



HAL
open science

Graph neural networks for laminar flow prediction around random two-dimensional shapes

J Chen, E Hachem, J Viquerat

► **To cite this version:**

J Chen, E Hachem, J Viquerat. Graph neural networks for laminar flow prediction around random two-dimensional shapes. 2021. hal-03432662

HAL Id: hal-03432662

<https://hal.science/hal-03432662v1>

Preprint submitted on 17 Nov 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

GRAPH NEURAL NETWORKS FOR LAMINAR FLOW PREDICTION AROUND RANDOM TWO-DIMENSIONAL SHAPES

A PREPRINT

J. Chen (陈峻峰)

MINES ParisTech, CEMEF
PSL - Research University
06904 Sophia Antipolis, France
junfeng.chen@mines-paristech.fr

E. Hachem

MINES Paristech, CEMEF
PSL - Research University
06904 Sophia Antipolis, France
elie.hachem@mines-paristech.fr

J. Viquerat*

MINES Paristech, CEMEF
PSL - Research University
06904 Sophia Antipolis, France
jonathan.viquerat@mines-paristech.fr

October 1, 2021

Abstract

In the recent years, the domain of fast flow field prediction has been vastly dominated by pixel-based convolutional neural networks. Yet, the recent advent of graph convolutional neural networks (GCNNs) have attracted a considerable attention in the computational fluid dynamics (CFD) community. In this contribution, we proposed a GCNN structure as a surrogate model for laminar flow prediction around two-dimensional (2D) obstacles. Unlike traditional convolution on image pixels, the graph convolution can be directly applied on body-fitted triangular meshes, hence yielding an easy coupling with CFD solvers. The proposed GCNN model is trained over a data set composed of CFD-computed laminar flows around 2,000 random 2D shapes. Accuracy levels are assessed on reconstructed velocity and pressure fields around out-of-training obstacles, and are compared with that of standard U-net architectures, especially in the boundary layer area.

Keywords Graph convolution · Triangular meshes · Computational fluid dynamics · Surrogate model

1 Introduction

During the last few years, the computational fluid dynamics (CFD) community has largely benefited from the fast-paced development of the machine learning (ML) field, and more specifically from that of the neural networks (NN) domain. In many cases, a data-driven surrogate model is constructed by training a neural network on a specific prediction task, in order to counter the high computational cost of CFD solvers. In this context, convolutional neural networks (CNNs) have been widely exploited. Multiple studies used CNNs for steady or unsteady flow prediction, with the flow fields represented

*Corresponding author

by image-like arrays [1, 2, 3, 4, 5, 6, 7]. CNNs are also used to directly perform lift or drag prediction when the geometric characteristics, Mach number and Reynolds number are at hand [8, 9]. Under the framework of reduced order modeling, CNNs have been applied to low-dimensional feature extraction of fluid systems [10, 11, 12]. The non-linearity of CNNs brings advantages over traditional linear methods like proper orthogonal decomposition. Recent research also involves the coupling of CNNs with PDE solvers. One interesting axe is to embed CNN corrections into a DNS solver on coarse grids to obtain more competitive results with remarkable computational speedup [13, 14]. CNNs-based surrogate models have also been applied to inverse problems to accelerate the optimization [15].

The use of convolution in neural networks was originally proposed in the context of image processing [16]. In practice, CNNs can be used to extract local features from grid-based data (include images) by sliding trainable kernels over an input image and performing convolution at each pixel. Yet, the use of uniform cartesian grids represent a hindrance for CFD applications, due to their poor intrinsic geometrical representation and large associated computational cost. For external and wall-bounded flows, body-fitted or unstructured meshes are usually adopted, in order to respect the geometric characteristics and to accurately capture the physics of the boundary layer. When coupled to CNN-based neural network models, data from the meshes have to be interpolated on *ad-hoc* cartesian grids, before being projected back on the mesh. Gao et al.[17] use elliptic coordinate transformation to transform the irregular physical domain to a regular reference domain. The coordinate transformation allows the authors to use a physics-informed CNN to solve partial differential equations in the reference domain. The physical solution in the irregular domain can then be obtained through an inverse transformation. Compared to fully-connected physics-informed neural networks, the proposed method converges faster and reaches a better accuracy.

A possible alternative consists in designing efficient convolution operations on graph structures, and to apply them directly on unstructured meshes. Extending convolutional neural networks to graphs is of great interest, as many CFD solvers are based on finite element/volume methods, which use unstructured meshes for discretization. In this regard, the design and use of trained graph neural network models can significantly enrich CFD workflows by being used for specific tasks during the resolution process, while still relying on the standard CFD mesh. Generally speaking, a graph $\mathbf{G}(V, E)$ is composed of a set of nodes V and a set of edges E , representing the connectivity of nodes. A node-level/edge-level feature vector represents the state of each node/edge. The goal of a graph convolution operation is to update the feature vectors through aggregating neighboring information. Two types of graph convolution are to be distinguished, namely spectral convolutions, and spatial convolutions. By means of graph Fourier transform, spectral convolution turns convolution into product between the entire graph’s feature matrix and the convolutional kernel [18]. In [19], the authors combine differentiable PDE solvers and graph convolutional networks (GCN) to perform flow prediction around 2D airfoils. An initial flow field is first obtained by CFD simulation on a coarse mesh, and is then passed to a GCN for a prediction on a refined mesh. This hybrid method provides better results than coarse CFD simulation alone, and is significantly less computationally-intensive than running a full CFD simulation on a refined mesh. However, spectral convolution techniques are based on a global Fourier transform of the entire graph, and therefore require a fixed graph structure. Hence, spectral convolution has thus limited application in real-world CFD sceneries, and are *de facto* excluded from contexts including mesh adaptation.

Spatial convolution is more similar to traditional convolution on cartesian data: a convolution kernel operates locally on the node and its neighbours, the kernel parameters being re-used across the entire mesh. A popular implementation of this concept is the message-passing neural network (MPNN), proposed by Gilmer *et al.* for quantum chemistry [20]. This method divides the convolution operation into two steps: first, neighbouring nodes and edge features are aggregated into a hidden node state, after what the hidden node state is used to update the node features. Both the aggregation and update steps exploit nonlinear differentiable functions that are represented by small, shallow neural networks, which can be regarded as the convolution kernels. In [21], a variant of MPNN called GraphSAGE [22] is applied to drag force prediction from given laminar velocity field around airfoils. On a dataset with 1550 different airfoils and 21 angles of attack, the graph convolution model has shown better performance than other machine learning methods. Battaglia *et al.* [23] propose an alternative method, called graph network (GN) blocks, for general multiple-step graph convolution. A GN block first passes messages from nodes to edges through an edge convolution kernel, before updating the edge features. The new

edge features are then aggregated to the nodes, through sum or other permutation invariant operations, as the edge messages. A node convolution kernel then takes the old node features and edges messages to obtain the new node features. In [24], the authors use graph networks-based simulators to learn physical systems of particles, including fluids, rigid solids and deformable materials. The proposed method shows high performance in reconstructing the results of mesh-free solvers. Pfaff *et al.* [25] apply graph networks to mesh-based simulations, including incompressible flow around cylinders and compressible flow around airfoils. The proposed neural network is trained to be an accurate incremental simulator, with capability in adapting the mesh. This method also generalizes well to mesh shapes and mesh sizes which are not present in the training set. A slightly different spatial convolution considers parameterized weight functions to aggregate the neighbouring information [26]. The features and relative position of each node pair are used to compute an edge weight through some trainable weight function, and a single weight function is applied to the entire graph inside one convolution. Since the weight function depends on node features, different aggregation patterns would be obtained at different locations of the graph. Xu *et al.* [27] applied this method to adjoint vector modelling for 2D flow around airfoils. The proposed weight function depends on the edge length and its angle to flow’s principle direction. By training a graph convolutional neural network with such weight functions, the authors get similar adjoint vectors to traditional adjoint methods. The predicted adjoint vectors are further used for shape optimization, with resulting shapes being very close to that obtained with direct adjoint methods.

