



A continuous convolutional trainable filter for modelling unstructured data

Dario Coscia¹ · Laura Meneghetti¹ · Nicola Demo¹ · Giovanni Stabile² · Gianluigi Rozza¹

Received: 26 October 2022 / Accepted: 11 February 2023 / Published online: 21 March 2023
© The Author(s) 2023

Abstract

Convolutional Neural Network (CNN) is one of the most important architectures in deep learning. The fundamental building block of a CNN is a trainable filter, represented as a discrete grid, used to perform convolution on discrete input data. In this work, we propose a continuous version of a trainable convolutional filter able to work also with unstructured data. This new framework allows exploring CNNs beyond discrete domains, enlarging the usage of this important learning technique for many more complex problems. Our experiments show that the continuous filter can achieve a level of accuracy comparable to the state-of-the-art discrete filter, and that it can be used in current deep learning architectures as a building block to solve problems with unstructured domains as well.

Keywords Convolutional neural network · Continuous filter · Deep learning · Unstructured data

1 Introduction

In the deep learning field, a convolutional neural network (CNN) [1] is one of the most important architectures, widely used in academia and industrial research. For an overview of the topic, the interested reader might refer to [2–6]. Despite the great success in many fields including, but not limited, to computer vision [7–9] or natural language processing [10,11], current CNNs are constrained to structural

data. Indeed, the basic building block of a CNN is a trainable filter, represented by a discrete grid, which performs cross-correlation, also known as convolution, on a discrete domain. Nevertheless, the idea behind convolution can be easily extended mathematically to unstructured domains, for reference see [12]. One possible approach for this kind of problem is the graph neural networks (GNN) [13,14], where a graph is built starting from the topology of the discretized space. This allows us to apply convolution even to unstructured data by looking at the graph edges, bypassing in this way the limitations of the standard CNNs approach. However, GNNs typically require huge computational resources, due to their implicit complexity.

Instead in this article, we present a methodology to apply CNNs to unstructured data by introducing a continuous extension of a convolutional filter, named *continuous filter*, without modeling the data using a graph.

The main idea, which is depicted graphically in Fig. 1, relies on approximating the continuous filter with a trainable function using a feed-forward neural network and perform standard continuous convolution between the input data and the continuous filter. Previous works have introduced different approaches to continuous convolution in various settings ranging from informatics and graph neural networks to physics and modeling quantum interactions, see for example [15–17]. Even so, the latter is difficult to generalize, and an analogy with a discrete CNN filter is not straightforward.

Laura Meneghetti, Nicola Demo and Giovanni Stabile have contributed equally to this work.

✉ Dario Coscia
dario.coscia@sissa.it

Laura Meneghetti
laura.meneghetti@sissa.it

Nicola Demo
nicola.demo@sissa.it

Giovanni Stabile
giovanni.stabile@uniurb.it

Gianluigi Rozza
gianluigi.rozza@sissa.it

¹ SISSA Mathematics Area, mathLab, via Bonomea, 265, Trieste 34136, Italy

² University of Urbino Carlo Bo, department of Pure and Applied Sciences, Informatics and Mathematics Section, piazza della Repubblica, 13, Urbino 61029, Italy

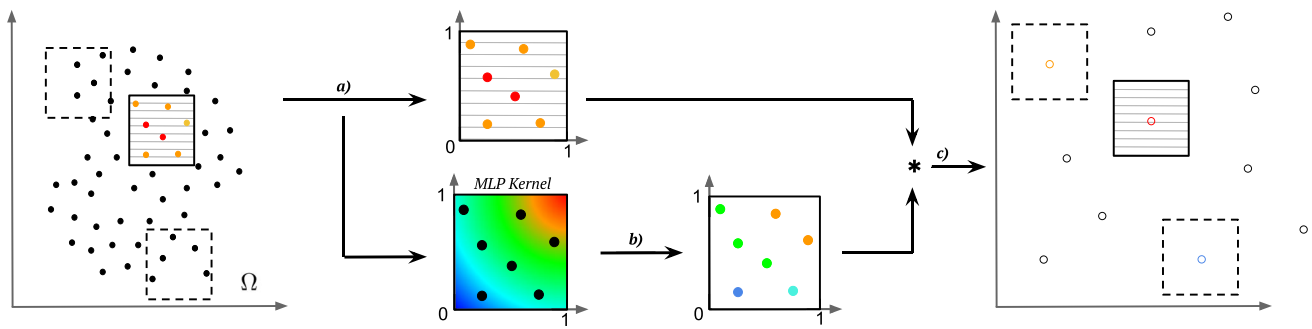


Fig. 1 Continuous convolutional filter process. The unstructured domain input points falling into the filter are mapped in the filter domain **a**. The filter values are approximated with a MLP kernel **b**. Finally, the convolution between the mapped values and the filter values is performed **c**

To our extent [18,19] are the closest works in literature to our approach, both approximating the trainable filter function with a feed-forward neural network and performing continuous convolution. However, [18] and [19] focus on filters with unbounded domains for convolution. In our work, we instead fix the dimension of the filter, as in state of the art discrete filters, and learn the approximation function on the filter domain. This introduces a neat analogy to discrete CNN filters. Furthermore, differently from [18,19], we also cover important properties of convolution, such as transposed convolution or different approaches to multichannel convolution. To summarize, in this work we aim to reproduce as closely as possible a discrete CNN filter but in a continuous not structured domain setting, in order to exploit the main deep learning architectures, based on CNNs, to solve problems in not discrete domains. To the best of the authors' knowledge, our approach to continuous convolution has not been explored in literature yet.

The main novelties of this work rely on:

- Building a new framework, based on continuous filters, for working with unstructured data (continuous filter).
- Defining a neat analogy between continuous (transposed) convolution and state of the art discrete (transposed) convolution in CNNs.
- Apply continuous convolutional layers in a CNN with partially-completed input.
- Exploiting general strategies to work with continuous convolutional autoencoders for dimensionality reduction and system output predictions at unseen time steps.

All this, we highlight, preserving the features of the standard CNNs, which make such an approach effective even dealing with large datasets. The present contribution is organised as follows: in Sect. 2, a small review of deep learning architectures useful for later analysis is done, as well as introducing the continuous filter for one-dimensional and multi-dimensional channels. In the same Section, we introduce the main idea to perform transposed continuous

convolution. Sect. 3 is focused on numerical results. First, we validate the proposed methodology on a discrete domain problem using a continuous CNN and compare it with its discrete representation. Second, we show that continuous convolution can also work with partially-completed images. Last, we present different deep learning architectures using continuous filters to solve the step Navier Stokes problem, and the multiphase problem. Finally, conclusions follow in Sect. 4.

2 Methodology

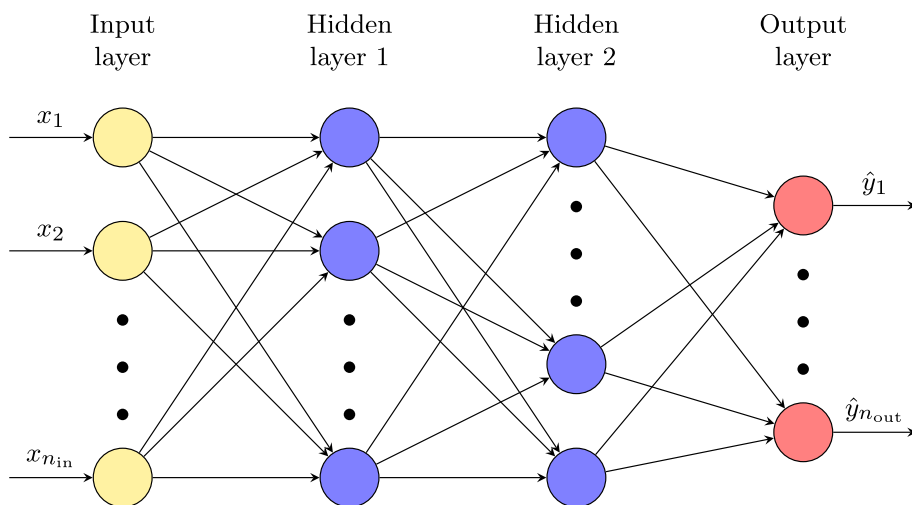
This Section focuses on the various methodologies we rely for building the continuous filter, as well as the introduction of the framework. First of all, we will describe briefly the feed-forward neural network and the discrete filter for a CNN in Sect. 2.1 and Sect. 2.2 respectively. As already mentioned in Sect. 1, one of the main novelty of the work is building a new framework based on continuous convolution. Hence, Sect. 2.3 concerns the introduction of our framework in different settings: single channel, multiple channel and transposed convolution using the continuous filter.

