

Use of Convolutional Neural Networks for Multifidelity CFD Problems

Original

Use of Convolutional Neural Networks for Multifidelity CFD Problems / Marci', S.G.I., Muscara, L., Folcarelli, L., Ferrero, A., Masseni, F., Pastrone, D.G.. - 69:(2026), pp. 66-71. (10th CEAS Aerospace Europe Conference and 28th AIDAA International Congress, 2025 Turin (ITA) December 1-4, 2025) [10.21741/9781644904251-12].

Availability:

This version is available at: 11583/3016391 since: 2026-10-07T09:48:04Z

Publisher:

Association of American Publishers

Published

DOI:10.21741/9781644904251-12

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Use of Convolutional Neural Networks for Multifidelity CFD Problems

Sebastiano Geremia Isidoro MARCI^{1,a*}, Luca MUSCARA^{1,b},
Lorenzo FOLCARELLI^{1,c}, Andrea FERRERO^{1,d}, Filippo MASSENI^{1,e}
and Dario Giuseppe PASTRONE^{1,f}

¹Politecnico di Torino, Dipartimento di Ingegneria Meccanica e Aerospaziale, Corso Duca degli Abruzzi 24, 10129 Torino, Italy

^asebastiano.marci@polito.it, ^bluca.muscara@polito.it, ^clorenzo.folcarelli@polito.it,

^dandrea_ferrero@polito.it, ^efilippo.masseni@polito.it, ^fdario.pastrone@polito.it

Keywords: Convolutional Neural Networks (CNN), U-Net, Computational Fluid Dynamics (CFD), Field Reconstruction, Super-Resolution, Multifidelity

Abstract. In recent years, there has been a growing interest in applying deep learning techniques to Computational Fluid Dynamics (CFD), particularly for accelerating simulations and enhancing data resolution. This work explores the use of Convolutional Neural Networks (CNNs) to reconstruct high-resolution CFD fields from low-resolution inputs. A U-Net architecture was selected due to its proven ability to capture both global context and fine-scale features through its encoder-decoder structure combined with skip connections, facilitating detailed reconstructions. The study focuses on two canonical CFD problems: a bump-in-a-duct configuration and a hydrogen combustion chamber. These cases were chosen to encompass both non-reactive and reactive flow regimes and to test the generalization capability of the model under different physical conditions. The network was trained to predict high-resolution pressure, axial velocity, and heat of reaction fields from their corresponding low-resolution versions, thus performing a super-resolution task. Due to the limited size of the available dataset, several regularization techniques were employed to improve model robustness, including Dropout layers and Batch Normalization. Additionally, different data normalization strategies were tested, such as Z-normalization and Min-Max scaling, to enhance training stability and ensure faster convergence. A multi-network approach was also developed to separately handle the prediction of pressure fields and that of axial velocity and heat release rate, mitigating the risk of feature mixing and allowing each network to specialize in learning specific patterns. The datasets were generated using two solvers: a custom Fortran-based code solving the non-dimensional Euler equations, and the commercial CFD software Ansys Fluent for more complex reactive flow simulations. This combination enabled the creation of a multifidelity database, covering a wide range of fluid dynamic behaviors and ensuring better generalization capabilities during training and validation phases. The training process involved careful hyperparameter optimization, early stopping strategies, and loss function selection to prevent overfitting and maximize prediction quality. Quantitative assessments demonstrated prediction errors on the order of a few percent, while qualitative analyses confirmed the ability of the models to accurately reconstruct the flow structures, even in complex reacting flows. The final phase of the project focused on identifying potential future developments of the proposed framework, particularly exploring the capability of the U-Net architecture to address a broader class of problems. One promising direction involves extending the methodology to correct two-dimensional solutions in order to approximate three-dimensional flow fields. This approach consists of training the network with the 2D planes extracted from the full 3D simulations. Furthermore, techniques such as data augmentation could be implemented to improve the model's generalization and robustness. The main contribution of this work includes the dataset generation, the design and implementation of the U-Net architecture, the training framework, and the



performance evaluation. The results highlight the promising role of CNN-based surrogate models in enhancing CFD simulation capabilities and accelerating future engineering workflows by enabling fast and accurate predictions across different fidelity levels.

Introduction

Deep learning approaches are increasingly used to complement conventional CFD [1] solvers, aiming to accelerate simulations, reduce computational costs, and improve data resolution [2]. Convolutional Neural Networks (CNNs), with their strong capability in pattern recognition and feature extraction, have shown potential in reconstructing high-fidelity flow fields from coarse data. This study applies a U-Net-based CNN architecture to perform a super-resolution task on CFD fields. The goal is to recover high-resolution distributions flow field from their corresponding low-resolution counterparts.




Figure 1 : The U-net structure.

Figure 2: Feature maps extracted during training.

The classic U-Net [3] (Fig. 1) follows a symmetric **encoder-decoder** structure with **skip connections**. It consists of two main parts:

- **Contracting path (encoder):** this path captures context through successive applications of *convolutions*, each followed by a non-linear *activation function* and a *max-pooling* operation for downsampling. With each downsampling step, the number of feature channels doubles.
- **Expanding path (decoder):** this path performs *upsampling* of the feature maps using interpolation methods (e.g., nearest-neighbor, or bilinear) followed by convolutional layers. Each step halves the number of channels and concatenates the result with the corresponding feature maps from the contracting path (skip connections). This fusion helps the network localize features accurately.