In the present contribution, a graph convolutional neural network(hereafter denoted by GCNN) is implemented for laminar flow prediction around random 2D shapes on triangular meshes. In section 2 the generation of the dataset used in this paper is presented. In section 3, insights about the convolution block and graph neural network architecture are provided. The training details and a global evaluation of the model are provided in section 4.1. In section 4.2 and 4.3, the flow predictions around a cylinder and an airfoil are compared to baseline CFD results. In particular, the velocity profiles and surface pressure distribution are carefully studied. In order to further show the potential of our method, drag forces are computed from the boundary layer information, and compared to the CFD results in section 4.4. A comparison of GCNNs with standard convolutional U-nets in terms of accuracy and computational cost is proposed in section 4.5. Finally, conclusions and future possible developments are proposed in section 5

2 Dataset construction

This section provides insights on the random shape dataset generation used to train networks in the next sections. This dataset was initially used in [9] (section 3.5), thus only the main lines are sketched here. For more details, the reader is referred to [9]. First, we describe the steps to generate arbitrary shapes by means of connected Bézier curves. Then, the solving of the Navier-Stokes equations with an immersed method is presented. Finally, the velocity and pressure fields are projected onto triangular meshes of a smaller domain for the training of neural networks.

2.1 Random shape generation

In the first step, n_s random points are drawn in $[0, 1]^2$, and translated so their center of mass is positioned in $(0, 0)$. An ascending trigonometric angle sort is then performed (see figure 1a), and the angles between consecutive random points are then computed. An average angle is then computed around each point (see figure 1b) using:

$$\theta_i^* = \alpha\theta_{i-1,i} + (1 - \alpha)\theta_{i,i+1},$$

with $\alpha \in [0, 1]$. The averaging parameter α allows to alter the sharpness of the curve locally, maximum smoothness being obtained for $\alpha = 0.5$. Then, each pair of points is joined using a cubic Bézier curve, defined by four points: the first and last points, p_i and p_{i+1} , are part of the curve, while the second and third ones, p_i^* and p_{i+1}^* , are control points that define the tangent of the curve at p_i and p_{i+1} . The tangents at p_i and p_{i+1} are respectively controlled by θ_i^* and θ_{i+1}^* (see figure 1c). A final sampling

of the successive Bézier curves leads to a boundary description of the shape (figure 1d). Using this method, a wide variety of shapes can be attained, as shown in figure 2.

2.2 Numerical resolution of Navier-Stokes equations

The flow motion of incompressible newtonian fluids is described by the Navier-Stokes (NS) equations:

$$\begin{cases} \rho (\partial_t \mathbf{v} + \mathbf{v} \cdot \nabla \mathbf{v}) - \nabla \cdot (2\eta \boldsymbol{\epsilon}(\mathbf{v}) - p \mathbf{I}) = \mathbf{f}, \\ \nabla \cdot \mathbf{v} = 0, \end{cases} \quad (1)$$

where $t \in [0, T]$ is the time, $\mathbf{v}(x, t)$ the velocity, $p(x, t)$ the pressure, ρ the fluid density, η the dynamic viscosity and \mathbf{I} the identity tensor. In order to efficiently construct the dataset, an immersed boundary method is used for resolution instead of the usual body-fitted method, avoiding a systematic re-meshing of the whole domain for each shape. This method rely on a unified fluid-solid eulerian formulation based on level-set description of the geometry [28], and leads to the following set of modified equations:

$$\begin{cases} \rho^* (\partial_t \mathbf{v} + \mathbf{v} \cdot \nabla \mathbf{v}) - \nabla \cdot (2\eta \boldsymbol{\epsilon}(\mathbf{v}) + \boldsymbol{\tau} - p \mathbf{I}) = \mathbf{f}, \\ \nabla \cdot \mathbf{v} = 0, \end{cases} \quad (2)$$

where we have introduced the following mixed quantities:

$$\begin{aligned} \boldsymbol{\tau} &= H(\alpha) \boldsymbol{\tau}_s, \\ \rho^* &= H(\alpha) \rho_s + (1 - H(\alpha)) \rho_f, \end{aligned}$$

where the subscripts f and s respectively refer to the fluid and the solid, and $H(\alpha)$ is the Heaviside function:


$$H(\alpha) = \begin{cases} 1 & \text{if } \alpha > 0, \\ 0 & \text{if } \alpha < 0. \end{cases} \quad (3)$$

The reader is referred to [29] for additional details about formulation (2). The shapes are placed at the origin of a rectangular domain $[-5, 10] \times [-5, 5]$ meters. The inlet velocity is 1m/s. With density $\rho = 1\text{kg/m}^3$ and kinematic viscosity $\nu = 0.1\text{m}^2/\text{s}$, the resulting Reynolds number is 10. At the top and bottom boundaries, $v_y = 0$ and $\frac{\partial v_x}{\partial y} = 0$ conditions are applied, while at the outlet we apply $\frac{\partial v_x}{\partial x} = \frac{\partial v_y}{\partial x} = 0$ and $p = 0$. Finally, no-slip condition $v_x = v_y = 0$ is imposed on the solid interface. The mesh size around the interface is set at 0.01m. Eventually, the modified equations (2) are cast into a stabilized finite element formulation, and solved using a variational multi-scale (VMS) solver [30, 31, 32, 33, 34].


2.3 Dataset for inference

The dataset² is composed of 2.000 shapes, along with their steady-state velocity and pressure fields at $Re = 10$. All the labels were computed following the methods exposed in section 2.2 using CimLib [29], an industrial-level CFD solver. As we are more concerned with the flow patterns near the obstacles, a smaller domain containing the obstacle is used for the neural networks's inference task. For each shape, a body-fitted triangular mesh is generated for the domain $[-2, 2]\text{m} \times [-2, 2]\text{m}$. The mesh size is 0.01m on the obstacle's surface, which is the same as the immersed meshes used for CFD simulation. The element size grows to 0.4m at the border of the domain. Depending on the complexity of the shapes, the number of nodes varies from 2,000 to 3,000 in such a mesh. The velocity and pressure fields from CFD simulations are projected onto these meshes through linear interpolation (see figure 3). For the learning phase, normalization over the entire data set is applied, mapping separately the two coordinates components, the two velocity components and the pressure into the $[0, 1]$ range. at the end of the prediction phase, the predicted fields are mapped back to their physical ranges, and


²This dataset was already exploited in previous studies, see [9, 4].




