

Article

Toward the Usage of Deep Learning Surrogate Models in Ground Vehicle Aerodynamics

Benet Eiximeno ^{1,2,*} , Arnau Miró ¹ , Ivette Rodríguez ²  and Oriol Lehmkuhl ¹ 

¹ Barcelona Supercomputing Center, 08034 Barcelona, Spain; arnau.miro@bsc.es (A.M.); oriol.lehmkuhl@bsc.es (O.L.)

² Turbulence and Aerodynamics Research Group, Universitat Politècnica de Catalunya (UPC), 08221 Terrassa, Spain; ivette.rodriguez@upc.edu

* Correspondence: benet.eiximeno@bsc.es

Abstract: This study introduces a deep learning surrogate model designed to predict the evolution of the mean pressure coefficient on the back face of a Windsor body across a range of yaw angles from 2.5° to 10°. Utilizing a variational autoencoder (VAE), the model effectively compresses snapshots of back pressure taken at yaw angles of 2.5°, 5°, and 10° into two latent vectors. These snapshots are derived from wall-modeled large eddy simulations (WMLESs) conducted at a Reynolds number of $Re_L = 2.9 \times 10^6$. The frequencies that dominate the latent vectors correspond closely with those observed in both the drag's temporal evolution and the dynamic mode decomposition. The projection of the mean pressure coefficient to the latent space yields an increasing linear evolution of the two latent variables with the yaw angle. The mean pressure coefficient distribution at a yaw angle of 7.5° is predicted with a mean error of $\bar{e} = 3.13\%$ when compared to the WMLESs results after obtaining the values of the latent space with linear interpolation.

Keywords: surrogate models; deep learning; variational autoencoder; computational fluid dynamics; wall-modeled large eddy simulations; car aerodynamics

MSC: 76G25



Citation: Eiximeno, B.; Miró, A.; Rodríguez, I.; Lehmkuhl, O. Toward the Usage of Deep Learning Surrogate Models in Ground Vehicle Aerodynamics. *Mathematics* **2024**, *12*, 998. <https://doi.org/10.3390/math12070998>

Academic Editor: Jonathan Blackledge

Received: 31 January 2024

Revised: 11 March 2024

Accepted: 25 March 2024

Published: 27 March 2024



Copyright: © 2024 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Road vehicles are constantly in a yawed free stream flow because they are either cornering, receiving a lateral wind gust, or overtaking other vehicles [1]. Hence, the parallel flow assumption is an unrealistic and invalid idealization of the conditions that the vehicle will find on the road. It is known that regardless of the vehicle shape, the drag force increases linearly for yaw angles in the range of $0^\circ \leq \delta \leq 15^\circ$ [2]. This drag augment is completely independent of the zero-yaw drag, thereby making it impossible to extrapolate the performance in crossflow conditions from the parallel flow case [2]. Thus, manufacturers need to assess the aerodynamic behaviors of their cars at several yaw angles to know how the flow changes with the direction of the incident velocity.

The aerodynamic evaluation of a new vehicle entails a combination of computational fluid dynamics simulations and wind tunnel testing [3]. Both methodologies are highly expensive, and their costs limit the amount of geometries and flow conditions that can be tested. Surrogate models can be used to overcome those limitations and obtain a more complete picture of the aerodynamic performance when the direction of the incident velocity changes.

Surrogate models are data-driven computational techniques used in various scientific and engineering fields to approximate complex systems or functions. These models serve as simpler substitutes for both experiments and more computationally intensive processes, thus providing quicker, yet sufficiently accurate results [4]. As they reduce the resource requirements for design exploration, they are mainly utilized to estimate the optimum

product solution or as instrumental tools to evaluate the performance in the initial stages of development [5,6].

There are some studies in the literature that present models that are able to evaluate the forces and moments evolutions with the yaw angle. For instance, Gong et al. [7] built a surrogate model based on the Kriging interpolation technique that finds the optimal wind deflector geometry in a tractor trailer to reduce drag in crosswind situations. Similarly, Ghoreyshi et al. [8] developed a model that is able to predict the forces and moments changes with the Mach number, angle of attack, and side slip angle of an aircraft, and Zhang et al. [9] presented a model to predict the derailment coefficient of a train in cross-flows. However, up to the authors' knowledge, there are no models that can evaluate the changes in the flow magnitudes with the yaw angle.

This work then aims to develop a surrogate model to compute the dependence of the flow field with the direction of the incident velocity. In particular, it is focused on the case of squared rear-end vehicles, because they are the most popular type of automobile. Not only is freight transport done on square-back trucks and lorries, but there is also an important customer interest in square- and estate-back cars, such as sport utility vehicles (SUVs) and station wagons [10,11]. Their growing popularity is based on their maximized cargo space; however, those geometries have been proven as the least aerodynamically efficient. Taking the DrivAer model [12] as an example, the drag coefficient of its estate-back version ($C_D = 0.314$) is significantly higher than the one of the same vehicle with a notchback ($C_D = 0.294$) or fastback ($C_D = 0.265$) rear end [13].

The chosen car model for the study is the Windsor body [14], which has a simplified geometry that mimics a squared rear-end vehicle (represented in Figure 1), and the target flow magnitude to predict with the surrogate model is the pressure on the back face of the automobile. The two reasons behind this choice are that all studies made on square-back cars (see for instance [15–23]) indicate that it is the principal contributor to the drag force and that there is a linear drag increase under crossflows [2] due to a lower base pressure footprint close to the leeward side of the vehicle [24].




Figure 1. Geometry of the Windsor body.

The objective of the surrogate model is to learn the temporal dynamics of the back pressure at the yaw angles of $\delta = [2.5^\circ, 5^\circ \text{ and } 10^\circ]$ and then be able to predict the mean back pressure at any yaw angle inside this range. Having such a fast prediction tool is relevant to understanding the evolution of the drag force and optimizing the vehicle shape to have a lower fuel consumption and greenhouse gases emissions.

The data have been obtained by means of wall-modeled large eddy simulations (WMLESs) at a Reynolds number of $Re_L = U_\infty L / \nu = 2.9 \times 10^6$, where L is the length of the car, U_∞ is the free stream velocity, and ν is the fluid kinematic viscosity. The choice of such a high-fidelity CFD methodology is because the AutoCFD3 workshop [25] proved it as the only computational method able to correctly predict the drag force and the back-pressure

dynamics on the Windsor body. In fact, having such a surrogate model gains importance, as no lower fidelity method can compute the correct mean back-pressure distribution.

Due to the complexity of the full-order model, the surrogate has been built on a reduced space, as suggested by [6,26]. Similar approaches have already been followed in the literature to optimize a UAV geometry [27], to predict the bidimensional flow around an airfoil at different Mach numbers and angles of attack [28], and to model the 2D mean flow at the hub height of a wind turbine given the inflow condition and the number of turbines in the wind farm [29].

In this work, the dimensionality reduction has been conducted by employing deep learning methodologies, as they outperform classical model order reduction techniques as the proper orthogonal decomposition [30], thanks to their ability to capture the nonlinearities of turbulent flow [31]. In particular, unsupervised learning algorithms, i.e., autoencoders (AEs), have shown excellent results for feature extraction in turbulent flows [32–34].

Among the different autoencoder architectures reported in the literature for model order reduction (for a list with recent relevant examples, please refer to Table 1), this work employs a convolutional neural network (CNN) variational autoencoder (VAE). The rationale behind this selection is the VAE capability to map the input data into an informative and smooth distribution of the latent space that allows for the generation of new samples.

Table 1. List of relevant examples of unsupervised learning algorithms for model order reduction in fluid mechanics.

Citation	Methodology
Eivazi et al. [32]	CNN-AE and LSTM for snapshots prediction
Murata et al. [35]	Introduction of modal decomposing CNN-AE
Fukami et al. [36]	Introduction of hierarchical AE for modal decomposition
Zhang [37]	Physics-assimilated CNN-AE
Akkari et al. [38]	VAE for model order reduction
Hines and Bekemeyer [39]	Graph NN AE for surface pressure prediction
Eivazi et al. [40]	beta-VAE for orthogonal modal decomposition
Wang et al. [34]	beta-VAE and transformers for snapshots prediction

The rest of the manuscript is organized as follows: Section 2 provides a description of the mathematical model and numerical implementation used to generate the training and validation data for the surrogate model; Section 3 presents a summary of the flow evolution with the yaw angle in the Windsor body, together with a validation of the simulations with the data from the experiments of Varney et al. [41] at the same Reynolds number; Section 4 presents a description of the implementation and training process of the neural network used for the surrogate model; and, finally, Section 5 shows the capabilities of the surrogate model.

2. Mathematical Description and Numerical Implementation

Wall-modeled large eddy simulations of the Windsor body (Figure 1) have been performed at a Reynolds number of $Re_L = LU_\infty/\nu = 2.9 \times 10^6$ and assuming incompressible flow at the yaw angles $\delta = [2.5^\circ, 5^\circ, 7.5^\circ, \text{ and } 10^\circ]$. The results at $\delta = 7.5^\circ$ have been used to validate the interpolation capabilities of the surrogate model.

