



HAL
open science

U-net architectures for fast prediction in fluid mechanics

Junfeng Chen, Jonathan Viquerat, Elie Hachem

► **To cite this version:**

Junfeng Chen, Jonathan Viquerat, Elie Hachem. U-net architectures for fast prediction in fluid mechanics. 2019. hal-02401465

HAL Id: hal-02401465

<https://hal.science/hal-02401465>

Preprint submitted on 10 Dec 2019

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.

U-net architectures for fast prediction in fluid mechanics

Junfeng Chen^a, Jonathan Viquerat^{a,*}, Elie Hachem^a

^a MINES ParisTech, PSL - Research University, CEMEF - Centre for material forming, CNRS UMR 7635, CS 10207 rue Claude Daunesse, 06904 Sophia-Antipolis Cedex, France.

Abstract

Machine learning is a popular tool that is being applied to many domains, from computer vision to natural language processing. It is not long ago that its use was extended to physics, but its capabilities remain to be accurately contoured. In this paper, we are interested in the prediction of 2D velocity and pressure fields around arbitrary shapes in laminar flows using supervised neural networks. To this end, a dataset composed of random shapes is built using Bézier curves, each shape being labeled with its pressure and velocity fields by solving Navier-Stokes equations using a CFD solver. Then, several U-net architectures are trained on the latter dataset, and their predictive efficiency is assessed on unseen shapes, using *ad hoc* error functions.

Keywords: machine learning, neural networks, convolutional networks, computational fluid dynamics, immersed method

1. Introduction

During the last few years, the CFD community has largely benefited from the fast-paced development of the machine learning (ML) field, and more specifically for that of the neural networks (NN) domain. In many cases, only a part of the resolution process is replaced with a trained NN, in order to reduce the computational cost. Examples for these applications are the prediction of closure terms in RANS [1] [2] or LES [3] computations. In other situations, a supervised network is directly trained to predict a flow profile: in [4], an auto-encoder NN architecture is used to obtain steady state flow predictions around elementary and real-life shapes; in [5], a fusion convolutional neural network (CNN) is trained to predict velocity snapshots around a cylinder in weakly turbulent flows, using the time history of pressure around the cylinder as an input.

The U-net is a particular CNN architecture, which was initially proposed by Ronneberger *et al* [6] for biomedical image segmentation. As shown in figure 1, U-nets are composed of (i) a contractive path based on the repetition of a convolution/max-pooling pattern, followed by (ii) an upscaling path, in which a pattern of deconvolution/concatenation/convolution is applied several times. The *skip connections* in U-nets, originally proposed by [7], are used to link the contracting path to the upscaling path, and help localize features of high resolution during the upsampling, thus improving the segmentation accuracy.

Based on the end-to-end property of U-nets, several variations have been proposed in the recent years. In [8], a stacked network architecture is proposed for the estimation human pose. The whole network is simply built by adding U-nets one after the other, the output of a U-net being the input of the following one. In [9], shortcuts were added from each U-net to all the following ones, in order to help the back-propagation of gradient. Multi-paths U-nets were introduced in [10], where separated contracting paths were used for images holding different informations. The different paths are merged at the bottom of the U-net, and share a single expanding path. In [11], the authors report the parallel training of independent U-nets, the final prediction being obtained as a weighted sum of the outputs of the different networks. Other methodological

* Corresponding author

Email address: jonathan.viquerat@mines-paristech.fr (Jonathan Viquerat)

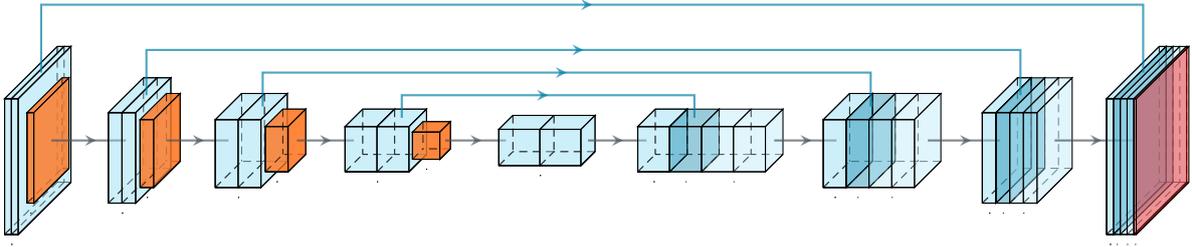


Figure 1: **Standard U-net architecture.** The contracting path is based on a pattern made of two convolutional layers (in light blue) followed by a max-pooling layer (in orange). At each occurrence of the pattern, the image size is divided by two, while the number of filters is doubled. In the upscaling path, a deconvolution step is first applied to the input, while the number of filters is halved. The output of this layer is then concatenated with its mirror counterpart in the contracting path using a shortcut connections (here noted with arrows). Finally, two convolution layers are applied. At the end of the last layer, a 1×1 convolution with 3 filters is applied to obtain an RGB image at the output (in red).

refinements demonstrated their potential, such as multi-scale inputs and supervision [12], for example. These different variations are discussed in details in section 2.

U-net architectures recently made their way to other physical fields, such as fluid dynamics. In [13], the authors exploit U-nets to infer the velocity and pressure of turbulent flow around airfoils computed in a Reynolds-Averaged Navier-Stokes framework. In [14], a traditional CNN is coupled to a Downsampled Skip-Connection Multi-Scale (DSC/MS) model to reconstruct the turbulent flow computed by direct numerical simulation. The DSC/MS model, which uses a downscaling-upsampling path like U-net, efficiently recovers the detailed velocity field.