(a) Draw n_s random points, translate them around $(0,0)$ and sort them by ascending trigonometric angle



(b) Compute angles between random points, and compute an average angle around each point θ_i^*



(c) Compute control points coordinates from averaged angles and generate cubic Bézier curve



(d) Sample all Bézier lines and export for mesh immersion

Figure 1: **Random shape generation with cubic Bézier curves.** The generation contains four steps, namely (1a) draw random points, (1b) compute angles, (1c) compute control points, (1d) sample.




Figure 2: **Shape examples drawn from the dataset.** A wide variety of shape is obtained using a restrained number of points ($n_s \in [4, 6]$), as well as a local curvature r and averaging parameter α .

non-dimensionalized by the inlet velocity and fluid density, both being equal to one. For additional details about the distribution of the elements in the dataset, the reader is referred to [9] (section 3.5). The whole data set is split into a training set with 1600 shapes, a validation set and a test set both with 200 shapes.

3 Network architecture

3.1 Convolution block

The basic component of the graph convolutional neural network is the convolution block. Here, the proposed convolution block has two components: a two-step graph convolutional layer adapted from the literature [23, 24, 25], followed by a two-step smoothing layer. In this section, the feature matrices on the nodes and the edges are respectively denoted by $\mathbf{X}_V \in \mathbf{R}^{N_V \times d_V}$ and $\mathbf{X}_E \in \mathbf{R}^{N_E \times d_E}$, with N_V and N_E being the number of nodes and edges. d_V and d_E are the dimensions of feature vectors on nodes and edges, which are denoted by $\mathbf{x}_v \in \mathbf{R}^{d_V}$ and $\mathbf{x}_e \in \mathbf{R}^{d_E}$.

Convolutional layer The convolutional layer propagates node-level messages to edges, then aggregates the new edge features and updates the node features according to the following rules:

$$\begin{aligned}
 \mathbf{x}'_e &= f_e \left(\frac{\mathbf{x}_{v_1} + \mathbf{x}_{v_2}}{2}, \frac{|\mathbf{x}_{v_1} - \mathbf{x}_{v_2}|}{2}, \mathbf{x}_e \right), \\
 \mathbf{x}'_v &= f_v \left(\mathbf{x}_v, \sum_{e_i \in N(v)} \mathbf{x}'_{e_i} \right),
 \end{aligned} \tag{4}$$

where v_1 and v_2 are the two nodes connected by edge e , and $N(v)$ is the set of neighbouring edges around node v . For both convolution kernels f_e and f_v , a fully-connected feedforward neural network with one hidden layer is used (see figure 4), the number of neurons in the hidden layer being set to 128. The number of neurons in the output layer remains modifiable, so the output dimension of the kernel can be customized. All the neurons are activated using the swish function[35], which performs better

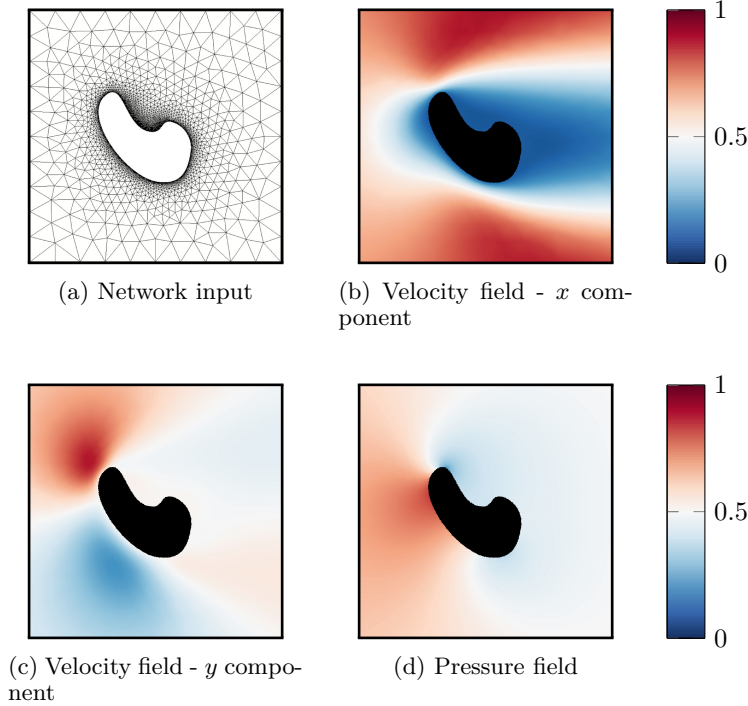


Figure 3: **Triangular mesh, velocity field and pressure field for a dataset element.** Body-fitted triangular meshes in a smaller domain (3a) are used for training the neural network, along with the interpolated velocity field (3b, 3c) and pressure field (3d).

than ReLU and hyperbolic tangent functions. The weights and bias of the kernels are initialized using the Glorot-normal distribution [36].

During the edge convolution step, symmetric node features $\frac{\mathbf{x}_{v_1} + \mathbf{x}_{v_2}}{2}$ and $\frac{|\mathbf{x}_{v_1} - \mathbf{x}_{v_2}|}{2}$ are preferred over the raw features \mathbf{x}_{v_1} and \mathbf{x}_{v_2} in order to preserve permutation invariance. The summation in the node convolution step is also invariant to permutation of neighbouring edges. As a consequence, the proposed convolution steps (4) are insensitive to nodes and edges indexing.

Smoothing layer The smoothing layer perform an average operation on the output graph. The averaging kernel implemented on triangular meshes is divided into two steps:

$$\begin{aligned} \mathbf{x}_e'' &= \frac{\mathbf{x}'_{v_1} + \mathbf{x}'_{v_2}}{2}, \\ \mathbf{x}_v'' &= \frac{1}{\text{card}(N(v))} \sum_{e_i \in N(v)} \mathbf{x}_{e_i}'' \end{aligned} \quad (5)$$

The smoothing layer is technically also a convolutional layer. The motivation of adding this layer is not for message propagation, but to reduce the spatial variability of the node features. It weighs down the feature map of a convolutional layer through neighbouring compensation. A sequence of convolution-smoothing blocks can hence regularize the output of the networks to reduce local minima. The smoothing layer is beneficial to further reducing the training loss and producing continuous flow fields. A similar idea is found in the regularized pooling operation used for convolutional layers on images [37]. The final output of a convolution block is the new edge feature matrix $\mathbf{X}'_E \in \mathbf{R}^{N_E \times d'_E}$ and the smoothed new node feature matrix $\mathbf{X}''_V \in \mathbf{R}^{N_V \times d'_V}$

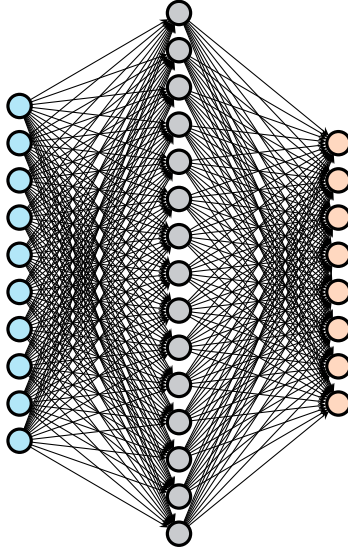


Figure 4: **Fully-connected feedforward neural network used as convolution kernel.** The vectors on the right hand side in equation (4) are concatenated as the inputs of the convolution kernel. In edge convolution, the input dimension is $2d_V + d_E$, while in node convolution, the input dimension is $d_V + d'_E$. The number of hidden nodes is always 128.