Large eddy simulations were performed here by solving the spatially filtered incompressible Navier–Stokes equations:

$$\frac{\partial \bar{u}_i}{\partial x_i} = 0 \quad (1)$$

$$\frac{\partial \bar{u}_i}{\partial t} + \frac{\partial \bar{u}_i \bar{u}_j}{\partial x_j} - \nu \frac{\partial^2 \bar{u}_i}{\partial x_j \partial x_j} + \rho^{-1} \frac{\partial \bar{p}}{\partial x_i} = -\frac{\partial \tau_{ij}}{\partial x_j} \quad (2)$$

where $(\bar{\cdot})$ represents the filtered variables. In the above equations, x_i are the spatial coordinates (or x , y , and z), u_i (or u , v , and w) stands for the velocity components, and p is the pressure. ν and ρ are the kinematic viscosity and the density of the fluid, respectively. The *rhs* term in Equation (2) defines the subgrid stresses; its anisotropic part is given as

$$\mathcal{T}_{ij} - \frac{1}{3}\mathcal{T}_{kk}\delta_{ij} = -2\nu_{sgs}\bar{\mathcal{S}}_{ij} \quad (3)$$

where the large-scale rate-of-strain tensor $\bar{\mathcal{S}}_{ij}$ is evaluated as $\bar{\mathcal{S}}_{ij} = \frac{1}{2}(\bar{g}_{ij} + \bar{g}_{ji})$, with $\bar{g}_{ij} = \partial\bar{u}_i/\partial x_j$ and δ_{ij} being the Kronecker delta. An appropriate expression for the subgrid scale (SGS) viscosity (ν_{sgs}) has to be introduced to close the above equations. In this work, the chosen SGS viscosity model is the local formulation of the integral-length-scale approximation (ILSA) used as in Lehmkuhl et al. [42]

The above equations were numerically solved using SOD2Ds (Spectral high-Order coDe 2 solve partial Differential equations) [43], thereby involving a low-dissipation spectral element method (SEM) code. SOD2Ds are based on a spectral-element version of Galerkin's finite element method continuous model, with a modified version of Guermond's entropy viscosity stabilization [44]. The aliasing effects of the reduced order integration caused by employing SEM integration for convective terms are countered with the skew-symmetric splitting presented by Kennedy and Gruber [45]. The time advancement algorithm is based on an explicit second-order Adams–Bashforth for the convection terms and an implicit second-order Crank–Nicolson for the diffusion terms. The fractional step method [46] is used to solve the coupled velocity pressure system using a conjugate gradient solver as an iterative solver of the final linear system.

The computations have been performed on the domain shown in Figure 2. A uniform velocity profile is fixed at the inlet boundary, with the velocity components depending on the yaw angle, δ , as $[u, v, w]_{\infty} = U_{\infty}[\cos(\delta), \sin(\delta), 0]$. Zero gradients are imposed at the outlet regions, and a slip boundary condition is imposed on the top wall. The back face of the car is treated as no-slip due to the geometric separation in the trailing edges; however, the rest of the body and the ground are considered as slip walls with a wall model implemented following the methodology from Owen et al. [47]. In this case, the Reichardt wall law [48] is used with an exchange location in the 5th node [49].

The cases have been solved on a computational grid made of fourth-order hexahedra, with 125 nodes each. The mesh has been generated following the guidelines of the finest grid provided in the 3rd Automotive CFD Workshop (AutoCFD3) [25], as it showed accurate results for the wall-modeled LES methodology. The resulting grid has a total amount of 135.2 million points, with most of them concentrating on the car and ground surfaces, as well as on the wake of the vehicle. The surface of the car has been refined to ensure that the nondimensional wall normal distance of the first grid point is in the range of $30 \leq y^+ \leq 60$, where $y^+ = u_{\tau}y/\nu$, with the skin friction velocity being $u_{\tau} = \sqrt{\tau_w/\rho}$, and where $\tau_w = \nabla u \cdot n$. A refinement region (depicted as A in Figure 3) extending from $x/L = -0.15$ to $x/L = 5$ and from $y/L = -0.66$ to $y/L = 0.66$ is added in the ground surface to ensure that this requirement is also met in the section that interacts with the flow under the body and with the wake of the car.

The wake of the car has also been specially refined to guarantee a correct resolution of the base pressure. The near-wake refinement (depicted as B in Figure 3) extends from $x/L = 1$ to $x/L = 2.8$ and from $y/L = -0.54$ to $y/L = 0.54$, with a maximum element length of $\Delta = 0.04L$. On the other hand, the far-wake refinement (depicted as C in Figure 3) extends from $x/L = 2.8$ to $x/L = 5.0$ and from $y/L = -0.66$ to $y/L = 0.66$, with a maximum element length of $\Delta = 0.08L$. Both regions grow from the ground ($z/L = 0$) and go up to $z/L = 0.4$.

The grid resolution in the near wake has been verified with an a posteriori analysis of the Kolmogorov length scales $\eta = \sqrt[4]{\nu^3/\varepsilon}$, where ε is the dissipation rate. The resulting ratio between the mesh size, Δ , in the near wake refinement region and the Kolmogorov scale, η , is within $2 \leq \Delta/\eta \leq 5$. Taking into account that the peak in the dissipation

spectrum is around $\Delta/\eta \approx 24$, the present grid resolution is fine enough to be considered a well-resolved LES in the near wake for the analyzed Reynolds number [50].

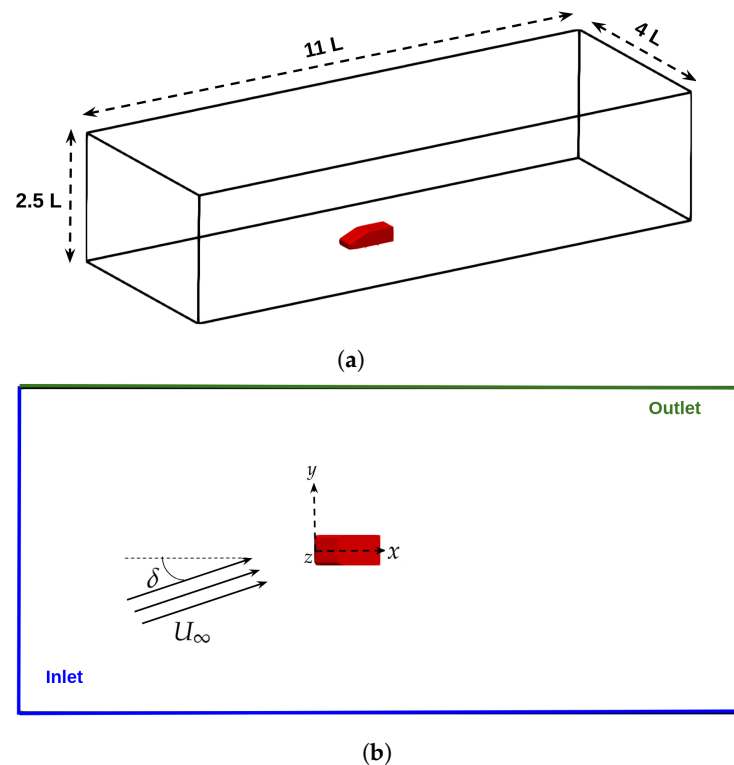


Figure 2. Isometric (a) and top (b) views of the domain used for the numerical simulations.




Figure 3. Schematic view of the three refinement regions of the mesh, with A being the surface refinement on the ground, B being the volume refinement in the near wake, and C being the volume refinement in the far wake.

All simulations have been run along 60 convective time units, $t = L/U_\infty$, once the initial transient has been washed out, thus collecting up to 660 snapshots of the back pressure per studied yaw angle for the training process of the surrogate model.

3. Flow Field Results

Figure 4 represents the instantaneous flow for the four yaw angles using isocontours of the Q criterion [51]. A relevant example of the differences between the four cases is that the A pillar vortex (depicted as A) follows the trajectory of the free stream velocity, and it interacts with a larger part of the upper trailing edge of the body as the yaw angle increases. Another difference that can be appreciated is that the vortex shed from the leeward corner of the bumper (depicted as a B) separates from the body and follows the direction of the free stream velocity at $\delta = 7.5^\circ$ and $\delta = 10^\circ$; however, it is embedded in the wake at $\delta = 2.5^\circ$ and $\delta = 5^\circ$.

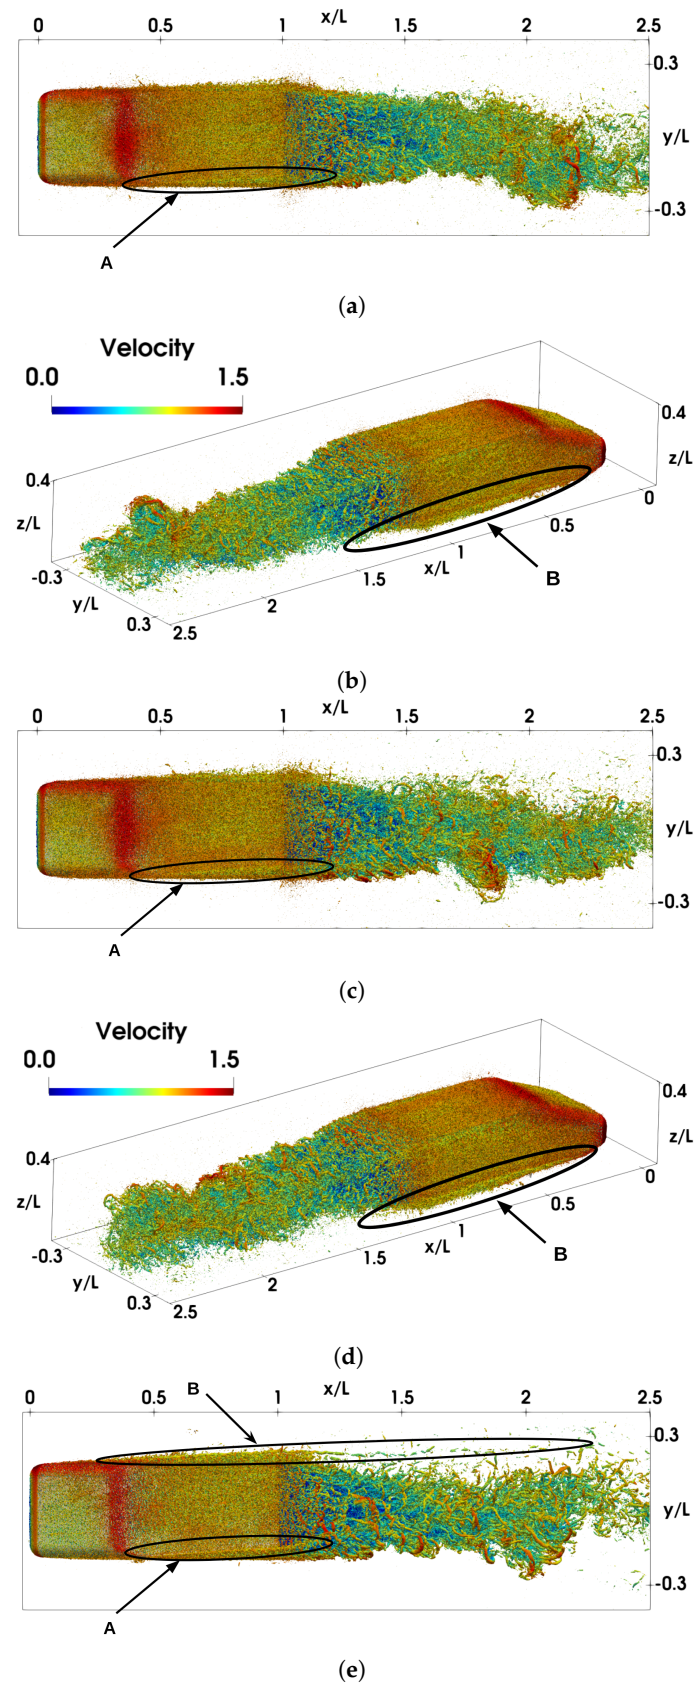


Figure 4. Cont.