In this paper, focus is made on the prediction of laminar flows around arbitrary 2D shapes. To this end, we assess the predictive capabilities of several U-net based architectures, and compare them to that of a 4-levels baseline U-net, presented in figure 1. In section 2, the functioning of the baseline U-net is briefly covered, and advanced architectures are introduced. Details about the dataset are given in section 3, along with a discussion on the metrics used to evaluate the results of the different architectures. In section 4, details are given about the training, and the performances of the different networks are compared on a test subset from the dataset. Finally, the predictions of the networks are assessed on unseen configurations, such as geometrical shapes or airfoils.

2. Existing U-net architectures

U-nets were introduced in 2015 [6] as an improvement over a sliding-window method for classical convolutional networks [15]. In classical convolutional networks, a contracting path similar to the first branch of the U-net is used to extract high-level features from an input image. However, as the abstraction level of the features increases, localization information about these features is lost, making it hard to directly reconstruct meaningful image as an output of the network. The idea behind U-net is to follow the contracting path by an upscaling path, in which the resolution of the feature maps is re-increased symmetrically to the contracting path using deconvolution layers. After each deconvolution, localization information is re-injected in the upscaling path by concatenating information from the symmetric level of the downscaling path using a shortcut connection. The concatenation is followed by a convolution operation that assembles the two informations (features coming from the previous level of the upscaling path, and localization coming from the shortcut connection) together.

Recently, improved architectures based on U-nets have been proposed. Newell *et al* [8] proposed a stacked architecture, similar to the one shown in figure 2b. The successive stacking of networks is said to consolidate the captured features, and to allow for the grasp of more complex spatial relationships. In [9], Tang *et al* introduced coupled U-nets (CU-nets), presented in figure 2d. Based on the current trend of dense connectivities in convolutional networks [16], the authors proposed to add a similar level of density to coupled U-nets, by introducing shortcuts between layers from same levels. They show that

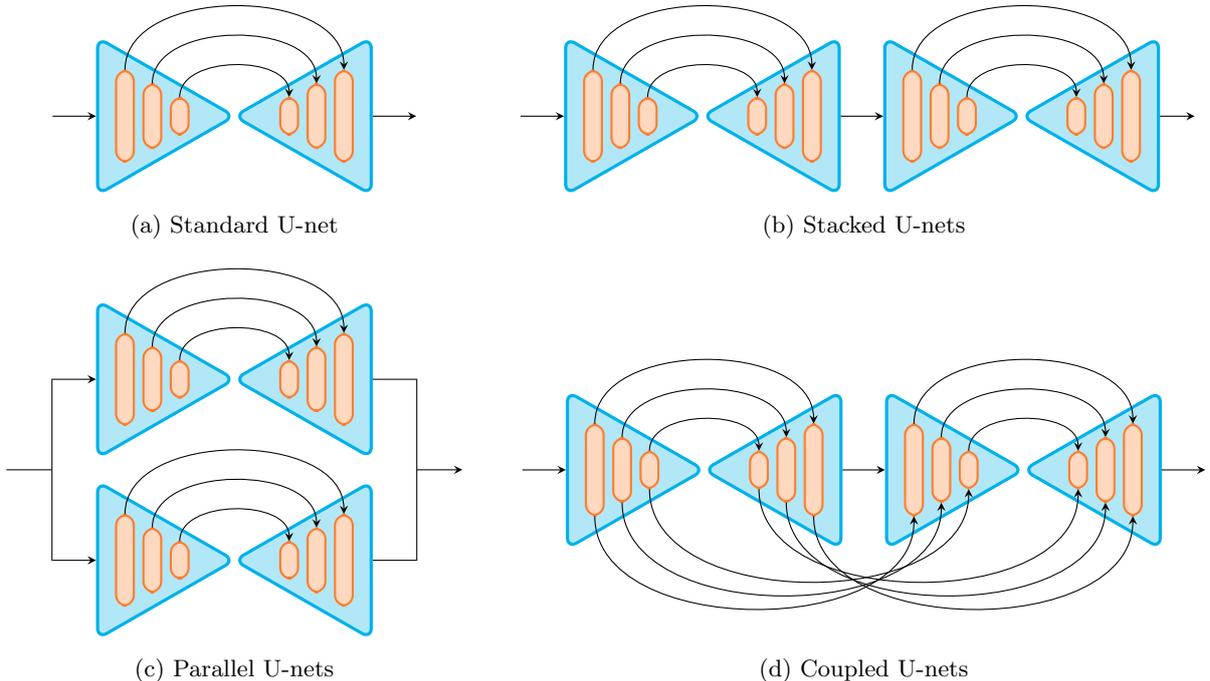


Figure 2: U-net architectures.

these shortcut help the information flow across the structure, while improving the efficiency of the gradient backpropagation. Finally, the separate training of U-nets in parallel is reported in [11]. In this contribution, the authors generate different versions of their initial dataset by adding different zero-mean gaussian noises. Then, three different U-nets are trained separately, each on a different version of the dataset. The final output is computed as a weighted combination of the outputs of the separate U-nets. Results show that this method reduces overfitting, thus inducing better performances on reduced datasets, and proved to be superior to a single U-net trained on any of the datasets (noisy or not).