3.2 Network architecture

The proposed architecture is presented in figure 5. As can be seen, the input is composed of three images respectively representing (i) the obstacle boundary, (ii) the x coordinates and (iii) the y coordinates. Eight convolution blocks/smoothing layers are stacked to form the graph convolutional neural network, followed by a 1×1 convolution as the output layer. Unless stated otherwise, the intermediate edge and node features dimensions in each graph convolution layer are respectively (4, 8, 16, 32, 64, 64, 32, 16) and (8, 16, 32, 64, 64, 32, 16, 8), as indicated in parenthesis in the graph convolution layers of figure 5, but a study of the impact of the features dimensions can be found in section 4.5. The final 1×1 convolutional layer transforms node features in \mathbf{R}^8 into vectors in \mathbf{R}^3 , which represent the predicted velocity and pressure. This symmetric architecture is similar to the encoder-decoder structure used in [4].

An important component of the proposed architecture is the skip connections from the input graph to the convolutional blocks. The coordinates of the nodes are concatenated to the node features after each smoothing layer. These skip connections provide spatial informations to the edge convolution steps in equation (4). The edge convolution hence takes into account the position of the edge center and relative position of the node pairs, allowing the convolution kernel to display different patterns at different locations in the domain.

4 Results

4.1 Training history

With architecture described in section 3, the proposed network has 217, 853 trainable parameters. The parameters are calibrated by minimizing a loss function defined on the predicted and the reference flow fields. The mean absolute error (MAE) is adopted as the loss function, and is defined as:

$$\mathcal{L} = \frac{1}{N_V} \sum_{i=1}^{N_V} \left(\left| u_i^{\text{pred}} - u_i^{\text{ref}} \right| + \left| v_i^{\text{pred}} - v_i^{\text{ref}} \right| + \left| p_i^{\text{pred}} - p_i^{\text{ref}} \right| \right),$$




Figure 5: **Proposed network architecture.** The provided input is composed of three images respectively representing (i) the obstacle boundary, (ii) the x coordinates and (iii) the y coordinates. The network has an encoder-decoder structure: in the encoder, the amount of features is doubled after each smoothing step, while in the decoder, it is halved after each smoothing step. For each graph convolution layer, the number of edge and node features are indicated in parenthesis. The network is terminated with a 1×1 convolutional layer to obtain the three output layers, holding respectively the u and v components of the velocity, and the pressure field p . The skip connections allow to inject coordinate informations from the input layer at each edge convolution step of the architecture.

with the considered fields being normalized to $[0, 1]$, as stated earlier. An in-house TensorFlow implementation is used for the network architecture and training. The network is trained using the Adam optimizer [38] for at most 1000 epochs, with mini-batches of size 64^3 to limit the required computational resources. The learning rate follows a pre-defined decay schedule described as:

$$l_r^e = \frac{l_r^0}{1 + \gamma e}, \quad (6)$$

where l_r^e is the learning rate used during epoch number e , and γ is the decay rate. In the experiments, both the initial learning rate and the decay rate were set equal to 0.002. To prevent overfitting, an early stopping criterion is applied: the training is terminated if the average MAE on the validation set is not improved in 60 epochs.

After 914 training epochs⁴, the training and validation losses converge respectively to 7.43×10^{-3} and 7.61×10^{-3} , while the MAE over the test set amounts to 7.70×10^{-3} , as shown in figure 6. From the MAE distribution over the test set, shown in figure 7, it appears that most elements in the test set are accurately predicted, with only a negligible amount of shapes exceeding a MAE of 0.01. This underlines the capabilities of the model to generalize on out-of-training data. Under the training settings described in section 4.1, the proposed model requires approximately 11.96 seconds per epoch on a Tesla V100 GPU card.

4.2 Field prediction on a cylinder

To test the trained GCNN model on an out-of-training input, we consider the velocity and pressure predictions performed on a 2D cylinder with radius 0.75m placed at the origin, and embedded in a body-fitted triangular mesh with 2447 nodes, with a 0.01m mesh size on the cylinder. The MAE of the predicted velocity and pressure fields is 4.3×10^{-3} , as the geometric characteristic of the cylinder are close to that of the Bézier shapes from the training set. In figure 8, we observe almost symmetric

³One mini-batch is a large graph containing 64 disjoint subgraphs. For more details please refer to the code page https://github.com/jviquerat/gnn_laminar_flow.

⁴This error level can be attained by following the weight initialization method in the published codes.




Figure 6: **Training history of the model.** In fewer than 1000 training epochs, the training loss and validation converge. No obvious overfitting is observed.




Figure 7: **Distribution of the MAE over the test set.** The test set contains 200 random Bézier shapes. Only 13% shapes are with MAE over 0.01.

velocity fields with respect to the x -axis for all three fields. Although the shapes in the training set are not necessarily symmetric, the trained GCNN is aware of the symmetry of flow around a cylinder. The velocity profiles and surface pressure distribution displayed in figure 9 further confirm the symmetric predictions around a cylinder.

4.3 Prediction on a NACA0012 airfoil

The NACA airfoils from the UIUC database are among the most studied geometries in the CFD community. In the context of this contribution, the prediction on airfoils is regarded as good test of the model generalization capabilities on out-of-training shapes, as their geometry is very different from that of Bézier random shapes. The demonstrated airfoil is NACA0012, with an angle of attack equal to 4° and a chord length equals to $1m$. The airfoil is embedded into a triangular mesh made of 1489 nodes. The referential flow fields are from a CFD resolution with the same configurations as the Bézier shapes and the cylinder. The MAE of the predicted fields obtained from the GCNN model is 1.04×10^{-2} , which corresponds to the worst prediction levels observed in the test set. Yet, the predicted fields presented in figure 10 display reasonable physical features. Although the model overestimates the velocity magnitude, it successfully detects the locations where the vertical velocity changes sign or attains local minimum and maximum, as can be further observed in the plots of figure 11. More, the surface pressure distribution is recovered by the model, although with an underestimated pressure magnitude on the leading edge and the trailing edge. The incorrect velocity and pressure magnitude both result from the dissimilarity between flow around the airfoil and flow around the Bézier shapes. Since the GCNN model in this contribution is trained under the supervised learning framework, it should be conservatively applied to geometries that are not too different from that of the training dataset.