The network ends with a 1×1 convolution to map the last layer feature vector to the desired number of output classes or regression values. This architecture allows the U-Net to combine coarse contextual information with fine spatial details.

Core Operations in U-Net

The U-Net architecture relies on a sequence of well-established operations in deep learning, we treated many of them in the previous section:

- **Convolution:** each convolutional layer applies learnable filters (typically 3x3) to the input feature maps. At the beginning of training the kernel are randomly weighted. The output of this operation is a set of output channels known as **feature maps**.
- **Max Pooling:** a downsampling operation where a window (commonly 2x2) slides over the input, and only the maximum value within the window is retained.
- **Upsampling:** In the decoder, spatial resolution is restored using learned upsampling, typically with *nearest neighbor* or *bilinear interpolation*, followed by convolutional layers to refine the features.
- **Concatenation (also known as Skip Connections):** Skip connections link encoder and decoder features at equal resolutions, preserving spatial details and aiding gradient flow. They prevent vanishing gradients, enhance training efficiency, and merge low-level and high-level information, enabling precise segmentation even with few labeled samples.

A python framework (Fig. 3) has been written in order to:

- Import CFD scenes from the solver and interpolate them onto a uniform grid.
- Crop a subregion of interest and apply a mask for out-of-domain points if needed.
- Normalize the dataset using different strategies.
- Train the U-net with two approaches:
 1. Run a single U-net model for all fields.
 2. Run separate U-net models for each field.
- Correct the low-fidelity simulation, denormalize the output, and interpolate back to the original grid.




Figure 3 : The workflow of the framework

Two test cases were selected to validate the framework: a non-reactive “bump-in-a-duct” configuration and a reactive hydrogen burner [4], representing a multifidelity scenario that combines Euler or RANS-based simulations.




Figure 4 : Bump - low fidelity mesh (3k nodes)




Figure 5 : Bump - high fidelity mesh (30k nodes)

Methodology

The database was generated using two distinct CFD solvers. The non-reactive dataset derives from a Fortran-based Euler code, while the reactive one comes from Ansys Fluent [5] simulations of a lean premixed hydrogen–air burner. The high-resolution fields were downsampled to generate corresponding low-fidelity maps, producing paired inputs and targets for network training. The U-Net architecture employed consists of an encoder–decoder structure with skip connections linking corresponding layers. These connections preserve spatial information that would otherwise be lost during the downsampling process. The encoder extracts global flow patterns, and the decoder reconstructs localized structures, such as shocks or flame fronts. Regularization was introduced via Dropout and Batch Normalization layers to enhance model generalization. A multi-network configuration was adopted to separately predict flow fields, avoiding feature maps interference.




Figure 6: Hyd. burner - Low fidelity mesh (40k nodes)




Figure 7 : Hyd. burner - high fidelity mesh (200k nodes)

Results and Discussion

Due to the limited size of the dataset, several combinations of batch size, dropout, and hyperparameters were tested, until good results have been reached. The trained networks achieved accurate reconstructions across both test cases. For the bump-in-a-duct case, the predicted wall-pressure distributions closely matched the high-fidelity results (Fig. 8). In the reactive burner case, the U-Net reproduced consistent with the high-fidelity reference solution (Fig. 9). The proposed framework successfully captured both global and local flow features, proving robust despite the dataset limitations. The analysis also highlighted the advantages of the multi-network configurations, which improved stability and accuracy for the different physical quantities.




Figure 8 : Reconstruction of the high-fidelity wall pressure in the bump-in-a-duct configuration.




Figure 9 : Reconstruction of the high-fidelity velocity field for the hydrogen burner case. Radial profiles are shown at three different axial positions.

Conclusions

A Convolutional Neural Network based on the U-Net architecture was developed and applied to the reconstruction of high-resolution CFD fields from low-resolution data under multifidelity conditions. The method effectively recovered detailed pressure, velocity, and heat of reaction (for the specific case of the combustor) distributions across different flow regimes. Future work will focus on extending the approach to three-dimensional field correction and on implementing data augmentation strategies to improve generalization.

References

- [1] O. Ronneberger, P. Fischer, T. Brox, U-Net: Convolutional networks for biomedical image segmentation, Medical Image Computing and Computer-Assisted Intervention (MICCAI), Springer, 2015, pp. 234–241. https://doi.org/10.1007/978-3-319-24574-4_28

- [2] N. Thuerey, K. Weißenow, L. Prantl, X. Hu, Deep learning methods for Reynolds-averaged Navier–Stokes simulations of airfoil flows, *AIAA Journal* 58 (1) (2020) 25–36.
<https://doi.org/10.2514/1.J058291>
- [3] L. Muscarà, M. Cisternino, A. Ferrero, A. Iob, F. Larocca, A comparison of local and global strategies for exploiting field inversion on separated flows at low Reynolds number, *Applied Sciences* 14 (18) (2024) 8382. <https://doi.org/10.3390/app14188382>
- [4] L. Folcarelli et al., Multifidelity modeling of a lean premixed swirl-stabilized hydrogen burner with axial air injection, *AIAA SciTech 2025 Forum*, Paper 0941, 2025.
<https://doi.org/10.2514/6.2025-0941>
- [5] ANSYS Inc., ANSYS Fluent Student, Release 2024 R1, Canonsburg, PA, USA, 2024.
Available at: <https://www.ansys.com/academic/students/ansys-fluent-student>