2.1 Feed-forward neural network

Feed-forward Neural Network, or *multi-layer perceptron* (MLP), is the most basic, yet one of the most important, building block of most current deep learning architectures [3,5,20]. Widely used in deep learning, MLPs have the ability to approximate any continuous function due to the universal approximation theorem [21–23].

More technically, given an input vector $\mathbf{x} \in \mathbb{R}^{n_{in}}$ and a function to approximate $\phi : \mathbb{R}^{n_{in}} \rightarrow \mathbb{R}^{n_{out}}$; the MLP approximation is done using a parameterised function class $\mathcal{F} = \{f_{\theta \in \Theta}\}$, where θ are trainable parameters of the network, belonging to the parameters' space Θ . A MLP can be represented as a directed acyclic graph, as depicted in Fig. 2. In particular, it is composed by an input layer, an output layer

Fig. 2 Schematic structure of Feed-Forward Neural Network



and a certain number of hidden layers, where the processing units of network, called *neurons*, perform the computation. Each layer i , with $i \in 0, \dots, M$, can be thought as a function $f^{(i)}$ belonging to \mathcal{F} , and the overall network function is given by the layers’ composition [24]:

$$f = f^{(M)} \circ f^{(M-1)} \circ \dots \circ f^{(1)} \circ f^{(0)}. \tag{1}$$

Hence, a single layer i , is a function $f^{(i)} : \mathbb{R}^{n_i} \rightarrow \mathbb{R}^{n_{i+1}}$, where n_i , represents the number of neurons in layer i , with $n_0 = n_{in}$ and $n_{M+1} = n_{out}$. Each layer i is composed by $\theta_i = (\mathbf{w}_{(i)}, \mathbf{b}_{(i)})$ parameters, where $\mathbf{w}_{(i)}$ is a real matrix $n_{i+1} \times n_i$, called *weight matrix*, and $\mathbf{b}_{(i)}$ is a real vector of dimension n_{i+1} , called *bias*. The output vector $\mathbf{h}_{(i+1)}$ of layer i , corresponding to input vector of layer $i + 1$ (except for the output layer), is then calculated using:

$$\mathbf{h}_{(i+1)} = f^{(i)}(\mathbf{h}_{(i)} | \theta_i) = \delta^{(i)}(\mathbf{w}_{(i)} \cdot \mathbf{h}_{(i)} + \mathbf{b}_{(i)}), \tag{2}$$

where $\mathbf{h}_{(0)} = \mathbf{x}$, and $\mathbf{h}_{(M+1)} = \hat{\mathbf{y}}$ is the output of the network. The function $\delta^{(i)} : \mathbb{R}^{n_i} \rightarrow \mathbb{R}^{n_{i+1}}$ is called *activation*, introducing non-linearity through the network; common choices are represented by the ReLU function, the sigmoid, the logistic function or the radial activation functions. By using Eq. 1 and 2, one can express mathematically a MLP architecture.

During the *training process*, in which a data-set $\mathcal{D} = \{(\mathbf{x}_i, \phi(\mathbf{x}_i))\}_{i=1}^n$ composed by n observation is fed into the network, the MLP parameters θ are modified in order to minimize a loss function $\mathcal{L}(\theta | \mathcal{D}, f)$. The choice of the loss function depends on the specific problem of application [3,5,25]. Hence, the learning phase can be summarised mathematically as:

$$\min_{\theta} \{ \mathcal{L}(\theta | \mathcal{D}, f) \}. \tag{3}$$

In practice, to solve the minimization problem, different optimization algorithms based on back-propagation can be used, see [26–28] for further reference. The optimization phase is done in multiple *training epochs*, i.e. a complete repetition of the parameter update involving the complete training data-set \mathcal{D} .

2.2 Discrete filter in convolutional neural networks

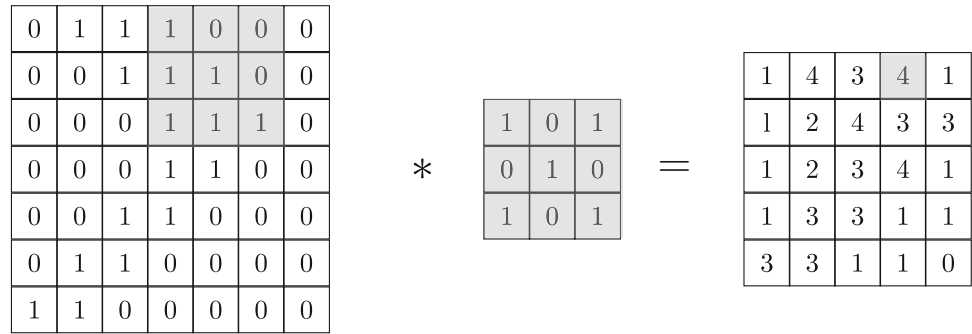
Convolutional Neural Network (CNN) is a class of deep learning architectures, vastly applied in computer vision [7–9,29]. Over the past years, different CNN architectures have been presented, for instance AlexNet [30], ResNet [31], Inception [32], VGGNet [33]. Differently from MLPs, in which affine transformations are performed for learning, a convolutional layer actually performs the convolution of the input data \mathcal{I} and the so called convolutive filter \mathcal{K} , such that

$$(\mathcal{I} * \mathcal{K})(\mathbf{x}) = \int_{-\infty}^{\infty} \mathcal{I}(\mathbf{x} + \boldsymbol{\tau}) \mathcal{K}(\boldsymbol{\tau}) d\boldsymbol{\tau}. \tag{4}$$

CNNs perform such convolution¹ in a discrete setting, using a tensorial representation of the two functions \mathcal{I} and \mathcal{K} instead of their continuous formulation. Thus, discrete correlation is computed as $(\mathcal{I} * \mathcal{K})(\mathbf{x}) = \sum_{\boldsymbol{\tau}=-\infty}^{\infty} \mathcal{I}(\mathbf{x} + \boldsymbol{\tau}) \mathcal{K}(\boldsymbol{\tau})$, with $\mathbf{x}, \boldsymbol{\tau} \in \mathbb{Z}^d$ (with d dimensions), where the latter infinite summation can be truncated by discarding the null products. In this way, it is not necessary to know the original function \mathcal{I} , but its evaluation at discrete coordinates. In this context, the filter \mathcal{K} can be represented as the tensor $\mathbf{K} \in \mathbb{R}^{N_1 \times \dots \times N_d}$ such that the element $\mathbf{K}_{i_1, \dots, i_d} \equiv \mathcal{K}(i_1, \dots, i_d)$ with $i_j \in$

¹ In many deep learning implementations the term convolution indicates what is known in mathematics as cross-correlation [3]. In this text, the term convolution will be used to indicate cross-correlation, thus adapting to the deep learning community convention.

Fig. 3 Discrete convolution operation on one dimensional tensor



$\{1, \dots, N_j\}, \forall j \in \{1, \dots, d\}$. Applying a similar representation also for the input, the convolution results in the sum of the element-wise multiplication between input and filter, as sketched in Fig. 3. The convolution is of course repeated for all the input components, by moving the filter across the input in a regularized fashion [3,5].

The filter components (the so-called weights) represent the trainable parameters of the convolutional layer, which are tuned during the training phase. In general, convolution reduces the size of a (multidimensional) array, performing *downsampling*. Conversely, the opposite transformation to downsampling, called *upsampling*, used by many deep learning architectures, e.g. autoencoders, uses *transposed convolution*. The interested reader might refer to [6,34] for more information regarding discrete (transposed) convolution.