4.4 Drag force prediction

In this section, the accuracy of the drag force prediction is used as a measure of the quality of the predicted boundary layer. The drag force at $Re=10$ is obtained by integrating the edge-wise contribution of pressure and viscous force on the obstacle, as shown in equation (7).

$$F_D = \mathbf{e}_x \cdot \oint_S (p \cdot \mathbf{n} + \mu \frac{\partial v_s}{\partial \mathbf{s}}) ds, \quad (7)$$

with \mathbf{e}_x is the unit vector in the x -direction, n and s respectively the inward- and outward-pointing unit normal vector. $\frac{\partial v_s}{\partial \mathbf{s}}$ is the normal gradient of the tangent velocity. The drag forces computed on the 200 shapes of the test set are compared with their reference drag values in figure 12. Overall, a good agreement is observed, with an average relative error of 3.43%, and a coefficient of determination equal to 0.987, which is in line with the results of the GraphSage method [21]. It must be noted that, in the literature, neural networks are usually directly trained on adequate databases to predict drag or lift coefficients. With the drag force computed by reconstructed flow fields on triangular meshes, it is possible to apply the GCNN model to shape optimization in future works.

4.5 Comparison with U-nets

U-nets are a popular neural network architecture for end-to-end image prediction. Initially introduced by Ronneberger *et al.* [39] in the field of biomedical image segmentation, they were recently exploited for flow prediction tasks [5, 6]. U-nets have an auto-encoder-like structure, but introduce skip connections between the corresponding convolutional and de-convolutional layers from the encoder and decoder paths in order to mix high-level features from the latent space with low-level ones from the encoder. A sketch is proposed in figure 13. By design, U-nets require grid-based data, such as images. As a consequence, flow fields as well as obstacle boundaries need to be projected onto a regular cartesian grid, implying geometric inaccuracies and a higher amount of data. Here, we consider a cartesian grid with a resolution equal to 0.01, in order to match the local resolution of the boundary layer of the graph network. The domain size is $[-2, 2] \times [-2, 2]$, hence each data instance is a 401×401 2D array. The U-net architecture is composed of 5 convolution-convolution-maxpooling blocks with 3×3 convolutional filters and 2×2 pooling size for the encoder part, and 5 deconvolution-convolution-convolution blocks with 2×2 deconvolutional filters and 3×3 convolutional filters for the decoder part.







Figure 8: **Flow and shape predictions around a cylinder using a GCNN model.** The normalized velocity and pressure fields outputted by the network are rescaled to their physical ranges and non-dimensionalized for better demonstration. The reference (8a, 8d, 8g) and predicted fields (8b, 8e, 8h) are very close, with no specific pattern in the MAE maps (8c, 8f, 8i).




(a) u profile at $x = -1$




(b) u profile at $x = 1$




(c) v profile at $x = -1$



(d) v profile at $x = 1$



(e) Pressure distribution on the lower surface



(f) Pressure distribution on the upper surface

Figure 9: **Velocity and pressure profiles around the cylinder.** The velocities are recorded at $x = -1$ (9a, 9c) and $x = 1$ (9b, 9d), and display the expected symmetry with respect to the x axis. The pressure is recorded on the lower (9e) and upper surfaces (9f) of the cylinder. The wiggles in the referential pressure profiles are caused by obstacle boundary resolution by the immersed boundary method, which employs an element size equal to 0.01m. These numerical artefacts could be reduced using smaller boundary elements, although we find they do not impact the neural networks training.

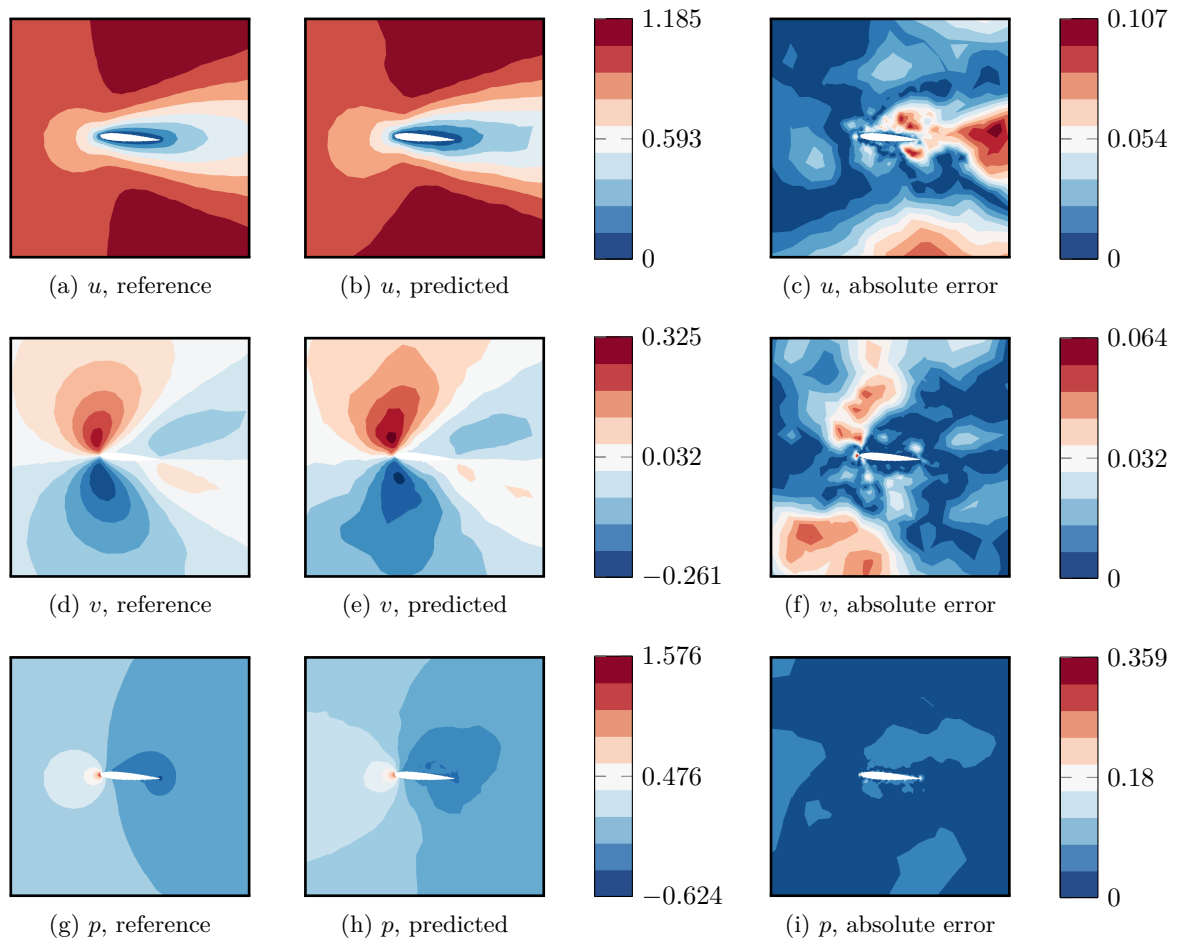





Figure 10: **Flow predictions on a NACA0012 airfoil using a GCNN model.** Although the field amplitudes are not correctly evaluated, the velocity (10a, 10b, 10c, 10d, 10e, 10f) and pressure maps (10g, 10h, 10i) display reasonably physical features, such as sign changes in the vertical velocity or localized high pressure area.