In the current paper, we introduce a parallel U-net (PU-net) architecture, shown in figure 2c. The input image is fed to two separate U-nets, which outputs are simply recombined using an average layer. At the difference of [11], the two networks are trained concurrently, using the same dataset.

3. 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 [17], thus only the main lines are sketched here. For more details, the reader is referred to [17]. First, we describe the steps to generate arbitrary shapes by means of connected Bezier curves. Then, the solving of the Navier-Stokes equations with an immersed method is presented. Finally, details about the dataset are given.

3.1. 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 3a), and the angles between consecutive random points are then computed. An average angle is then computed around each point (see figure 3b) 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 3c). A final sampling of the successive Bézier curves leads to a boundary description of the shape (figure 3d). Using this method, a wide variety of shapes can be attained, as shown in figure 4.

3.2. 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 [18], 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 [19] for additional details about formulation (2). Eventually, the modified equations (2) are cast into a stabilized finite element formulation, and solved using a variational multi-scale (VMS) solver [19].

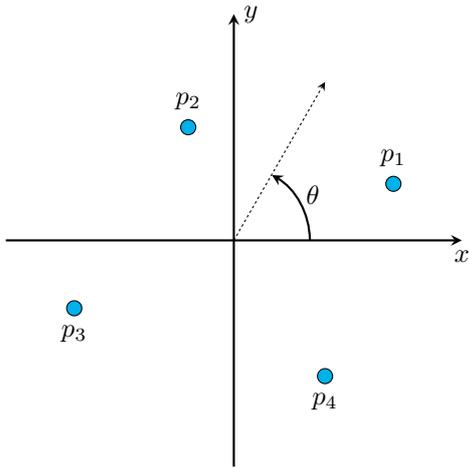
3.3. Dataset

The dataset (DS) is composed of 12.000 shapes, along with their steady-state velocity and pressure fields at $Re = 10$ (see figure 5). All the labels were computed using cimLib [19], following the methods exposed in section 3.2. In the following, the DS is systematically divided into three sets: 9600 shapes for the training set, 1200 shapes for the validation set, and 1200 shapes for the test set.

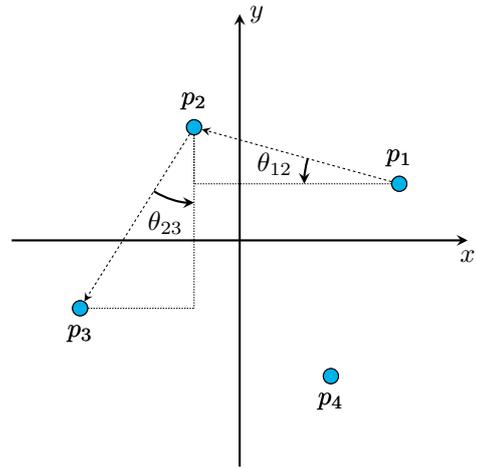
4. Networks training

4.1. Networks

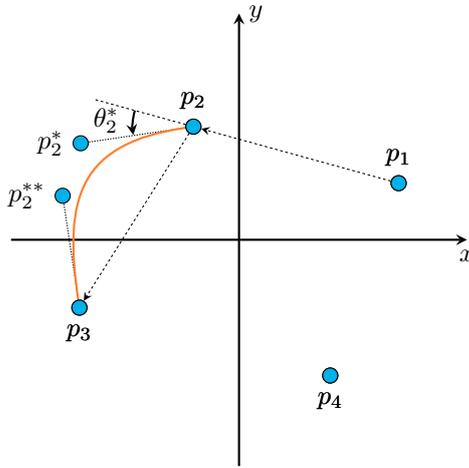
The baseline U-net exploited in this paper is composed of 4 blocks in the contractive path, a bottleneck block, and 4 blocks in the upscaling path (see figure 1). The initial number of kernels is set to 32 and is doubled at each pooling operation, reaching 512 in the bottleneck. It is then divided by two at each upscaling step, dropping back to 32 in the final block. The upscaling path is followed by a last 1×1 convolution with 3 filters, thus reconstructing an RGB image. All convolution kernels are of size 3×3 , deconvolution kernels and strides of size 2×2 , and pooling sizes and strides also of size 2×2 . All activation functions are rectified linear units (ReLU).



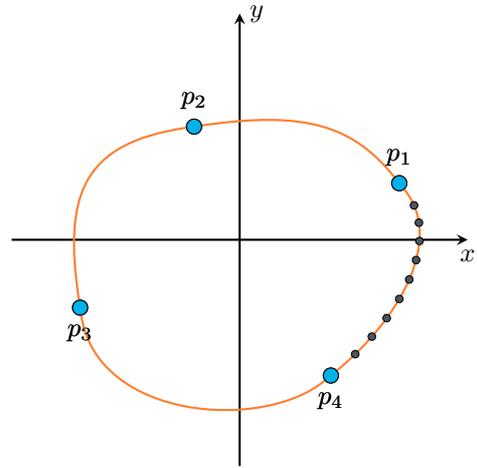
(a) Draw n_s random points, translate them around $(0, 0)$ and sort them by ascending trigonometric angle



(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 3: **Random shape generation with cubic Bézier curves.**

