are then retransformed,

$$\mathbf{X}^*(\mathbf{x}, \mathbf{t}) = \mathcal{F}^{-1}(\underline{\mathbf{X}}^*(\mathbf{k}, \mathbf{t})).$$
(19)

We emphasize that this is not related to the method of slicing, but rather is more efficient and accurate than shifting the snapshots and performing a linear interpolation if s is not an integer. This makes s an angular rather than a translational shift. Since this angle is discontinuous when progressing from $-\pi$ to π , the STN experiences difficulties capturing this behavior. To counteract, the STN can be set to predict two output parameters, s_1 and s_2 , which are then merged following

$$s = \arctan 2(s_1, s_2) \frac{L_x}{2\pi}.$$
(20)

Neither of these parameters are discontinuous themselves, since they represent the sine and cosine of this shift angle. Indeed, this convention also lets us explore the connection between the above network structure and the method of slicing as described by Budanur *et al.* [4], see section 3.1.

The transformed state is then manipulated according to a discrete symmetry group that is hardcoded into the network, according to the global discrete symmetries in the system of equations,

$$[\mathbf{X}'_0 \dots \mathbf{X}'_l \dots \mathbf{X}'_n] = [\mathbf{G}^0 \mathbf{X}^* \dots \mathbf{G}^l \mathbf{X}^* \dots \mathbf{G}^n \mathbf{X}^*].$$
(21)

For each of these discrete symmetry operations a separate branch is created to which the autoencoder is applied,

$$[\hat{\mathbf{X}}_0'\dots\hat{\mathbf{X}}_l'\dots\hat{\mathbf{X}}_n'] = [D(E(\hat{\mathbf{X}}_0'))\dots D(E(\hat{\mathbf{X}}_l'))\dots D(E(\hat{\mathbf{X}}_n'))].$$
(22)

A network with split branches each containing the same subnetwork is known as a Siamese network. We note that the weights of the autoencoder are identical across each of these parallel branches. In the case of reflectional symmetries, two branches are created, where \mathbf{G}^0 is the identity and \mathbf{G}^1 is the reflection operation $(u, v)(x, y) \rightarrow (u, -v)(x, -y)$. We then obtain a modal reconstruction of each of the modified samples. As a

next step, we identify the branch, l, that has the lowest reconstruction error with respect to its input to the autoencoder, i.e.

$$l = \arg\min_{i} (|\mathbf{X}'_{i} - \hat{\mathbf{X}}'_{i}|_{2}).$$
⁽²³⁾

Lastly, the reconstruction with the lowest error is retransformed into the input space. In other words, we apply the inverse of the discrete symmetry operation that was used for the creation of the chosen branch,

$$\hat{\mathbf{X}}^* = \mathbf{G}^{-l} \hat{\mathbf{X}}_l^{\prime}. \tag{24}$$

In the special case of the reflectional symmetry each transform is its own inverse. The inverse of the continuous symmetry group for the parameter s

$$\hat{\mathbf{X}} = \mathbf{F}^{-1}(\mathbf{s})[\hat{\mathbf{X}}^*] \tag{25}$$

is then applied. For the previously described case (the group of continuous translations) this becomes

$$\hat{\mathbf{X}} = \mathcal{F}^{-1}(e^{-i\mathbf{k}2\pi\frac{-s}{L}}\mathcal{F}(\hat{\mathbf{X}}^*)) \quad \forall k \in \mathbf{k}.$$
(26)

As a result, we achieved a reconstruction of the original sample in input space – with all invariances encoded into the same latent space. For a full description of the input the shift parameter s and the branch l are now part of the latent space as well. Note that for training the weights we calculate the error $|\mathbf{X} - \hat{\mathbf{X}}|_2$ and backpropagate it only through the chosen branch and the *STN*. In the case of s-nlPCA, the autoencoder branches are nonlinear, whereas in s-PCA the branches are set to be linear. To make s-PCA more akin to traditional PCA we train the network with linear autoencoders and perform PCA on the modified dataset comprised of samples \mathbf{X}'_l . Hence, the symmetry-aware modes are automatically orthogonal, and we deduce an energetic relevance measure for each mode from the associated eigenvalues.

When comparing the computational effort of this model to that of a conventional AE design we do not notice a significant increase in cost. While the network calculates the forward pass of each sample in parallel, which is comparable to the cost of a matrix multiplication, the more costly backpropagation is performed on only one branch (with the minimal error). As a result, the parallel branches of the Siamese network do not substantially increase the cost. The *STN*, however, requires a small additional training effort, equivalent to an encoder for a two-dimensional latent space.

3.1. STNs and continuous translations

Given a dataset, $\mathbf{X}(\mathbf{x}, \mathbf{t})$, with periodic boundary conditions in \mathbf{x} on the domain $[-\pi, \pi]$, the method of slicing aims at aligning each snapshot by setting the phase of the first complex Fourier coefficient, \underline{X}_1 , to 0. All subsequent modes are then shifted by this change in phase. We employ the DFT and describe the method as follows

$$\underline{X}(k,\mathbf{t}) = \sum_{n=0}^{N_x - 1} \mathbf{X}(x_n, \mathbf{t}) e^{-i\frac{2\pi}{N_x}kn} = (\mathbf{C}_k - i\mathbf{S}_k)\mathbf{X}(\mathbf{x}, \mathbf{t}) \quad \forall k \in \mathbf{k},$$
(27)

where **k** is the set of wavenumbers in the frequency domain and N_x the dimension of **x**. We define

$$\mathbf{C}_{k} = \begin{bmatrix} 1 & \cos(\frac{2\pi}{N_{x}}k) & \dots & \cos(\frac{2\pi}{N_{x}}(N_{x}-1)k) \end{bmatrix}$$
(28)

$$\mathbf{S}_k = \begin{bmatrix} 0 & \sin(\frac{2\pi}{N_x}k) & \dots & \sin(\frac{2\pi}{N_x}(N_x - 1)k) \end{bmatrix}.$$
(29)

The shift coefficient s and the subsequently shifted Fourier coefficients, $\underline{\mathbf{X}}^*$, are then defined as

$$s = \arctan2(-\mathbf{S}_1\mathbf{X}(\mathbf{x}, \mathbf{t}), \mathbf{C}_1\mathbf{X}(\mathbf{x}, \mathbf{t}))$$
(30)

$$\underline{X}^{*}(k,\mathbf{t}) = \underline{X}(k,\mathbf{t})e^{-ik2\pi\frac{d}{L_{x}}} \quad \forall k \in \mathbf{k}.$$
(31)

The relation between the method of slicing and the STN becomes apparent when the STN is set to be a linear network without biases. Recall that with eq. (15) and eq. (20) this linear STN takes on the form

$$\mathbf{s} = \mathbf{W}_{STN} \mathbf{X},\tag{32}$$

where $\mathbf{s} = [s_1 \ s_2]$ and $\mathbf{W}_{STN} \in \mathbb{R}^{2 \times N_{\xi}}$ are the weights of the *STN*. Thus, the rows of the weights of a linear *STN*, \mathbf{W}_{s_1} and \mathbf{W}_{s_2} , henceforth referred to as shift-modes, can be viewed as mode-like objects with an associated spatial structure.

Further, we introduce spatial smoothing by convolving the weights of the STN with Gaussian kernels after each training step. For one-dimensional problems we convolve the weights with

$$\mathbf{B} = \frac{1}{4} \begin{bmatrix} 1 & 2 & 1 \end{bmatrix},\tag{33}$$

whereas for two-dimensional cases we use

$$\mathbf{B} = \frac{1}{16} \begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{bmatrix}.$$
 (34)

Both expressions represent approximations of a Gaussian filter with variance $\sigma = 1/\sqrt{2}$.