(a) u profile at $x = -1$




(b) u profile at $x = 1$




(c) v profile at $x = -1$



(d) v profile at $x = 1$



(e) Pressure distribution on the lower surface



(f) Pressure distribution on the upper surface

Figure 11: **Velocity and pressure profiles around the NACA0012 airfoil.** The velocities are recorded at $x = -1$ (11a, 11c) and $x = 1$ (11b, 11d), and display significant errors around the trailing edge. With significant discrepancy of velocity magnitude, the physical pattern of the velocity profiles is still well recovered. The pressure is recorded on the lower (11e) and upper surfaces (11f) of the airfoil, with remarkable error at the leading and trailing edges.

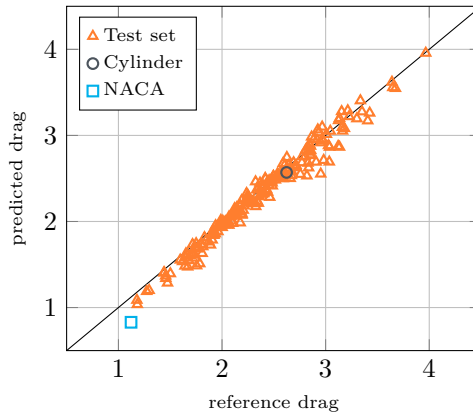


Figure 12: **Drag force computed from the boundary layer reconstruction obtained with the graph neural network.** The included data contains the 200 Bézier shapes from the test set, along with the cylinder and the NACA0012 airfoil considered in the previous sections. Excellent agreement is observed on most shapes from the test set, with an average relative error of 3.43%. While the cylinder drag is accurately computed, the airfoil drag is off by 26.1% from its reference value.




Figure 13: **Sketch of auto-encoder and U-net architectures.** Standard auto-encoders (13a) are composed of an encoder and a decoder paths, and can be exploited either for end-to-end regression tasks (in a supervised way, with labels), or for the inference of latent space representations (in an unsupervised way, without labels). U-net auto-encoders (13b) are a specific class of AE, in which skip connections are added from the encoder branch to the decoder one in order to mix high-level features from the latent space with low-level one from the contractive path. They usually present a superior level of performance on regression tasks.

In the encoder branch, the amount of filters is doubled after each block, while in the coder, the amount of filters is halved after each block. The U-net’s architecture is hence symmetric, and scalable through a modification of the initial number of filters in the first block. The relation between the number of initial filters and the number of trainable parameters is shown in table 1. Similarly, in a GCNN, one can modify the node feature dimension of the first convolutional block while proportionally changing the other intermediate feature dimensions.

To perform a performance comparison between U-net and GCNN architectures, the U-nets are implemented and trained in the Keras API with the following configurations: batch size equals to 64, optimized by the Adam optimizer with an initial learning rate equals to 0.001, a learning rate decay equals to 0.0002, early stopping patience is 20 epochs. A sample prediction obtained with a U-net holding 1 944 763 free parameters is shown in figure 15. The prediction on this sample is globally of high quality with an MAE equals to 7.09×10^{-3} , but contrarily to GCNN predictions, most of the discrepancies are concentrated on the border of the obstacle. Indeed, the U-net not only predicts the velocity and pressure fields, but also has to reconstruct the obstacle boundary, which cannot be strictly restored due to the staircasing effect of the cartesian grid. To perform a fair quantitative comparison

1 st layer	Parameters	1 st layer	Parameters
7	1 489 162	2	61 190
8	1 944 763	4	113 411
9	2 461 080	6	165 632
10	3 038 113	8	217 853

(a) U-net

1 st layer	Parameters
2	61 190
4	113 411
6	165 632
8	217 853

(b) GCNN

Table 1: **Trainable parameters of U-nets and GCNNs.** For simplification, both architectures use symmetric and scalable configurations. Hence, their complexity can be represented by the number of convolutional filters or node feature’s dimension in the 1st convolutional layer.

with GCNN prediction, the results obtained with the U-net model were interpolated on the triangular meshes. The average MAE on the test set is used for comparison, and is displayed in figure 16a as a function of the number of trainable parameters, both for GCNN and U-net. As can be observed, the GCNN require fewer trainable parameters than U-net (approximately 1 order of magnitude) to reach a similar accuracy, highlighting the cost-efficient of GCNN in field reconstruction compared to grid-based techniques. Indeed, the GCNN is able to exploit local grid refinement or coarsening to save degrees of freedom, which is a central strength of finite-element and finite-volume CFD methods. In terms of memory consumption, the GCNN requires 1.05 GB to import the data set (including the nodes coordinates and connectivity, and the velocity and pressure fields), which must be compared to the 13.5 GB required by the U-net architecture for the same task. Hence, the data used for GCNN is more likely to fit into standard GPU cards.

Yet, the main observed drawback of GCNNs is the larger training time required to reached a converged model. Compared to a traditional convolution, the 4-step graph convolutional block presented in equations (4) and (5) requires heavy matrix operations, which significantly increases the time required for a forward calculation. In order to implement automatic differentiation, the generated intermediate variables are not deleted from the memory until one whole back-propagation of the gradients in all the blocks is finished. The required memory to store the intermediate variables are compared in 16c when batch size is 64. In spite of the small input data size, GCNNs still require large memory for the intermediate variables in the training process. The limitation of memory thus does not support using large batch size, which also slows the training. In figure 16b and 16d, the training time of both methods are compared when batch size is 64. Considering the input size as well as the network’s complexity, training GCNNs is more expensive than training U-nets.

5 Conclusion

In the present contribution, a graph neural network architecture for 2D incompressible laminar flow prediction was presented. Compared to traditional convolution neural networks, the graph based method works directly on body-fitted triangular meshes, making it a method of choice for coupling with finite element numerical simulations. The proposed convolution method preserves permutation invariance, and is therefore insensitive to node and edge indexing. The GCNN network was trained on a dataset of 2000 laminar flows around random 2D shapes, generated using Bézier curves, resulting in an averaged MAE on the test set at 7.70×10^{-3} , with no evidence of overfitting. The resulting model was proved to efficiently reconstruct velocity and pressure fields around a cylinder, respecting the symmetry induced by the input geometry. In particular, the predicted flow remained highly accurate even in the boundary layer of the obstacle. The predicted flow around a significantly different shape, namely a NACA0012 airfoil, presented higher MAE levels and ill-estimated amplitudes, while still respecting the physics of the reference flow. The predicted velocity and pressure fields were further used to compute the drag force by integration of the fields in the boundary layer, resulting in an average relative error equal to 3.43% on the test set. Finally, the GCNN model performance was compared to that of a standard U-net model. The GCNN architecture proved its superiority in terms