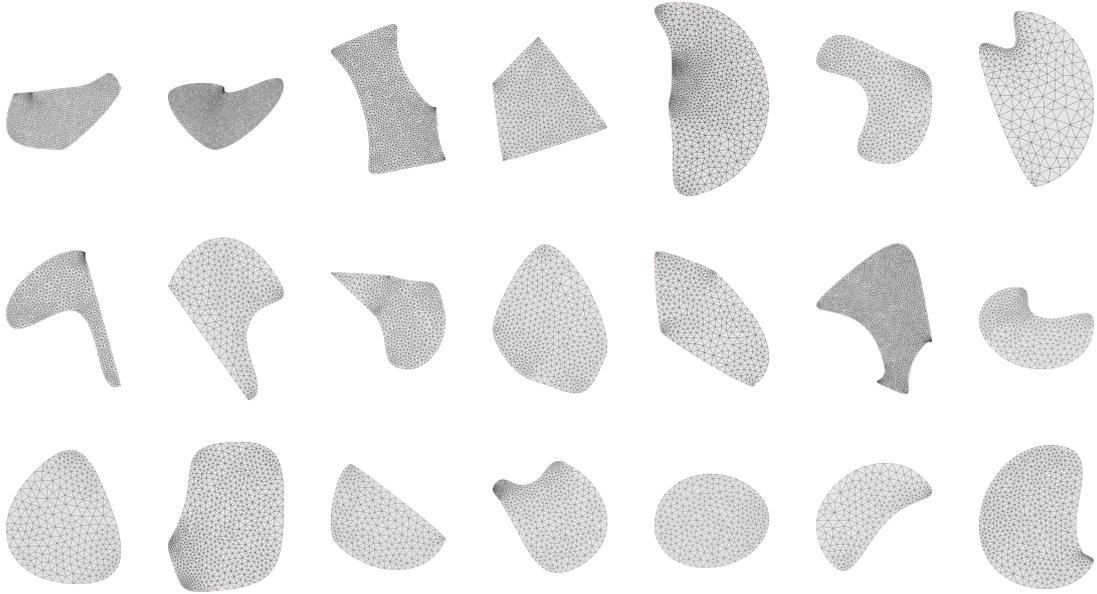


Figure 4: **Shape examples drawn from the dataset.** A wide variety of shape is obtained using a restrained number of points ($n_s \in [4, 6]$), as well as a local curvature r and averaging parameter α .

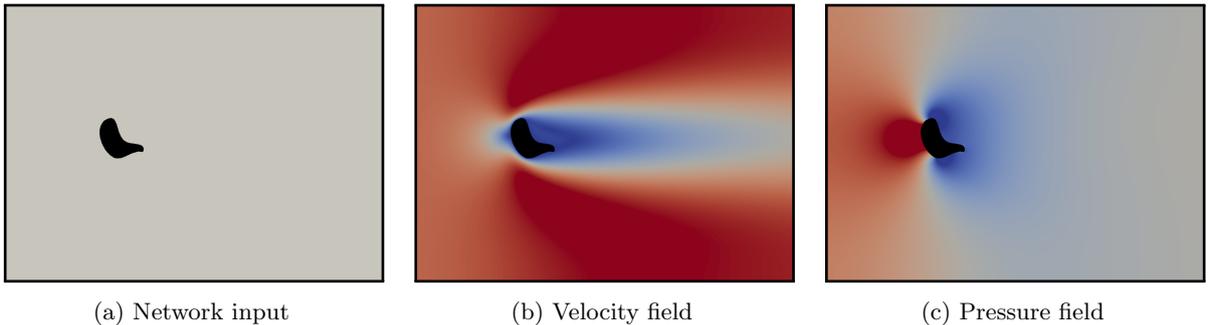


Figure 5: **Network input, velocity field and pressure field for a dataset element.** The shape is shown in its computational domain (left), along with the computed velocity field (center) and pressure field (right). Fields are scaled to $[0, 1]$ for velocity and $[-1, 1]$ for pressure.

4.2. Training

The learning process in neural networks consists in adjusting all the biases and weights of the network in order to reduce the value of a well-chosen loss function. For regression cases, it is common to choose the mean squared error (MSE). With the loss function at hand, the optimization of the weights and biases is performed with a stochastic gradient descent that exploits the back-propagation algorithm [20]. The gradient descent is usually performed multiple times on small random subsets (called mini-batches), instead of considering a single evaluation of the gradient on the whole dataset. This method represents a compromise between an accurate gradient computation and a low computational cost. In this paper, the batch size is set to 8. The whole dataset, split in mini-batches, is re-used multiple times in random orders, a full re-use of it being called an epoch. This procedure is repeated as long as the validation accuracy (the accuracy of the prediction computed on the validation subset) increases, and is stopped when it starts decreasing, meaning that the network is overfitting (*i.e.* learning non-generic features) on the training subset. In this paper, the input images are downscaled to a constant size of 107×78 pixels to limit the training cost. To ensure fair

comparisons, all random seeds are set once and for all for the training of the different considered networks.

The amount of ready-to-use neural networks libraries has exploded in the recent years, most of them exploiting C++ or Python. For supervised learning, they usually include a wide range of choices regarding layer types, activation functions, losses, optimizers and so on. In this paper, we chose to use Keras [21] for its high level of abstraction and the ease of use provided by the Python language.

