


Contents lists available at [ScienceDirect](https://www.sciencedirect.com)

## Comput. Methods Appl. Mech. Engrg.

journal homepage: [www.elsevier.com/locate/cma](http://www.elsevier.com/locate/cma)

# Thermodynamics-informed super-resolution of scarce temporal dynamics data

Carlos Bermejo-Barbanoj<sup>a</sup>, Beatriz Moya<sup>b</sup>, Alberto Badías<sup>c</sup>, Francisco Chinesta<sup>b,d</sup>, Elías Cueto<sup>a,\*</sup>

<sup>a</sup> *ESI Group-UZ Chair of the National Strategy on Artificial Intelligence, Aragon Institute of Engineering Research (I3A), Universidad de Zaragoza, Zaragoza, Spain*

<sup>b</sup> *CNRS@CREATEJ LTD., Singapore*

<sup>c</sup> *ETSIAE, Universidad Politécnica de Madrid, Madrid, Spain*

<sup>d</sup> *ESI Chair and PIMM Lab, ENSAM ParisTech, 155 Boulevard de l'Hôpital, 75013 Paris, France*

## ARTICLE INFO

### Keywords:

Superresolution  
Deep learning  
Adversarial autoencoders  
Thermodynamics  
Model reduction  
Structure preserving

## ABSTRACT

We present a method to increase the resolution of measurements of a physical system and subsequently predict its time evolution using thermodynamics-aware neural networks. Our method uses adversarial autoencoders, which reduce the dimensionality of the full order model to a set of latent variables that are enforced to match a prior, for example a normal distribution. Adversarial autoencoders are seen as generative models, and they can be trained to generate high-resolution samples from low-resolution inputs, meaning they can address the so-called super-resolution problem.

Then, a second neural network is trained to learn the physical structure of the latent variables and predict their temporal evolution. This neural network is known as a structure-preserving neural network. It learns the metriplectic-structure of the system and applies a physical bias to ensure that the first and second principles of thermodynamics are fulfilled.

The integrated trajectories are decoded to their original dimensionality, as well as to the higher dimensionality space produced by the adversarial autoencoder and they are compared to the ground truth solution. The method is tested with two examples of flow over a cylinder, where the fluid properties are varied between both examples.

## 1. Introduction

Resolution Augmentation techniques, frequently known as super-resolution, refer to a series of techniques that aim to enhance the level of detail of data, often an image, through computational techniques. Their main goal is to produce a high-resolution version of a low-resolution input, improving overall detail and, in some cases, revealing smaller details that may not appear in the original input. Although these techniques have been extensively studied for years by the computer vision community, the growing advances in machine learning have supposed an important boost to this field, as they allow to generate better quality high-resolution samples while improving efficiency. Deep learning based approaches, such as Convolutional Neural Networks (CNNs) [1] and Generative Adversarial Networks (GANs) [2–4] have proved as an efficient way to augment the resolution of an image, as they are able to learn a mapping from low-resolution to high-resolution images, surpassing more traditional techniques like interpolation-based methods.

\* Corresponding author.

*E-mail addresses:* [cbarbanoj@unizar.es](mailto:cbarbanoj@unizar.es) (C. Bermejo-Barbanoj), [beatriz.moya@cnsrcreate.sg](mailto:beatriz.moya@cnsrcreate.sg) (B. Moya), [abadias@unizar.es](mailto:abadias@unizar.es) (A. Badías), [francisco.chinesta@ec-nantes.fr](mailto:francisco.chinesta@ec-nantes.fr) (F. Chinesta), [ecueto@unizar.es](mailto:ecueto@unizar.es) (E. Cueto).

<https://doi.org/10.1016/j.cma.2024.117210>

Received 27 February 2024; Received in revised form 31 May 2024; Accepted 3 July 2024

Available online 19 July 2024

0045-7825/© 2024 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

One field that could benefit from these recent advances in super-resolution are predictive digital twins for physical systems. When developing a digital twin of a system, sensors are commonly employed to capture information of the fields of interest. Although sensors provide accurate measurements, usually their placement is limited to concrete areas due to physical limitations, resulting in sparse spatial measurements. Super-resolution techniques can be applied to this partial information to generate dense output fields [5–7].

While the Big Data paradigm has become famous in news media of all kinds, the reality is that big data is rarely available in engineering applications. Sensors are often expensive, and the storage, curation and subsequent handling of large amounts of data is no easy task. The result is that we are often faced with situations where less information is available than we would like or need. Given this situation, the application of super-resolution techniques to the world of time series forecasting becomes an urgent necessity.

The most widespread super-resolution techniques (mainly in the world of computational imaging) use black-box techniques to generate the missing information. Logically, this provides much better results than simple interpolation. However, such techniques show severe limitations, and recently new approaches to the problem have been tried, if we stick to the case of prediction of physical phenomena. Since we deal with physical phenomena, a logical and a priori very attractive option is based on taking advantage of the scientific knowledge developed over centuries of research to complement the missing information. Thus, for example, if we are faced with a fluid mechanics problem, the imposition of the Navier–Stokes equations provides valuable information to achieve a successful super-resolution [8].

Predictive digital twins constitute a natural field of application of these techniques [9,10]. They aim to predict the time evolution of the real physical system they represent. However, these systems commonly exhibit complex behaviours which makes their real-time prediction difficult. In order to obtain a complete analysis of the phenomena that describe the behaviour of those systems, physical simulations must be performed using computational tools such Finite Element Method (FEM) for solid mechanics and Computational Fluid Mechanics (CFD) for fluid mechanics. These tools require the discretization of the domain into fine meshes, in most cases with millions of degrees of freedom. As a result, generally simulations are very computationally expensive, often making it impossible to obtain almost real-time predictions of the time evolution of the system. One approach to overcome this problem is to use model order reduction (MOR) methods, as often the solution of the system is contained in a lower-dimensional space, as stated in the manifold hypothesis [11]. Basic approaches like Proper Orthogonal Decomposition (POD) [12] rely on linear transformations to project the information to a low dimensional space, but they usually fail to model complex nonlinear phenomena. The ROM community has developed some techniques to overcome this limitation and obtain nonlinear mappings, like Local Linear Embedding (LLE) [13] and kernel Principal Component Analysis (k-PCA) [14], but in recent years, deep learning based methods [15] have been gaining popularity, with Autoencoders [16] being the most common approach. Some works within the ROM community have addressed the multi-fidelity problem, for instance the Non-Intrusive Reduced Basis (NIRB) [17,18]. Those methods could benefit from the advantages of super-resolution techniques to obtain high-fidelity data or to enhance the outputs. Multi-scale problems could also benefit from the interaction of both methods, as the NIRB could handle the large scale information, while the super-resolution could refine the fine details, leading to high resolution results. In the present work we focus on autoencoders, as they have proven their capabilities to produce highly nonlinear manifolds for a wide range of applications that include physical simulations [19,20]. Moreover, some autoencoder architectures exhibit generative capabilities, which makes them a feasible option to generate high-resolution outputs from low-resolution data.

In order to predict the time evolution of the analysed systems, deep learning approaches can be also applied. While classically these approaches have been seen as black boxes, as they require large amounts of data and fail to generalize, leading to unreliable predictions, in the recent years there has been a growing interest in physics-consistent deep learning. These techniques consist in adding some physical knowledge of the system to neural network to guarantee the physical consistency of the solution, minimizing the amount of data needed and improving generalization capabilities. Some works in this field are based on solving the PDEs that govern the problem, which leads to very accurate results [21,22]. The main drawback of these methods is that they require some knowledge of the governing equations of the phenomena, and in practical applications they are often not fully known. An alternative approach is thus to enforce more general physics, or physics of a higher epistemic level. In this last case, thermodynamics comes into play as a natural choice when more detailed information is missing. Some approaches have been done by imposing the so-called GENERIC (General Equation for Non-Equilibrium Reversible-Irreversible Coupling) metriplectic structure [23,24] of the problem, by means of the so-called Structure Preserving Neural Networks [25] and Thermodynamics-Informed Graph Neural Networks [26]. These neural networks lead to a thermodynamically-consistent prediction that can be applied to both conservative and dissipative systems. Recently, new insights in the way we can impose the fulfilment of the first and second laws of thermodynamics to the learning process have been included in [27,28]. Previous works [29] have proved the efficiency of combining model order reduction by autoencoders and time evolution prediction by structure-preserving neural networks, leading to fast and accurate predictions.

The aim of this work is to develop a method to augment the resolution of the low-resolution fields of the state variables of a system and consequently to predict a physically-consistent evolution of this system at the high-resolution regime. The proposed methodology is very general as the used formulation to predict the time evolution of the system is valid for a wide variety of dynamical systems, although we focus on fluid mechanics. The resulting high-resolution reconstruction of the system dynamics is guaranteed to fulfil the first and second principles of thermodynamics (energy conservation and non-negative entropy production).

In this way, both the super-resolution of the state variables of our system and the prediction of the time evolution of their dynamics will be carried out under the perspective of the same formalism, the so-called GENERIC equation, whose usefulness and physical correctness for a multitude of phenomena has already been demonstrated in previous works [30].

The structure of the paper is as follows. A description of the problem setup is presented in Section 2. The methodology is presented in Section 3, where both the model order reduction autoencoder and the GENERIC formalism to predict the evolution are described. In Section 4 two examples are analysed: the flow past a cylinder in a Newtonian and non-Newtonian setting. Finally, the conclusions of the paper are discussed at Section 5.