2.3 Continuous convolutional filter

In contrast to discrete convolution as described in the previous Section, continuous two-dimensional convolution is mathematically defined as:

$$\mathcal{I}_{\text{out}}(x, y) = \int_{\mathcal{X}} \int_{\mathcal{Y}} \mathcal{I}(x + \tau_x, y + \tau_y) \cdot \mathcal{K}(\tau_x, \tau_y) d\tau_x d\tau_y, \tag{5}$$

where $\mathcal{K} : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$ is the *continuous filter* function, and $\mathcal{I} : \Omega \subset \mathbb{R}^2 \rightarrow \mathbb{R}$ is the input function. The continuous filter function is approximated using a MLP, thus trainable during the training phase. In order to maintain the parallelism with discrete convolution in CNNs, the definition adopted for continuous convolution differs from the mathematical one for which $\mathcal{X} = \mathcal{Y} = \mathbb{R}$. In fact, the continuous filter presented is defined on a close domain, smaller than the input function domain, as in the case of the discrete filter. The integral in Eq. 5 can be evaluated using different techniques from numerical analysis [35]. In our implementation, for simplicity, the double integral is approximated by a double sum on the grid nodes inside the filter domain. Given $\{(x_i, y_i)\}_{i=1}^n$ points of the input function mapped on the $\mathcal{X} \times \mathcal{Y}$ filter domain, we

approximate Eq. 5 as:

$$\mathcal{I}_{\text{out}}(\tilde{x}_i, \tilde{y}_i) = \sum_{x_i \in \mathcal{X}} \sum_{y_i \in \mathcal{Y}} \mathcal{I}(x_i + \tau_x, y_i + \tau_y) \cdot \mathcal{K}(x_i, y_i), \tag{6}$$

where $(\tau_x, \tau_y) \in \mathcal{S}$, with \mathcal{S} the set of available strides, corresponds to the current stride position of the filter, and $(\tilde{x}_i, \tilde{y}_i)$ points are obtained by taking the centroid of the filter position mapped on the Ω domain. It is important to remark, that the discretization strategy for Ω is irrelevant for the filter usage, since it can work even for holed or not connected domains. Finally, the strides positions \mathcal{S} , constituting the different positions of the filter on Ω , do not have to cover the whole Ω domain, thus representing a generalization of the discrete filter stride.

Having described the overall mathematical idea behind one channel continuous convolutions, in Algorithm 1 a possible simple implementation is reported. When implementing the continuous convolutional layer, we identified opportunities to optimize it (e.g., stacking mapped input to call just once the MLP or calling the mapping function on the entire data before entering in the `for` loop, allowing just one point search per forward pass). The pseudo-code in Algorithm 1 does not report, for simplicity, these optimizations used in building the layer. Moreover, the mapping function, which is responsible for finding the points inside the filter, can be implemented via linear search or advanced algorithms based on efficient data structures. The usage of efficient data structures may decrease the overall time complexity of the algorithm, leading to an efficient implementation of the convolutional layer. Note that the implementation used in this article relies on linear search for the mapping function.

The latter is easily extendable for multichannel convolution. Here two approaches can be adopted: multiple two-dimensional filters, or multiple three-dimensional filters. In the first case, the framework is analogous to the one of the discrete convolution in multi dimensional channels, see Sect. 2.2. We suggest to use this if the channels represent independent quantities (e.g., pressure and velocity in a fluid, or image channels). The second possibility is to define a three-dimensional filter, meaning a function

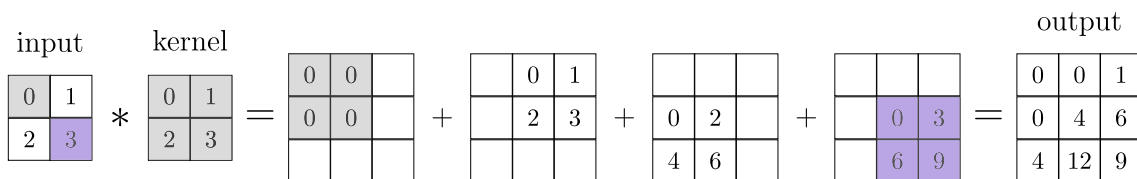


Fig. 4 Example of transposed convolution in discrete setting

Algorithm 1 One channel continuous convolution algorithm

```

Require:  $\mathcal{I}$  : input function
Require:  $\mathcal{K}$  : MLP approximation function
Require:  $\mathcal{S}$  : set of stride positions
Require:  $\Omega$  : domain discretization
 $\mathcal{I}_{out} \leftarrow \text{Initialize}(\Omega, \mathcal{S})$            ▷ Initialize output coordinates
for  $(\tau_x, \tau_y) \in \mathcal{S}$  do
   $(\mathbf{x}, \mathbf{y}) \leftarrow \text{Map}(\Omega, (\tau_x, \tau_y))$            ▷  $\Omega \rightarrow \mathcal{X} \times \mathcal{Y}$ 
   $\mathbf{v}_{\mathcal{K}} \leftarrow \mathcal{K}(\mathbf{x}, \mathbf{y})$                        ▷ MLP evaluation
   $\mathcal{I}_{out} \leftarrow \text{Convolve}(\mathcal{I}(\mathbf{x} + \tau_x, \mathbf{y} + \tau_y), \mathbf{v}_{\mathcal{K}})$    ▷ Convolution
end for
return  $\mathcal{I}_{out}$ 
    
```

$\mathcal{K} : \mathcal{X} \times \mathcal{Y} \times \mathcal{P} \rightarrow \mathbb{R}$, which takes as extra argument a channel parameter $p \in \mathcal{P}$. We suggest to use this if the channels represent correlated quantities (e.g., velocity along two directions in a fluid). While the multiple two-dimensional filters strategy needs Output Channels \times Input Channels independent filters, i.e. neural networks; the multiple three-dimensional filters strategy only needs Output Channels independent filters, since the multichannel in the input is already handled by the extra dimension in the filter. In our implementation we considered and tested both strategies.

Finally, we highlight that such filter allows to trivially extend the transposed convolution in the continuous setting. In several architectures, like the autoencoder employed in the numerical experiments of this contribution, it is indeed necessary to upsample the output of a given layer. With discrete filters, such operation is obtained by simply multiplying any elements in the input tensor by all the elements of the filter in a element-wise fashion. Figure 4 sketches a simple example of transposed convolution with 2D tensors. In case of overlapping elements in the output — such condition depends by the stride used in the transposed convolution — these elements are typically summed, but also other choices can be employed, like averaging.

In the continuous counterpart, the transposed operation is semantically the same, but the elements within the filter are not anymore finite. It is necessary indeed, in this latter case, to select the points within the filter domain, sample the kernel function at those points and finally weight them with the elements of the input tensor to produce the output.

3 Numerical results

In this Section different experiments and their results, using our continuous convolutional filter, are presented. In Sect. 3.2 a comparison between the state-of-the-art discrete filter and the continuous filter is done, while in Sect. 3.2.1 the advantages of using a continuous filter for partially-complete images are shown. Following, in Sect. 3.3 we validate the usage of transposed convolution using a Navier Stokes problem. Finally, in Sect. 3.4, the continuous filter is used to build an autoencoder architecture for reconstructing not seen time snapshot for the multiphase problem.

3.1 Software

In order to implement and construct the continuous convolutional filter, as well as performing all the experiments in this Section, we employed PyTorch [36] due to its versatility and its wide use in the deep learning community. Moreover, the open source C++ finite volume library OpenFOAM [37] is used for the mesh creation for the problems presented in Sects. 3.3 and 3.4, and the FEniCSx [38] for solving the Navier Stokes problem in Sect. 3.3 using finite elements method. Finally, the EZyRB Python library [39] is used for performing a comparison with proper orthogonal decomposition in Sect. 3.4.

3.2 Convolutional neural networks comparison on MNIST dataset

The first problem presented for validating the methodology is the classification task in a supervised learning setting. Training and testing is done on the MNIST dataset [40], as it represents a standard benchmark for classification. In particular we compare two CNNs, differing only for the first convolutional layer in which discrete and continuous filters are interchanged. The MNIST dataset used during test is composed by 60000 training images, and 10000 testing images. The images are saved as $1 \times 28 \times 28$ matrix I , where the first dimension represents the number of channels, while second and third dimension represent height and width respectively. When passing through the continuous filter the image is trans-

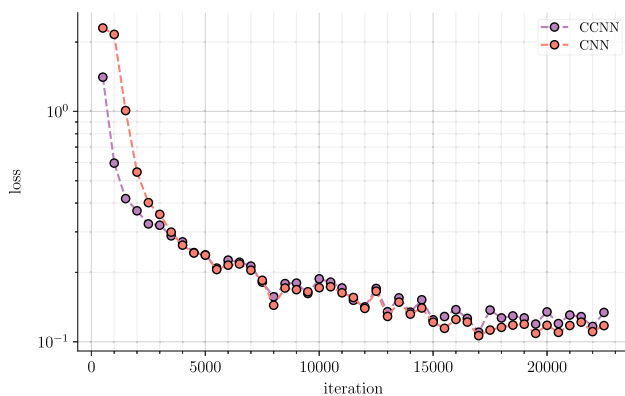


Fig. 5 Comparison of CCNN and CNN with respect to loss value and the number of iterations

formed using a bed of nail representation

$$I_{\text{cont}}(x, y) = \sum_{i,j} \delta(x - i, y - j) I[i, j], \quad (7)$$

where the sum is done on the matrix indices, as also done in [16]. This representation allows to treat a matrix as a continuous function, thus permitting the use of the continuous filter. The opposite transformation is done by arranging the function value in a matrix whose pixel positions are given by the (x, y) coordinates of the function I_{cont} .