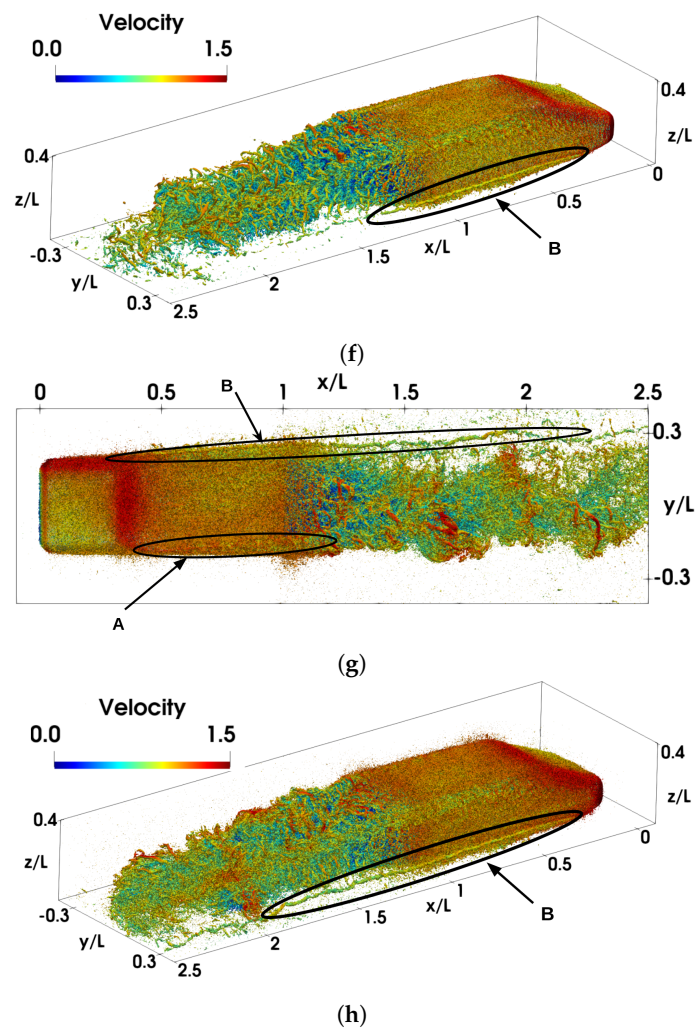


Figure 4. Instantaneous representation of the flow using Q criterion [51] isocontours colored with the velocity magnitude at $Q^* = Q/U_\infty^2 = 500$ for $\delta = 2.5^\circ$ (a,b) and $\delta = 5^\circ$ (c,d). The highlighted areas with A and B represent the windward A pillar vortex and the bumper leeward vortex, respectively. Instantaneous representation of the flow using Q criterion [51] isocontours colored with the velocity magnitude at $Q^* = Q/U_\infty^2 = 500$ for $\delta = 7.5^\circ$ (e,f) and $\delta = 10^\circ$ (g,h). The highlighted areas with A and B represent the windward A pillar vortex and the bumper leeward vortex, respectively.

Apart from affecting the A pillar and bumper vortices, the yaw angle has a significant impact on the recirculation region. In the case of square-back bluff bodies in a yawed free stream flow, the wake is dominated by two hairpin vortices: one on the leeward side and one on the windward side [15]. Figure 5 compares the stream lines on a horizontal plane located at $z/L = 0.186$ for the four simulated yaw angles. It shows that the vortex on the leeward side (depicted as L) dominated the recirculation and gained intensity over the windward vortex (depicted as W) as the yaw angle increased. Moreover, the core of both the vortices and the saddle point of the recirculation (depicted as S) got closer to the car as the yaw angle increased. Figure 6 plots the recirculation length in the symmetry plane of the car at $z/L = 0.186$ and the cross-stream velocity in a line crossing the wake at $x/L = 1.156$ at the same height to confirm the trend identified by Lorite-Díez et al. [52] in a square-back Ahmed body, i.e., the recirculation length decreases with the yaw angle, and the recirculation bubble is deflected towards the leeward side of the wake. In Figure 6, the experimental results from Varney et al. [41] at the same Reynolds number are also plotted. The streamwise velocity was coincident with the experiments inside the recirculation region, and the cross-stream velocity also presented a reasonable agreement with the experimental results.

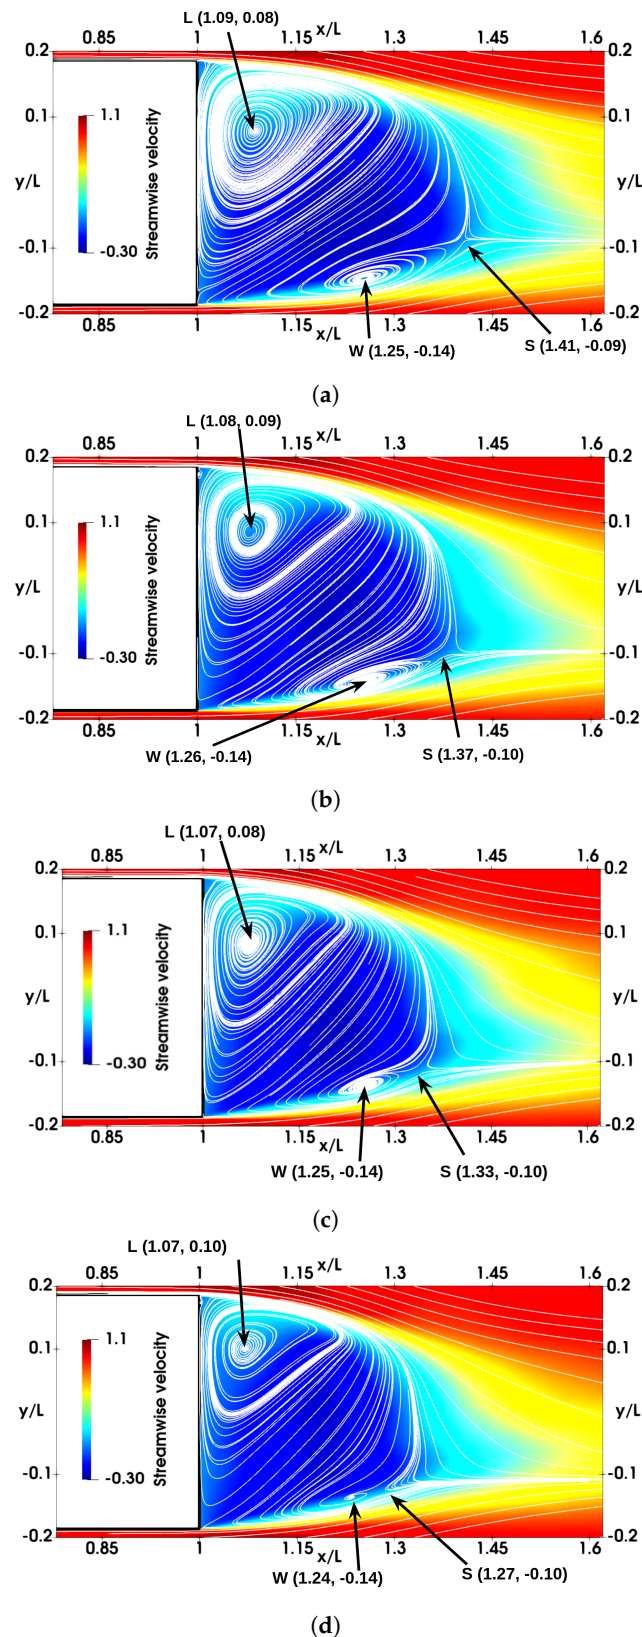


Figure 5. Representation of the recirculation with streamlines of the velocity vector at a horizontal plane at $z/L = 0.186$ for $\delta = 2.5^\circ$ (a), $\delta = 5^\circ$ (b), $\delta = 7.5^\circ$ (c), and $\delta = 10^\circ$ (d). The core of the leeward and windward vortices are depicted as L and W, respectively. The saddle point of the recirculation is depicted as S. The background plane is colored with the streamwise component of the velocity.




Figure 6. Streamwise velocity on a line extending from the back face of the car at the symmetry plane of the car at $z/L = 0.186$ (left) and cross-stream velocity on a line crossing the wake of the car at $x/L = 1.156$ and at $z/L = 0.186$ (right). The data depicted as crosses are those obtained experimentally by Varney et al. [41].