## 2. Problem statement

In this work we propose a framework to estimate the temporal evolution of a physical system from data, and to augment its spatial resolution, given the assumption of scarce data. We apply superresolution techniques based in the employ of deep learning and making use of the so-called dynamical system equivalence of scientific machine learning [31]. We assume a dynamical system governed by a set of state variables  $\mathbf{x} \in \mathcal{M} \subseteq \mathbb{R}^D$ , with  $\mathcal{M}$  the state space of these variables, assumed to evolve on a differentiable manifold in  $\mathbb{R}^D$ , thanks to the widespread manifold hypothesis [11]. The full-order model of a physical phenomenon can be expressed as a system of differential equations that give the temporal evolution of a set of state variables  $\mathbf{x}$ ,

$$\dot{\mathbf{x}} = \frac{d\mathbf{x}}{dt} = \mathbf{F}(\mathbf{x}, t), \quad t \text{ in } I = (0, T], \quad \mathbf{x}(0) = \mathbf{x}_0, \quad (1)$$

where  $t$  refers to the time coordinate in the time interval  $I$  and  $\mathbf{F}(\mathbf{x}, t)$  is an a priori unknown nonlinear function that represents the flow map of the governing variables. The identification of this function  $\mathbf{F}$  from data is precisely the objective of this work, where we assume that we work in a scarce data scenario.

Since in the most common applications of such techniques (such as the aforementioned digital twins) there is the additional circumstance of strong real-time constraints, it is also assumed that there is a need to work on reduced models of the physics under study. The dimensionality reduction procedure looks for a simpler representation of the full-order state vector represented by  $\mathbf{x}$ , through a set of reduced (also denoted as latent in the literature of machine learning) variables  $\mathbf{z} \in \mathcal{N} \subseteq \mathbb{R}^d$  contained in a manifold with a dimensionality lower than the original space  $\mathcal{M}$ . The mapping between both spaces is denoted by  $\phi : \mathcal{M} \subseteq \mathbb{R}^D \rightarrow \mathbb{R}^d$ , where  $d \ll D$ . An inverse mapping  $\phi^{-1}$  allows to undo the transformation, recovering the information in the full-order space

The goal of this paper is two-fold. First, to find a mapping  $\phi$  for a dynamical system governed by Eq. (1) that allows us to predict its temporal evolution under stringent real-time constraints on a reduced-order manifold, and then to augment the spatial dimensionality of the data, back to the original, full-order state space manifold. The mapping  $\phi$  allows to learn the underlying physics of the system in the reduced space  $\mathcal{N}$  and then predict its temporal evolution. In order to obtain a physically consistent prediction of the system, the solution must fulfil the laws of thermodynamics, which are enforced by assuming that their evolution occurs under the GENERIC framework. The second objective is to achieve this while simultaneously augmenting the spatial resolution of the data to generate a dense solution field, thus obtaining a solution space with a higher dimensionality than the original space  $\mathcal{M}$ .

## 3. Methodology

The proposed framework splits the problem in two main steps. First, the low-resolution, full order model is encoded (projected) onto a reduced-order manifold (or latent space) with an autoencoder, thus achieving a nonlinear mapping  $\phi$ . The autoencoder learns a coded representation of the physical system, which allows to work with the data in a compact form. Moreover, the autoencoder is trained to generate high resolution fields of the state variables from the low resolution input data.

Then, a structure-preserving neural network is trained with (low resolution, but full order) simulation data so as to obtain a temporal prediction of the evolution of the dynamics of the system. This network predicts the time evolution of the system by using the GENERIC formalism. Finally, these latent variables are projected back by the decoder to both the original manifold of the low-resolution full order model and a higher resolution manifold. A general scheme of this procedure is shown in Fig. 1. In this work, the full order model data has been generated in silico, although this procedure could be applied to measurements coming from a real physical system.

### 3.1. Model reduction with adversarial autoencoders

An autoencoder is a type of neural network that learns an efficient codification or embedding of data. This results in a dimensionality reduction of the input information into a set of latent variables, which ideally contain the same information as the original data. The classical autoencoder architecture is composed by two basic elements: an encoder,  $\varepsilon_\phi$ , that maps the high-dimensional information into a low-dimensional code, and a decoder,  $\delta_\phi$ , that applies the inverse operation, recovering the information into the original full-order manifold.

$$\varepsilon_\phi : \mathbb{R}^D \rightarrow \mathbb{R}^d, \quad \mathbf{z} = \varepsilon_\phi(\mathbf{x}), \quad (2)$$

$$\delta_\phi : \mathbb{R}^d \rightarrow \mathbb{R}^D, \quad \hat{\mathbf{x}} = \delta_\phi(\mathbf{z}). \quad (3)$$

In this work, we use adversarial autoencoders (AAEs) [32]. This kind of autoencoder enforces the latent vector to follow a desired distribution or prior, similarly to variational autoencoders (VAEs) [33], but instead of predicting the mean and standard deviation to enforce that the latent code follows the prior, it is enforced by using an additional network called discriminator. This allows the latent variables to follow not only normal distributions, like VAEs, but also more complex ones.

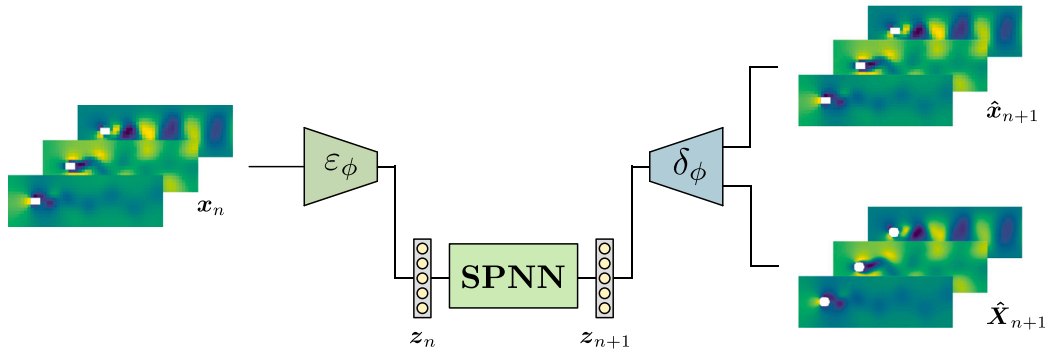


Fig. 1. Scheme of the proposed framework. First, an encoder is used to reduce the dimensionality of the problem, obtaining a set of reduced variables or latent code. Then, a structure-preserving neural network (SPNN) is trained to integrate the time evolution of the reduced variables of the system. Thus, given the state of the system at time instant  $n$ , the net obtains the state at time instant  $n + \Delta n$ . Finally the decoder is used to recover the data to its original dimensionality and to generate the output in a resolution that is higher than the input one.

The discriminator is a simple neural network, usually a Multilayer Perceptron (MLP), that takes as input the latent code generated by the autoencoder and a random sample that follows the prior. It compares both to determine how close the latent code is to the prior. As the training process advances, the latent code produced by the AAE is closer to the prior, which means that the discriminator finds harder to discern if the sample comes from the prior or from the autoencoder.

AAEs are seen as a mix between VAEs and Generative Adversarial Networks (GANs) [2], as they enforce the latent code to follow a prior but make use of a discriminator to ensure that this prior is matched. Like VAEs and GANs, AAEs are considered as generative models. This results in a very useful feature for the proposed task, as they can be used to generate a high-resolution output from a low-resolution input. The resolution augmentation has been achieved by training the decoder of the AAE to output the low-dimensional data (same as input, as in classical autoencoders) and also the high-dimensional data, supervising the training with the ground truth information. The AAE scheme can be seen in Fig. 2.

The loss function of the Autoencoder is composed therefore by three terms:

- **Low-Resolution data loss:** The output of the autoencoder,  $\hat{x}$ , must match the ground truth, in this case the input of the network, the low resolution pressure and velocity fields,  $x$ . The accuracy of the network is evaluated using the mean squared error:

$$\mathcal{L}_{\text{mse, LR}}^{\text{AAE}} = \frac{1}{n_{\text{snap}}} \sum_{i=0}^{n_{\text{snap}}} (x_i - \hat{x}_i)^2, \tag{4}$$

where the subscript  $i$  refers to the snapshot number,  $i = 1, \dots, n_{\text{snap}}$ .

- **High-Resolution data loss:** The output of the autoencoder,  $\hat{X}$ , must match the ground truth, the high-resolution pressure and velocity fields obtained from the in-silico simulations,  $X$ . As with the low resolution data, the accuracy of the autoencoder is evaluated using the mean squared error:

$$\mathcal{L}_{\text{mse, HR}}^{\text{AAE}} = \frac{1}{n_{\text{snap}}} \sum_{i=0}^{n_{\text{snap}}} (X_i - \hat{X}_i)^2. \tag{5}$$

- **Adversarial loss:** The third term of the loss function is the contribution of the discriminator,  $\mathcal{L}_{\text{adv}}^{\text{AAE}}$ . This term measures the likelihood between the proposed distribution and the distribution obtained by the encoder.