With these conventions, the weights of the network predicting s_1 and s_2 converge to a sine and cosine with wavenumber k = 1 in the translated direction when weight smoothing is applied, see section 4.1 and section 4.3. That is

$$\mathbf{W}_{s_1} \approx \left[\cos(\phi) \quad \cos(\frac{2\pi}{N} + \phi) \quad \dots \quad \cos(\frac{2\pi}{N}(N-1) + \phi)\right] \tag{35}$$

$$\mathbf{W}_{s_2} \approx \left[\sin(\phi) \quad \sin(\frac{2\pi}{N} + \phi) \quad \dots \quad \sin(\frac{2\pi}{N}(N-1) + \phi)\right]. \tag{36}$$

In essence, the STN is then a DFT calculating the first odd and even Fourier coefficient with a random phase shift ϕ . A linear STN is thus equivalent to the method of slicing, up to a random phase shift ϕ and computational errors. In addition, the shift coefficients s of a linear network are (up to a constant phase shift) entirely identical to the ones for a nonlinear STN, as we will discuss in section 4.1. However, the weights of our network only converge to a pure sine and cosine function in the translated, homogeneous direction. No constraint is placed on the modulation of the weights in the non-translated directions, but such a modulation may be desirable under certain circumstances, see Marensi *et al.* [16].

3.2. Discrete symmetries and Siamese networks

To decompose the dataset into relevant parts, conditional methods have been employed in the past, see e.g. Schmidt & Schmid [27]. In these cases, additional conditions, such as pressure or dissipation signals, are evaluated *a priori*, and only snapshots satisfying prescribed criteria are included in the dataset on which PCA is performed. Conceivably these methods can also be applied to filter the dataset for discrete symmetrical states; However, they would require additional work and may decrease the number of samples in the dataset by a degree that compromises the statistical convergence of the covariances for a meaningful PCA. Our proposed network performs the same task when using linear AEs, without requiring additional input while avoiding any reduction in the dataset. For the example of a diffuser, see section 4.2, where the flow exhibits the discrete group of reflections we will show that the signal of l, i.e. the choice of the parallel branches used for model reduction, largely coincides with the y-position of the problem's barycenter, which expresses the symmetry of the system.



Figure 3: Temporal evolution of Burgers' equation for the chosen initial condition (see text).



Figure 4: Visualization of the symmetry-reducing autoencoder architecture used for Burgers' equation. A sample **X** is passed to the network. The input is then fed to a secondary network, called an *STN*, to generate a shift parameter s ($\mathbf{X} \rightarrow s$). The input is then shifted in x with $\mathbf{T}(s)$ ($\mathbf{X} \rightarrow \mathbf{X}^*$). We next apply an autoencoder on this transformed snapshot ($\mathbf{X}^* \rightarrow \mathbf{\hat{X}}^*$). The resulting reconstruction is transformed back to the original input space by applying the inverse shift, $\mathbf{T}(-s)$ ($\mathbf{\hat{X}}^* \rightarrow \mathbf{X}'$). Both the autoencoder and the STN are trained by backpropragating the reconstruction error $|\mathbf{\hat{X}} - \mathbf{X}|_2$.

4. Results

4.1. Burgers' equation

As a first example, we will investigate a dataset generated by approximating the inviscid Burgers' equation discretized with a first-order upwind scheme in space, where viscosity is introduced implicitly by the discretisation scheme. The domain $\mathbf{x} \in [0, 2\pi]$ with periodic boundary conditions is resolved with $N_x = 64$ grid points, and the initial condition

$$u'(x,0) = \begin{cases} 1 & \text{for } x \le \pi \\ 2 & \text{for } \pi < x \end{cases}$$
(37)

is integrated using the first-order Euler time scheme with $\Delta t_{int} = 0.0001$ until t = 10. We extract 1000 samples with $\Delta t_{rec} = 0.01$. Since the solution asymptotically converges for large times, we define $\mathbf{u}(\mathbf{x}, \mathbf{t}) = \mathbf{u}'(\mathbf{x}, \mathbf{t}) - \lim_{t\to\infty} \mathbf{u}'(\mathbf{x}, t) = \mathbf{u}'(\mathbf{x}, \mathbf{t}) - 1.5$. We observe from fig. 3 a solution that is both advecting in x and non stationary in the moving reference frame. The classical model reduction method of PCA struggles with both of these characteristics, leading to high reconstruction errors and poorly interpretable modes.

Since the dataset only exhibits continuous translations, the network from fig. 2 simplifies to the one showcased in fig. 4. Instead of requiring Siamese branches for discrete symmetry groups, we only need to



Figure 5: L_2 reconstruction error for PCA, nlPCA, s-PCA and s-nlPCA for Burgers' equation over the number of modes used for the reconstruction.



Figure 6: Shift modes $\mathbf{W}_{s,1}$ and $\mathbf{W}_{s,2}$ and $\sin x$ and $\cos x$ (left). Predicted shift s using a linear and nonlinear STNs (right).

align the flow in the continuous direction. The AE components, i.e. E and D, can be chosen to be either linear or nonlinear. The exact architectures together with the hyperparameters used for training the linear (table A.1) and nonlinear (table A.2) AEs are given in the appendix.

We compare the normalized L_2 reconstruction error, $|\hat{\mathbf{X}} - \mathbf{X}|_2/|\mathbf{X}|_2$, over the number of modes used for all four methods. In fig. 5 we observe that all other methods surpass the reconstruction capabilities of PCA. In addition, a plateau in the reconstruction error at 6×10^{-7} is reached for both s-PCA and nlPCA, indicating that the most relevant features have been successfully captured. For s-PCA and nlPCA this limit is approximately reached with N = 8, while conventional PCA fails to converge to this level even with the maximum number of considered modes, N = 16. Using s-nlPCA, the same limit is achieved with only two modes. We notice that the L_2 -error for nlPCA using only one mode is significantly higher than using two modes, which is a consequence of the solution essentially being a modified travelling wave, which requires a complex mode pair to be properly represented; prescribing only one latent variable yields a nonphysical approximation of this complex mode by the network. In the case of s-nlPCA, on the other hand, a nearly complete reconstruction of the input is achieved with merely one mode, since the translation is handled by the symmetry-aware approach and only the dissipation needs to be approximated.

To highlight the similarities between the STN and the method of slicing, we produce the weights of the STN for a linear configuration without biases, see fig. 6. As mentioned above, setting the STN to be nonlinear does not bear any decrease in error and generates the same shift for each timestep, up to a constant phase difference, see fig. 6. In fig. 6 the weights of the network, i.e. the shift-modes, converge to $\sin(x + \phi)$ and $\cos(x + \phi)$, with ϕ as a random phase shift. Hence, in this case the method of slicing and the STN are equivalent. The shifts generated by a nonlinear STN largely coincide with the shifts from linear case – up to a constant phase shift. Consequently, the method of slicing appears to be the optimal way of aligning samples in this case. Another interesting point is the slope of the shift s for the considered methods. In all cases, $\frac{ds}{dt} \approx \frac{3}{2}$, which is exactly the shock speed of the Riemann problem for the inviscid solution.



Figure 7: Normalized modal energies of the PCA and s-PCA eigenvalue problem for Burgers' equation (left). PCA (top right) and s-PCA (bottom right) modes for Burgers' equation for $i \in \{1, 2, 3, 4\}$.