Figure 7 shows that the changes in the recirculation came together with a decrease in the base pressure coefficient, where $C_{pb} = (p - p_\infty) / (0.5\rho U_\infty^2)$ [53]. It is important to see how as the yaw angle augmented, a larger and more intense suction region developed on the leeward side of the vehicle's rear face. This led to an expansion of the area with higher Reynolds stresses from the windward to the leeward side [24]. Figure 8 plots the evolution of the pressure coefficient on the symmetry plane of the vehicle and a horizontal plane at $z/L = 0.249$. The pressure coefficient on the upper and lower faces of the car remained constant along the range of yaw angles studied. The most relevant suction peak was located on the lower face of the car at the transition between the nose and the floor. The upper side presented a wide suction region in the transition between the windscreen and the roof of the car. However, there was an increase in the pressure difference between the windward and leeward sides. The suction of the windward A pillar was reduced, and there was a clear augment in the suction at the leeward A pillar. Both results agree with the data from Varney et al. [41], who also showed that the changes in the yaw angle are too small to have an effect in the vertical direction, but they are relevant in affecting the lateral and longitudinal motions.

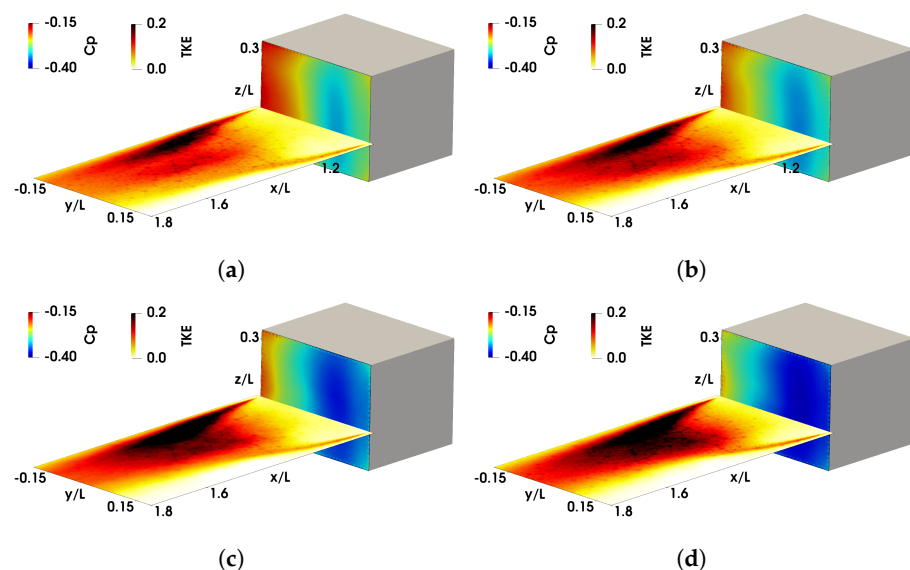


Figure 7. Pressure coefficient on the back face of the car, together with the turbulent kinetic energy on a plane located at $z/L = 0.15$ for $\delta = 2.5^\circ$ (a), $\delta = 5^\circ$ (b), $\delta = 7.5^\circ$ (c), and $\delta = 10^\circ$ (d).

The greater suction in the back face of the car and the bigger pressure difference between the windward and leeward faces led to a linear increase of the drag and side forces [2,54], which is depicted in Figure 9. Both force coefficients agree with the ones obtained experimentally by Varney et al. [41]. The drag coefficient mismatch between their results and the ones computed in the present study relies on the lower blockage ratio of the computational domain (1.07%) compared with the one of the wind tunnel, which increased from 4.5% for $\delta = 2.5^\circ$ to 6.2% for $\delta = 10^\circ$ [55].

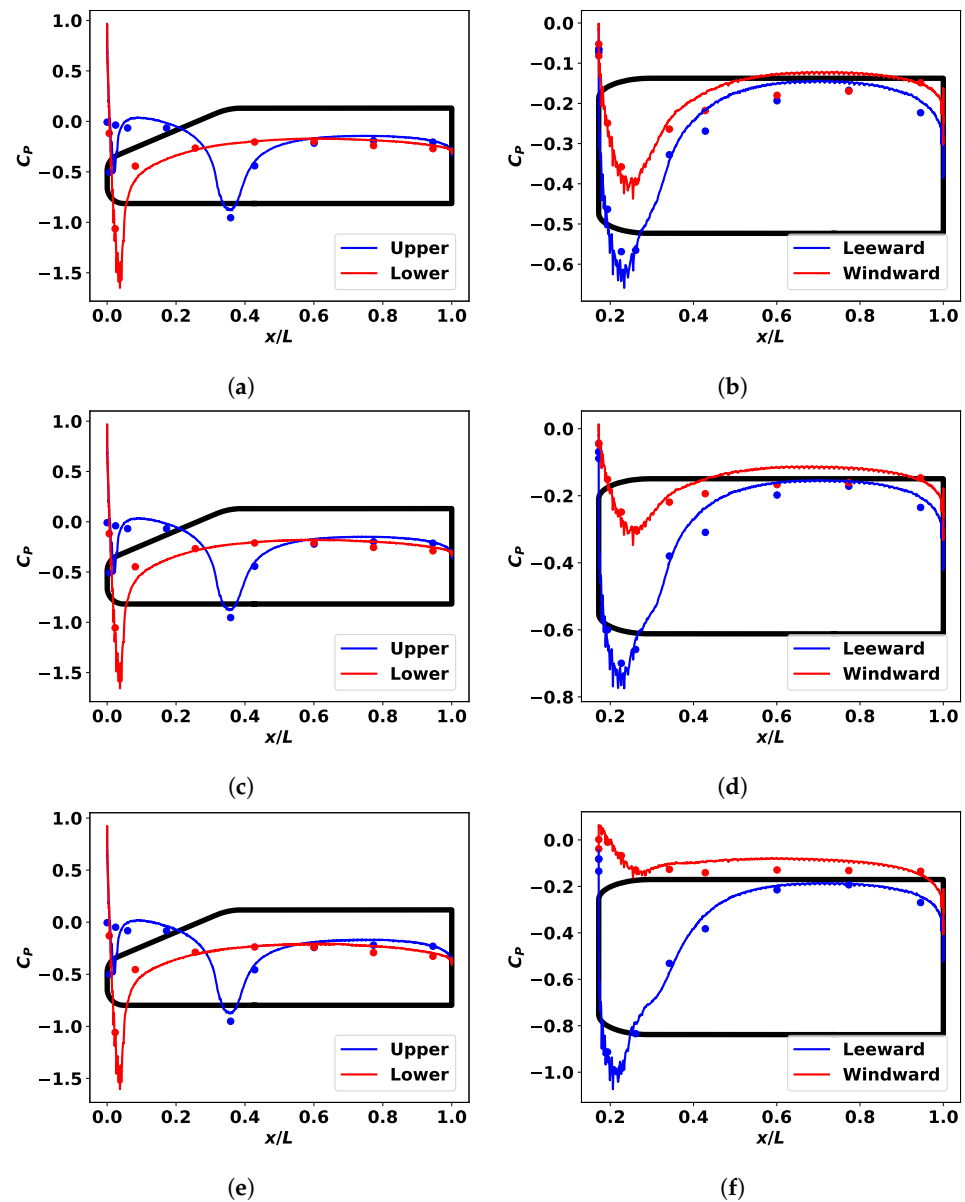


Figure 8. Pressure coefficient value at the symmetry plane of the car for $\delta = 2.5^\circ, 5^\circ, 10^\circ$ (a,c,e) and pressure coefficient on a horizontal plane at $z/L = 0.249$ for the same angles (b,d,f). The data represented in points are those obtained experimentally by Varney et al. [41].




Figure 9. Comparison of the drag, C_D (a), and side C_S (b), force coefficients with the ones obtained experimentally by Varney et al. [41].

4. Deep Learning Surrogate Model Definition

To build the surrogate model, the most important features in the back pressure at $\delta = [2.5^\circ, 5^\circ \text{ and } 10^\circ]$ have been learned by a variational autoencoder [56]. An autoencoder (AE) is a deep neural network (DNN) with an architecture that is suitable for unsupervised feature extraction. Autoencoders are built in two parts: First, an encoder \mathcal{E} maps the input data \mathbf{x} to a low-dimensional latent space \mathbf{z} with \mathbf{d} dimensions and $\mathbf{z} \in \mathbb{R}^d$. Then, a decoder \mathcal{D} projects the low-dimensional data back to the reference space $\tilde{\mathbf{x}}$. The model is trained by optimizing its parameters \mathbf{w} to minimize the reconstruction loss \mathcal{L}_{rec} between the input and the reconstructed data.

Kingma and Welling [56] proposed an evolution of vanilla AEs, known as variational autoencoders (VAEs), to map the input data into an informative and smooth distribution of the latent space. This solution overcomes the limitations of vanilla AEs when generating new data due to their fixed and noninformative latent space representation. In this study, a Gaussian distribution has been used; therefore, the input data are projected through \mathcal{E} to a mean μ and a standard deviation σ . The latent space distribution is sampled by a normally distributed random number ϵ , $\mathbf{z} = \mu + \sigma \odot \epsilon$. The loss function of a variational autoencoder includes a penalization from latent space information with the Kullback–Leibler (KL) divergence, D_{KL} , which is defined as follows:

$$\mathcal{L}(\mathbf{x}) = \mathcal{L}_{rec} - \frac{1}{2} \sum_{i=1}^d \left(1 + \log(\sigma_i^2) - \mu_i^2 - \sigma_i^2 \right) \quad (4)$$

The VAE architecture used in this work is presented in Figure 10. The input data have been formed by snapshots of the pressure coefficient on the back face of the car, with $n_y = 352$ points in the horizontal direction and $n_z = 256$ points in the vertical direction. The baseline architecture was taken from previous works on autoencoders to compress fluid flow data [34,36,40,57]. All of them captured the flow features with five convolutional layers. For the current application, applying a filter with a kernel size that reduces the dimension of the input data yielded more robust results than using pooling layers for dimensionality reduction. Hence, to halve the output data dimension in each layer, a kernel size of (4×4) with a padding of 1 and a (2×2) stride was used. The first layer has a filter size of 16, and doubles at each layer to preserve the patterns in the data. After the last convolutional layer, the extracted features are flattened and fed to a fully connected layer of length $l = 128$ so that the dimension is reduced smoothly. The final step of the encoder is mapping the features onto two fully connected layers in parallel of dimension \mathbf{d} , with one of them representing the mean μ and the other the variance σ of the latent distribution. In the current study, \mathbf{d} was set to $\mathbf{d} = 2$. The sizes of the filter, the fully connected layer, and the latent space were the minimum ones to obtain coherent reconstructions of the mean back pressure coefficient.