The final loss function of the autoencoder is composed of a weighted sum of all the terms. A hyperparameter,  $\lambda_{\text{adv}}^{\text{AAE}}$  is added to control its influence to the total loss function,

$$\mathcal{L}^{\text{AAE}} = \mathcal{L}_{\text{rec}}^{\text{AAE}} + \lambda_{\text{adv}}^{\text{AAE}} \cdot \mathcal{L}_{\text{adv}}^{\text{AAE}}, \tag{6}$$

where the  $\mathcal{L}_{\text{rec}}^{\text{AAE}}$  term represents the reconstruction capabilities of the autoencoder and is composed by the terms associated to the low-resolution and high-resolution fields,

$$\mathcal{L}_{\text{rec}}^{\text{AAE}} = \mathcal{L}_{\text{mse, LR}}^{\text{AAE}} + \mathcal{L}_{\text{mse, HR}}^{\text{AAE}}. \tag{7}$$

### 3.2. Learning the dynamical evolution of the system by structure-preserving neural network

One of our main interests is to develop a framework that satisfies a priori, by construction, known principles of physics about the phenomenon at hand. This is crucial, as we want our framework to provide credible, robust and accurate predictions to help in

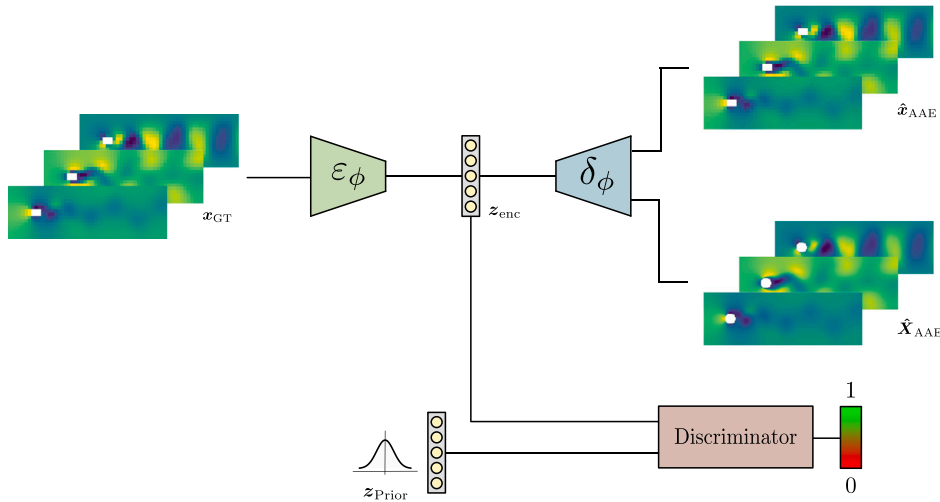


Fig. 2. Scheme of the adversarial autoencoder (AAE). The autoencoder takes a snapshot of the simulation as input and learns an encoded representation of the data. The decoder recovers the data to its original dimensionality,  $\hat{x}_{AAE}$ , and it is trained to generate a higher-resolution output,  $\hat{x}_{AAE}$ , achieving “superresolution” of the analysed simulation. The discriminator takes as input a sample that follows the proposed distribution or prior (in this work a normal distribution) and the latent code generated by the autoencoder. The discriminator output is in the range between 0 and 1. An output close to 0 means that the latent code does not match the prior distribution, while if the output is close to 1 the autoencoder is able to produce a latent code that fits into the prior distribution.

fields like decision-making, and this can only be achieved with predictions that fulfil the basic principles of physics. In our approach, this is achieved by using physics principles as inductive bias. An inductive bias is a set of assumptions about the data that prioritize one solution over the rest — precisely, the one fulfilling known physical principles —, preventing the learning process from finding a local minimum of the loss function.

Maybe the most popular method in our community at this moment is the so-called Physics-Informed Neural Networks (PINN) [21], in which we enforce the fulfilment of a particular partial differential equation that governs our system. However, there are some situations where the governing equations are not well known or they cannot be applied easily. In other situations, there are models that are well known but nevertheless provide unconvincing results in predicting the evolution of the system. In this case, a very attractive option is to learn only the “ignorance” about the physical behaviour, so that the prediction is the sum of the evolution predicted by the model and the prediction of the learnt ignorance model about the system. This is the approach that has been followed, for example, in [9,34].

For that reason, we want to guarantee the physical meaning of the solution, but without enforcing any particular physical equation.

For this purpose, we use a structure-preserving neural network (SPNN) [25]. Structure-preserving neural networks refer to a class of methods that are constructed to satisfy some high-level epistemic properties of the problem, for example, the principles of thermodynamics. SPNN can be applied to conservative and dissipative problems, ensuring that the principles of thermodynamics are satisfied by construction. This property allows us to use the thermodynamics laws as an inductive bias [35], ensuring the physical consistency of the results.

### 3.2.1. GENERIC formalism

To guarantee the physical meaning of the solution, we enforce the “General Equation for Non-Equilibrium Reversible-Irreversible Coupling”, usually referred as GENERIC formalism [23,24]. This formalism is a generalization of the classic Hamiltonian formulation to dissipative systems. This approach assumes the reversible or conservative contribution to be of Hamiltonian form, thus requiring an energy function and a Poisson bracket. The irreversible contribution to the energetic balance is generated by the non-equilibrium entropy and an irreversible or friction bracket [36].

The GENERIC formulation of time evolution for non-equilibrium systems, parameterized by a set of state variables able to describe the evolution of the energy of the system,  $z$ —the choice is thus not unique—, is given by:

$$\frac{dz}{dt} = \{z, E\} + [z, S], \tag{8}$$

where the so-called Poisson bracket  $\{\cdot, \cdot\}$  and dissipative bracket  $[\cdot, \cdot]$  have been used. For practical use, the bracket notation is often reformulated using two linear operators:

$$L : T^*\mathcal{M} \rightarrow T\mathcal{M}, \quad M : T^*\mathcal{M} \rightarrow T\mathcal{M}, \tag{9}$$

where  $T^*\mathcal{M}$  and  $T\mathcal{M}$  represent, respectively, the cotangent and tangent bundles of the state space  $\mathcal{M}$ . The operator  $L$  represents the Poisson bracket and must be skew-symmetric, while the operator  $M$ , the friction matrix, describes the irreversible part of the

system and must be positive semidefinite to make sure that the dissipation rate is positive. For phenomena involving plasticity, for instance, this approach may not be valid, and a more general form of the dissipative term should be considered. A more general one is developed in [37,38], among other references.

Replacing the original bracket formulation in Eq. (8) with their respective operators, the time evolution equation for the state variables  $\mathbf{z}$  is derived,

$$\frac{d\mathbf{z}}{dt} = \mathbf{L} \frac{\partial E}{\partial \mathbf{z}} + \mathbf{M} \frac{\partial S}{\partial \mathbf{z}}. \quad (10)$$

The equation is completed by adding the so-called degeneracy conditions:

$$\{S, \mathbf{z}\} = [E, \mathbf{z}] = \mathbf{0}. \quad (11)$$

The first expression states that the entropy is a degenerate functional of the Poisson bracket, and shows the reversible nature of the Hamiltonian contribution to the dynamics. The second expression states that the energy is a degenerate functional of the friction matrix, so the total energy of the system is conserved. These conditions can be reformulated into a matrix form in terms of the previously defined  $\mathbf{L}$  and  $\mathbf{S}$  operators, which results in the following degeneracy conditions:

$$\mathbf{L} \frac{\partial S}{\partial \mathbf{z}} = \mathbf{M} \frac{\partial E}{\partial \mathbf{z}} = \mathbf{0}. \quad (12)$$

The degeneracy conditions, in addition to the non-negativeness of the irreversible bracket, guarantees that the first (energy conservation) and the second (entropy inequality) laws of thermodynamics are fulfilled.

$$\frac{dE}{dt} = 0, \quad \frac{dS}{dt} \geq 0. \quad (13)$$

### 3.2.2. Structure-preserving neural networks

The structure-preserving neural networks impose the GENERIC formalism to guarantee the thermodynamical consistency of the solution. In order to work with the data coming from the simulation, the GENERIC formalism is discretized along time intervals  $\Delta t$ ,

$$\frac{\mathbf{z}_{n+1} - \mathbf{z}_n}{\Delta t} = \mathbf{L}_n \left( \frac{\partial E}{\partial \mathbf{z}} \right)_n + \mathbf{M}_n \left( \frac{\partial S}{\partial \mathbf{z}} \right)_n, \quad (14)$$

where we employ the subscript  $n$  to refer to time instant  $t = n\Delta t$  and, therefore,  $n + 1$  to refer to  $t + \Delta t = (n + 1)\Delta t$ .