Due to construction, the continuous filter is expected to perform equivalently to a discrete filter with structured data, thus classification on MNIST can be used as a validation methodology for the continuous filter framework. The general network used in training and testing is composed by two convolutional layers followed by three fully-connected layers of size 150, 24, and 10 respectively, with hyperbolic tangent activation after the first two layers. The first convolutional layer has filter size and stride both equal to 4, with one input and output channel, and zero padding. The second convolutional layer has filter size and stride both equal to 1, with one input channel and four output channels, and zero padding. For the continuous filter we used a MLP with two layers of size 12, 12 with ReLU activation for the inner network. Through the experiments only the first convolutional layer is modified in the general network by changing discrete with continuous filters (and vice-versa), but without changing any hyper-parameter, e.g. stride, filter size or padding. Let us indicate with CNN the general network using a discrete filter in the first layer, and CCNN the continuous version of the CNN, i.e. the one in which the first convolutional layer discrete filter is replaced with a continuous filter. Both networks are trained, minimizing the cross entropy loss [3], for 22500 iterations on the training set using the Stochastic Gradient Descent optimizer [41] with learning rate and momentum set to 0.001 and 0.9 respectively, and batch size equal to 8.

In Fig. 5 the training loss for the two networks is reported at different iterations. The result shows a faster convergence of the CCNN network with respect to the CNN network. Nevertheless, at plateau, both networks reach similar level of training loss, as expected. Furthermore, both networks have a similar decay shape of the loss, evidencing that the training process on a continuous filter using discrete data is comparable to a discrete filter.

In Table 1 the training accuracy, the testing accuracy, the size of the networks used during training and the training time are reported. The accuracy is defined as the number of correctly classified predictions divided by the total number of predictions. It is possible to note from the table that the training and testing accuracy is approximately the same between the two networks, thus confirming heuristically the hypothesis that both networks behave the same on structured data. The discrete filter tends to have a small percentage advantage 0.2% on training data, and 0.5% on test data over the continuous filter. Nevertheless, fine-tuning the inner network for the continuous filter might increase the test accuracy as well. Furthermore, we stress that the true power of continuous filters is employed when unstructured data are used, and the test only wants to confirm the validity of the methodology. The CCNN network presents a small overhead in terms of the size of the network due to the MLP used for approximating the continuous filter function. Nevertheless, the sizes of the two networks are still comparable in order of magnitude. Finally, due to expensive routines, e.g. the mapping function, the CCNN is slower compared to a standard CNN, resulting in higher training time. To summarize, in terms of accuracy using a CNN or a CCNN does not affect the overall performance. Even so, CCNN can be slightly slower than normal CNN due to continuous filter architecture construction, but CCNN can work with unstructured domains. In the following experiments, the power of continuous filter will be shown by working with unstructured domains.

3.2.1 Partially-completed images

Continuous convolutional neural networks have the ability to work also with partially completed images, unlike normal convolutional neural networks. Recalling the bed of nails representation reported in Eq. 7, some information can be removed from the image I by avoiding some indices (i^*, j^*) into the summation. This procedure corresponds to not evaluating the continuous representation of I in $(x = i^*, y = j^*)$, thus removing information from the image. Notice that the (i^*, j^*) positions are not filled with zeros, but they are actually not considered by the continuous filter. This procedure can not be done on a regular discrete filter since it would require to remove the (i^*, j^*) entry of the matrix I , thus losing the structure in the data, and making not possible to perform matrix operations. On the contrary, by using a con-

Table 1 Results obtained for CNN and CCNN networks trained on MNIST dataset. Train time and test time reported in seconds

Network	Train accuracy (%)	Test accuracy (%)	Size	Train time	Test time
CCNN	96.6	96.1	33593	630	21.1
CNN	96.8	96.6	33449	40	1.1

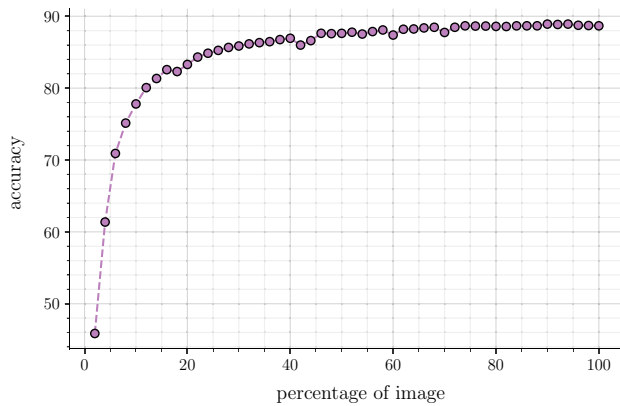


Fig. 6 Variation in testing accuracy at difference percentage of information on input images

tinuous representation of I , removing pixels corresponds to not evaluating the continuous function on them, and the filter can still be used since it can deal with unstructured data by construction.

The second problem presented deals with the classification task on a supervised learning setting using partially-completed images. In order to understand if a continuous filter can work with missing pixels images, i.e. the filter is able to generalize with loss of information, a simple two-layer network is used. In the first layer a continuous filter of size and stride both equal to 4 is posed, followed by the ReLU activation function and a linear layer (single MLP layer) of size 10 used to classify the output of the convolution. By adopting this simple network it is possible to understand how a single continuous filter is able to classify images as the pixel information reduces. Indeed, the network is not built for performance in accuracy, rather it is used to show that the continuous filter can work and achieve effective results also with unstructured missing pixels images. Training and testing is done on the MNIST dataset presented in the previous Section. The network is trained, minimizing the cross entropy loss for 22500 iterations on the training set using the Stochastic Gradient Descent optimizer with learning rate set to 0.001, and batch size equal to 8.

Let us define $\mathcal{P} \in [0, 100]$ the percentage of information in the image, i.e. the percentage of pixels of the original image. The pixels are removed using a uniform two dimensional distribution, selecting the $(x = i^*, y = j^*)$ coordinates where not evaluating the continuous representation of I . The experiment is performed by training, for each

different percentage in \mathcal{P} , the proposed network,² and then validating the network on the testing data. In Fig. 6 the results of the experiment are depicted, where the network accuracy is reported as a function of the percentage of image. The results show that the network is able to learn with missing value pixels. Starting from the baseline given by a percentage of 100, the network is robust in learning, obtaining comparable results to the baseline till roughly a percentage of 75. As the percentage decreases the accuracy also decreases, which is expected since information is hidden to the network. Nevertheless, with an image percentage of 20 the network decreases in accuracy only for a small 8% value, before a fast drop in accuracy due to little information fed into the network. Overall, it is evident that by using missing value images the continuous filter is able to generalize and learn, due to its ability of learning using unstructured domains.