4.3. Model evaluation

The mean squared error is often used as a loss function for the train of neural networks. It is defined as:

$$\varepsilon_{\text{mse}} = \frac{1}{n(P)} \sum_{p=1}^{n(P)} \left(\frac{1}{3} \sum_{j=1}^3 (y_{pj} - \hat{y}_{pj})^2 \right), \quad (4)$$

where P is the set of all pixels in an image, $n(P)$ is the number of pixels in the image, j is the index looping over the three channels of an RGB image, \hat{y}_{pj} is the predicted value of channel j at pixel p , and y_{pj} is its corresponding label. Still, it can be hard to compare the quality of different predictions by relying only on the mean squared error, as it is directly related to the amplitude of the solution. To this end, a complementary pixel-wise relative error function is proposed for the assessment of the predictions quality:

$$\varepsilon_{\alpha} = \frac{1}{n(P)} \sum_{p=1}^{n(P)} \mathbb{1}(\varepsilon_{\alpha}^p > 0), \quad \text{with } \varepsilon_{\alpha}^p = \mathbb{1} \left(\left| \frac{y_p - \hat{y}_p}{y_p + \eta} \right| > \alpha \right) \quad \text{and } y_p = \sum_{j=1}^3 y_{pj}. \quad (5)$$

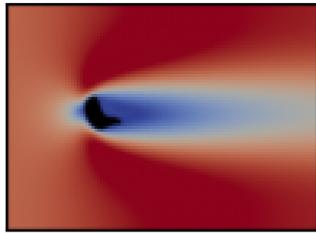
In the latter expression, η is a small positive number introduced to ensure a non-zero denominator in the case where $y_{pj} = 0$, and α is a tolerance for the relative error. Hence, this metric returns the proportion of pixels which prediction accuracy is lower than a given threshold level α , without being proportional to the solution amplitude. Regarding the perturbation induced by η , we argue that only the pixels whose absolute value is close to η will be impacted. By setting $\eta = 1 \times 10^{-6}$, the variations observed in ε_{α} were systematically lower than 0.1%, which represents a negligible error, as will be shown in later sections. Additionally, for a single image, it is possible to plot the 2D map of ε_{α}^p to locate the areas of low prediction accuracy.

5. Results

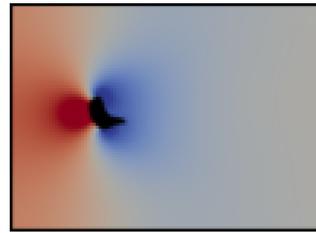
In this section, a prediction from the test subset obtained with the baseline U-net is first presented and analyzed in details with the two metrics presented in section 4.3. Then, the different U-net architectures are statistically compared on the test subset. In a third time, the predictive efficiency of the different networks is assessed on several unseen shapes, including NACA airfoils.

A prediction example drawn from the test subset is shown in figure 6, along with its corresponding exact solutions and error maps. As can be seen, the predictions are physically sound, and, to the eye, fairly close to their targets. Looking at the pressure predictions, we observe that the network successfully predicted two minimal pressure areas at the top and bottom right of the shape. The maximal errors are located in the direct vicinity of the shape, and in areas of high gradients. By comparing the two error maps $\varepsilon_{0.01}^p$ and $\varepsilon_{0.05}^p$, one sees that most pixels are predicted within a 5% relative error range, and that the remaining pixels mostly concentrate at the surface of the shape. These error levels could be reduced by reducing the inference region of the network.

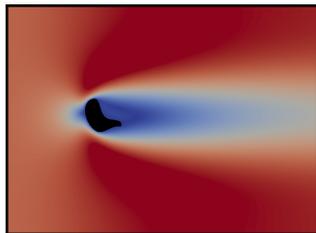
As both ε_{mse} and ε_{α} are related to a single image, we compute their average, standard deviation and maximal values to represent these estimators over the entire test subset. The different errors, shown in figure 7, are all normalized to their baseline U-net counterpart for better comparison. On the top row, the error levels for the velocity predictions indicate that, in average, advanced compositions of U-nets provide better predictions in average, up to 45% of improvement for the PU-net in the MSE sense. Regarding standard deviation and maximum error, SU-nets present a significant advantage, with a drop close to 40% for standard deviation and 50% for maximum error using MSE. These results suggest that SU-nets are able to catch outliers with more efficiency than CU- and PU-nets.



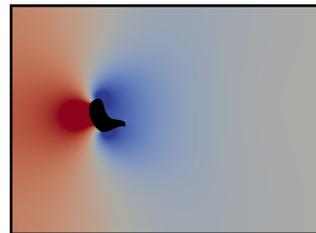
(a) Velocity prediction



(b) Pressure prediction



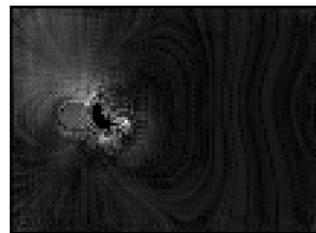
(c) Exact velocity



(d) Exact pressure



(e) Absolute error for velocity prediction



(f) Absolute error for pressure prediction



(g) Pixels with failed velocity prediction for $\alpha = 0.01$



(h) Pixels with failed pressure prediction for $\alpha = 0.01$



(i) Pixels with failed velocity prediction for $\alpha = 0.05$



(j) Pixels with failed pressure prediction for $\alpha = 0.05$

Figure 6: Velocity (top left) and pressure (top right) predictions on a shape from the test subset, along with their exact solutions and errors to CFD solution. Fields are scaled to $[0, 1]$ for velocity and $[-1, 1]$ for pressure, while absolute errors are normalized to their maximal value and scaled to $[0, 1]$.