4.1.1. Linear modal representations

Having already shown the dominance of s-PCA over PCA when considering the reconstruction error, we expect to extract more physically meaningful modes using s-PCA. In the latter case, a Fourier-like basis is obtained in the *x*-direction, due to the periodic boundary conditions, and consequently large modal energies ensue even for higher modes. We emphasize that the identified basis is not a pure Fourier basis, since such a basis is only attained for statistically stationary solutions. Since s-PCA does not need to approximate the convective features, we expect the modes to largely model the diffusion of the initial condition. We present the normalized modal energy

$$e_i = \frac{\lambda_i}{\sum_i \lambda_i} \tag{38}$$

in fig. 7 for both the PCA and s-PCA eigenvalue problem alongside their respective modal shapes. We recall that the eigenvalues for s-PCA are computed by performing a classical PCA on the modified dataset, obtained by compiling the inputs to the autoencoders into the symmetry-aware space. From fig. 7 we can see that the energy spectrum for standard PCA is broader when compare to s-PCA. While we observe a steady decrease in modal energy, the relative energy of e.g. mode 30 is still rather high at $e_{PCA,30} \approx 1 \times 10^{-6}$. On the other hand the slope of the energy decrease for s-PCA is noticeably steeper, and $e_{s-PCA,30} \approx 1 \times 10^{-12}$, suggesting a far small contribution of this mode to the overall solution. While the first mode in both scenarios seems to have similar energy, $e_{PCA,1} = 0.4813$ and $e_{s-PCA,1} = 0.9776$, the energy of the first s-PCA mode is more than twice the value of the corresponding PCA mode.

As expected, the modes for PCA in this translationally invariant problem take on Fourier-like shapes. Some minor deviation from a proper periodic base can be observed due to diffusion. We can clearly see that for modes ϕ_3 and ϕ_4 we have a decreasing amplitude for increasing x. The modal shapes for s-PCA are not as smooth as for PCA. Since in the case of s-PCA we only have to approximate the dissipation within the problem, the modes are conditioned accordingly. However, since the initial solution is discontinuous, which gradually progresses towards a smooth solution, this transition has to be modeled in the linear framework. The same issue can be observed for the original PCA modes. For example the first mode shows a discontinuous slope at $x = \pi$, which coincides with the location of the original shock. However, the discontinuity gets smoothed out due to the translational invariance overshadowing such issues. The appearance of the discontinuity in s-PCA is an indicator for the removal of the artificial smoothing introduced by the Fourier basis and, as a result, for a better approximation of the input field.

4.1.2. Equivalent PCA energy for different reconstructions

As stated above, it is challenging to visualize nlPCA and consequently s-nlPCA modes in a satisfactory manner. We can, however, project the nonlinearly reconstructed solutions onto the original PCA modes in order to procure an equivalent PCA energy (section 2.3), see fig. 8. We show this equivalent energy of the reconstructed solutions using one s-PCA mode. Alternatively, two modes are used for nlPCA, owing to the



Figure 8: Equivalent PCA energy of Burgers' equation for an s-PCA reconstruction using one mode (left), an nlPCA reconstruction using two modes (middle) and an s-nlPCA reconstruction using one mode (right).

complex-valued nature of a travelling wave, and one mode for s-nlPCA.

When projecting the s-PCA reconstruction onto the original PCA modes, we see that one s-PCA mode captures the energy of the first 15 PCA modes almost exactly, and there exists a rather good agreement for the rest of the modes. Both nonlinear methods, i.e. nlPCA and s-nlPCA, capture large portions of the linear energy content. With two nlPCA modes nearly 18 linear modes can be captured, and an agreeable approximation of the remaining modes is observed. More apparent is the exact capturing of the first 20 linear modes when using s-nlPCA with only one mode. When comparing to the linear reconstruction, we can see a clear increase in detained energy when introducing nonlinearities into the model reduction. Even in the absence of convection, the structure of the shock still changes due to diffusion. Hence, we cannot expect to model this behaviour with simply one linear mode. The nonlinear model, s-nlPCA, however, can well approximate this behavior, once the convective component has been removed.

4.2. Flow through a step diffuser

As a more complex example the sudden expansion of the flow behind a step diffuser is investigated. The dataset is generated using Freefem++, solving the incompressible Navier-Stokes equations on an unstructured grid, see Hecht [9]. The flow field is composed of the inflow $(\mathbf{x}_{in} \in [-3, 0), \mathbf{y}_{in} \in [-0.5, 0.5])$ and the main body $(\mathbf{x}_m \in [0, 15], \mathbf{y}_m \in [-1.5, 1.5])$, and simulations are run at $Re_h = \frac{U_m h}{v} = 150$, where U_m represents the maximum velocity of the parabolic inflow velocity profile and h is the channel height.

The symmetric setup of the problem introduces the discrete group of reflections, \mathbf{R} : $(u, v)(x, y) \rightarrow (u, -v)(x, -y)$. As such, the flow exhibits a nonlinearly bistable behavior as indicated by the bifurcation diagram in fig. 9. An initial symmetric solution is found using a Newton method for Re= 70 by prescribing a uniform $(\mathbf{u}, \mathbf{v}) = (1, 0)$. To obtain an asymmetric solution an initial guess was set to $(\mathbf{u}, \mathbf{v}) = (0.7, 0.15)$ and the Reynolds number was increased until the Newton method did no longer converge to a symmetric solution. The Reynolds number was further increased and the symmetric and asymmetric solutions for the previous Reynolds number were used as initial guesses for subsequent Reynolds numbers. The bifurcation diagram is then generated by calculating

$$y_{cu} = \frac{\int u(x, y)y dx dy}{\int u(x, y)dx dy}$$
(39)

for each of the discovered solutions. fig. 9 shows that the flow remains symmetric for Re < 79. Beyond this value the symmetric state becomes linearly unstable. In addition, two nonlinearly stable attractors, which become increasingly asymmetric for larger Reynolds number, are found. These attractors are mutually invariant under the group of discrete reflections. Since these attractors are both nonlinearly stable we introduce a momentum forcing of the form $n(t)\mathbf{f}(\mathbf{x},\mathbf{y})$ to the inflow region. Here, n(t) is a Gaussian white noise, and $\mathbf{f}(\mathbf{x},\mathbf{y})$ corresponds to the unstable adjoint global mode of the symmetric base-flow. Under this forcing, the evolution equations are integrated in time using $\Delta t_{int} = 0.0075$, and 36862 samples are extracted using $\Delta t_{rec} = 7.5$. In order to increase robustness of the PCA, we concatenate our dataset with



Figure 9: Bifurcation diagram of the flow through a sudden diffuser (left). The flow has one linearly stable attractor until $Re \approx 79$. Past this value, it splits up into one linearly unstable symmetric state (dashed) and two nonlinearly stable attractors (solid) that are reflectionally invariant. A symmetric, unstable solution A at Re = 150 (top right), an asymmetric nonlinearly stable solution B at Re = 150 (middle right) and a disturbed asymmetric solution B - A (bottom right) are shown.