In this scheme the time derivative is substituted by a forward-Euler scheme with time increments  $\Delta t$ . The accuracy and stability of different time discretizations of the GENERIC equation have been deeply analysed in [39,40]. The Poisson and friction operators are discretized as  $\mathbf{L}_n$  and  $\mathbf{M}_n$ . Similarly, energy and entropy gradients are discretized as  $\left( \frac{\partial E}{\partial \mathbf{z}} \right)_n$  and  $\left( \frac{\partial S}{\partial \mathbf{z}} \right)_n$ . Eq. (14) can be rewritten to the proposed integration scheme to predict the temporal evolution of the system:

$$\mathbf{z}_{n+1} = \mathbf{z}_n + \Delta t \cdot \left( \mathbf{L}_n \left( \frac{\partial E}{\partial \mathbf{z}} \right)_n + \mathbf{M}_n \left( \frac{\partial S}{\partial \mathbf{z}} \right)_n \right). \quad (15)$$

Additionally, discretized degeneracy conditions are added to ensure the thermodynamical consistency of the prediction:

$$\mathbf{L}_n \left( \frac{\partial S}{\partial \mathbf{z}} \right)_n = \mathbf{0}, \quad \mathbf{M}_n \left( \frac{\partial E}{\partial \mathbf{z}} \right)_n = \mathbf{0}. \quad (16)$$

The GENERIC structure is imposed to the encoded space learnt by the adversarial autoencoder, similarly to [29,41]. The SPNN is a feed-forward neural network composed by a set of fully connected layers. The input of the net is the encoded state vector at a given timestep  $\mathbf{z}_n^{\text{AAE}}$ . The output from the net is a vector containing the predicted  $\mathbf{L}_n$  and  $\mathbf{M}_n$  matrices, as well as the predicted energy and entropy gradients,  $\left( \frac{\partial E}{\partial \mathbf{z}} \right)_n$  and  $\left( \frac{\partial S}{\partial \mathbf{z}} \right)_n$ .

Actually, to enforce the skew-symmetry and positive semi-definiteness of matrices  $\mathbf{L}$  and  $\mathbf{M}$ , the output of the network is a pair of matrices  $\mathbf{l}$  and  $\mathbf{m}$ , reshaped in lower-triangular matrices,

$$\mathbf{L} = \mathbf{l} - \mathbf{l}^\top, \quad \mathbf{M} = \mathbf{m} \mathbf{m}^\top. \quad (17)$$

Then, using the integration scheme showed in Eq. (15), the reduced space state vector at the next time step is obtained  $\mathbf{z}_{n+1}^{\text{SPNN}}$  (see Fig. 3).

The loss function used to train the SPNN is composed by two different terms:

- **Data loss:** The output of the integration scheme,  $\mathbf{z}_{n+1}^{\text{SPNN}}$ , must match the ground truth, in this case the encoded state vector,  $\mathbf{z}_{n+1}^{\text{AAE}}$ , predicted by the autoencoder. The accuracy of the network is evaluated using the mean squared error:

$$\mathcal{L}_{\text{mse}}^{\text{SPNN}} = \frac{1}{n_{\text{snap}}} \sum_{n=0}^{n_{\text{snap}}} (\mathbf{z}_{n+1}^{\text{AAE}} - \mathbf{z}_{n+1}^{\text{SPNN}})^2. \quad (18)$$

- **Degeneracy conditions loss:** The loss function includes the fulfilment of the degeneracy conditions, ensuring the thermodynamical consistency of the solution. They are measured as the sum of the squared values of both conditions:

$$\mathcal{L}_{\text{degen}}^{\text{SPNN}} = \frac{1}{n_{\text{snap}}} \sum_{n=0}^{n_{\text{snap}}} \left( \left( \mathbf{L}_n \left( \frac{\partial S}{\partial \mathbf{z}} \right)_n \right)^2 + \left( \mathbf{M}_n \left( \frac{\partial E}{\partial \mathbf{z}} \right)_n \right)^2 \right). \quad (19)$$






Fig. 3. Scheme of the structure-preserving neural network (SPNN). The SPNN is trained to predict the full time evolution of the latent variables generated by the AAE by applying the GENERIC structure of the underlying physics of the problem. The network takes the current snapshot as input and outputs the L and M matrices, as well as the energy and entropy gradients. Then, they are integrated following the GENERIC formalism, as shown in Eq. (15), and the latent variables of the next snapshot are obtained. This process can be done iteratively, obtaining the rollout prediction of the full simulation.

The final loss function is composed of a weighted sum of both terms. A hyperparameter  $\lambda_{mse}^{SPNN}$  is added to control its influence and balance both of them,

$$\mathcal{L}^{SPNN} = \lambda_{mse}^{SPNN} \cdot \mathcal{L}_{mse}^{SPNN} + \mathcal{L}_{degen}^{SPNN}. \tag{20}$$

### 4. Results

#### 4.1. Example 1: Flow past a cylinder of a Newtonian fluid

##### 4.1.1. Database generation

The first example consists in an unsteady flow past a cylindrical obstacle. The geometry of the obstacle is fixed for all examples and the flow conditions are varied by modifying the freestream velocity, which results in a variable Reynolds regime and generates a Kármán vortex street that exhibits a periodic behaviour during the steady state. The state variables for the flow past a cylinder are the velocity and pressure fields,

$$S = \{ \mathbf{x} = (\mathbf{u}, P) \in \mathbb{R}^2 \times \mathbb{R} \}. \tag{21}$$

The ground truth simulations are computed solving the 2D Navier–Stokes equations using OpenFOAM software [42]. No-slip condition is applied in the cylinder obstacle. The fluid is assumed to have a Newtonian behaviour with density of  $\rho = 1$  and dynamic viscosity  $\mu = 10^{-3}$ . The freestream velocity is contained within the interval  $\mathbf{u} \in [0.8, 3.4]$ , resulting in a total of  $N_{sim} = 27$  cases. Each case is discretized in  $n_{snap} = 800$  time increments of  $\Delta t = 0.005$ .

The input of the autoencoder are the low resolution velocity and pressure fields, with size  $3 \times 16 \times 48$ , while the output are the velocity (two components) and pressure (a scalar) fields at the original resolution and a higher one, with sizes  $3 \times 16 \times 48$  and  $3 \times 64 \times 192$ , respectively. Both the encoder and decoder use convolutional layers with  $N_{ch} = 64$  channels and a kernel size of  $k = 3$ , following a ResNet-like structure [43]. The number of latent variables at the bottleneck is set to  $d = 5$ . The activation function used is the Leaky-ReLU with a negative slope of 0.1, except for the last layer of both the encoder and decoder, where linear activations are used. The adversarial hyperparameter weight is set to  $\lambda_{adv}^{AAE} = 10^{-3}$ . The optimizer used is Adam [44] with a learning rate set to  $l_r^{AAE} = 10^{-4}$  with decreasing order of magnitude on epochs 600 and 1200, a weight decay set to  $w_d^{AAE} = 10^{-4}$ , and a total number of  $N_{epochs} = 1800$  epochs. Latent variables obtained at the bottleneck are then used as input variables for the structure preserving neural network that, as explained before, operated in the latent manifold. The training and validation loss curves for the autoencoder are shown in Fig. 4(a)

The SPNN input size coincides with the AAE latent dimension,  $N_{in}^{SPNN} = d = 5$ , while the output size is  $N_{out}^{SPNN} = d \cdot (d + 2) = 35$ , see Eq. (17). The number of hidden layers of the SPNN is  $N_{hl}^{SPNN} = 5$  with 100 neurons each one, Leaky-ReLU activations and linear for the last layer. The data weight hyperparameter is set to  $\lambda_{data}^{SPNN} = 10^2$ . The SPNN is trained for  $N_{epochs} = 4500$  epochs using the Adam optimizer. The learning rate is set to  $l_r^{SPNN} = 10^{-3}$ , decreasing one order of magnitude on epoch 1500 and 3000. The weight decay is set to  $w_d^{SPNN} = 10^{-4}$  and noise variance is added to the train set,  $\sigma_{noise}^2 = 10^{-6}$ . The Fig. 4(b) shows the training and validation curves for the SPNN.

##### 4.1.2. Results

Fig. 5 shows the prediction achieved for the pressure and velocity fields predicted by the autoencoder in low and high resolution, as well as the absolute error for each field. The AAE prediction shows good agreement with the reconstructed low resolution fields and those generated in high resolution. Fig. 6 shows a box plot of the data error for the train and test sets, obtaining a mean error lower than 3% for the pressure and velocity fields in both low and high resolution.

In order to prove the convenience of our proposed method we compare it to classical resolution augmentation techniques. The AAE is compared with a common technique to augment resolution in the computer vision field: the bicubic interpolation. Comparison between both is shown in Table 1. The results of the AAE outperform the bicubic interpolation, while being considerably faster, leading to a 37× speed increment.




Fig. 4. Training and validation loss curves for the Adversarial Autoencoder (Left) and the SPNN (right).