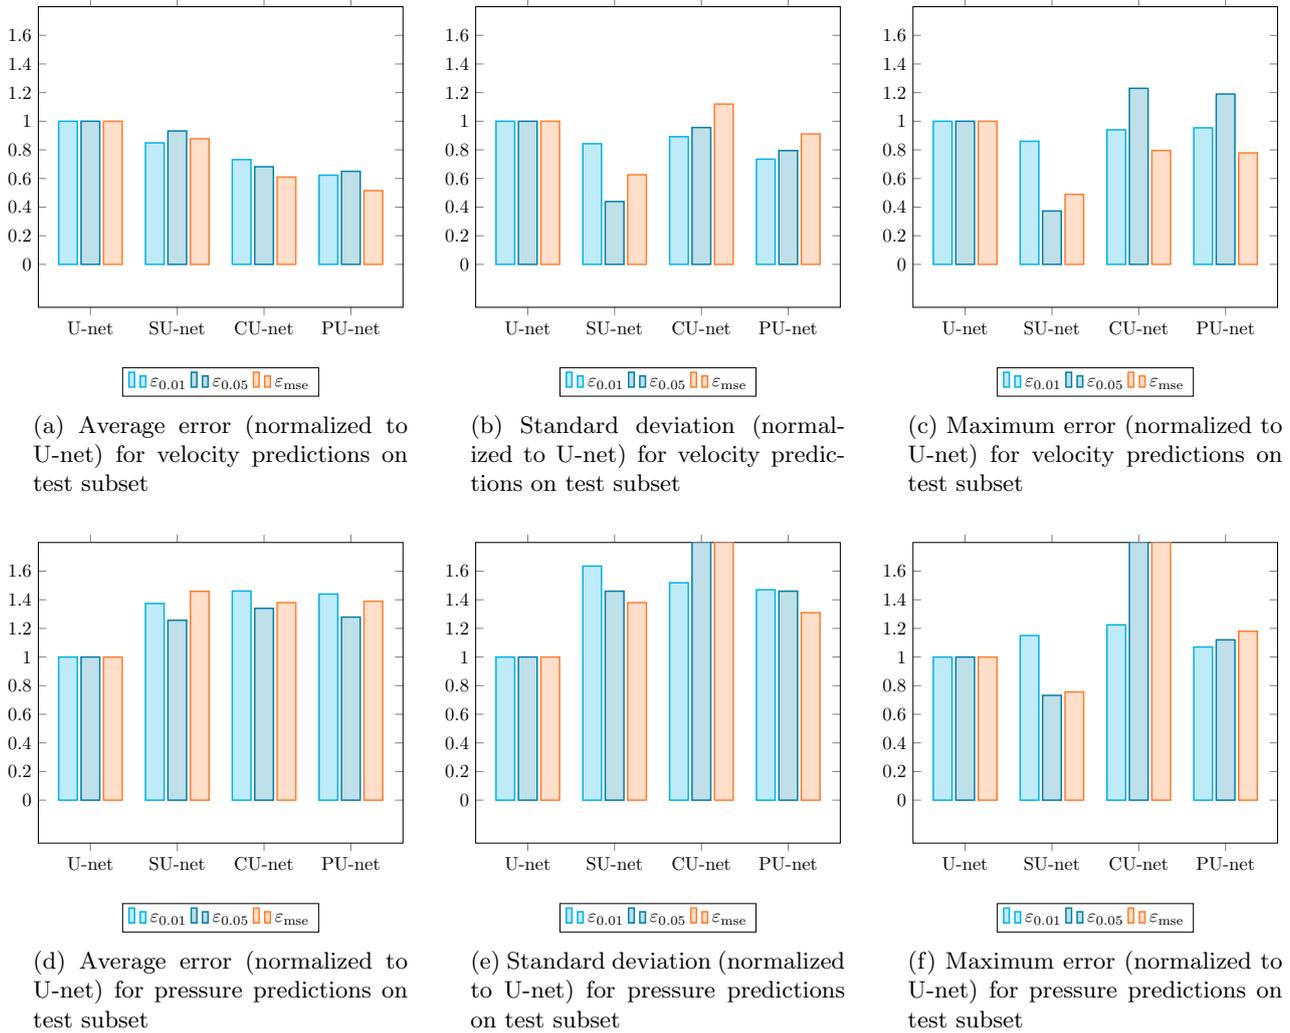


Figure 7: Average errors, deviations and maximal errors obtained on the test subset using different U-net-based architectures for velocity (top row) and pressure fields (bottom row).

Although one could expect to observe similar results for pressure predictions, the performances of SU-, CU- and PU-nets are clearly below that of the baseline U-net for average error, standard deviation and maximum error. The fact that pressure predictions cannot be improved using more complex networks can arise from an insufficient pre-processing of the target data, causing the learning problem to be ill-conditioned, as suggested in [13]. Still, the improvements obtained on the velocity predictions when using more complex networks may not be worthwhile in regard of the overhead of computational effort they require. Indeed, the baseline U-net shown in figure 1 contains 7.7 millions parameters, and requires 3 hours of training on a regular computer (not using GPU cards), while the SU-, CU- and PU-nets contain between 15 and 18 million parameters, leading to training times of 6 to 7 hours. For that reason, we suggest that, although improved U-net-based architectures provide better results in segmentation tasks, such superiority may not reflect in regression tasks such as field map predictions.