Figure 10: Visualization of the symmetry-reducing autoencoder architecture used for the sudden expansion problem. A sample **X** is passed to the network, and the predefined reflectional symmetry is applied, $(\mathbf{X} \to \mathbf{X}')$. The network splits into two branches where each branch is either the original input or the reflected state. The same autoencoder is then applied to all branches in parallel, $(\mathbf{X}' \to \hat{\mathbf{X}}')$. The resulting reconstructions are merged by selecting the branch with the minimal L_2 -error to its respective input, $l(\hat{\mathbf{X}}' \to \hat{\mathbf{X}}'_l)$. This state is finally retransformed by applying the reflectional symmetry again if the reflected branch was chosen, $(\hat{\mathbf{X}}'_l \to \hat{\mathbf{X}})$. The autoencoder is trained by backpropragating the reconstruction error $|\hat{\mathbf{X}} - \mathbf{X}|_2$.

the symmetric transformations of each sample, an approach advocated by Sirovich [28]. In the following we utilize the disturbance velocities, $\mathbf{u}' = \mathbf{u} - \mathbf{u}_{base}$ and $\mathbf{v}' = \mathbf{v} - \mathbf{v}_{base}$, where values with the subscript *base* are those of the symmetric, unstable solution. The symmetric unstable solution A, one of the aforementioned attractors B and the disturbed state B - A are shown in fig. 9. To facilitate the application to autoencoders, the solution was interpolated onto a rectangular grid with $\Delta x = \Delta y = 0.05$, which still contains both the inflow and the main body.

Since we only have the discrete reflectional symmetry group, our network from fig. 2 simplifies to the one illustrated in fig. 10. Instead of having to align the flow in a continuous direction we only need to create two Siamese branches to handle the reflectional symmetries. The AE can either be linear or nonlinear. See table A.3 and table A.4 for linear and nonlinear AE architectures as well as the hyperparameters used for training.

The L_2 reconstruction errors for all considered methods, i.e. PCA, nlPCA, s-PCA and s-nlPCA, are compiled in fig. 11. There is a clear progressive reduction in L_2 reconstruction error when moving from



Figure 11: L_2 reconstruction error for PCA, nIPCA, s-PCA and s-nIPCA for the flow through a step diffuser versus the number of modes used for the reconstruction.



Figure 12: Branch chosen by the s-PCA network resulting in the lowest reconstruction error, l, and normalized center of pressure in the y-direction, $y_{cp,n}$, over time t.

PCA to s-PCA to nlPCA to s-nlPCA. Moreover, we can investigate how this reduction in error is related to other methods, such as conditional PCA. For example, we may perform conditional PCA on our dataset by transforming the dataset depending on the normalized center of pressure in the y-direction

$$y_{cp,n} = \frac{y_{cp}}{\max_t(y_{cp})},\tag{40}$$

where we either execute or omit the invariance operation, depending on the sign of the pressure center. In fig. 12 we show the branch chosen by the network, l, that results in the lowest reconstruction error, where -1 is branch 1 and 1 is branch 2. In the same graph we plot the normalized center of pressure in the y-direction. We see how the development of the chosen branch, l, largely coincides with the signal of the normalized pressure center $y_{cp,n}$. The difference in sign is only due to initialization. Our neural network then automatically performs a conditional PCA, without the need for a priori data manipulation or the definition of a global, conditioning quantity such as the pressure center.

However, due our limitation to only two possible transformations, rather than an infinite range of them, the difference in error is not as pronounced as for the previous case of Burgers' equation. Nonetheless, this exercise serves as a proof of concept for the removal of discrete symmetry groups from our dataset and paves the way for more complex discrete groups. Additionally, the interpretability of linear modes, which we will focus on next, is enhanced.

4.2.1. Linear modal representations

In fig. 13 the normalized modal energy

$$e_i = \frac{\lambda_i}{\sum_i \lambda_i} \tag{41}$$

for PCA and s-PCA is shown. The complexity of the flow through a step diffuser exceeds the simpler flow



Figure 13: Normalized modal energies of the PCA and s-PCA eigenvalue problem for flow through a step diffuser.



Figure 14: u-Modes, $i \in [1, 4]$, created with PCA (left) and s-PCA (right).

behavior governed by Burgers' equation, as the modal energy for mode 30 is still rather high at around 1×10^{-3} . The use of s-PCA lowers the energy of higher modes, though not to the degree we observed for Burgers' equation. In the case of s-PCA the modal energy of the first mode is increased, while the modal energy of the second mode is diminished to nearly the energy level of the original third mode. A crossover point at i = 10 is reached, where the modal energies of the s-PCA modes consistently fall below those of PCA, indicating a superior representation of key features with the first ten modes.

We proceed by comparing the modes created by PCA and s-PCA in fig. 14. The PCA modes adhere to the condition, given in Sirovich [28], of becoming symmetric and antisymmetric. Additionally, the third and fourth PCA modes form an antisymmetric travelling wave, which can be deduced from their nearly identical modal energy in fig. 13. Due to the instability of the symmetric solution, however, the samples are asymmetric at most times.png. When moving to s-PCA, the use of the network eliminates the constraint on the modes to be either purely symmetric or purely antisymmetric, and symmetric and antisymmetric modes merge into composite modes. Evidence for this phenomenon is found in the fact that the equivalent mode for PCA mode 2 does not appear in the mode set for s-PCA. More specifically, the first and second PCA modes appear to combine into the first s-PCA mode. We will investigate this connection further in fig. 15. The modes that describe a travelling wave solution are now modes 2 and 3, and this wave is no longer antisymmetric, but asymmetric. Mode 4 merely describes the flow in the inlet, further confirming that the bulk of the principle dynamics has been captured by the first three modes.

In fig. 15 we show the first mode generated by s-PCA together with

$$\tilde{\mathbf{\Phi}} = \frac{\lambda_1 \mathbf{\Phi}_1 + \lambda_2 \mathbf{\Phi}_2}{|\lambda_1 \mathbf{\Phi}_1 + \lambda_2 \mathbf{\Phi}_2|}.$$
(42)



Figure 15: **u**-Mode 1 of s-PCA (left) and $\hat{\Phi}$, generated by combining the first two PCA modes (right).



Figure 16: Equivalent PCA energy for flow through a step diffuser using an s-PCA reconstruction (left), an nlPCA reconstruction (middle) and an s-nlPCA reconstruction (right) based on four modes each.

We see how the combination of the first two PCA modes, $\hat{\Phi}$, closely resembles the first s-PCA mode, even though a few minor discrepancies remain. We note also the resemblance of both modal shapes to the disturbed state in fig. 9; the same resemblance for higher modes, e.g. mode 3 and mode 10, ceases to exist.

4.2.2. Equivalent PCA energy for different reconstructions

In fig. 16 the equivalent PCA energy of s-PCA, nlPCA and s-nlPCA using four modes, calculated by projecting the reconstructions onto the original PCA modes, see section 2.3, are shown. First, we note that s-PCA using four modes can capture five PCA modes exactly. In addition, we can detect an increase in the equivalent PCA energy of modes $i \in \{10, 11, 15\}$ above the ground level of the remaining modes. Interestingly, these latter modes are symmetric, whilst modes $i \in [1, 9] \setminus 2$ are antisymmetric. This again points towards s-PCA merging antisymmetric and symmetric modes into composite modes. Both nlPCA and s-nlPCA capture nearly five linear modes exactly and raise the energy of some of the remaining higher modes – s-nlPCA more so than nlPCA. Remarkably, for nlPCA the energies for linear modes $i \in \{10, 11, 15\}$ are augmented above the base level compared to the other modes, corroborating that the nonlinearities in the autoencoder merge symmetric and antisymmetric modes as well. While s-nlPCA does increase the energy of the symmetric modes, the principal improvement is directed towards modes 6 and 7, which now nearly match the linear energy. From this we deduce that symmetry reduction frees up capabilities in the network – otherwise spent to combine symmetric and antisymmetric modes – that are employed towards a better approximation of higher modes.