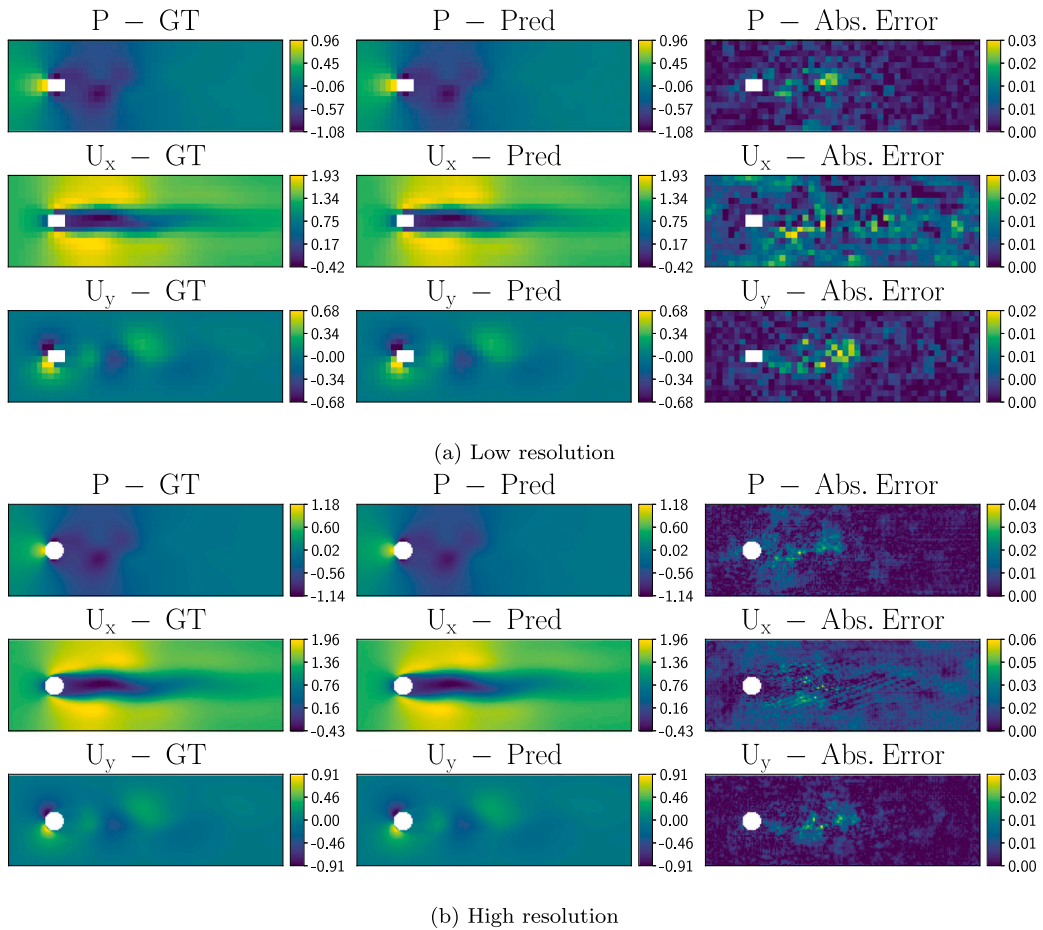
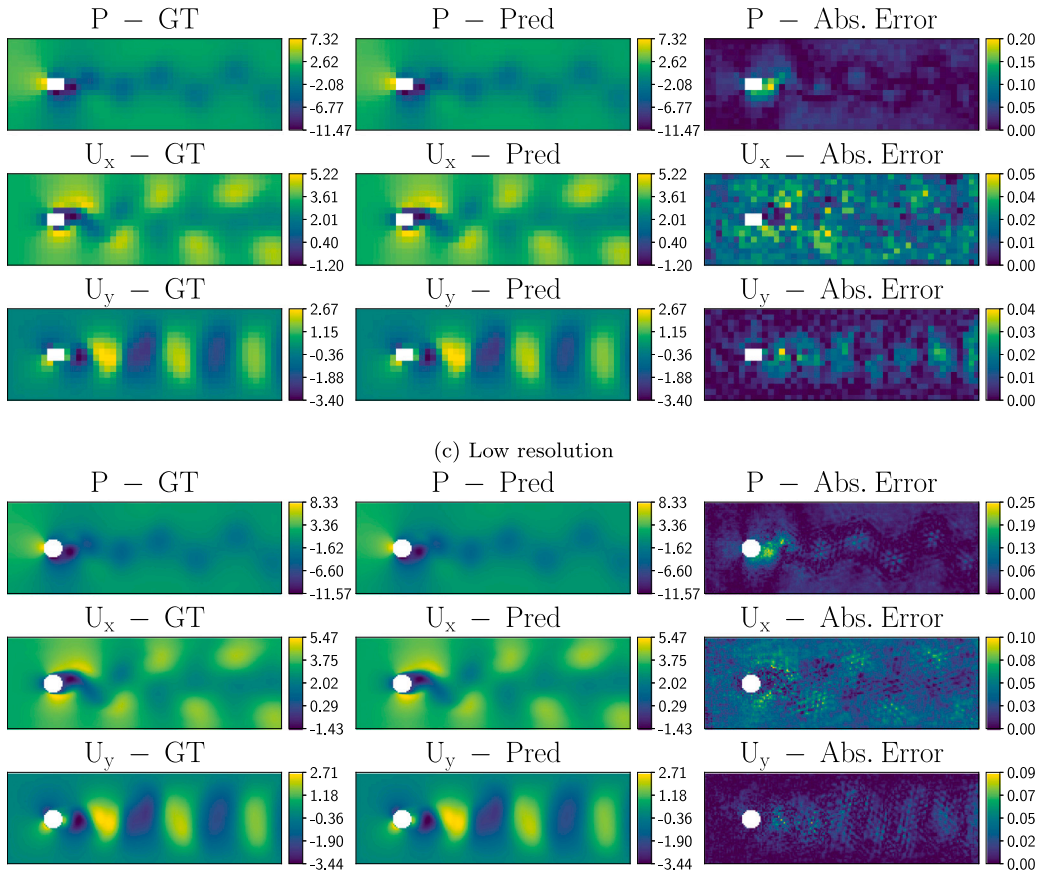


Fig. 5. Results of the prediction made by the Adversarial Autoencoder (AAE). 5a: Low resolution Ground Truth (GT), AAE prediction and absolute error for  $P$ ,  $U_x$  and  $U_y$ . The ground truth is the input of the AAE and it predicts the low resolution fields and the high resolution fields shown in Fig. 5b. 5b: High resolution GT, AAE prediction and absolute error fields for the same snapshot shown in Fig. 5a. Fig. 5c: Low resolution fields for a different snapshots. Fig. 5d: High resolution GT, AAE pred. and abs. error fields for the snapshot shown in Fig. 5c.





(d) High resolution

Fig. 5. (continued).





Fig. 6. Box plots for the relative L2 error of the autoencoder for all the snapshots of the Newtonian fluid for both train and test cases, in low (left) and high (right) resolution. The state variables represented are pressure (P) and velocity ( $U_x$  and  $U_y$ )

Fig. 7 shows the comparison between the AAE latent variables and the rollout prediction made by the SPNN, for the case with input velocity  $u = 3.4$ —the worst case scenario among all considered—. The SPNN is able to integrate the latent variables in the reduced space in good agreement with the original AAE encoding, considered as the ground truth for the SPNN.


The SPNN is also compared with a black box (BB) neural network. The black box neural network predicts the increment of the latent variables  $z$  and uses a forward Euler integration scheme to obtain the next snapshot of the simulation. The black box is trained using the same hyperparameters as the SPNN, except for the output size, which is the same dimension as the input, for this example

**Table 1**  
Comparison between the proposed method and bicubic interpolation. For every variable, mean relative error is shown, and the reconstruction time for the whole dataset is computed.

	AAE	Bicubic interpolation
$P$ (-)	0.0247	0.0621
$U_x$ (-)	0.0105	0.0285
$U_y$ (-)	0.0196	0.0524
Time (s)	711.05	26303.31



**Fig. 7.** Results of the SPNN integration with respect to the ground truth (GT) latent variables obtained by the AAE for two different simulations of the Newtonian fluid. The dashed line corresponds to the SPNN prediction while the continuous line is the ground truth. To facilitate the identification of ground truth and prediction values, each latent variable is represented by a distinct colour. This simplifies the comparison between the ground truth and the SPNN prediction. Left: simulation at  $U_{in}=1.4$ . Right: simulation at  $U_{in}=3.4$ .



**Fig. 8.** Evolution of the relative error during the whole integration of the velocity for all the simulations. Even if the error of the SPNN increases as a consequence of the explicit Euler integration scheme employed, the black box approach fails during the rollout and diverges from the solution. The SPNN error accumulates during the rollout prediction but its able to converge to a meaningful solution.

$N_{in}^{BB} = N_{out}^{BB} = 5$ . Results are compared in Fig. 8, which shows the velocity accumulated error mean and standard deviation with a confidence interval of 95% for every simulation at each snapshot. The SPNN (Fig. 8(a)) error raises as the prediction advances, which was expected, as the forward Euler is a first order integration scheme, but the network is able to converge to the solution. Meanwhile, the black box neural network (Fig. 8(b)) is not able to integrate the predicted trajectory and diverges from the ground truth, proving that the thermodynamic bias guides the network to converge to a meaningful solution.

The ground truth simulations were performed on a MacBook Pro M1 Pro. Each simulation took around 20 min to complete. The AAE and the SPNN were trained using the Pytorch framework. The computer used to train both networks was a Linux-based machine equipped with a Intel i9-13900K CPU and a NVIDIA RTX 4090 GPU. The AAE training time was approximately 4 h, while the SPNN took around 20 min to train. While working on inference, the prediction for the latent variables can be obtained in 1–2 s in a MacBook Pro M1 Pro, while rendering the video for the complete simulation takes around 15 s, achieving a considerable speedup when compared with the computational cost of running the high-fidelity simulation.




Fig. 9. Adversarial Autoencoder (left) and the SPNN (right) training and validation loss curves for the non-Newtonian example.