3.3 Navier stokes problem

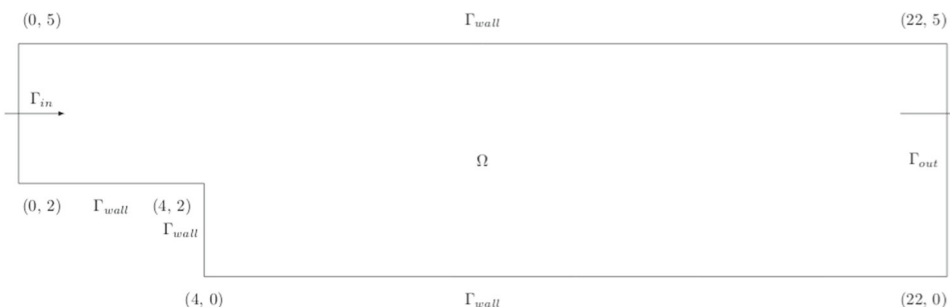
The second experiment shows the capability of the continuous filter in a unsupervised learning framework using autoencoders [42]. In detail, a typical CFD parametric benchmark is considered, where Navier Stokes equations are used to describe the flow in a back-step domain (Fig. 7). More specifically, equations and boundary conditions are set as follows:

$$\begin{aligned}
 \nu \Delta \mathbf{u} + (\mathbf{u} \cdot \nabla) \mathbf{u} + \nabla p &= 0 && \text{in } \Omega, \\
 \nabla \cdot \mathbf{u} &= 0 && \text{in } \Omega, \\
 \mathbf{u} &= \mu \left\{ \frac{1}{2.25} (x_1 - 2)(5 - x_1), 0 \right\} && \text{on } \Gamma_{\text{in}}, \\
 \mathbf{u} &= 0 && \text{on } \Gamma_{\text{wall}}, \\
 \nu \frac{\partial \mathbf{u}}{\partial \mathbf{n}} - p \mathbf{n} &= 0 && \text{on } \Gamma_{\text{out}},
 \end{aligned} \tag{8}$$

where $\mathbf{u} \equiv \mathbf{u}(\mathbf{x}, \mu)$ and $p \equiv p(\mathbf{x}, \mu)$ are the spatial and parameter dependent velocity and normalized pressure fields, respectively, with $\mathbf{x} = (x_0, x_1) \in \Omega$ and $\mu \in [1, 80] \subset \mathbb{R}$. The letter \mathbf{n} denotes the boundary normals, whereas the symbol ν indicates the kinematic viscosity, here imposed to 1.0. Thus, we are imposing the so-called no-slip condition on the walls of the domain, a parametric horizontal velocity profile at the inflow and the directional do-nothing condition in the outflow.

² By decreasing the number of informative pixels, the output size of the continuous convolutional filter decreases as well. Hence, the input of following linear layer is different for different percentage of images. Nevertheless, the overall structure of the network remains the same.

Fig. 7 The domain for the Navier Stokes experiment



To numerically solve such (nonlinear) problem, we have discretized the domain by means of a non-overlapping triangulation, onto which we compute the weak solution by exploiting a finite element framework. Regarding the test functions, we have chosen a second order polynomial for the velocity space and a first order one for the pressure space, i.e. (P2, P1). The (parameter dependent) solutions belong then to an irregular discrete space, forbidding to apply standard CNNs to this kind of data. Thanks to the continuous filter proposed in the present work, we are able to perform convolution also in this setting, avoiding also post-processing steps — e.g. projecting or interpolating the solutions on a regular grid.

An autoencoder (AE) is a type of deep learning architecture, which uses a MLP to encode the input into a lower dimensional meaningful representation, and then decode it back to obtain a reconstructed input as similar as possible to the original one. Convolutional autoencoder (CAE) is a type of AE where convolutional layers are used. In particular, CAE have been widely applied in fluid dynamics [43,44] as reduced order methods. However, as previously mentioned, pre-processing and post-processing steps are needed since discrete filters do not work with unstructured domains. Therefore, a continuous convolutional autoencoder (CCAЕ) is presented and used to extrapolate the latent representation of the Navier Stokes solution, without the need to pre-process or post-process it, thus keeping the unstructured domain representation. In particular the proposed autoencoder is composed by an encoder, represented by a single channel continuous convolutional layer followed by a single fully-connected layer of size 90 (latent representation size) and GELU [45] activation function after it; and a decoder, represented by a fully-connected layer of size 840 followed by a single channel transposed continuous convolutional layer. The (transposed) continuous convolutional layers have filter size and stride both equal to [0.75, 0.18], while the inner network is approximated using a MLP of size 40, 40 and GELU activation function. The CCAЕ is trained, minimizing the l_1 loss for 150 epochs using the Adam optimizer [46] with learning rate set to 0.001. The training data and testing data are divided by a 20 – 80% rule on the total data set which is composed of 500 solution samples at different inlet velocities

of the fluid. Notice that the choice of a small percentage of data in the training set is made to guarantee that the network can generalize easily even with a low number of samples used during training.

Overall, the CCAЕ has an l_2 percentage error of 3.9% both on training and testing. In Fig. 8 samples of reconstructed solutions using the CCAЕ are reported for different inlet velocities. The reconstructed samples show that the network has correctly generalized to unseen solutions, thus it can correctly extrapolate the latent dimension of them.

A further meaningful test is to compare the CCAЕ with a simple AE that only uses a MLP for encoding and decoding. In fact, a MLP preserves the structure in the data as CCAЕ, avoiding pre- and post-processing. Nevertheless, we show that the MLP AE is characterized by a higher error rate than the CCAЕ. The test is done using the same CCAЕ structure defined above, while the MLP AE network is obtained by removing the continuous convolutional layer and the transposed continuous convolutional layer, thus it is represented by a MLP of input and output size of 1639 (dimension of the Navier–Stokes solution vector), and hidden size of 90 with GELU activation function after the first layer. Both networks are trained minimizing the l_1 loss, using the Adam optimizer with learning rate set to 0.001. In order to see the ability to correctly learn a meaningful latent representation of the solution vector, both networks are trained multiple times. For each training, a different percentage of training data with respect to the overall data set size is used, while testing is done on the remaining solution samples. Hence, this test allows to compare the ability of both networks to generalize to unseen solutions as the number of data in training increases. The test is repeated 5 times with different parameters initialization for both networks, and results are averaged to reduce noise. The results of the experiment (Fig. 9) highlight that both networks tend to reduce the overall test error as the number of training data increases, as expected. However, the CCAЕ tends to outperform the MLP AE constantly for different percentages of training data, which is shown by a lower l_2 error curve.

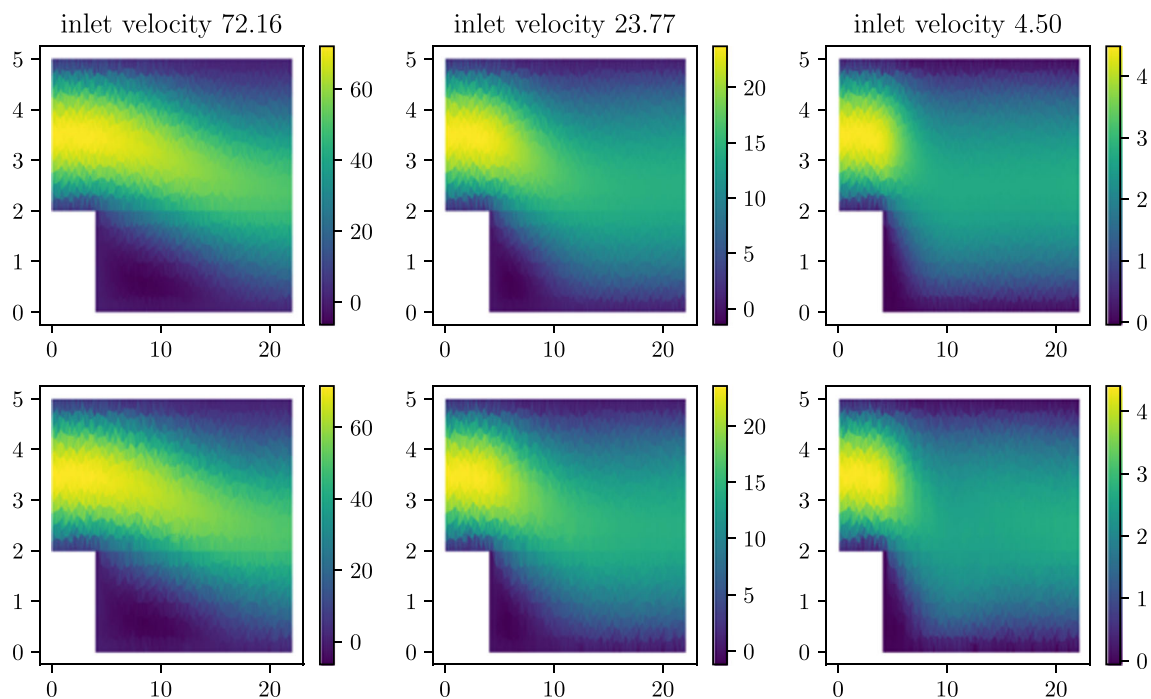


Fig. 8 Samples of CCAE reconstructed input (below) compared to the originals (above)

3.4 Multiphase state problem

The last problem presented shows how a continuous filter can be used in a more advanced deep learning framework. Specifically, we present a deep learning architecture for performing inference on unseen snapshots in a computational fluid dynamics simulation using the multiphase problem. This problem is represented by two fluids occupying a certain fraction of the volume domain separated by a sharp interface. The two fluids evolve in time resulting in a wave moving in a specific direction. In order to perform the mesh we followed [47], which gives also a mathematical formulation of the problem using the unsteady Navier Stokes equations. The mesh contains 18750 grid points for a single snapshot, while 200 different time shots of the dynamics are saved.