We now focus on the prediction of field maps in realistic, unseen configurations such as geometric shapes or NACA airfoils, using the baseline U-net architecture of figure 1. As shown in figure 8, predictions provided by the trained U-net are still physically sound, and field details such as multiple high and low pressure areas in the vicinity of the obstacles are well reproduced, to the limit of the resolution of the network output. MSE

error levels for these predictions are in the same order of magnitude as those obtained on the test subset, *i.e.* between 1×10^{-5} and 5×10^{-5} for both pressure and velocity. Further testing would require to reduce the predicted domain around the shapes to obtain better details in its vicinity, or to output predictions with higher resolution, although it would require more computational power.

6. Conclusion

In the current contribution, we presented the prediction of velocity and pressure field maps around arbitrary 2D shapes in laminar flows exploiting U-net-based architectures. First, details were given about the considered U-net architectures, and how they recently started to be exploited for regression tasks. Insights were then given on the dataset construction and the networks training, followed by a discussion on the metrics used to measure the quality of the predictions. A prediction example was then analyzed on a shape drawn from the test subset, that showed good agreement with the exact solution for both velocity and pressure. Then, the predictions were assessed statistically on the entire test subset, showing that, although they perform better on segmentation tasks, advanced architectures such as stacked, coupled and parallel U-nets do not provide a clear advantage on regression tasks such as field map predictions. Finally, we showed that such predictions could be made on unseen shapes with reasonable accuracy, and that the predictions were physically sound, thanks to the rich geometrical features present in the dataset.

Although this application is focused on a fairly basic case, these results give hope for the possibility of tackling more complex problems, such as turbulent viscosity maps at higher Reynolds numbers, or 3D fields, for example. Additionally, such tools can be exploited in related domains as surrogate models, for optimization purposes.

Acknowledgements

The authors would like to acknowledge the support of the MINDS project from the Carnot M.I.N.E.S. institute.

- [1] J. Ling, A. Kurzwski, J. Templeton, Reynolds averaged turbulence modelling using deep neural networks with embedded invariance, *Journal of Fluid Mechanics* 807 (2016) 155–166.
- [2] B. D. Tracey, K. Duraisamy, J. J. Alonso, A Machine Learning Strategy to Assist Turbulence Model Development, in: 53rd AIAA Aerospace Sciences Meeting, 2015, pp. 1–23.
URL <http://arc.aiaa.org/doi/10.2514/6.2015-1287>
- [3] A. D. Beck, D. G. Flad, C.-D. Munz, Deep Neural Networks for Data-Driven Turbulence Models, arXiv.
URL <http://arxiv.org/abs/1806.04482>
- [4] X. Guo, W. Li, F. Iorio, Convolutional Neural Networks for Steady Flow Approximation, *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD '16* (2016) 481–490.
URL <http://dl.acm.org/citation.cfm?doid=2939672.2939738>
- [5] X. Jin, P. Cheng, W. L. Chen, 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).
- [6] O. Ronneberger, P. Fischer, T. Brox, U-net: Convolutional networks for biomedical image segmentation, arXiv.
- [7] J. Long, E. Shelhamer, T. Darrell, Fully Convolutional Networks for Semantic Segmentation, in: *IEEE Conference on Computer Vision and Pattern Recognition*, 2015.
URL <https://ieeexplore.ieee.org/document/7298965>
- [8] A. Newell, K. Yang, J. Deng, Stacked hourglass networks for human pose estimation, *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 9912 LNCS (2016) 483–499.
- [9] Z. Tang, X. Peng, S. Geng, Y. Zhu, D. N. Metaxas, CU-Net: Coupled U-Nets.
URL <http://arxiv.org/abs/1808.06521>
- [10] J. Dolz, I. Ben Ayed, C. Desrosiers, Dense multi-path u-net for ischemic stroke lesion segmentation in multiple image modalities, *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 11383 LNCS (2019) 271–282.
- [11] D. Li, J. Wu, Y. He, X. Yao, W. Yuan, D. Chen, H.-C. Park, S. Yu, J. L. Prince, X. Li, Parallel deep neural networks for endoscopic OCT image segmentation, *Biomedical Optics Express* 10 (3) (2019) 1126.
- [12] J. Stawiaski, A Multiscale Patch Based Convolutional Network for Brain Tumor Segmentation (2017) 1–11.
URL <http://arxiv.org/abs/1710.02316>
- [13] N. Thuerey, K. Weissenow, L. Prantl, X. Hu, Deep Learning Methods for Reynolds-Averaged Navier-Stokes Simulations of Airfoil Flows.
URL <http://arxiv.org/abs/1810.08217>