4.3. Kolmogorov flow

We shall now investigate the more complex scenario of Kolmogorov flow, an example that involves both continuous and discrete symmetry groups. We used the code by Wan *et al.* [29], which simulates Kolmogorov flow first described by Meshalkin & Sinai [19]: it is a spatially periodic, forced Navier-Stokes solution for periodic boundary conditions in both x and y. The equation is often used to mimic turbulent behaviour in two dimensions. We choose a domain size of 32×32 , a forcing wavenumber of k = 4 in the y-direction and Re = 35, which leads to a weakly chaotic regime ($Re_{crit} = 33.6$ [5]). We collect 300000 samples by integrating a random initial condition with $\Delta t_{int} = 0.005$ and recording vorticity fields in intervals $\Delta t_{rec} = 0.1$. Due to the large size of the dataset and the number of discrete symmetries, data manipulation based on symmetry transformations, as was done in section 4.2, was not performed. We gather the vorticity ($\boldsymbol{\omega}$) snapshots,



Figure 17: Snapshots of vorticity, $\boldsymbol{\omega}$, for Kolmogorov flow at Re = 35 at two different times, exhibiting low (left) and high (right) dissipation.



Figure 18: Visualization of the symmetry-reducing autoencoder architecture used for Kolmogorov flow. A sample \mathbf{X} is passed to the network. A secondary network, an STN extracts a continuous transformation parameter s from the sample $(\mathbf{X} \to s)$. The original sample is then shifted with $\mathbf{T}(s)$ and s is passed on, $(\mathbf{X} \to \mathbf{X}^*)$. Each transformation of the set of predefined discrete transformations, \mathbf{G} , is then applied, $(\mathbf{X}^* \to \mathbf{X}')$. The network splits into 16 branches where each branch is a candidate for a discrete transform of the modified input. The same autoencoder is then applied to all branches in parallel, $(\mathbf{X}' \to \mathbf{\hat{X}}')$. The resulting reconstructions are subsequently merged by choosing the branch, l, with the least L_2 -error to its respective \mathbf{X}'_l , $(\mathbf{\hat{X}}' \to \mathbf{\hat{X}}'_l)$. We next apply the inverse discrete transform of \mathbf{G}^l , \mathbf{G}^{-l} $(\mathbf{\hat{X}}'_l \to \mathbf{\hat{X}}^*)$. The sample is un-shifted by $\mathbf{T}(-s)$ to bring the reconstruction to the original input space, $(\mathbf{\hat{X}}^* \to \mathbf{\hat{X}})$. We train the autoencoder as well as the STN by backpropragating the reconstruction error $|\mathbf{\hat{X}} - \mathbf{X}|_2$.

for a low- and high-dissipation event, shown in fig. 17. Note that, due to the doubly periodic boundaries, vorticity can straightforwardly be transformed back into the velocity components \mathbf{u} and \mathbf{v} .

Due to the periodic boundary conditions and the wavenumber of k = 4 for the forcing in y, multiple symmetry groups for the exact solution exist: the group of continuous translations in x, i.e. $\mathbf{T}(s): \omega(x, y) \rightarrow \omega(x + s, y)$, the 8-cyclic shift and reflect symmetry $\mathbf{S}: \omega(x, y) \rightarrow \omega(-x, y + \pi/4)$, and the 2-cyclic rotational symmetry through $\pi, \mathbf{P}: \omega(x, y) \rightarrow \omega(-x, -y)$, see e.g. Chandler & Kerswell [5]. In what follows, we combine the two discrete symmetry groups into a group $\mathbf{G} = [\mathbf{P}_0 \mathbf{S}_0, \dots, \mathbf{P}_0 \mathbf{S}_7, \mathbf{P}_1 \mathbf{S}_0, \dots, \mathbf{P}_1 \mathbf{S}_7]$. Hence, our network is now a combination of an *STN* that compensates for the translational invariance, and 16 Siamese branches that remove the discrete group \mathbf{G} from our dataset. From fig. 18 we see that for Kolmogorov flow the network needs to cope with the full complexity shown in fig. 2. In a first step the *STN* determines by how much each sample has to be shifted in the x-direction. Afterwards, the network is split into 16 parallel branches,



Figure 19: L_2 reconstruction error for PCA, nlPCA, s-PCA and s-nlPCA for Kolmogorov flow versus number of modes used for the reconstruction.



Figure 20: Number of times each branch was chosen as the one generating the minimum error, l (left) and shifts generated by the method of slicing and the STN over an arbitrary time span (right).

each representing a candidate transformation of G with the same autoencoder. From these branches we select the one with the lowest reconstruction error with respect to its input, before applying the inverse of the discrete and shifting transforms to arrive back in input space.

In table A.5 and table A.6 linear and nonlinear AE architectures as well as hyperparameters used for training are collected. For the Kolmogorov-flow dataset a hybrid AE employing both convolutional layers as well as fully connected layers have been used. The normalized reconstruction errors, $|\mathbf{X} - \mathbf{X}|_2/|\mathbf{X}|_2$, for all considered methods, are visualized in fig. 19. This figure shows that all alternative methods outperform PCA by a wide margin. While in PCA the error starts to decrease in a meaningful way for N > 32, s-PCA exhibits a substantial decrease at around N = 8 and shows some sign of levelling off at N = 128, which is absent for PCA. The error for nlPCA is initially higher than for s-PCA, but crosses below the latter at N = 2, to then decrease with a similar slope. Finally, the error for nlPCA tapers off at N = 64, at seemingly the same error s-PCA is converging to. Convergence of s-nlPCA also fades at N = 64, however, its error is lower everywhere, including the ultimate error at N = 128.

In fig. 20 we show the number of times each branch was chosen as the one with the lowest reconstruction error, *l*. Additionally, the connection between the method of slicing and *STN* is investigated. As discussed above, the *STN* does not need to be nonlinear to properly discover the required shift to align the samples. For this reason, we can compare the shift created by the method of slicing and the equivalent shift generated by our *STN*. For the discrete-symmetry-reducing step there exist 16 symmetry transformations, $G = [P_0S_0, \ldots, P_0S_7, P_1S_0, \ldots, P_1S_7]$, resulting in 16 Siamese branches for the autoencoder, which allows us to compute the number of times, *N*, each branch was chosen as the minimum-error option. This provides insight into the statistical convergence of our dataset in terms of the discrete invariant transformations. We observe that the part of **G** connected to P_0 , i.e. without the application of the rotation through π , remains largely unpopulated, since occurrences there are statistically insignificant. In addition, only four shift-and-reflect branches of P_1 are populated above N = 100. This comes as no surprise since it is deceptive to assume that weak turbulence, as encountered in our example, is necessarily ergodic, see e.g. Holmes *et al.*



Figure 21: Shift modes $\mathbf{W}_{s,1}$ (left) and $\mathbf{W}_{s,2}$ (right).



Figure 22: Normalized modal energies of the PCA and s-PCA eigenvalue problem for Kolmogorov flow at Re = 35.