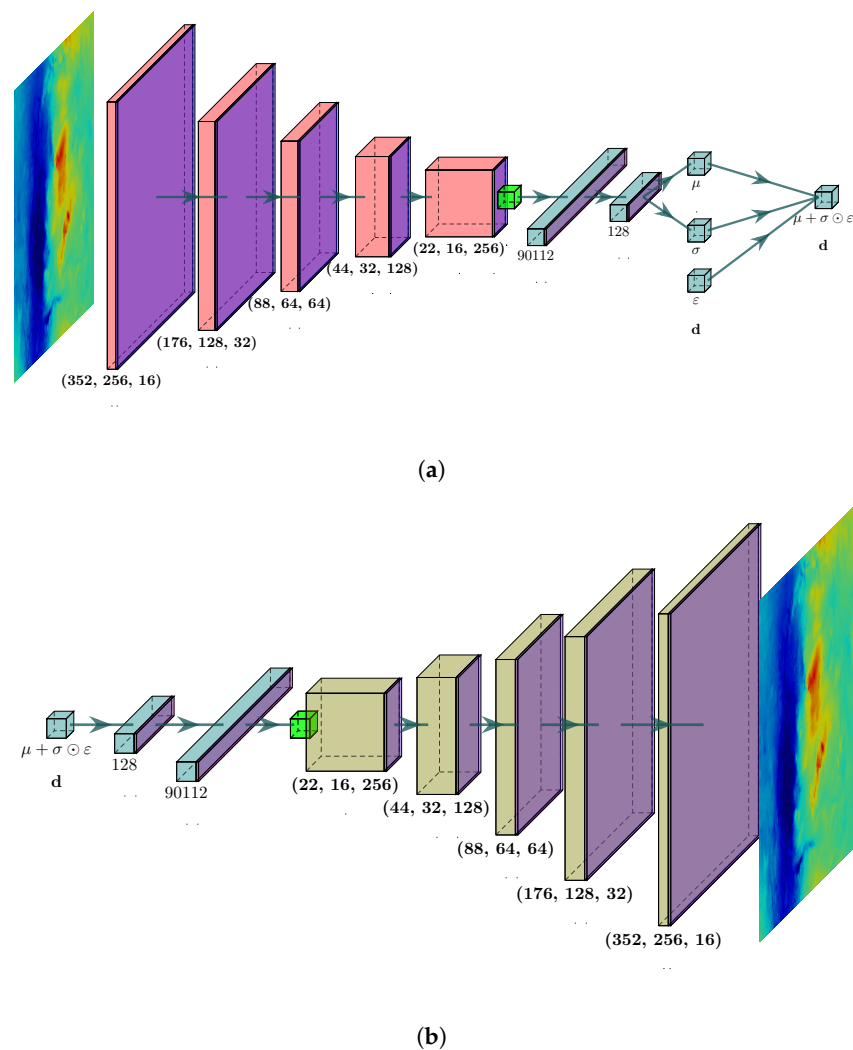


Figure 10. Encoder (a) and decoder (b) architectures. The output size of each layer is indicated under the corresponding block in the format (height × width × channels). The color coding for each layer is as follows: 2D convolution layer, 2D transpose convolution layer, fully connected layer, reshape, and tanh activation function.

The decoder has been designed symmetrically, where fully connected layers are fed from the latent distribution, and then the data are reshaped to the same output of the last convolutional layer of the encoder. Five transposed convolutional layers were used to reconstruct the original data by increasing the spatial dimension progressively. All layers were followed by a batch norm layer [58] before applying the nonlinear activation function, which in this study is the hyperbolic tangent (tanh) function, as it was the best performing one for the studied data. All the layers in the model were initialized with the Xavier uniform scheme [59].

The reconstruction loss \mathcal{L}_{rec} was taken as the mean squared error (MSE), and the model parameters \mathbf{w} were updated using the Adam algorithm [60]. To ensure the convergence of the loss function, the model was trained for 300 epochs, with a learning rate lr scheduled to follow a decay from 5×10^{-4} as

$$lr_i = \frac{lr_{i-1}}{1 + 1 \times 10^{-4} \cdot i} \quad (5)$$

where i is the current epoch number. To prevent \mathbf{w} from overfitting, an early-stopping scheduler was set to the total loss $\mathcal{L}(\mathbf{x})$ when it no longer improved [61].

A percentage of 80% of the total 1980 collected snapshots for $\delta = [2.5^\circ, 5^\circ \text{ and } 10^\circ]$ have been used to train the autoencoder, while the remaining 20% were reserved as a validation set to assess possible overfitting of the model to the seen data. To ensure that the model did not have a bias to any of the angles, the training and validation datasets have been created with the same number of snapshots from each yaw angle case. Apart from that, the mean value of the back pressure at the three angles was also added to the training set to ensure that the VAE could reproduce it correctly. The reader is referred to Algorithm 1 for a detailed description of the training process of the variational autoencoder. The evolution of the training and validation losses during the training process of the VAE is depicted in Figure 11. It is important to point out that each wall-modeled LES simulation took 960 h in 2300 CPU Intel Xeon Platinum 8160, and the training of the model took 3 h and 51 min on a GPU AMD Radeon Instinct M150. Once the model is trained, the inference to obtain the mean base pressure distribution at another yaw angle is immediate.

Algorithm 1: Training of a CNN variational autoencoder

```

1 Initialize parameters  $\mathbf{w}$  with Xavier uniform scheme [59];
2  $LearningRate = 5 \times 10^{-4}$ ;
3  $iepoth = 0$ ;
4 repeat
5   Update learning rate (Equation (5));
6   for  $batch$  in  $DataTraining$  do
7      $\mu, \sigma, \varepsilon \leftarrow \mathcal{E}(batch)$ ;           // Encode the data batch to a Gaussian distribution
8      $\mathbf{z} \leftarrow \mu + \sigma \odot \varepsilon$ ;           // Compute the latent vectors
9      $recBatch \leftarrow \mathcal{D}(\mathbf{z})$ ;           // Reconstruct the batch using the decoder
10     $\mathcal{L}_{rec} \leftarrow MSE(batch, recBatch)$ ;
11    Compute the training loss  $\mathcal{L}_{tr}$  (Equation (4))
12    Update the parameters  $\mathbf{w}$  using the Adam algorithm [60]
13  end
14   $\bar{\mathcal{L}}_{tr} \leftarrow$  mean loss of all training batches;
15  for  $batch$  in  $DataValidation$  do
16     $\mu, \sigma, \varepsilon \leftarrow \mathcal{E}(batch)$ 
17     $\mathbf{z} \leftarrow \mu + \sigma \odot \varepsilon$ 
18     $recBatch \leftarrow \mathcal{D}(\mathbf{z})$ 
19     $\mathcal{L}_{rec} \leftarrow MSE(batch, recBatch)$ 
20    Compute the validation loss  $\mathcal{L}_{va}$  (Equation (4))
21  end
22   $\bar{\mathcal{L}}_{va} \leftarrow$  mean loss of all validation batches;
23  if  $EarlyStopp(\bar{\mathcal{L}}_{va}, prevLoss, \bar{\mathcal{L}}_{tr})$  then
24    break;           // Early stopping to prevent overfitting [61]
25  end
26   $prevLoss \leftarrow \bar{\mathcal{L}}_{tr}$ ;
27   $iepoth \leftarrow iepoth + 1$ ;
28 until  $iepoth = finalEpoch$ ;

```

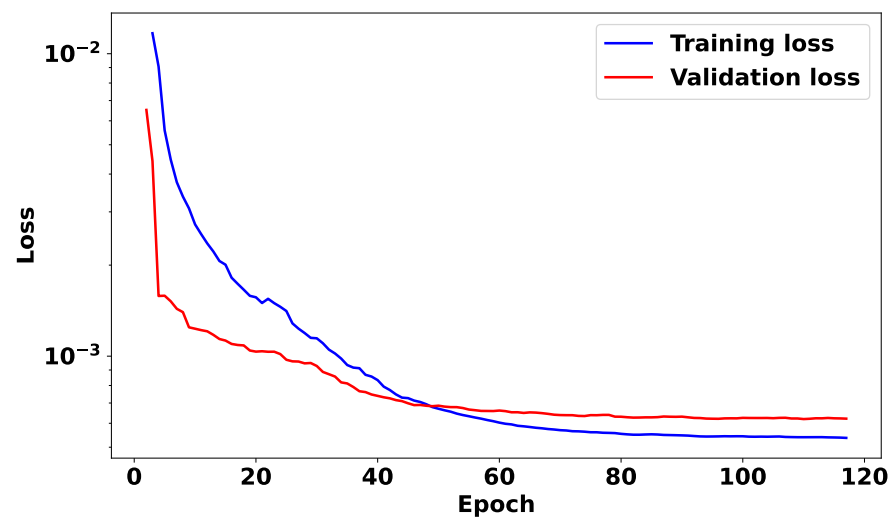


Figure 11. Evolution of the training and validation losses. Early stopper was activated at the 117th epoch.