As depicted in Fig. 10, the overall deep learning architecture used (MP-Net) can be summarised into two building networks: an autoencoder (AE_{net}) to find the latent representation z ; and a MLP ($TimeNet$) to map the time t instance to the corresponding latent representation of the wave at time t . The AE_{net} is composed by an encoder f_{θ} and a decoder g_{ϕ} . The encoder is a single continuous convolutional block with one input channel and output channel, stride and filter dimension both equal to [0.4, 0.2], followed by a linear layer of dimension (150, 30) with ELU [48] activation function.

The decoder, represented in Fig. 11, is composed by a linear layer of dimension (30, 150), followed by ELU activation function. The output of the ELU layer is passed to two independent layers: firstly, a linear layer with dimension

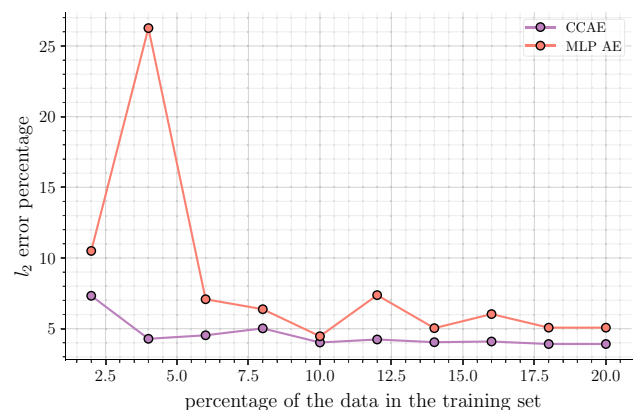


Fig. 9 Variation in test error for different percentage of data in the training set

(150, 18750); secondly, a transpose continuous filter layer of one input and output channel, stride and kernel size both equal to 0.4, followed by a sigmoid activation function. The results of these two independent layers are multiplied, followed by an adaptive sigmoid activation function [49], which gives the output of the autoencoder. Both continuous filters, direct and transposed, are internally approximated by a MLP of dimension 10, 40, 80 with an adaptive sigmoid activation function. The choice of multiplying the outputs of the two independent networks is due to the fact that a simple MLP decoder tends to over-fit, resulting in multiple waves in the output; while the continuous transpose filter tends to create a noisy output, resulting in a poor boundary approximation.

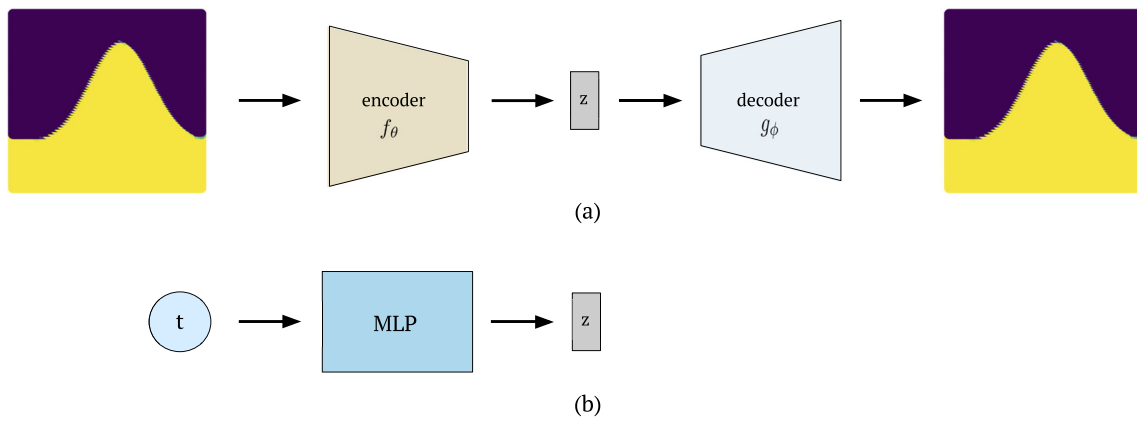


Fig. 10 Multiphase network with AE_{net} a and $TimeNet$ b

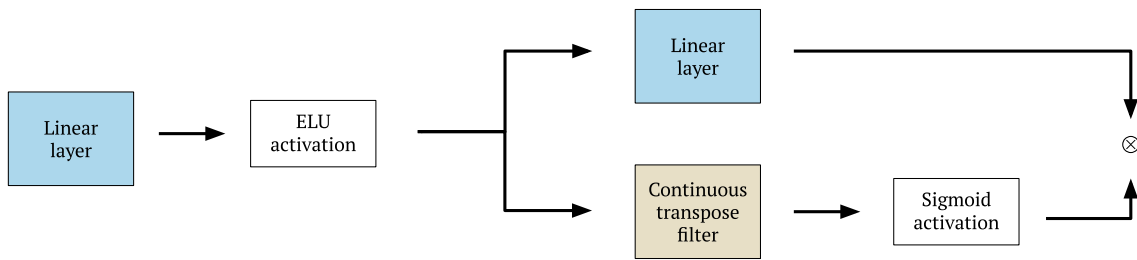


Fig. 11 Decoder g_ϕ structure for the MP-Net

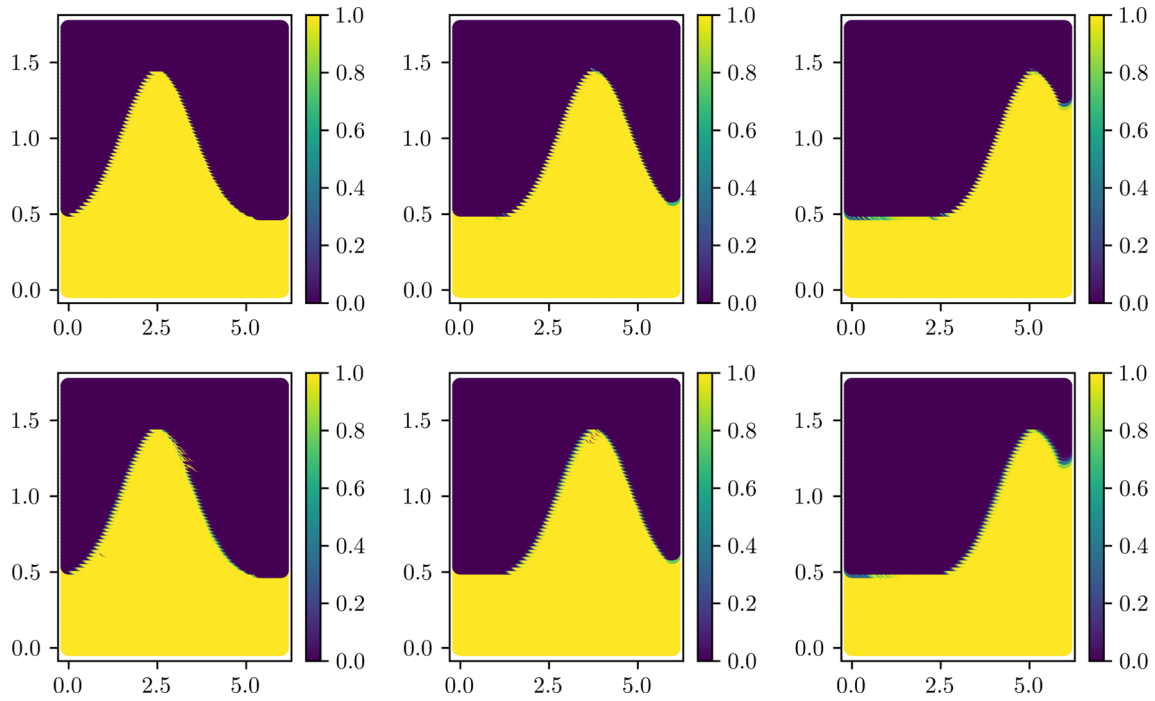


Fig. 12 Samples of predicted snapshots (below) comparing to the originals (above) at different time steps