4.2. Example 2: Flow past a cylinder of a non-Newtonian fluid

4.2.1. Database generation

The second example is generated by using the same geometry than in the previous example, but the fluid is replaced by a non-Newtonian fluid. As in the previous case, the flow conditions are obtained by varying the initial velocity of the flow, which results in different Reynolds numbers. The state variables for the non-Newtonian flow past a cylinder are the velocity, shear rate and pressure fields, although good prediction results can be achieved by using only the pressure and velocity fields obtained from the solver:

$$S = \{x = (u, \dot{\gamma}, P) \in \mathbb{R}^2 \times \mathbb{R} \times \mathbb{R}\}. \tag{22}$$

Ground truth simulations are obtained by solving the 2D Navier–Stokes equations using OpenFOAM [42]. A no-slip condition is applied at the wall of the cylinder. In this example, a non-Newtonian fluid behaviour is applied using the Herschel-Bulkey model in OpenFOAM, defined by the following parameters:  $\rho = 1, \nu_0 = 0.00125, \tau_0 = 0.00125, k = 0.015625, n = 1.88$ . The freestream velocity is contained within the interval  $u \in [1.0, 2.0]$ , with speed increments of 0.1, which results in a total of  $N_{sim} = 11$  cases. Each simulation is discretized in  $n_{snap} = 600$  time increments of  $\Delta t = 0.005$ .

The input of the autoencoder are the low resolution velocity and pressure fields, with size  $3 \times 16 \times 48$ , while the output are the velocity and pressure fields at the original resolution and a higher one, with sizes  $3 \times 16 \times 48$  and  $3 \times 64 \times 192$ . Both the encoder and decoder use convolutional layers with  $N_{ch} = 64$  channels at a kernel size of  $k = 3$ , following a ResNet-like structure [43]. The number of latent variables at the bottleneck is set to  $d = 6$ . The activation function used is the leaky-ReLU with a negative slope of 0.1, except for the last layer of both the encoder and decoder, where linear activations are used. The adversarial hyperparameter weight is set to  $\lambda_{adv}^{AAE} = 10^{-3}$ . The optimizer used is Adam [44] with a learning rate set to  $\lambda_r^{AAE} = 10^{-4}$  with decreasing order of magnitude on epochs 600 and 1200, a weight decay set to  $w_d^{AAE} = 10^{-6}$ , and a total number of  $N_{epochs} = 1800$  epochs. Latent variables obtained at the bottleneck are then used as input variables for the structure preserving neural network. Fig. 9(a) shows the training and validation loss for the adversarial autoencoder.

The SPNN input size coincides with the AAE latent dimension,  $N_{in}^{SPNN} = d = 6$ , while the output size is  $N_{out}^{SPNN} = d \cdot (d + 2) = 48$ . The number of hidden layers of the SPNN is  $N_{hl}^{SPNN} = 5$  with 120 neurons each one, Leaky-ReLU activations and linear for the last layer. The data weight hyperparameter is set to  $\lambda_{data}^{SPNN} = 10^2$ . The SPNN is trained for  $N_{epochs} = 6000$  epochs using the Adam optimizer, and a batch size of  $B_{size} = 128$ . The learning rate is set to  $\lambda_r^{SPNN} = 10^{-3}$ , decreasing one order of magnitude on epoch 2000 and 4000. The weight decay is set to  $w_d^{SPNN} = 10^{-4}$  and noise variance is added to the train set,  $\sigma_{noise}^2 = 10^{-5}$ . Fig. 9(b) shows the training and validation loss for the SPNN.

4.2.2. Results

Fig. 10 shows the prediction results for the pressure and velocity fields obtained by the autoencoder for both, low and high resolution, as well as the absolute error for each field. The low and high resolution fields reconstructed by the autoencoder show good agreement with the ground truth fields obtained from the CFD simulation. A box plot containing the error of the state variables for the train and test cases is shown in Fig. 11, achieving a mean error lower than 3% for the pressure and velocity fields in both low and high resolution.

As in the previous case, the autoencoder has been compared with the bicubic interpolation technique. Comparison is shown in Table 2. As expected, considering the results obtained for the Newtonian fluid case, the AAE clearly outperforms the bicubic interpolation, specially if the time difference between both methods is considered, with the AAE being almost 45 times faster than the bicubic interpolation.

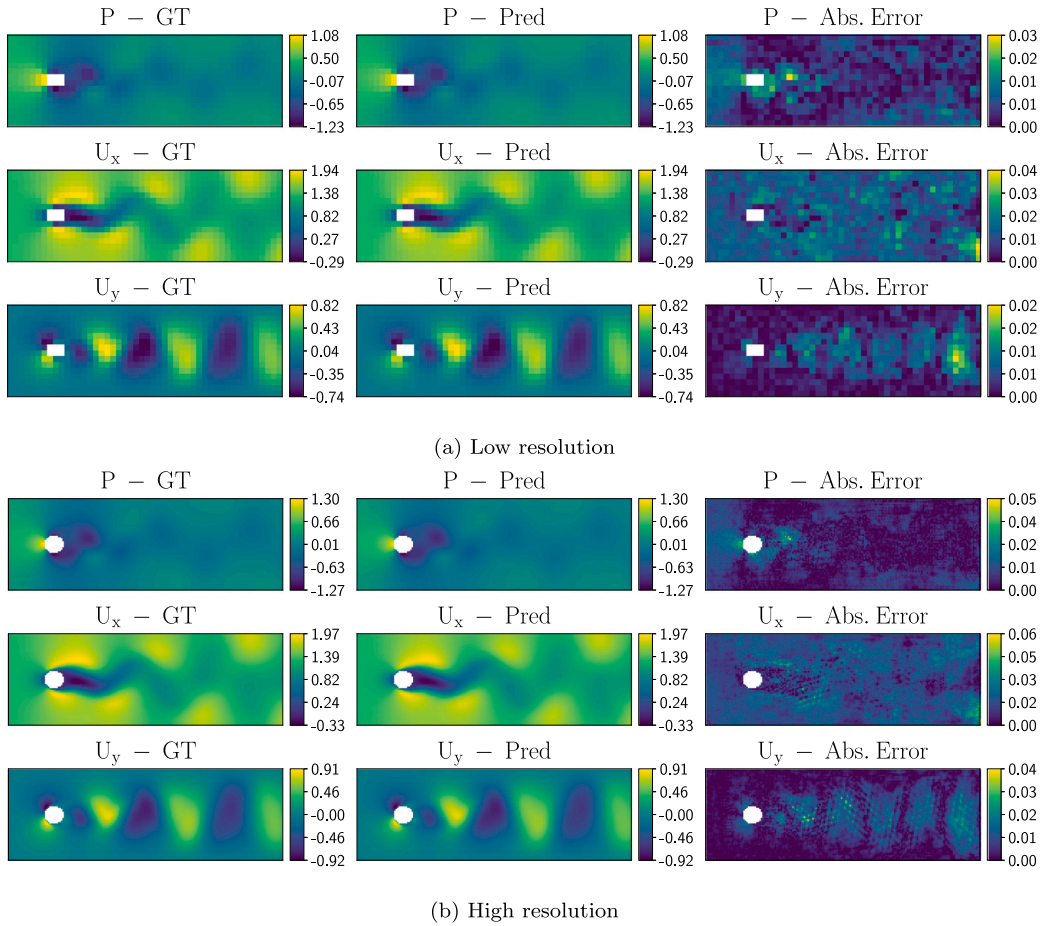


Fig. 10. Results of the AAE for two different snapshots. 10a: Low resolution Ground Truth (GT), AAE prediction and absolute error fields for  $P$ ,  $U_x$  and  $U_y$  fields for the first snapshot. 10b: High resolution GT, AAE prediction and absolute error fields for the same snapshot shown in Fig. 10a. 10c: Low resolution fields for the second snapshot. 10d: High resolution fields for the snapshot shown in 10c.

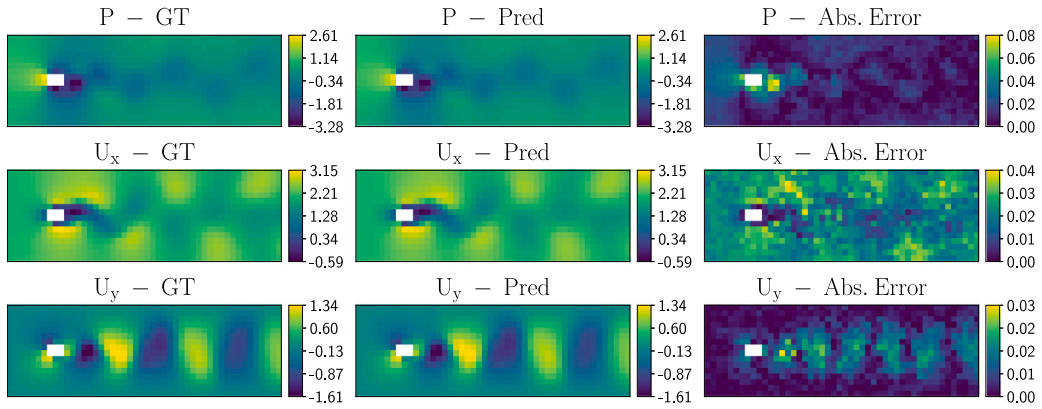
**Table 2**  
 Comparison between the AAE and bicubic interpolation. For every variable, relative error is shown, and time reconstruction time for all dataset is computed.