5. Deep Learning Surrogate Model Properties and Capabilities

The evaluation of the trained variational autoencoder as a surrogate model for the yaw angle dependence of the back pressure of the Windsor body entails three different parts: reproducing the training data, interpolating the mean back pressure at another yaw angle ($\delta = 7.5^\circ$), and extracting the temporal dynamics of a dataset.

Figure 12 illustrates the reconstruction of the mean back-pressure coefficient at $\delta = [2.5^\circ, 5^\circ, 10^\circ]$, together with the original field. The mean relative error on all points is defined as

$$e = \frac{|C_{PbVAE} - C_{Pb}|}{|C_{Pb}|} \quad (6)$$

where the original and the reconstruction was lower than $\bar{e} \leq 1\%$ for all three angles. The error was concentrated on two symmetric vertical lines close to $y/L = -0.09$ and $y/L = 0.09$. Those lines are seen in Figure 12 as small discontinuities in the pressure field resulting from the encoding and decoding process in the VAE.

The value of the two latent vectors obtained when projecting the mean back pressure through the encoder had a strong linear dependence on the yaw angle (Figure 13). Both of the latent vectors had a linear increasing tendency with the yaw angle as the one seen for the drag and side force coefficients (Figure 9). Thus, it is possible to interpolate the value of the latent vectors to any angle between $2.5^\circ \leq \delta \leq 10^\circ$ with a linear regression. The correlation between the regression and the original data was calculated at $R^2 = 0.9510$ for the first latent vector \mathbf{z}_1 and $R^2 = 0.997$ for the second latent vector \mathbf{z}_2 . Moreover, the reconstruction of the base pressure distribution was unaffected when the regression results were taken as the \mathbf{z}_i for $\delta = 2.5^\circ$, $\delta = 5^\circ$, and $\delta = 10^\circ$.

To illustrate the capabilities of the surrogate model, in Figure 14, the reconstruction of the mean back pressure after decoding the \mathbf{z}_1 and \mathbf{z}_2 values computed with the linear regressions for $\delta = 7.5^\circ$ is shown. The mean relative error between the interpolated field and the WMLES results was $\bar{e} = 3.13\%$. This result proves that the model was not overfitted to the training angles and could generate new samples for the unseen angles. The base pressure distribution obtained with the surrogate model yielded a contribution to the total drag force of $\bar{C}_{Db} = 0.284$. This result is fully in agreement with that contribution obtained with the baseline CFD simulation at $\bar{C}_{Db} = 0.276$. Both results, i.e., the back-pressure distribution and base pressure, are fully satisfactory, especially when considering the low cost of the evaluation of the model for a new yaw angle and the high accuracy with respect to the baseline CFD results.

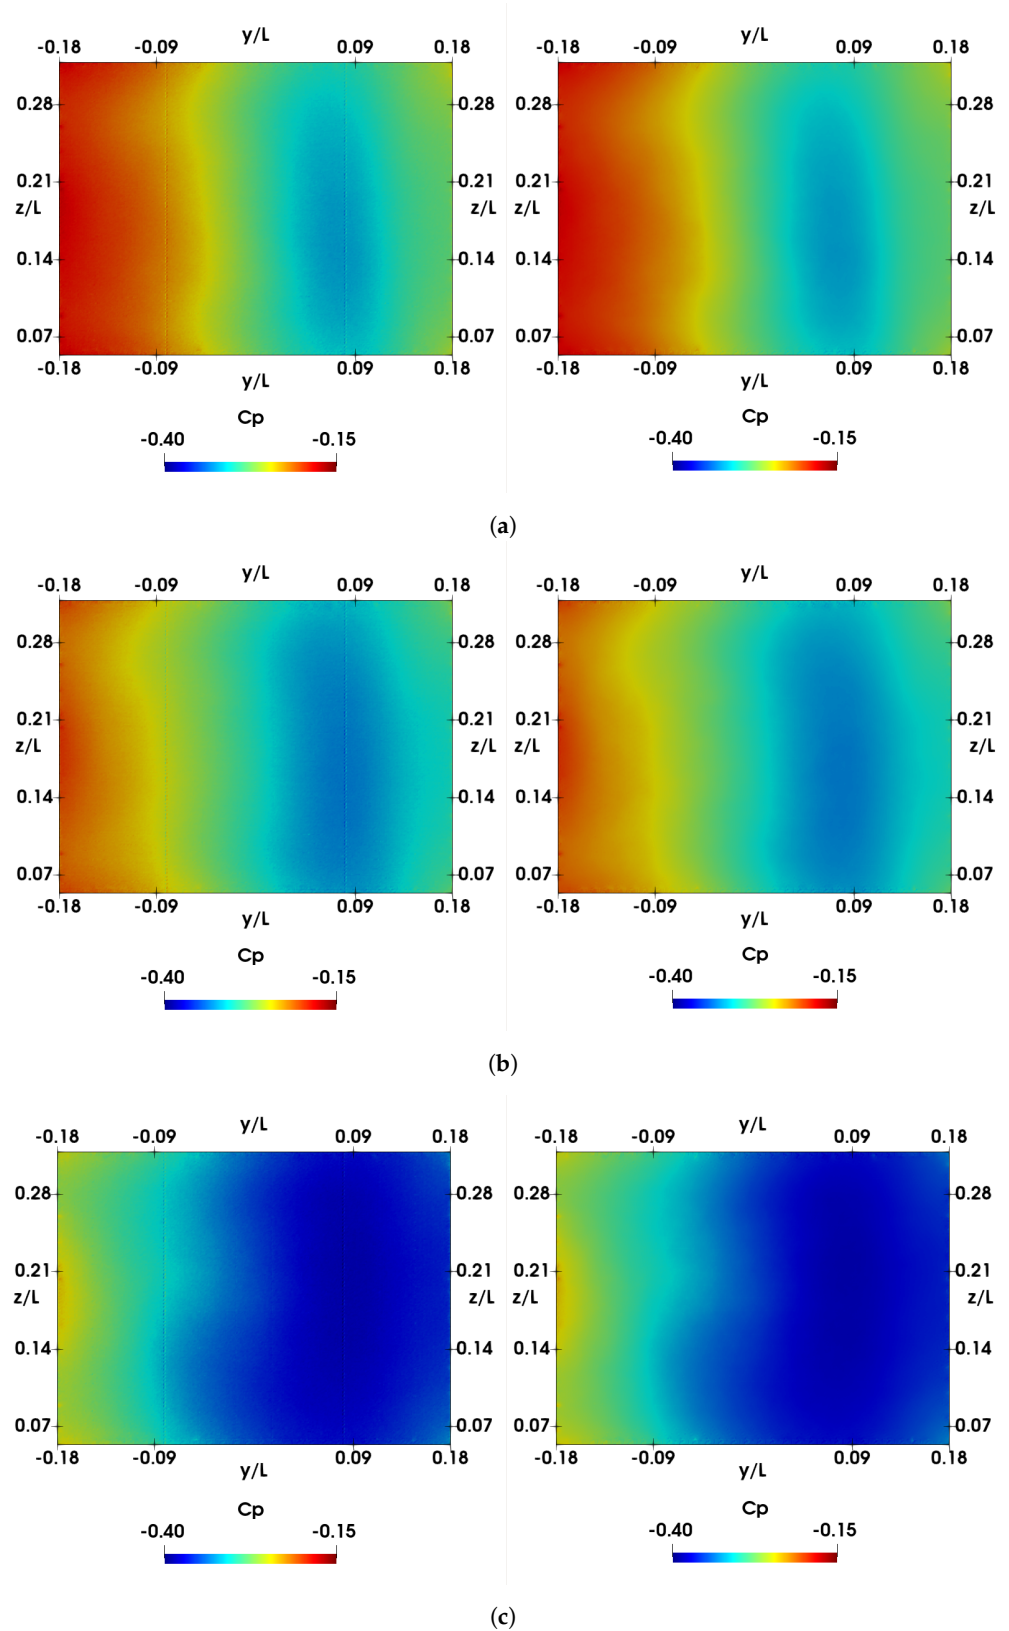


Figure 12. Comparison between the reconstruction (left) and original (right) mean pressure coefficient at $\delta = 2.5^\circ$ (a), $\delta = 5^\circ$ (b), and $\delta = 10^\circ$ (c).




Figure 13. Linear regression of the mean values of each latent vector. $Z_1[\delta] = 0.211\delta + 4.67$ with a correlation of $R^2 = 0.9510$, and $Z_2[\delta] = 0.373\delta + 0.374$ with a correlation of $R^2 = 0.997$.




Figure 14. Comparison between the predicted C_p with the variational autoencoder (left) and the wall-modeled LES results (right) at $\delta = 7.5^\circ$

After showing that the evolution of the latent space with the yaw angle agreed with the C_D and C_S tendency with the yaw angle, the temporal evolution of the latent vectors was assessed. For this study, the 660 snapshots of the back pressure at each yaw angle $\delta = [2.5^\circ, 5^\circ \text{ and } 10^\circ]$ were projected through the encoder while keeping its time correlation. Thus, a coherent time evolution of the latent space was obtained. Figures 15–17 present the frequency spectra of the temporal evolution of the latent space for $\delta = 2.5^\circ$, $\delta = 5^\circ$, and $\delta = 10^\circ$, respectively. All spectra are plotted together with the frequency–amplitude distribution of the dynamic mode decomposition (DMD) [62] analysis of the back pressure snapshots.

DMD is a frequency-based modal decomposition that yields modes that characterize fluid elements with prominent dynamic activity within the observed data set (the reader is referred to Appendix A for the details of the DMD implementation used in this work). Note that the higher PSD maxima in the latent vector spectra also correspond to relevant maxima in the DMD frequency–amplitude distribution, thereby giving evidence of the flow dynamics preservation in the latent space.




Figure 15. Frequency spectra of the latent space at $\delta = 2.5^\circ$ (a) and frequency–amplitude distribution of the DMD analysis (b). Frequencies have been nondimensionalized with the height of the vehicle, and the spectra have been shifted for clarity. The dominant DMD frequencies are highlighted in red.