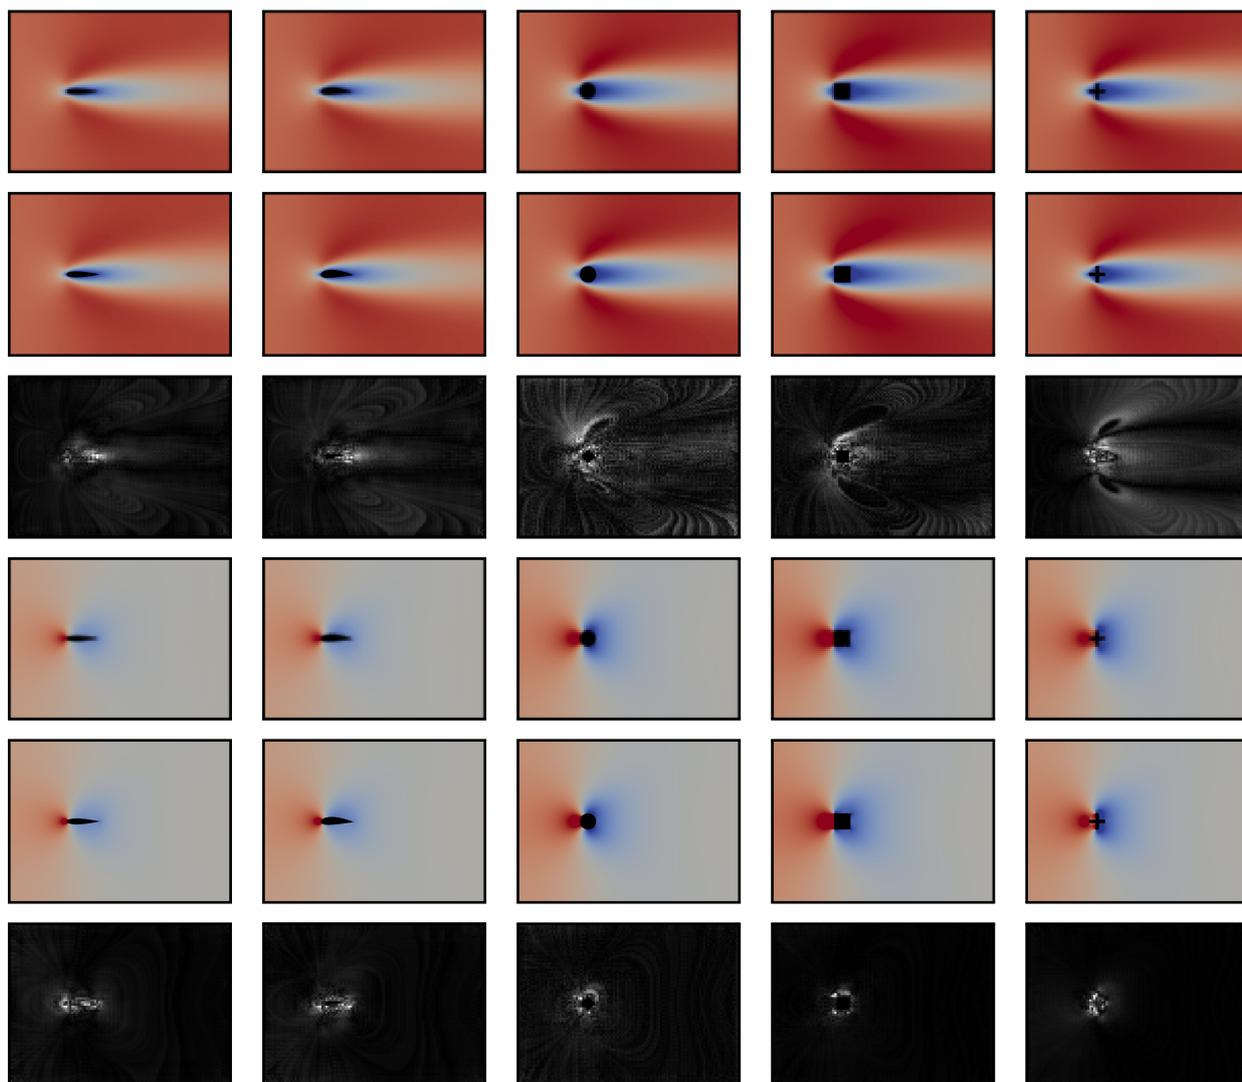


Figure 8: Velocity (first row) and pressure (fourth row) predictions compared to their exact counterparts (second and fifth rows), along with their associated absolute error maps (third and sixth rows) for different real-life shapes.. From left to right: NACA 0018, NACA 4424, cylinder, square and cross.

- [14] K. Fukami, K. Fukagata, K. Taira, Super-resolution reconstruction of turbulent flows with machine learning (2018) 1–14. URL <http://arxiv.org/abs/1811.11328>
- [15] D. Cireşan, A. Giusti, L. M. Gambardella, J. Schmidhuber, Deep Neural Networks Segment Neuronal Membranes in Electron Microscopy Images, Nips (2012) 2843–2851. URL <http://papers.nips.cc/paper/4741-deep-neural-networks-segment-neuronal-membranes-in-electron-microscopy-images.pdf>
- [16] G. Huang, Z. Liu, L. Van Der Maaten, K. Q. Weinberger, Densely connected convolutional networks, Proceedings - 30th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017 2017-Janua (2017) 2261–2269.
- [17] J. Viquerat, E. Hachem, A supervised neural network for drag prediction of arbitrary 2d shapes in low reynolds number flows (2019). [arXiv:1907.05090](https://arxiv.org/abs/1907.05090).
- [18] J. Bruchon, H. Digonnet, 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) (2009) 980–1008.
- [19] E. Hachem, S. Feghali, R. Codina, T. Coupez, Immersed stress method for fluid structure interaction using anisotropic mesh adaptation, International Journal for Numerical Methods in Engineering 94 (2013) 805–825.
- [20] A. C. Ian Goodfellow, Yoshua Bengio, The Deep Learning Book, MIT Press, 2017.
- [21] F. Chollet, Deep Learning with Python, Manning, 2018.