	AAE	Bicubic interpolation
$P$ (-)	0.0257	0.0557
$U_x$ (-)	0.0113	0.0228
$U_y$ (-)	0.0240	0.0639
Time (s)	197.57	8818.12

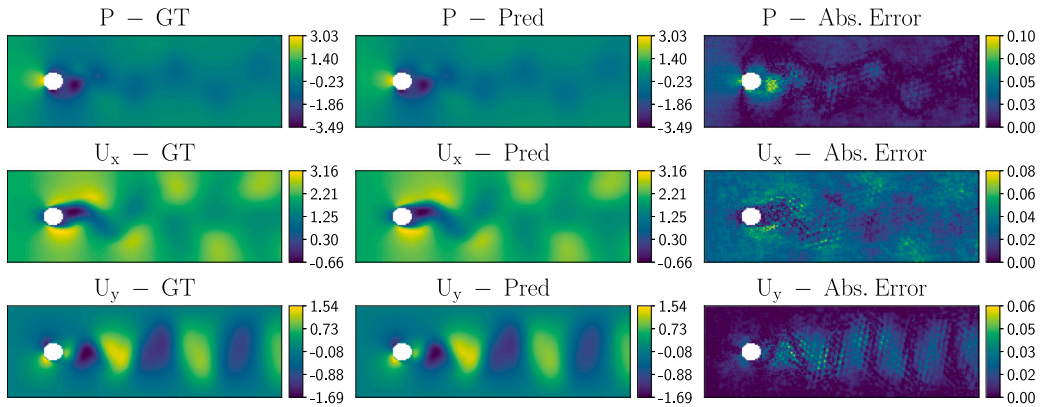
Fig. 12 shows the comparison between the ground truth latent variables, the ones obtained by the AAE and the prediction made by the SPNN. As in the previous case, the SPNN is able to integrate the latent variables in the reduced space successfully with respect to the original AAE encoding.

As previously, the SPNN is compared with a black box neural network, which is trained using the same hyperparameters as the SPNN, except for the outputs size, for this example  $N_{in}^{BB} = N_{out}^{BB} = 6$ , the same as the bottleneck of the AAE and the input of the SPNN. Results are compared in Fig. 13, which shows the velocity accumulated error mean and standard deviation with a confidence interval of 95% for every simulation at each snapshot. The SPNN error (Fig. 13(a)) remains lower than the black box neural network (Fig. 13(b)), proving that the thermodynamic bias helps the network to converge to the correct solution.

The ground truth simulations were performed on a MacBook Pro M1 Pro, with each simulation taking around 30 min to complete. Both, the AAE and the SPNN were trained on a Linux-based machine equipped with a Intel i9-13900K CPU and a NVIDIA RTX 4090 GPU using the Pytorch framework. Training the AAE took 2.5 h in the NVIDIA GPU, while the SPNN training time was around 20 min. Regarding the inference time, each latent variables prediction can be obtained in less than 1 s on a MacBook Pro M1 Pro, although rendering the video with the complete prediction takes about 10 s, which is considerably faster than that computing the high-fidelity model.



(c) Low resolution



(d) High resolution

Fig. 10. (continued).

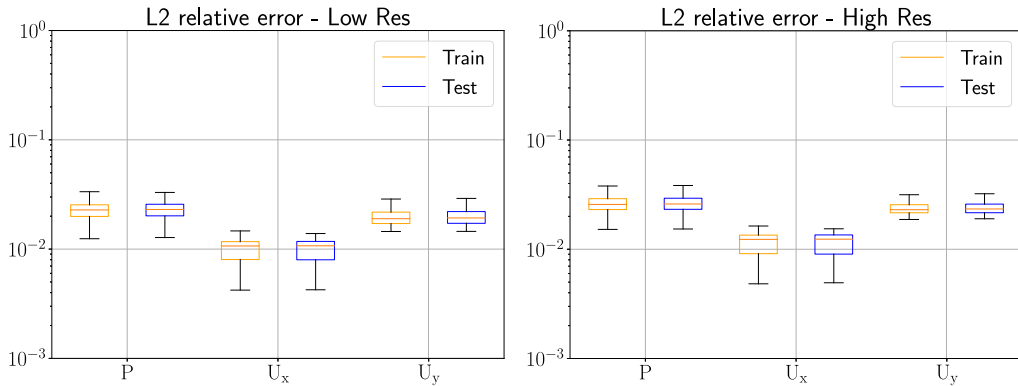


Fig. 11. Box plots for the relative L2 error of the autoencoder for all the snapshots of the non-Newtonian fluid for both train and test cases, in low (left) and high (right) resolution.

5. Conclusions

In this work we have presented a new methodology to increase the spatial resolution of predictions obtained by learned simulators, while ensuring a thermodynamics-aware prediction, satisfying the basic principles of thermodynamics. The proposed AAE architecture is able to encode the information to a reduced-order space and to produce high-resolution output fields from low resolution input thanks to its generative capabilities. The AAE has been compared to a classical resolution augmentation technique: the bicubic interpolation. Not only the AAE outperforms the bicubic interpolation, but it is also considerably faster, making it






Fig. 12. Results of the SPNN integration with respect to the latent variables obtained by the AAE (ground truth) for two different simulations for the non-Newtonian fluid. The dashed line corresponds to the SPNN prediction while the continuous line is the ground truth. Each latent variable is represented in a different colour to simplify the identification of the ground truth and its corresponding predicted value. Left: simulation at  $U_{in}=1.4$ . Right: simulation at  $U_{in}=1.7$ .




Fig. 13. Evolution of the relative error during the whole integration of the velocity for all the simulations. The SPNN error remains lower than the black box neural network during all the prediction rollout, proving that the physical bias helps to reach the correct solution.

feasible for quasi-real-time or even real-time applications. Additionally, AAE resolution augmentation technique can be applied to a wider range of geometries than bicubic interpolation. The structure-preserving neural network is able to estimate the evolution of the encoded variables in the reduced space and then the decoder re-projects the SPNN prediction to the original and higher resolution spaces. The SPNN is compared to a black-box approach, outperforming it thanks to the GENERIC formalism, as it adds physical constraints to the prediction that act as an inductive bias. The results show good agreement between our predictions and the synthetic ground truth obtained by CFD for the two examples analysed. However, there are some limitations in the current work that could be improved in the future:

- **Database:** The present work makes use of a synthetic database generated by a CFD tool. However, real data coming from sensors could be used to train a system to work with real-world digital twins. Additionally, the database could be augmented with different geometry cases, improving the generalization to unseen geometries.
- **Integration scheme:** In this work, an Euler integration scheme is used. This integration scheme is simple, and higher order integration schemes like the midpoint rule, Heun’s method or a Runge–Kutta method [28,45] could improve the accuracy of the SPNN. This would also allow the network to work with bigger time increments. However, increasing the complexity of the integration scheme would require more forward passes of the neural network for each time step, slowing the training process.
- **Net architecture:** Graph Neural Networks (GNNs) [35,46] could be used to take advantage of their unstructured data, in comparison to convolutional neural networks, that require grid-structured information. Thus, GNNs could be applied to real-world applications, e.g., a digital-twin of a system whose sensors are not evenly distributed.



## CRedit authorship contribution statement

**Carlos Bermejo-Barbanoj:** Writing – original draft, Visualization, Validation, Software, Investigation, Formal analysis, Data curation. **Beatriz Moya:** Writing – review & editing, Supervision, Software, Methodology, Investigation, Formal analysis, Conceptualization. **Alberto Badías:** Writing – review & editing, Supervision, Investigation, Formal analysis, Conceptualization. **Francisco Chinesta:** Writing – review & editing, Resources, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Conceptualization. **Elías Cueto:** Writing – original draft, Supervision, Resources, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Conceptualization.

## Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Elias Cueto reports financial support was provided by US Army Combat Capabilities Development Command Army Research Laboratory. Elias Cueto reports financial support was provided by Ministry of Economic Affairs and Digital Transformation. Elias Cueto reports financial support was provided by ESI Group. Francisco Chinesta reports financial support was provided by ESI Group. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Data availability

Data will be made available on request.

## Acknowledgements

This work was supported by the Spanish Ministry of Science and Innovation, AEI/10.13039/501100011033, through Grant number PID2020-113463RB-C31 and by the Ministry for Digital Transformation and the Civil Service, through the ENIA 2022 Chairs for the creation of university-industry chairs in AI, through Grant TSI-100930-2023-1.

This material is also based upon work supported in part by the Army Research Laboratory and the Army Research Office under contract/grant number W911NF2210271.

This research is also part of the DesCartes programme and is supported by the National Research Foundation, Prime Minister Office, Singapore under its Campus for Research Excellence and Technological Enterprise (CREATE) programme.

The authors also acknowledge the support of ESI Group through the chairs at the University of Zaragoza and at ENSAM Institute of Technology.