Figure 16. Frequency spectra of the latent space at $\delta = 5^\circ$ (a) and frequency–amplitude distribution of the DMD analysis (b). Frequencies have been nondimensionalized with the height of the vehicle, and the spectra have been shifted for clarity. The dominant DMD frequencies are highlighted in red.




Figure 17. Frequency spectra of the latent space at $\delta = 10^\circ$ (a) and frequency–amplitude distribution of the DMD analysis (b). Frequencies have been nondimensionalized with the height of the vehicle, and the spectra have been shifted for clarity. The dominant DMD frequencies are highlighted in red.

For the case of $\delta = 2.5^\circ$ (Figure 15), the DMD mode with the highest amplitude happened at $fH/U_\infty = 0.068$ and matches the frequency with the maximum PSD in the first latent vector z_1 and the frequency with the second-highest PSD value of z_2 . Volpe et al. [19] identified a similar low frequency ($0.07 \leq fH/U_\infty \leq 0.08$) due to the recirculation bubble impingement in the wake of a square-back Ahmed body under parallel flow conditions. This motion is driven by the interaction of the vortices shed by the vertical and horizontal walls and is still relevant at small yaw angles [15].

In addition to this main frequency, the two additional relative maxima in the DMD frequency–amplitude distribution, $fH/U_\infty = 0.095$ and $fH/U_\infty = 0.121$, had their footprint in the latent space. The most energetic fluctuation of z_2 happened at $fH/U_\infty = 0.091$, and the second most energetic fluctuation of z_1 happened at $fH/U_\infty = 0.115$. The latter frequency is similar to that reported by Booyesen et al. [15] for the vortex shedding from the vertical walls of an Ahmed body ($0.12 \leq fH/U_\infty \leq 0.14$).

A similar trend is appreciated for the case of $\delta = 5^\circ$ (Figure 16). In this case, the frequency of the most relevant DMD mode happened at $fH/U_\infty = 0.107$ as the PSD peak in the second latent vector. The most energetic fluctuation of the latent vector z_1 , $fH/U_\infty = 0.060$ also coincided with a relevant maximum in the spectrum of the second latent vector and the DMD frequency–amplitude distribution.

Finally, for $\delta = 10^\circ$ (Figure 17), the DMD mode with a higher amplitude, $fH/U_\infty = 0.130$, gained relevance over the rest of the modes, and both latent vectors had the maximum energy peak at this frequency. This agrees with the results from Booyesen et al. [15], who identified that the energy of the flow oscillations increases with the yaw angle in square-back bluff bodies. This nondimensional frequency matches the one found by Volpe et al. [19] as the shedding frequency from the lateral walls of a square-back Ahmed body, $fH/U_\infty = 0.130$.

In all three cases, the results from the latent space and the DMD agree with the experimental characterizations of Booyesen et al. [15] and Volpe et al. [19] for a square-back Ahmed body at similar Reynolds numbers. Moreover, the identified frequencies also play a

relevant role in the spectra of the drag force fluctuations (Figure 18), as the energy peaks occurred at the same frequencies as in the DMD and the latent vectors.




Figure 18. Spectra of the drag force for all the training yaw angles. Frequencies have been nondimensionalized with the height of the vehicle, and spectra have been shifted for clarity.

These comparisons in the frequency domain yield the conclusion that the spectra of the latent space can have similar interpretations to the results from linear decompositions of the DMD, thus meaning that variational autoencoders are not only able to reproduce and generate new samples of a fluid flow but also that the physical phenomena seen in the reference space have an identifiable footprint on the latent space. DMD could also be used for the interpolation of the data, but it would be necessary to find the evolution of the spatial correlations with the yaw angles instead of working only with two scalar variables.

It is expected that this model could be extended to a broader yaw angle range if simulation data are available. One should also anticipate a linear evolution of the latent variables in the yaw angle range where the drag coefficient increases linearly, but it is yet to be demonstrated that a change in the tendency of the C_D leads to a change in the dependence of the latent variables with the yaw angle.

The current surrogate model might be generalized to other flow conditions, e.g., different Reynolds numbers, if additional back pressure data were provided. For Reynolds numbers higher than $Re_L \geq 2.9 \times 10^6$, it should be possible to obtain the base pressure coefficient from the current model, as the drag force coefficient is invariant in this range [63]. However, for its use at lower Reynolds numbers, where the back pressure is more Reynolds-dependent, transfer learning might be used in order to retrain the model at the

6. Summary and Conclusions

In the present study, a probabilistic deep neural network architecture based on a variational autoencoder [56] has been proposed as a surrogate model for the base pressure coefficient changes with the yaw angle. In particular, the work focused on the Windsor body, which is a simplified square-back car geometry. Instantaneous snapshots of the back pressure were encoded through five convolutional layers to two latent vectors. Each layer halves the dimensions of the input field while doubling the number of filters. The decoding process followed a symmetric architecture (Figure 10).

The training snapshots have been obtained with wall-modeled large eddy simulations at $Re_L = 2.9 \times 10^6$ for the yaw angles of $\delta = [2.5^\circ, 5^\circ, 10^\circ]$. The simulations present the tendencies previously identified in the literature [15,52,55] regarding the intensification of the hairpin vortex located on the leeward side of the wake (Figure 5), as well as the suction increase with the yaw angle in the leeward side of the back face (Figure 7). Moreover, the results show good agreement with the experimental data obtained by Varney et al. [41] (Figures 6 and 8) in terms of both the wake velocity and the pressure coefficient on the body.

The compression of the mean base pressure at $\delta = [2.5^\circ, 5^\circ, 10^\circ]$ yielded a linear evolution of the two latent variables with the yaw angle (Figure 13). This finding enabled their interpolation in the range of $2.5^\circ \leq \delta \leq 10^\circ$ with linear regression and decoding the base pressure distribution from the interpolated latent variables. The pointwise mean error between the predicted reconstruction and the WMLES results at $\delta = 7.5^\circ$ was $\bar{e} \leq 3.13\%$ (Figure 14).

In addition to presenting a linear growth with the yaw angle, when a set of time-correlated snapshots were encoded, the two latent vectors had a time evolution that was coherent with the temporal sequence of the drag force and with the results obtained from the dynamic mode decomposition analysis [62].

The three dominant frequencies resulting from the DMD analysis for $\delta = 2.5^\circ$ were $fH/U_\infty = 0.068, 0.095$ and 0.121 . All three of them played a relevant role in the temporal evolution of the latent vectors, as the most energetic peak of the \mathbf{z}_1 spectrum occurred at $fH/U_\infty = 0.067$ and the second most relevant peak at $fH/U_\infty = 0.115$. In the case of \mathbf{z}_2 , its peaks occurred at $fH/U_\infty = 0.091$ and $fH/U_\infty = 0.067$. In the case of $\delta = 5^\circ$, the DMD reported two dominant modes at $fH/U_\infty = 0.107$ and 0.060 , while the spectrum \mathbf{z}_1 had its maximum at $fH/U_\infty = 0.060$, and \mathbf{z}_2 had its peak value at $fH/U_\infty = 0.107$. Finally, at $\delta = 10^\circ$, the most relevant DMD modes occurred at $fH/U_\infty = 0.130$ and 0.083 . In this case, both latent vectors had the energy peak at $fH/U_\infty = 0.130$, and \mathbf{z}_1 had a second peak at $fH/U_\infty = 0.087$.

All in all, this work gives additional evidence of the capabilities of variational autoencoders for the surrogate modeling of turbulent flows [34,40], thereby showing their ability regarding data compression on a highly turbulent case, which is relevant for the automotive industry. It also provides clear proof that the latent space is governed by the phenomena seen in the reference space. This work encourages the application of similar techniques for flow prediction in the wake of the body and the usage of transfer learning methods to extrapolate this model to different flow conditions with respect to changes in the Reynolds number.

Author Contributions: Conceptualization, B.E. and O.L.; methodology, B.E., O.L., and I.R.; software, O.L., B.E., and A.M.; data curation, B.E.; formal analysis, B.E., O.L., and A.M.; investigation, all authors; writing—original draft preparation, B.E., I.R., and A.M.; writing—review and editing, all authors. All authors have read and agreed to the published version of the manuscript.

Funding: The research leading to this work has been partially funded by the European Project NextSim, which has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 956104 and cofounded by the Spanish Agencia Estatal de Investigación (AEI) under grant agreement PCI2021-121962. This work was partially financially supported by the Ministerio de Economía, Industria, y Competitividad, as well as the Secretaría de Estado de Investigación, Desarrollo, e Innovación in Spain (refs PID2020-116937RB-C21 and PID2020-116937RB-C22). B. Eiximeno's work was funded by a contract from the Subprograma de Ayudas Predoctorales given by the Ministerio de Ciencia e Innovación (PRE2021-096927). Oriol Lehmkuhl has been partially supported by a Ramon y Cajal postdoctoral contract (Ref: RYC2018-025949-I). The authors acknowledge the support of the Departament de Recerca i Universitats de la Generalitat de Catalunya through the research group Large-scale Computational Fluid Dynamics (Code: 2021 SGR 00902) and the Turbulence and Aerodynamics Research Group (Code: 2021 SGR 01051).

Data Availability Statement: The data that support the findings of this study are available from the corresponding author upon request.

Acknowledgments: We acknowledge the Barcelona Supercomputing Center for awarding us access to the MareNostrum IV machine based in Barcelona, Spain.

Conflicts of Interest: The authors declare no conflicts of interest.