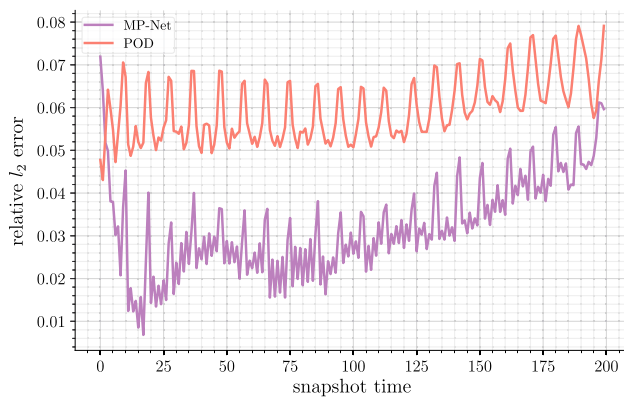


Fig. 13 Variation of l_2 relative error for different snapshot for POD and MP-Net

By passing the continuous transpose filter representation to a sigmoid activation, we consider the output of the filter as a probability weight. Thus, each coordinate of the linear layer output is multiplied with this weight, with the intention of zeroing the noisy extra waves wrongly produced by a linear layer decoder, and keeping the wave of interest. Finally, the *TimeNet* is a simple MLP of size 40, 80 with a ReLU activation function.

Training is done on 90 snapshots equally spaced in time, while testing is done on the remaining snapshots. The training process is divided in two parts as well: first training AE_{net} for obtaining a latent representation z of the input I through encoding $z = f_{\theta}(I)$ and decoding $g_{\phi}(z)$; second, once AE_{net} is completed, training *TimeNet* to map a time step t to its corresponding latent representation z by using a MLP, minimizing the mean square error between z and the *TimeNet* output. AE_{net} is trained using Adam optimizer [46] with initial learning rate 0.001 and using a exponential learning rate scheduler with multiplicative factor of learning rate decay equal to 0.99. The training is done on 300 epochs with data distributed in batches of 5, minimizing the mean square error loss. *TimeNet* is trained using Adam optimizer with learning rate 0.001, minimizing the mean square loss for 10000 epochs with a batch size of 5.

During inference, given only a time t the snapshot is constructed. In particular, for each t a latent representation z is obtained using *TimeNet*, thus the former is passed to g_{ϕ} decoder for getting the snapshot prediction. Notice that at inference time the encoder is not used since the snapshot samples are not known during inference. Samples of predicted snapshots at different times are reported in Fig. 12. Overall, on the 110 unseen snapshots the network has an l_2 relative error of 2.7 on training, and 3.6 on testing. The obtained test error is notable considering the fact that the network is equation-agnostic, and only 90 samples are used for training. As a comparison baseline, the Proper Orthogonal Decomposition (POD), a possible technique for model order reduction,

has been employed. By setting the same latent dimension size and snapshot used for MP-Net, and by using a radial basis function interpolator, the POD leads to a relative l_2 test error of 6.1 percentage.

In Fig. 13, we report the variation of l_2 relative error for all the snapshots (training + testing) for POD and MP-Net. The graph highlights that MP-Net outperforms POD for almost all snapshots, showing that MP-Net is able to find a better latent representation than a classical POD.

4 Conclusion

We propose a continuous trainable filter to perform convolution on unstructured data. We demonstrate the validity of the methodology on different machine learning frameworks, such as supervised and self-supervised learning, and on different type of data e.g., images (structured) and meshes (unstructured). We show that a continuous filter is able to reproduce the same results as the state-of-the-art discrete filter on structured data, and it can be used efficiently in deep learning architectures with unstructured data to reproduce unseen snapshots in a fluid dynamics simulation. In addition, we show that by using continuous convolutional autoencoders higher performances can be achieved with respect to the more simpler feed-forward autoencoder, when unstructured data are used. Possible further investigation include finding the most performing inner network architecture in the continuous convolutional filter, or using the latter for transfer learning representations: indeed, by being independent on the input image size, incompatible datasets can easily be merged. Furthermore, possible researches using the continuous filter for image inpainting algorithms could be done, considering the shown ability of the filter to work with partially completed images. Moreover, extending the framework to a 3-dimensional convolution might open the possibility to apply the continuous convolution to more complex manifolds. Finally, the use of the presented continuous filter in physics informed neural network architectures might be investigated, especially for problems where data are not structured. The architecture could be also applied to projection-based [50,51] and non-intrusive reduced order models [52] extending, for example, the work done in [43] to unstructured meshes. To summarize, the presented methodology, combined with the proposed research directions, opens the possibility to use the continuous convolutional filter in very complex settings, spacing in several application fields of science and industry.

Acknowledgements This work is partially supported by European Union Funding for Research and Innovation—Horizon 2020 Program—in the framework of European Research Council Executive Agency: H2020 ERC CoG 2015 AROMA-CFD project 681447 “Advanced Reduced Order Methods with Applications in Computational Fluid

Dynamics” P.I. Professor Gianluigi Rozza, by European Union Funding for Research and Innovation—Horizon Europe Program—in the framework of European Research Council Executive Agency: ERC POC 2022 ARGOS project 101069319 “Advanced Reduced order modelling: Online computational web server for complex parametric Systems” P.I. Professor Gianluigi Rozza, by European High-Performance Computing Joint Undertaking project Eflows4HPC GA N. 955558, by PRIN “Numerical Analysis for Full and Reduced Order Methods for Partial Differential Equations” (NA-FROM-PDEs) project.