## References

- [1] Y. Lecun, L. Bottou, Y. Bengio, P. Haffner, Gradient-based learning applied to document recognition, *Proc. IEEE* 86 (11) (1998) 2278–2324.
- [2] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, Yoshua Bengio, Generative adversarial nets, in: Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, K.Q. Weinberger (Eds.), in: *Advances in Neural Information Processing Systems*, vol. 27, Curran Associates, Inc, 2014.
- [3] Christian Ledig, Lucas Theis, Ferenc Huszar, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, Wenzhe Shi, Photo-realistic single image super-resolution using a generative adversarial network, 2017.
- [4] Xintao Wang, Ke Yu, Shixiang Wu, Jinjin Gu, Yihao Liu, Chao Dong, Chen Change Loy, Yu Qiao, Xiaoou Tang, ESRGAN: Enhanced super-resolution generative adversarial networks, 2018.
- [5] Mathis Bode, Michael Gauding, Zeyu Lian, Dominik Denker, Marco Davidovic, Konstantin Kleinheinz, Jenia Jitsev, Heinz Pitsch, Using physics-informed enhanced super-resolution generative adversarial networks for subfilter modeling in turbulent reactive flows, *Proc. Combust. Inst.* 38 (2) (2021) 2617–2625.
- [6] Kai Fukami, Koji Fukagata, Kunihiko Taira, Super-resolution analysis via machine learning: A survey for fluid flows, *Theor. Comput. Fluid Dyn.* 37 (4) (2023) 421–444.
- [7] Linqi Yu, Mustafa Z. Yousif, Meng Zhang, Sergio Hoyas, Ricardo Vinuesa, Hee-Chang Lim, Three-dimensional ESRGAN for super-resolution reconstruction of turbulent flows with tricubic interpolation-based transfer learning, *Phys. Fluids* 34 (12) (2022) 125126.
- [8] Daniel Kelshaw, Georgios Rigas, Luca Magri, Physics-informed CNNs for super-resolution of sparse observations on dynamical systems, 2022.
- [9] Francisco Chinesta, Elias Cueto, Emmanuelle Abisset-Chavanne, Jean Louis Duval, Fouad El Khaldi, Virtual, digital and hybrid twins: A new paradigm in data-based engineering and engineered data, *Arch. Comput. Methods Eng.* 27 (2020) 105–134.
- [10] Adil Rasheed, Omer San, Trond Kvamsdal, Digital twin: Values, challenges and enablers, 2019, arXiv preprint arXiv:1910.01719.
- [11] Charles Fefferman, Sanjoy Mitter, Hariharan Narayanan, Testing the manifold hypothesis, *J. Amer. Math. Soc.* 29 (4) (2016) 983–1049.
- [12] S. Niroomandi, I. Alfaro, E. Cueto, F. Chinesta, Real-time deformable models of non-linear tissues by model reduction techniques, *Comput. Methods Programs Biomed.* 91 (3) (2008) 223–231.
- [13] Alberto Badías, Sarah Curtit, David González, Iciar Alfaro, Francisco Chinesta, Elías Cueto, An augmented reality platform for interactive aerodynamic design and analysis, *Internat. J. Numer. Methods Engrg.* 120 (1) (2019) 125–138.
- [14] Beatriz Moya, Iciar Alfaro, David Gonzalez, Francisco Chinesta, Elías Cueto, Physically sound, self-learning digital twins for sloshing fluids, *PLoS One* 15 (6) (2020) e0234569.
- [15] Zulkeefal Dar, Joan Baiges, Ramon Codina, Artificial neural network based correction for reduced order models in computational fluid mechanics, *Comput. Methods Appl. Mech. Engrg.* 415 (2023) 116232.
- [16] Ian J. Goodfellow, Yoshua Bengio, Aaron Courville, *Deep Learning*, MIT Press, Cambridge, MA, USA, 2016, <http://www.deeplearningbook.org>.
- [17] Rachida Chakir, Benjamin Streichenberger, P. Chatellier, A non-intrusive reduced basis method for urban flows simulation, 2021.
- [18] Elise Grosjean, Yvon Maday, Error estimate of the non-intrusive reduced basis (NIRB) two-grid method with parabolic equations, *SMAI J. Comput. Math.* 9 (2023) 227–256.

- [19] Hamidreza Eivazi, Soledad Le Clainche, Sergio Hoyas, Ricardo Vinuesa, Towards extraction of orthogonal and parsimonious non-linear modes from turbulent flows, *Expert Syst. Appl.* 202 (2022) 117038.
- [20] Yuning Wang, Alberto Solera-Rico, Carlos Sanmiguel Vila, Ricardo Vinuesa, Towards optimal  $\beta$ -variational autoencoders combined with transformers for reduced-order modelling of turbulent flows, *Int. J. Heat Fluid Flow* 105 (2024) 109254.
- [21] M. Raissi, P. Perdikaris, G.E. Karniadakis, Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations, *J. Comput. Phys.* 378 (2019) 686–707.
- [22] Chayan Banerjee, Kien Nguyen, Clinton Fookes, George Karniadakis, Physics-informed computer vision: A review and perspectives, 2023.
- [23] Miroslav Grmela, Hans Christian Öttinger, Dynamics and thermodynamics of complex fluids. I. Development of a general formalism, *Phys. Rev. E* 56 (1997) 6620–6632.
- [24] Hans Christian Öttinger, Miroslav Grmela, Dynamics and thermodynamics of complex fluids. II. Illustrations of a general formalism, *Phys. Rev. E* 56 (1997) 6633–6655.
- [25] Quercus Hernández, Alberto Badías, David González, Francisco Chinesta, Elías Cueto, Structure-preserving neural networks, *J. Comput. Phys.* 426 (2021) 109950.
- [26] Quercus Hernandez, Alberto Badías, Francisco Chinesta, Elías Cueto, Thermodynamics-informed graph neural networks, *IEEE Trans. Artif. Intell.* (2022) 1.
- [27] Kookjin Lee, Nathaniel A. Trask, Panos Stinis, Machine learning structure preserving brackets for forecasting irreversible processes, 2021.
- [28] Zhen Zhang, Yeonjong Shin, George Em Karniadakis, GFNNs: GENERIC formalism informed neural networks for deterministic and stochastic dynamical systems, *Phil. Trans. R. Soc. A* 380 (2229) (2022) 20210207.
- [29] Quercus Hernandez, Alberto Badías, David González, Francisco Chinesta, Elías Cueto, Deep learning of thermodynamics-aware reduced-order models from data, *Comput. Methods Appl. Mech. Engrg.* 379 (2021) 113763.
- [30] Michal Pavelka, Václav Klika, Miroslav Grmela, Multiscale Thermo-Dynamics: Introduction to GENERIC, Walter de Gruyter GmbH & Co KG, 2018.
- [31] E. Weinan, A proposal on machine learning via dynamical systems, *Commun. Math. Stat.* 5 (1) (2017) 1–11.
- [32] Alireza Makhzani, Jonathon Shlens, Navdeep Jaitly, Ian Goodfellow, Adversarial autoencoders, in: *International Conference on Learning Representations*, 2016.
- [33] Diederik P. Kingma, Max Welling, Auto-encoding variational bayes, 2013, arXiv preprint arXiv:1312.6114.
- [34] Beatriz Moya, Alberto Badías, Iciar Alfaro, Francisco Chinesta, Elías Cueto, Digital twins that learn and correct themselves, *Internat. J. Numer. Methods Engrg.* 123 (13) (2022) 3034–3044.
- [35] Peter W. Battaglia, Jessica B. Hamrick, Victor Bapst, Alvaro Sanchez-Gonzalez, Vinicius Zambaldi, Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan Faulkner, Caglar Gulcehre, Francis Song, Andrew Ballard, Justin Gilmer, George Dahl, Ashish Vaswani, Kelsey Allen, Charles Nash, Victoria Langston, Chris Dyer, Nicolas Heess, Daan Wierstra, Pushmeet Kohli, Matt Botvinick, Oriol Vinyals, Yujia Li, Razvan Pascanu, Relational inductive biases, deep learning, and graph networks, 2018.
- [36] Philip J. Morrison, A paradigm for joined Hamiltonian and dissipative systems, *Physica D* 18 (1) (1986) 410–419.
- [37] Alexander Mielke, On thermodynamically consistent models and gradient structures for thermoplasticity, *GAMM-Mitt.* 34 (1) (2011) 51–58.
- [38] Alexander Mielke, Formulation of thermoelastic dissipative material behavior using GENERIC, *Contin. Mech. Thermodyn.* 23 (3) (2011) 233–256.
- [39] Ignacio Romero, Algorithms for coupled problems that preserve symmetries and the laws of thermodynamics: Part I: Monolithic integrators and their application to finite strain thermoelasticity, *Comput. Methods Appl. Mech. Engrg.* 199 (25–28) (2010) 1841–1858.
- [40] Ignacio Romero, Algorithms for coupled problems that preserve symmetries and the laws of thermodynamics: Part II: Fractional step methods, *Comput. Methods Appl. Mech. Engrg.* 199 (33–36) (2010) 2235–2248.
- [41] Beatriz Moya, Alberto Badías, David González, Francisco Chinesta, Elías Cueto, Physics perception in sloshing scenes with guaranteed thermodynamic consistency, *IEEE Trans. Pattern Anal. Mach. Intell.* 45 (2) (2023) 2136–2150.
- [42] H.G. Weller, G. Tabor, H. Jasak, C. Fureby, A tensorial approach to computational continuum mechanics using object-oriented techniques, *Comput. Phys.* 12 (6) (1998) 620–631.
- [43] Kaiming He, Xiangyu Zhang, Shaoqing Ren, Jian Sun, Deep residual learning for image recognition, in: *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, 2016, pp. 770–778.
- [44] Diederik P. Kingma, Jimmy Ba, Adam: A method for stochastic optimization, 2017.
- [45] Yi-Jen Wang, Chin-Teng Lin, Runge-Kutta neural network for identification of dynamical systems in high accuracy, *IEEE Trans. Neural Netw.* 9 (2) (1998) 294–307.
- [46] Michael M. Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, Pierre Vandergheynst, Geometric deep learning: Going beyond Euclidean data, *IEEE Signal Process. Mag.* 34 (4) (2017) 18–42.