[10]. If PCA were to be employed to reduce our dataset, it would be non-robust in its outcome and result in a largely unrepresentative basis; in this case, only an exhaustive combinatorial search would identify proper bases and yield meaningful results. Regarding robustness, our method does not require ergodicity with respect to the symmetry groups, but, rather, only ergodicity within the symmetry-reduced space. For the continuous symmetry group, similar to the Burgers' equation example, the predicted shift of each sample largely coincides with the shift given by the method of slicing, up to an arbitrary phase shift which does not impact the model reduction capabilities. For the one-dimensional case of Burgers' equation, we reported a convergence of the weights of the STN to trigonometric functions in the translated direction. We anticipate the same behavior here, even though a prediction on the shape of the weights in the homogeneous direction y cannot be made. According to Marensi et al. [16], a dependence of the transformation parameter on the non-translated, inhomogeneous directions may be required, if discontinuities of the predicted shift are to be avoided. For Kolmogorov flow, this dependence is not necessary, since the amplitude of the first complex Fourier coefficient never reaches zero. Investigating the weights reveals that such a dependence can theoretically be learned, without any further user input. The weights, i.e. the shift-modes, are compiled in fig. 21. As expected, the weights take on a shape similar to $\cos x$ and $\sin x$. However, we do notice a modulation of the weights in the y-direction. It appears reasonable that this modulation in y assumes a wavenumber of k = 4, aligning with the forcing wavenumber of Kolmogorov flow in our example. While this modulation is not strictly necessary in our example, it illustrates that the autonomous detection of such an intrinsic modulation is feasible. To aid convergence, we could conceivably prescribe the shape of $\cos x$ and $\sin x$ and let the network only discover the shape in y. We did however refrain from this step, since the network showed satisfactory convergence without it. As a next step, we will investigate the use of a linear autoencoder in our symmetry-aware framework and compare the results to PCA.

4.3.1. Linear modal representations

In fig. 22 the normalized modal energy

$$e_i = \frac{\lambda_i}{\sum_i \lambda_i} \tag{43}$$

for PCA and s-PCA, is compared. For PCA the modal energies remain rather large for all modes shown



Figure 23: Vorticity modes for Kolmogorov flow using PCA (top) and s-PCA (bottom) for $i \in \{1, 2, 3, 4\}$.

here, never decreasing below 7×10^{-4} . We can observe typical paired energies for a number of modes, such as modes 1 and 2 and modes 4 and 5, indicating that the accociated modes describe a travelling-wave structure in a Fourier basis. The modal energies of s-PCA modes decrease below the corresponding value of PCA modes for i > 1. The first s-PCA mode now carries over 80% of the energy content of the whole solution compared to 20% for PCA. All consecutive modes are less important than their PCA counterparts, and s-PCA mode 35, for example, carries as little energy as the 100th PCA mode.

The modal shapes for both PCA and s-PCA are compared in fig. 23, in an attempt to explain the difference in reconstruction error. As expected, the PCA modes assume a Fourier-like shape in the *x*-direction. As discussed above, modes 1 and 2 form a mode pair that together reflects a translational dynamic. Mode 3 shows little modularity in *x* but retains a wavenumber k = 4 in its *y*-direction. The same pattern can be observed in modes 1 and 2. This wavenumber stems from the stationary forcing applied to the Navier-Stokes equations to generate Kolmogorov flow. Mode 4 appears to be a higher harmonic of the previous modes. It is obvious that this basis is not very representative of the low-dissipation snapshot in fig. 17, in whose vicinity the solution resides for a majority of time. When moving to s-PCA, a different behavior emerges. Mode 1 now represents a distorted version of the equilibrium state at this Reynolds number; see e.g. Platt *et al.* [24] for this equilibrium state. Modes 2 and 3 now approximate a vortex pair with centers at $(x_{c,1}, y_{c,1}) \approx (\pi/2, 5\pi/4)$ and $(x_{c,2}, y_{c,2}) \approx (3\pi/2, \pi/4)$. This approximation is relatively crude, owing to the linear nature of the s-PCA. Mode 4 models the structures shed by the vortex pairs. We shall now investigate how well each of the non-PCA reconstructions capture the linear PCA modes.

4.3.2. Equivalent PCA energy for different reconstructions

To gather additional information, we once again rely on the equivalent PCA energy, see section 2.3, for s-PCA, nlPCA and s-nlPCA, using two modes, see fig. 24. When projecting an s-PCA reconstruction with only two modes onto the PCA modes, we capture nearly 20 PCA modes exactly. For the remaining PCA modes, we observe substantial amounts of energy using the s-PCA reconstruction. With two nlPCA modes we can capture the energy content of 13 linear modes almost exactly, with higher modes retaining large amounts of energy. Introducing symmetry reduction results in an increase in the number of captured linear modes to around 35. Additionally, the retained energy of higher modes is closer to the actual energy. This underlines the drastic increase in compression performance of s-nlPCA, compared to PCA.

5. Conclusion

In this work we present a novel, symmetry-aware machine learning approach for model reduction of physical problems. For both linear and nonlinear methods our symmetry-aware embedding was able to



Figure 24: Equivalent PCA energy of the Kolmogorov flow for the s-PCA reconstruction (left), the nlPCA reconstruction (middle) and the s-nlPCA reconstruction (right) using two modes each.

reduce the L_2 reconstruction error by large margins when compared to more conventional methods. Hence, a comparable accuracy of reconstructions could be achieved with far fewer modes, which makes each individual mode more relevant and worth investigating. Our method was tested on multiple setups, exhibiting different types of symmetry groups. The first test case, Burgers' equation, only supports the continuous group of translations, $\mathbf{T}(s) : u(x) \to u(x + s)$. Our second test case, the expansive flow in a step diffuser only sustains a pair of reflectional symmetries, $\mathbf{R} : (u, v)(x, y) \to (u, -v)(x, -y)$. The last case, Kolmogorov flow, is characterized by a more complex scenario: a combination of the translational invariance $\mathbf{T}(s)$ in x, along with two discrete symmetry groups, namely a rotation through π , $\mathbf{P} : \omega(x, y) \to \omega(-x, -y)$ and a shift-and-reflect symmetry, $\mathbf{S} : \omega(x, y) \to \omega(-x, y + \pi/4)$.

Our method is based on transforming our input dataset into an invariant subspace and performing model reduction in that space. Continuous symmetries are being removed using an STN-like method. Here, a subnetwork detects the continuous transformation parameter s, the translational shift, which is then applied through a prescribed transformation function. We have demonstrated that the STN does not necessarily need to be nonlinear in order to discover the proper shifts for both Burgers' equation and Kolmogorov flow. Consequently, the spatial shapes of our STN can be examined, since the network can be represented by a matrix $\mathbf{W}_{STN} \in \mathbb{R}^{2 \times N_{\xi}}$, with N_{ξ} as the number of grid points. These shapes resemble trigonometric functions (distinguished by a random phase shift) in the alignment direction. Furthermore, a modulation in the inhomogeneous direction, which does not change the alignment of the samples, can be used in cases such as turbulent channel flow, see e.g. Marensi *et al.* [16], to avoid infinite shifts due to a vanishing Fourier amplitude for the first complex coefficient [4]. In a conventional application of the method of slicing this modulation has to be user-supplied and carefully tailored, whereas in our network it is learned as part of the network training.

To remove the discrete invariances from the dataset a Siamese network is formed to autoencode all potential transformation functions simultaneously. After the autoencoding process, we prune our network to one single branch by selecting the branch with the lowest reconstruction error relative to its respective input. Thus, when a sample reappears in a different symmetry state, it is automatically reoriented, since the AE has been previously conditioned to compress a deviating orientation. In the simple case of the expansion in a step diffuser, the chosen branch correlates with the location of the pressure center in y. We have thus created a method that behaves like a conditional-filter method based on the pressure signal, though without the need for *a priori* data modifications or the calculation of global quantities.

Within our symmetry-aware framework, both linear and nonlinear methods can be embedded. Applying a linear AE, we arrive at a method termed symmetry-aware linear PCA (s-PCA). Here, we do not achieve compression rates typical of nonlinear methods, but we gain insight into the dynamics of the system, as the modal structures are both extractable and interpretable.