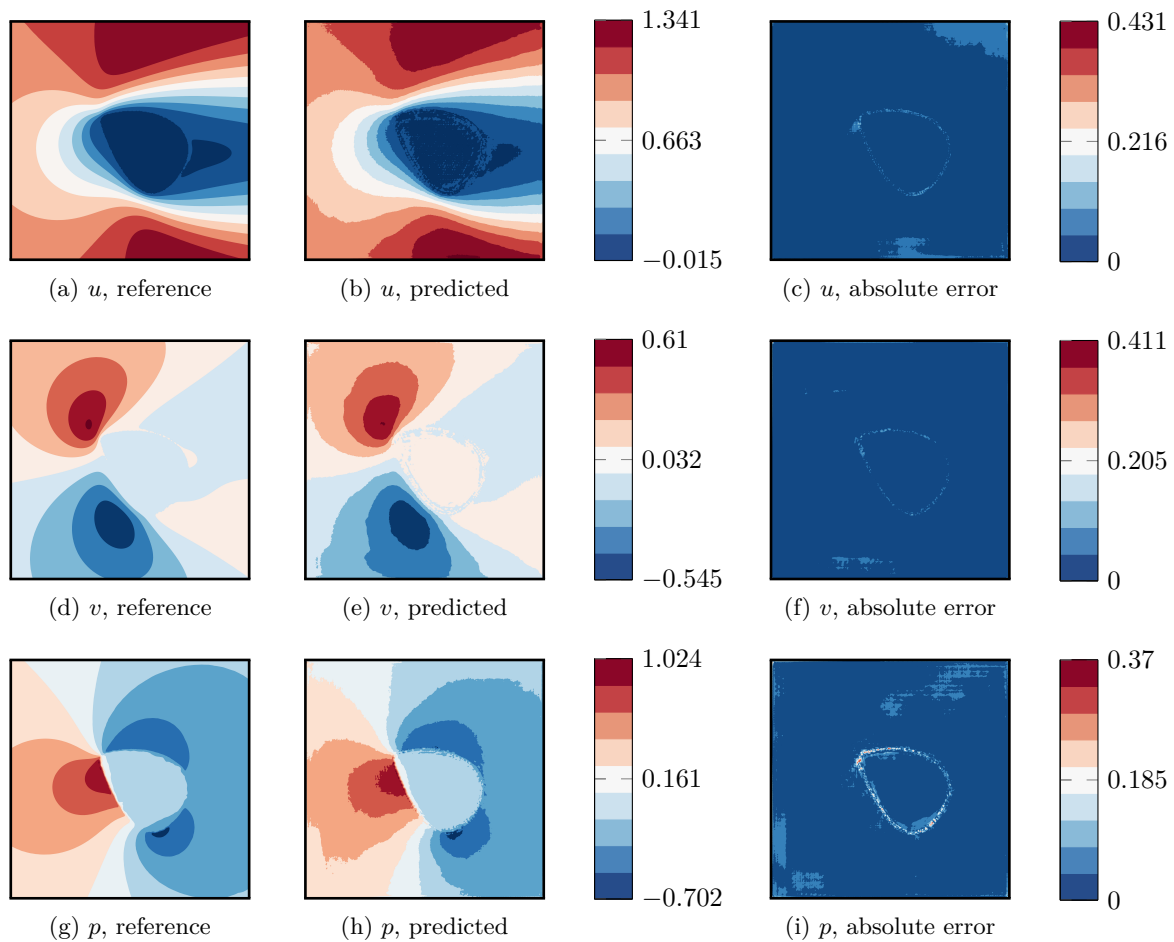


Figure 14: **Flow predictions on a Bézier shape from the test set using a U-net model.** The horizontal velocity (14a, 14b, 14c), vertical velocity (14d, 14e, 14f) and pressure fields (14g, 14h, 14i) display similar error levels. The trained U-net has 8 filters in its first convolutional layer, leading to 1 944 763 trainable parameters in total. The error is mostly concentrated in the border of the obstacle.


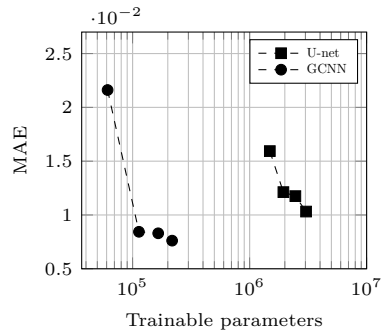



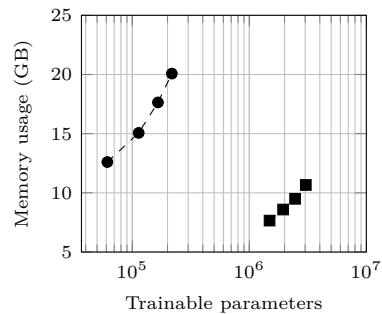
Figure 15: **Flow predictions on a Bézier shape from the test set using a GCNN model.** The horizontal velocity (15a, 15b, 15c), vertical velocity (15d, 15e, 15f) and pressure fields (15g, 15h, 15i) display similar error levels. The concerned GCNN has 217 853 trainable parameters in total. The maximum absolute error is much smaller than the U-net method, with no specific pattern displayed.



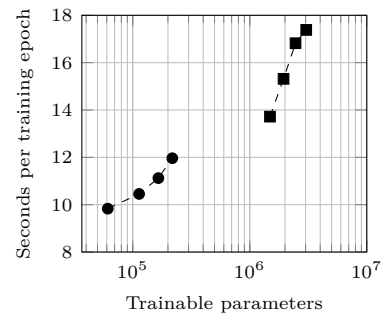
(a) MAE on the test set.



(b) Total training time.



(c) Required memory for gradient backpropagation.



(d) Training time per epoch.

Figure 16: **Performance comparison between U-net and GCNN** in terms of test set MAE (16a), total training time (16b), memory usage (16c) and per-epoch training time (16d). As can be seen, GCNNs require much fewer parameters to reach the same accuracy of U-nets. However, training GCNNs requires more memory to perform gradient backpropagation and is more time-consuming considering their smaller model size.

of accuracy and model complexity, although its required training time was significantly longer. Yet, the versatility of GCNN and its capabilities to work on body-fitted meshes with arbitrarily complex geometric characteristics make it a method of choice for the generation of accurate, fast prediction models coupled to computational fluid dynamics.

Acknowledgements

This work is supported by the Carnot M.I.N.E.S. Institute through the M.I.N.D.S. project.

Author declarations

Conflicts of interest

The authors have no conflicts to disclose.

A Data availability statement

The code of this project is available on the following github repository: https://github.com/jviquerat/gnn_laminar_flow⁵.