Funding Open access funding provided by Scuola Internazionale Superiore di Studi Avanzati - SISSA within the CRUI-CARE Agreement.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- LeCun Y (1989) Generalization and network design strategies. *Connect Perspect* 19(143–155):18
- Li Z, Liu F, Yang W, Peng S, Zhou J (2021) A survey of convolutional neural networks: Analysis, applications, and prospects. *IEEE Transactions on Neural Networks and Learning Systems* 1–21. <https://doi.org/10.1109/TNNLS.2021.3084827>
- Goodfellow I, Bengio Y, Courville A (2016) *Deep Learning*. MIT Press, Cambridge, MA. <http://www.deeplearningbook.org>
- Alzubaidi L, Zhang J, Humaidi AJ, Al-Dujaili A, Duan Y, Al-Shamma O, Santamaría J, Fadhel MA, Al-Amidie M, Farhan L (2021) Review of deep learning: concepts, cnn architectures, challenges, applications, future directions. *J Big Data* 8(1):1–74. <https://doi.org/10.1186/s40537-021-00444-8>
- Calin O (2020) *Deep learning architectures: a mathematical approach*. Springer series in the data sciences. Switzerland, Springer
- Zhang A, Lipton ZC, Li M, Smola AJ (2011) Dive into deep learning. arXiv preprint [arXiv:2106.11342](https://arxiv.org/abs/2106.11342)
- Krizhevsky A, Sutskever I, Hinton GE (2012) Imagenet classification with deep convolutional neural networks. *Adv Neural Inf Process Syst* 25:1097–1105
- Shanmugamani R, Moore SM (2018) *Deep learning for computer vision: expert techniques to train advanced neural networks using TensorFlow and Keras*. Packt Publishing Ltd, Birmingham, UK
- Jiang X, Hadid A, Pang Y, Granger E, Feng X (2019) *Deep learning in object detection and recognition*. Springer, Singapore
- Young T, Hazarika D, Poria S, Cambria E (2018) Recent trends in deep learning based natural language processing. *IEEE Comput Intell Mag* 13(3):55–75
- Deng L, Liu Y (2018) *Deep learning in natural language processing*. Springer, Singapore
- Heil C (2019) *Convolution and the Fourier Transform*. Springer, Cham, pp 327–386
- Kipf T.N, Welling M (2017) Semi-supervised classification with graph convolutional networks. In: International conference on learning representations. <https://openreview.net/forum?id=SJU4ayYgl>
- Wu Z, Pan S, Chen F, Long G, Zhang C, Yu PS (2021) A comprehensive survey on graph neural networks. *IEEE Trans Neural Netw Learn Syst* 32(1):4–24. <https://doi.org/10.1109/TNNLS.2020.2978386>
- Schütt K.T, Kindermans P.-J, Sauceda H.E, Chmiela S, Tkatchenko A, Müller K.-R (2017) Schnet: A continuous-filter convolutional neural network for modeling quantum interactions. In: Proceedings of the 31st International conference on neural information processing systems. NIPS’17, pp. 992–1002. Curran Associates Inc., Red Hook, NY, USA. <https://doi.org/10.5555/3294771.3294866>
- Shocher A, Feinstein B, Haim N, Irani M (2020) From discrete to continuous convolution layers. arXiv preprint [arXiv:2006.11120](https://arxiv.org/abs/2006.11120)
- Boscaini D, Masci J, Rodolà E, Bronstein M (2016) Learning shape correspondence with anisotropic convolutional neural networks. In: Lee D, Sugiyama M, Luxburg U, Guyon I, Garnett R (eds) *Advances in neural information processing systems*, vol 29. Curran Associates Inc, Barcelona, Spain
- Wang S, Suo S, Ma W.-C, Pokrovsky A, Urtasun R (2018) Deep parametric continuous convolutional neural networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 2589–2597. <https://doi.org/10.1109/CVPR.2018.00274>
- Romero D.W, Knigge D.M, Gu A, Bekkers E.J, Gavves E, Tomczak J.M, Hoogendoorn M (2022) Towards a General Purpose CNN for Long Range Dependencies in ND. arXiv. <https://doi.org/10.48550/ARXIV.2206.03398>. <https://arxiv.org/abs/2206.03398>
- Fine TL (2006) *Feedforward neural network methodology*. Springer, New York
- Hornik K, Stinchcombe M, White H (1989) Multilayer feedforward networks are universal approximators. *Neural Netw* 2(5):359–366
- Cybenko G (1989) Approximation by superpositions of a sigmoidal function. *Math Control Signals Syst* 2(4):303–314. <https://doi.org/10.1007/BF02551274>
- Leshno M, Lin VY, Pinkus A, Schocken S (1993) Multilayer feedforward networks with a nonpolynomial activation function can approximate any function. *Neural Netw* 6(6):861–867. [https://doi.org/10.1016/S0893-6080\(05\)80131-5](https://doi.org/10.1016/S0893-6080(05)80131-5)
- Deisenroth MP, Faisal AA, Ong CS (2020) *Mathematics for Machine Learning*. Cambridge University Press, Cambridge
- Kriesel D (2007) *A Brief Introduction to Neural Networks*, Zeta2 edn
- Rojas R (1996) *The Backpropagation Algorithm*, pp. 149–182. Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-642-61068-4_7
- Sun S, Cao Z, Zhu H, Zhao J (2019) A survey of optimization methods from a machine learning perspective. arXiv. <https://doi.org/10.48550/ARXIV.1906.06821>. <https://arxiv.org/abs/1906.06821>
- Zaki MJ, Meira W Jr (2020) *Data mining and machine learning: fundamental concepts and algorithms*. Cambridge University Press, Cambridge, UK
- Rawat W, Wang Z (2017) Deep convolutional neural networks for image classification: a comprehensive review. *Neural Comput* 29(9):2352–2449
- Krizhevsky A, Sutskever I, Hinton GE (2012) Imagenet classification with deep convolutional neural networks. In: Pereira F, Burges CJ, Bottou L, Weinberger KQ (eds) *Advances in neural information processing systems*, vol 25. Curran Associates Inc, Red Hook, NY, USA
- He K, Zhang X, Ren S, Sun J (2015) Deep residual learning for image recognition. arXiv. <https://doi.org/10.48550/ARXIV.1512.03385>. <https://arxiv.org/abs/1512.03385>
- Szegedy C, Liu W, Jia Y, Sermanet P, Reed S, Anguelov D, Erhan D, Vanhoucke V, Rabinovich A (2014) Going deeper with convolu-

- tions. arXiv . <https://doi.org/10.48550/ARXIV.1409.4842>. <https://arxiv.org/abs/1409.4842>
33. Simonyan K, Zisserman A (2014) Very Deep Convolutional Networks for Large-Scale Image Recognition. arXiv. <https://doi.org/10.48550/ARXIV.1409.1556>. <https://arxiv.org/abs/1409.1556>
 34. Dumoulin V, Visin F (2016) A guide to convolution arithmetic for deep learning. arXiv. <https://doi.org/10.48550/ARXIV.1603.07285>. <https://arxiv.org/abs/1603.07285>
 35. Stroud A.H (1976) Methods of numerical integration Philip J. Davis and Philip Rabinowitz). SIAM Rev. (3), 528–529
 36. Paszke A, Gross S, Massa F, Lerer A, Bradbury J, Chanan G, Killeen T, Lin Z, Gimelshein N, Antiga L, Desmaison A, Kopf A, Yang E, DeVito Z, Raison M, Tejani A, Chilamkurthy S, Steiner B, Fang L, Bai J, Chintala S (2019) Pytorch: an imperative style, high-performance deep learning library. Advances in neural information processing systems 32
 37. OpenFOAM. (Accessed on 07/09/2022). <https://www.openfoam.com/>
 38. Logg A, Mardal K-A, Wells GN et al (2012) Automated solution of differential equations by the finite element method. Springer, Berlin. <https://doi.org/10.1007/978-3-642-23099-8>
 39. Demo N, Tezzele M, Rozza G (2018) EZyRB: Easy Reduced Basis method. 0.21105/joss.00661. <https://github.com/mathLab/EZyRB>
 40. Cohen G, Afshar S, Tapson J, van Schaik A (2017) EMNIST: an extension of MNIST to handwritten letters. CoRR [arXiv:1702.05373](https://arxiv.org/abs/1702.05373)
 41. Ruder S (2016) An overview of gradient descent optimization algorithms. arXiv. <https://doi.org/10.48550/ARXIV.1609.04747>. <https://arxiv.org/abs/1609.04747>
 42. Bank D, Koenigstein N, Giryes R (2020) Autoencoders. arXiv. <https://doi.org/10.48550/ARXIV.2003.05991>. <https://arxiv.org/abs/2003.05991>
 43. Romor F, Stabile G, Rozza G (2022) Non-linear manifold ROM with convolutional autoencoders and reduced over-collocation method. arXiv. <https://doi.org/10.48550/ARXIV.2203.00360>. <https://arxiv.org/abs/2203.00360>
 44. Fu R, Xiao D, Navon I.M, Wang C (2021) A data driven reduced order model of fluid flow by Auto-Encoder and self-attention deep learning methods. arXiv . <https://doi.org/10.48550/ARXIV.2109.02126>. <https://arxiv.org/abs/2109.02126>
 45. Hendrycks D, Gimpel K (2016) Gaussian Error Linear Units (GELUs). arXiv . <https://doi.org/10.48550/ARXIV.1606.08415>. <https://arxiv.org/abs/1606.08415>
 46. Kingma D.P, Ba J (2014) Adam: A method for stochastic optimization. arXiv. <https://doi.org/10.48550/ARXIV.1412.6980>. <https://arxiv.org/abs/1412.6980>
 47. Papapicco D, Demo N, Girfoglio M, Stabile G, Rozza G (2022) The neural network shifted-proper orthogonal decomposition: a machine learning approach for non-linear reduction of hyperbolic equations. Comput Methods Appl Mech Eng 392:114687. <https://doi.org/10.1016/j.cma.2022.114687>
 48. Clevert D.-A, Unterthiner T, Hochreiter S (2015) Fast and accurate deep network learning by exponential linear units (ELUs). arXiv. <https://doi.org/10.48550/ARXIV.1511.07289>. <https://arxiv.org/abs/1511.07289>
 49. Jagtap AD, Kawaguchi K, Karniadakis GE (2020) Adaptive activation functions accelerate convergence in deep and physics-informed neural networks. J Comput Phys 404:109136. <https://doi.org/10.1016/j.jcp.2019.109136>
 50. Stabile G, Zancanaro M, Rozza G (2020) Efficient geometrical parametrization for finite-volume based reduced order methods. Int J Numer Methods Eng 121(12):2655–2682
 51. Georgaka S, Stabile G, Rozza G, Bluck MJ (2019) Parametric POD-Galerkin Model Order Reduction for Unsteady-State Heat Transfer Problems. Commun Comput Phys 27(1):1–32
 52. Tezzele M, Demo N, Stabile G, Mola A, Rozza G (2020) Enhancing CFD predictions in shape design problems by model and parameter space reduction. Advanced Modeling and Simulation in Engineering Sciences 7(1) [arXiv:2001.05237](https://arxiv.org/abs/2001.05237). <https://doi.org/10.1186/s40323-020-00177-y>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.