When removing the advection from the dataset generated by Burgers' equation with s-PCA, the linear modes only need to model the dissipation of the system. For the expansion in a step diffuser we removed the constraint that the modal shapes have to adhere either to symmetric or antisymmetric structures, generated by the global invariances. As a consequence, our modes represent the flow field in a more accurate asymmetric

manner in accordance with the asymmetric shapes of the underlying stable solutions. Removing all symmetry groups for Kolmogorov flow changes the modes from an approximate Fourier basis, representative of the boundary conditions and the forcing wavenumber, to more interpretable structures. As a result, the first mode closely approximates the equilibrium solution, while higher modes recover the dynamics of a vortex pair in the flow.

Using nonlinear methods leads to the method of symmetry-aware nonlinear PCA (s-nlPCA). The main gain here is a substantial increase in compression over conventional nonlinear PCA (nlPCA). For Burgers' equation, we could show that nearly all features are captured by one mode. Using four s-nlPCA modes for the dataset generated by the step diffuser and two modes for Kolmogorov flow yields reconstructions equivalent (in captured energy) to PCA reconstructions using 20 and 35 modes, respectively.

Whether the learned modulation of our *STN* weights in the non-aligned direction, can indeed circumvent the issue of infinite shifts for vanishing first Fourier coefficients remains to be investigated in a future effort. In addition, s-nlPCA, with its high compression rate, should provide sufficient compression in order to properly cluster data in the latent space and thus get a characterization of different states. Dynamic modelling of states using approximate Koopman operators, see Lusch *et al.* [15] or Otto & Rowley [22], also rely on AEs to encode measured states into a set of observables which can be linearly advanced in time. Symmetryaware structures should also prove useful in these applications, since the dynamics does not depend on the orientation of a sample in space.

References

- BALDI, P. & HORNIK, K. 1989 Neural networks and principal component analysis: Learning from examples without local minima. Neural Networks 2 (1), 53–58.
- [2] BOURGEOIS, J. A., NOACK, B. R. & MARTINUZZI, R. J. 2013 Generalized phase average with applications to sensor-based flow estimation of the wall-mounted square cylinder wake. *Journal of Fluid Mechanics* 736, 316–350.
- [3] BUDANUR, N. B. 2015 Exact coherent structures in spatiotemporal chaos: from qualitative description to quantitative predictions. PhD thesis, Georgia Institute of Technology.
- [4] BUDANUR, N. B., CVITANOVIĆ, P. DAVIDCHACK, R. L. & SIMINOS, E. 2015 Reduction of SO(2) symmetry for spatially extended dynamical systems. *Phys. Rev. Lett.* 114, 084102.
- [5] CHANDLER, G. J. & KERSWELL, R. R. 2013 Invariant recurrent solutions embedded in a turbulent two-dimensional Kolmogorov flow. *Journal of Fluid Mechanics* 722, 554–595.
- [6] CYBENKO, G. 1989 Approximation by superpositions of a sigmoidal function. Mathematics of control, signals and systems 2 (4), 303–314.
- [7] FREUND, J. B. & COLONIUS, T. 2009 Turbulence and sound-field POD analysis of a turbulent jet. International Journal of Aeroacoustics 8 (4), 337–354.
- [8] FUKAMI, K., NAKAMURA, T. & FUKAGATA, K. 2020 Convolutional neural network based hierarchical autoencoder for nonlinear mode decomposition of fluid field data. *Physics of Fluids* 32 (9), 095110.
- [9] HECHT, F. 2012 New development in FreeFem++. J. Numer. Math. 20 (3-4), 251–265.
- [10] HOLMES, P., LUMLEY, J. L., BERKOOZ, G. & ROWLEY, C. W. 2012 Turbulence, Coherent Structures, Dynamical Systems and Symmetry, 2nd edn. Cambridge Monographs on Mechanics. Cambridge University Press.
- [11] IOLLO, A. & LOMBARDI, D. 2014 Advection modes by optimal mass transfer. Phys. Rev. E 89, 022923.
- [12] JADERBERG, M., SIMONYAN, S., ZISSERMAN, A. & KAVUKCUOGLU, K. 2015 Spatial Transformer Networks. In NIPS.
- [13] LINOT, A. J. & GRAHAM, M. D. 2020 Deep learning to discover and predict dynamics on an inertial manifold. *Physical Review E* 101 (6).
- [14] LUMLEY, J. L. 1967 The structure of inhomogeneous turbulent flows. Atmospheric Turbulence and Radio Wave Propagation.
- [15] LUSCH, B., KUTZ, J. N. & BRUNTON, S. L. 2018 Deep learning for universal linear embeddings of nonlinear dynamics. *Nature Communications* 9 (1).
- [16] MARENSI, E., YALNIZ, G., HOF, B. & BURAK B., N. 2021 Symmetry-reduced Dynamic Mode Decomposition of Near-wall Turbulence. arXiv e-prints p. arXiv:2101.07516, arXiv: 2101.07516.
- [17] MEHR, É., LIEUTIER, A., BERMUDEZ, F. S., GUITTENY, V., THOME, N. & CORD, N. 2018 Manifold learning in quotient spaces. In CVPR, pp. 9165–9174.
- [18] MENDIBLE, A., BRUNTON, S. L., ARAVKIN, A. Y., LOWRIE, W. & KUTZ, J. N. 2020 Dimensionality reduction and reducedorder modeling for traveling wave physics. *Theoretical and Computational Fluid Dynamics* 34 (4), 385–400.
- [19] MESHALKIN, L. D. & SINAI, IA. G. 1961 Investigation of the stability of a stationary solution of a system of equations for the plane movement of an incompressible viscous liquid. *Journal of Applied Mathematics and Mechanics* 25 (6), 1700–1705.
- [20] MILANO, M. & KOUMOUTSAKOS, P. 2002 Neural network modeling for near wall turbulent flow. Journal of Computational Physics 182, 1–26.

- [21] MURATA, T., FUKAMI, K. & FUKAGATA, K. 2020 Nonlinear mode decomposition with convolutional neural networks for fluid dynamics. Journal of Fluid Mechanics 882, A13.
- [22] OTTO, S. & ROWLEY, C. 2019 Linearly recurrent autoencoder networks for learning dynamics. SIAM J. Appl. Dyn. Sys. 18, 558–593.
- [23] PAGE, J., BRENNER, M. P. & KERSWELL, R. R. 2021 Revealing the state space of turbulence using machine learning. *Physical Review Fluids* 6 (3).
- [24] PLATT, N., SIROVICH, L. & FITZMAURICE, N. 1991 An investigation of chaotic Kolmogorov flows. *Physics of Fluids A:* Fluid Dynamics **3**.
- [25] PLAUT, E. 2018 From Principal Subspaces to Principal Components with Linear Autoencoders. arXiv e-prints p. arXiv:1804.10253, arXiv: 1804.10253.
- [26] ROWLEY, C. W., KEVREKIDIS, I. G., MARSDEN, J. E. & LUST, L. 2003 Reduction and reconstruction for self-similar dynamical systems. *Nonlinearity* 16 (4), 1257–1275.
- [27] SCHMIDT, O. T. & SCHMID, P. J. 2019 A conditional space-time POD formalism for intermittent and rare events: example of acoustic bursts in turbulent jets. *Journal of Fluid Mechanics* 867, R2.
- [28] SIROVICH, L. 1987 Turbulence and the dynamics of coherent structures part ii: Symmetries and transformations. Quarterly of Applied Mathematics 45 (3), 573–582.
- [29] WAN, Z. Y., VLACHAS, P., KOUMOUTSAKOS, P. & SAPSIS, T. 2018 Data-assisted reduced-order modeling of extreme events in complex dynamical systems. PLOS ONE 13 (5), 1–22.