References

- [1] X. Guo, W. Li, and F. Iorio. Convolutional neural networks for steady flow approximation. In *22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 481–490, 2016.
- [2] X. Jin, P. Cheng, W. L. Chen, and H. Li. Prediction model of velocity field around circular cylinder over various reynolds numbers by fusion convolutional neural networks based on pressure on the cylinder. *Physics of Fluids*, 30(4), 2018.
- [3] S. Lee and D. You. Data-driven prediction of unsteady flow over a circular cylinder using deep learning. *Journal of Fluid Mechanics*, 879:217–254, 2019.
- [4] Junfeng Chen, Jonathan Viquerat, Frederic Heymes, and Elie Hachem. A twin-decoder structure for incompressible laminar flow reconstruction with uncertainty estimation around 2d obstacles, 2021.
- [5] N. Thuerey, K. Weißenow, L. Prantl, and X. Hu. Deep learning methods for reynolds-averaged navier–stokes simulations of airfoil flows. *AIAA Journal*, 58(1):25–36, 2020.
- [6] K. Fukami, K. Fukagata, and K. Taira. Super-resolution reconstruction of turbulent flows with machine learning. *arXiv preprints*, 2018.
- [7] Aakash Patil, Jonathan Viquerat, George El Haber, and Elie Hachem. Robust deep learning for emulating turbulent viscosities, 2021.
- [8] Y. Zhang, W. J. Sung, and D. N. Mavris. Application of convolutional neural network to predict airfoil lift coefficient. In *2018 AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, 2018.
- [9] J. Viquerat and E. Hachem. A supervised neural network for drag prediction of arbitrary 2d shapes in laminar flows at low reynolds number. *Computers & Fluids*, 210:104645, 2020.
- [10] K. Lee and K. T. Carlberg. Model reduction of dynamical systems on nonlinear manifolds using deep convolutional autoencoders. *Journal of Computational Physics*, 404:108973, 2020.
- [11] S. R. Bukka, A. R. Magee, and R. K. Jaiman. Deep convolutional recurrent autoencoders for flow field prediction, 2020.

⁵The code will be released upon publication of the present manuscript

- [12] F. J. Gonzalez and M. Balajewicz. Deep convolutional recurrent autoencoders for learning low-dimensional feature dynamics of fluid systems, 2018.
- [13] Jaideep Pathak, Mustafa Mustafa, Karthik Kashinath, Emmanuel Motheau, Thorsten Kurth, and Marcus Day. Using machine learning to augment coarse-grid computational fluid dynamics simulations. arXiv preprint arXiv:2010.00072, 2020.
- [14] Dmitrii Kochkov, Jamie A Smith, Ayya Alieva, Qing Wang, Michael P Brenner, and Stephan Hoyer. Machine learning–accelerated computational fluid dynamics. Proceedings of the National Academy of Sciences, 118(21), 2021.
- [15] Mark C Messner. Convolutional neural network surrogate models for the mechanical properties of periodic structures. Journal of Mechanical Design, 142(2):024503, 2020.
- [16] Yann LeCun and Yoshua Bengio. Convolutional networks for images, speech, and time series. In Michael A. Arbib, editor, Handbook of Brain Theory and Neural Networks, page 3361. MIT Press, 1995.
- [17] Han Gao, Luning Sun, and Jian-Xun Wang. Phygeonet: Physics-informed geometry-adaptive convolutional neural networks for solving parameterized steady-state pdes on irregular domain. Journal of Computational Physics, 428:110079, 2021.
- [18] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks, 2017.
- [19] Filipe De Avila Belbute-Peres, Thomas Economou, and Zico Kolter. Combining differentiable PDE solvers and graph neural networks for fluid flow prediction. In Hal Daumé III and Aarti Singh, editors, Proceedings of the 37th International Conference on Machine Learning, volume 119 of Proceedings of Machine Learning Research, pages 2402–2411. PMLR, 13–18 Jul 2020.
- [20] Justin Gilmer, Samuel S. Schoenholz, Patrick F. Riley, Oriol Vinyals, and George E. Dahl. Neural message passing for quantum chemistry. In Doina Precup and Yee Whye Teh, editors, Proceedings of the 34th International Conference on Machine Learning, volume 70 of Proceedings of Machine Learning Research, pages 1263–1272. PMLR, 06–11 Aug 2017.
- [21] Francis Ogoke, Kazem Meidani, Amirreza Hashemi, and Amir Barati Farimani. Graph convolutional neural networks for body force prediction, 2020.
- [22] Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, Advances in Neural Information Processing Systems, volume 30. Curran Associates, Inc., 2017.
- [23] 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, and Razvan Pascanu. Relational inductive biases, deep learning, and graph networks, 2018.
- [24] Alvaro Sanchez-Gonzalez, Jonathan Godwin, Tobias Pfaff, Rex Ying, Jure Leskovec, and Peter W. Battaglia. Learning to simulate complex physics with graph networks, 2020.
- [25] Tobias Pfaff, Meire Fortunato, Alvaro Sanchez-Gonzalez, and Peter Battaglia. Learning mesh-based simulation with graph networks. In International Conference on Learning Representations, 2021.
- [26] F. Monti, D. Boscaini, J. Masci, E. Rodolà, J. Svoboda, and M. M. Bronstein. Geometric deep learning on graphs and manifolds using mixture model cnns. In 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 5425–5434, 2017.
- [27] Mengfei Xu, Shufang Song, Xuxiang Sun, and Weiwei Zhang. A convolutional strategy on unstructured mesh for the adjoint vector modeling. Physics of Fluids, 33(3):036115, 2021.
- [28] J. Bruchon, H. Digonnet, and T. Coupez. Using a signed distance function for the simulation of metal forming processes: Formulation of the contact condition and mesh adaptation. International Journal for Numerical Methods in Engineering, 78(8):980–1008, 2009.

- [29] E. Hachem, S. Feghali, R. Codina, and T. Coupez. Immersed stress method for fluid structure interaction using anisotropic mesh adaptation. International Journal for Numerical Methods in Engineering, 94:805–825, 2013.
- [30] T.E. Tezduyar, S. Mittal, S.E. Ray, and R. Shih. Incompressible flow computations with stabilized bilinear and linear equal-order-interpolation velocity-pressure elements. Computer Methods in Applied Mechanics and Engineering, 95(2):221–242, 1992.
- [31] Y. Bazilevs, V.M. Calo, J.A. Cottrell, T.J.R. Hughes, A. Reali, and G. Scovazzi. Variational multiscale residual-based turbulence modeling for large eddy simulation of incompressible flows. Computer Methods in Applied Mechanics and Engineering, 197(1):173–201, 2007.
- [32] K. Takizawa, T. E. Tezduyar, and Y. Otoguro. Stabilization and discontinuity-capturing parameters for space–time flow computations with finite element and isogeometric discretizations. Computational Mechanics, 62(5):1169–1186, 2018.
- [33] Y. Otoguro, K. Takizawa, T. E. Tezduyar, K. Nagaoka, R. Avsar, and Y. Zhang. Space–time vms flow analysis of a turbocharger turbine with isogeometric discretization: computations with time-dependent and steady-inflow representations of the intake/exhaust cycle. Computational Mechanics, 64(5):1403–1419, 2019.
- [34] Y. Otoguro, K. Takizawa, and T. E. Tezduyar. Element length calculation in b-spline meshes for complex geometries. Computational Mechanics, 65:1085–1103, 2020.
- [35] Prajit Ramachandran, Barret Zoph, and Quoc V. Le. Searching for activation functions, 2018.
- [36] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In Yee Whye Teh and Mike Titterton, editors, Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics, volume 9 of Proceedings of Machine Learning Research, pages 249–256, Chia Laguna Resort, Sardinia, Italy, 13–15 May 2010. PMLR.
- [37] Takato Otsuzuki, Hideaki Hayashi, Yuchen Zheng, and Seiichi Uchida. Regularized pooling. In Artificial Neural Networks and Machine Learning – ICANN 2020, pages 241–254, Cham, 2020. Springer International Publishing.
- [38] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017.
- [39] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015, pages 234–241, 2015.