Appendix A. Network architectures and hyperparameters

Collected below are the architectures and hyperparameters used to generate the reconstructions in this paper for different testcases. The architectures are compiled in tables showing the encoder and decoder levels with descending layers being further removed from the input and latent space, respectively. The column 'datasize' describes the output size of each of the layers. The layers used are

- Reshape: Reshapes the layer inputs to the sizes prescribed in 'data size'.
- Fully connected: Fully connected layer, i.e. all input nodes are connected to an output node of the layer with a weight and bias. The used activation functions are noted in brackets. If no bias was used (for PCA and s-PCA), this is also noted in brackets.
- Convolution: Performs a convolution, with the parameters in brackets in the format (kernel size x, kernel size y, number of filters). Additionally, we specify the used activation in brackets. Note that due to periodic boundaries, convolutions were performed with circular padding. Using a 3×3 kernel this means that the first column is padded by the last column and similar for other boundary points.
- Average Pooling: Pool a number of nodes by the kernel size given in brackets by averaging all values within the kernel size.
- Upsampling: Upsample nodes with nearest neighbor upsampling with the kernel size given in brackets.

In what follows, the autoencoding architectures together with the hyperparameters used for training the networks for PCA, s-PCA, nlPCA and s-nlPCA are given for the datasets generated by Burgers' equation, flow through a step diffuser and Kolmogorov.

Architecture					
Encoder		Decoder			
Layer Data size		Layer Data			
Input (Batch, 64)		Latent space	(Batch, N)		
Fully connected (linear, no bias)	(Batch, N)	Fully connected (linear, no bias)	(Batch, 64)		
Latent space	(Batch, N)	Output	(Batch, 64)		

Hyperparameters				
Parameter	Value	Parameter	Value	
Maximum epochs	2000	Early stopping patience	500	
Learning rate	10^{-2}	Early stopping delta	10^{-9}	
Learning rate decay steps	10000	Batch size	500	
Learning rate decay rate	0.1	Optimizer	Adam (Keras default values)	

Table A.1: Autoencoding architecture and hyperparameters used to train the autoencoding architecture for PCA and s-PCA for the Burgers' equation dataset.

Architecture					
Encoder		Decoder			
Layer Data size		Layer	Data size		
Input	(Batch, 64)	Latent space	(Batch, N)		
Fully connected (swish)	(Batch, 512)	Fully connected (swish)	(Batch, 512)		
Fully connected (linear)	(Batch, N)	Fully connected (linear)	(Batch, 64)		
Latent space	(Batch, N)	Output	(Batch, 64)		

Hyperparameters				
Parameter	Value	Parameter	Value	
Maximum epochs	15000	Early stopping patience	5000	
Learning rate	10^{-3}	Early stopping delta	10^{-9}	
Learning rate decay steps	100000	Batch size	500	
Learning rate decay rate	0.1	Optimizer	Adam (Keras default values)	

Table A.2: Autoencoding architecture and hyperparameters used to train the autoencoding architecture for nlPCA and s-nlPCA for the Burgers' equation dataset.

Architecture						
Encoder		Decoder				
Layer Data size		Layer	Data size			
Input	(Batch, 19180, 2)	Latent space	(Batch, N)			
Reshape	(Batch, 38360)	Fully connected (linear, no bias)	(Batch, 38360)			
Fully connected (linear, no bias)	(Batch, N)	Reshape	(Batch, 19180, 2)			
Latent space	(Batch, N)	Output	(Batch, 19180, 2)			

Hyperparameters				
Parameter	Value	Parameter	Value	
Maximum epochs	1000	Early stopping patience	50	
Learning rate	10^{-3}	Early stopping delta	10^{-7}	
Learning rate decay steps	1000	Batch size	100	
Learning rate decay rate	0.2	Optimizer	Adam (Keras default values)	

Table A.3: Autoencoding architecture and hyperparameters used to train the autoencoding architecture for PCA and s-PCA for the step diffuser dataset.

Architecture					
Encode	er	Decoder			
Layer	Data size	Layer	Data size		
Input	(Batch, 19180, 2)	Latent space	(Batch, N)		
Reshape	(Batch, 38360)	Fully connected (swish)	(Batch, 512)		
Fully connected (swish)	(Batch, 512)	Fully connected (linear)	(Batch, 38360)		
Fully connected (linear)	(Batch, N)	Reshape	(Batch, 19180, 2)		
Latent space	(Batch, N)	Output	(Batch, 19180, 2)		

Hyperparameters				
Parameter	Value	Parameter	Value	
Maximum epochs	1000	Early stopping patience	50	
Learning rate	10^{-3}	Early stopping delta	10^{-7}	
Learning rate decay steps	1000	Batch size	100	
Learning rate decay rate	0.2	Optimizer	Adam (Keras default values)	

Table A.4: Autoencoding architecture and hyperparameters used to train the autoencoding architecture for nlPCA and s-nlPCA for the step diffuser dataset.

Architecture					
Encoder		Decoder			
Layer Data size		Layer	Data size		
Input	(Batch, 32, 32)	Latent space	(Batch, N)		
Reshape	(Batch, 1024)	Fully connected (linear, no bias)	(Batch, 1024)		
Fully connected (linear, no bias)	(Batch, N)	Reshape	(Batch, 32, 32)		
Latent space	(Batch, N)	Output	(Batch, 32, 32)		

Hyperparameters				
Parameter	Value	Parameter	Value	
Maximum epochs	10	Early stopping delta	10^{-7}	
Learning rate	10^{-3}	Batch size	500	
Early stopping patience	5	Optimizer	Adam (Keras default values)	

Table A.5: Autoencoding architecture and hyperparameters used to train the autoencoding architecture for PCA and s-PCA for the Kolmogorov flow dataset.

Architecture					
Encoder		Decoder			
Layer	Data size	Layer	Data size		
Input	(Batch, 32, 32)	Latent space	(Batch, N)		
Reshape	(Batch, 32, 32, 1)	Fully connected (swish)	(Batch, 2048)		
Convolution $((3, 3, 16), \text{swish})$	(Batch, 32, 32, 16)	Fully connected (linear)	(Batch, 4096)		
Average pooling $(2, 2)$	(Batch, 16, 16, 16)	Reshape	(Batch, 8, 8, 64)		
Convolution $((3, 3, 32), \text{swish})$	(Batch, 16, 16, 32)	Upsampling $(2, 2)$	(Batch, 16, 16, 64)		
Average pooling $(2, 2)$	(Batch, 8, 8, 32)	Convolution $((3, 3, 32), \text{swish})$	(Batch, 16, 16, 32)		
Convolution $((3, 3, 64), \text{swish})$	(Batch, 8, 8, 64)	Upsampling $(2, 2)$	(Batch, 32, 32, 32)		
Reshape	(Batch, 4096)	Convolution $((3, 3, 16), \text{swish})$	(Batch, 32, 32, 16)		
Fully connected (swish)	(Batch, 2048)	Convolution $((3, 3, 1), \text{linear})$	(Batch, 32, 32, 1)		
Fully connected (linear)	(Batch, N)	Reshape	(Batch, 32, 32)		
Latent space	(Batch, N)	Output	(Batch, 32, 32)		

Hyperparameters					
Parameter Value Parameter Value					
Maximum epochs	10	Early stopping delta	10^{-7}		
Learning rate	10^{-3}	Batch size	500		
Early stopping patience	5	Optimizer	Adam (Keras default values)		

 $\label{eq:alpha} Table A.6: Autoencoding architecture and hyperparameters used to train the autoencoding architecture for nlPCA and s-nlPCA for the Kolmogorov flow dataset.$