Appendix A. Dynamic Mode Decomposition Implementation

Dynamic mode decomposition (DMD) [62] is a frequency-based modal decomposition. The modes in DMD are seen as structures that represent a linear tangent estimate of the base flow. These modes characterize fluid elements showcasing prominent dynamic activity within the observed data set. An interesting aspect of DMD modes is that they are part of a dynamic system, thereby indicating that they represent not merely a spatial correlation with a time signal but rather a coherent structure characterized by a specific frequency, amplitude, and damping ratio. DMD operates as a data-centric method, thus disregarding any information other than the flow field data. These data need to be presented in a snapshot sequence denoted by the matrix \mathcal{D}_1^N :

$$\mathcal{D}_1^N = [d_1, d_2, d_3, \dots, d_N] \quad (\text{A1})$$

here, d_i represents the i th flow field. DMD assumes the presence of a linear mapping A linking one field d_i to its subsequent field d_{i+1} :

$$d_{i+1} = Ad_i \quad (\text{A2})$$

When A is presumed to be constant, the sequence of flow fields can be conceptualized as a Krylov sequence:

$$\mathcal{D}_1^N = [d_1, Ad_1, A^2d_1, \dots, A^{N-1}d_1] \quad (\text{A3})$$

In this study, the DMD computations have been conducted using pyLOM [64], which is a high-performance computing tool for modal decompositions already used by Miró et al. [65] and Eiximeno et al. [66] for DMD analysis. In pyLOM, the computation of A follows the algorithm proposed by Tu et al. [67]. Subsequently, the attributes of the dynamical process described by A were determined using eigenvalue decomposition. The concluding phase of dynamic mode decomposition involves calculating the amplitude of each mode, thus employing the method introduced by Jovanovic et al. [68].

References

1. D'Hooge, A.; Palin, R.; Rebbeck, L.; Gargoloff, J.; Duncan, B. Alternative Simulation Methods for Assessing Aerodynamic Drag in Realistic Crosswind. *SAE Int. J. Passeng. Cars Mech. Syst.* **2014**, *7*, 617–625. [CrossRef]
2. Howell, J. Aerodynamic Drag of Passenger Cars at Yaw. *SAE Int. J. Passeng. Cars Mech. Syst.* **2015**, *8*, 306–316. [CrossRef]
3. Zhang, X.; Toet, W.; Zerihan, J. Ground Effect Aerodynamics of Race Cars. *Appl. Mech. Rev.* **2006**, *59*, 33–49. [CrossRef]
4. Sun, G.; Wang, S. A review of the artificial neural network surrogate modeling in aerodynamic design. *Proc. Inst. Mech. Eng. Part G J. Aerosp. Eng.* **2019**, *233*, 5863–5872. [CrossRef]
5. Kuya, Y.; Takeda, K.; Zhang, X.; Forrester, A.I.J. Multifidelity Surrogate Modeling of Experimental and Computational Aerodynamic Data Sets. *ALAA J.* **2011**, *49*, 289–298. [CrossRef]
6. Yondo, R.; Andrés, E.; Valero, E. A review on design of experiments and surrogate models in aircraft real-time and many-query aerodynamic analyses. *Prog. Aerosp. Sci.* **2018**, *96*, 23–61. [CrossRef]
7. Gong, X.; Gu, Z.; Li, Z. Surrogate model for aerodynamic shape optimization of a tractor-trailer in crosswinds. *Proc. Inst. Mech. Eng. Part D J. Automob. Eng.* **2012**, *226*, 1325–1339. [CrossRef]
8. Ghoreyshi, M.; Cummings, R.M. Unsteady aerodynamics modeling for aircraft maneuvers: A new approach using time-dependent surrogate modeling. *Aerosp. Sci. Technol.* **2014**, *39*, 222–242. [CrossRef]
9. Zhang, L.; Li, T.; Zhang, J.; Piao, R. Optimization on the Crosswind Stability of Trains Using Neural Network Surrogate Model. *Chin. J. Mech. Eng.* **2021**, *34*, 86. [CrossRef]
10. Bekker, H. 2023 (January) Europe: Car Sales and Market Analysis. 2023. Available online: <https://www.best-selling-cars.com/europe/2023-january-europe-car-sales-and-market-analysis/> (accessed on 9 January 2024)
11. GoodCarBadCar. 2023 US SUV Sales Figures by Model. 2023. Available online: <https://www.goodcarbadcar.net/2023-us-suv-sales-figures-by-model/> (accessed on 9 January 2024)

12. Heft, A.I.; Indinger, T.; Adams, N.A. Introduction of a New Realistic Generic Car Model for Aerodynamic Investigations. In Proceedings of the SAE 2012 World Congress & Exhibition. Technische Universität München, 2012, Vehicle Aerodynamics, 2012-SP-2333, Detroit, MI, USA, 23–26 April 2012; p. 14. [\[CrossRef\]](#)
13. Experimental and Numerical Investigation of the DrivAer Model. In Proceedings of the Fluids Engineering Division Summer Meeting, Rio Grande, PR, USA, 8–12 July 2012. [\[CrossRef\]](#)
14. Littlewood, R.; Passmore, M. *The Optimization of Roof Trailing Edge Geometry of a Simple Square-Back*; Technical Report 2010-01-0510; SAE International: Warrendale, PA, USA, 2010. [\[CrossRef\]](#)
15. Booyesen, A.; Das, P.; Ghaemi, S. Large-scale 3D-PTV measurement of Ahmed-body wake in crossflow. *Exp. Therm. Fluid Sci.* **2021**, *132*, 617–625. [\[CrossRef\]](#)
16. Fan, Y.; Parezanović, V.; Cadot, O. Wake transitions and steady z-instability of an Ahmed body in varying flow conditions. *J. Fluid Mech.* **2022**, *942*, A22. [\[CrossRef\]](#)
17. Grandemange, M.; Gohlke, M.; Cadot, O. Turbulent wake past a three-dimensional blunt body. Part 1. Global modes and bi-stability. *J. Fluid Mech.* **2013**, *722*, 51–84. [\[CrossRef\]](#)
18. Grandemange, M.; Gohlke, M.; Cadot, O. Turbulent wake past a three-dimensional blunt body. Part 2. Experimental sensitivity analysis. *J. Fluid Mech.* **2014**, *752*, 439–461. [\[CrossRef\]](#)
19. Volpe, R.; Devinant, P.; Kourta, A. Experimental characterization of the unsteady natural wake of the full-scale square back Ahmed body: Flow bi-stability and spectral analysis. *Exp. Fluids* **2015**, *56*, 99. [\[CrossRef\]](#)
20. Bruneau, C.H.; Creusé, E.; Gilliéron, P.; Mortazavi, I. Effect of the vortex dynamics on the drag coefficient of a square back Ahmed body: Application to the flow control. *Eur. J. Mech. Fluids* **2014**, *45*, 1–11. [\[CrossRef\]](#)
21. Haffner, Y.; Li, R.; Meldi, M.; Borée, J. Drag reduction of a square-back bluff body under constant cross-wind conditions using asymmetric shear layer forcing. *Int. J. Heat Fluid Flow* **2022**, *96*, 109003. [\[CrossRef\]](#)
22. Evrard, A.; Cadot, O.; Herbert, V.; Ricot, D.; Vigneron, R.; Détery, J. Fluid force and symmetry breaking modes of a 3D bluff body with a base cavity. *J. Fluids Struct.* **2016**, *61*, 99–114. [\[CrossRef\]](#)
23. Perry, A.K.; Pavia, G.; Passmore, M. Influence of short rear end tapers on the wake of a simplified square-back vehicle: Wake topology and rear drag. *Exp. Fluids* **2016**, *57*, 169. [\[CrossRef\]](#)
24. Li, R.; Borée, J.; Noack, B.R.; Cordier, L.; Harambat, F. Drag reduction mechanisms of a car model at moderate yaw by bi-frequency forcing. *Phys. Rev. Fluids* **2019**, *4*, 034604. [\[CrossRef\]](#)
25. 3rd Edition of the AutoCFD, the Automotive Industry CFD Workshop. Workshop, 2022. Available online: <https://autocfd.eng.ox.ac.uk/> (accessed on 10 November 2022).
26. Bertram, A.; Othmer, C.; Zimmermann, R., Towards Real-time Vehicle Aerodynamic Design via Multi-fidelity Data-driven Reduced Order Modeling. In Proceedings of the 2018 AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, Washington, DC, USA, 8–12 January 2018. [\[CrossRef\]](#)
27. Karali, H.; Demirezen, M.U.; Yukselen, M.A.; Inalhan, G. Design of a Deep Learning Based Nonlinear Aerodynamic Surrogate Model for UAVs. In Proceedings of the AIAA Scitech 2020 Forum, Orlando, FL, USA, 6–10 January 2020. [\[CrossRef\]](#)
28. Dupuis, R.; Jouhaud, J.C.; Sagaut, P. Surrogate Modeling of Aerodynamic Simulations for Multiple Operating Conditions Using Machine Learning. *AIAA J.* **2018**, *3622*–3635. [\[CrossRef\]](#)
29. Zhang, J.; Zhao, X. Machine-Learning-Based Surrogate Modeling of Aerodynamic Flow Around Distributed Structures. *AIAA J.* **2021**, *59*, 868–879. [\[CrossRef\]](#)
30. Lumley, J.L. Rational Approach to Relations between Motions of Differing Scales in Turbulent Flows. *Phys. Fluids* **1981**, *10*, 1405. [\[CrossRef\]](#)
31. Brunton, S.L.; Noack, B.R.; Koumoutsakos, P. Machine Learning for Fluid Mechanics. *Annu. Rev. Fluid Mech.* **2020**, *52*, 477–508. [\[CrossRef\]](#)
32. Eivazi, H.; Veisi, H.; Naderi, M.H.; Esfahanian, V. Deep neural networks for nonlinear model order reduction of unsteady flows. *Phys. Fluids* **2020**, *32*, 105104. [\[CrossRef\]](#)
33. Fukami, K.; Fukagata, K.; Taira, K. Super-resolution reconstruction of turbulent flows with machine learning. *J. Fluid Mech.* **2019**, *870*, 106–120. [\[CrossRef\]](#)
34. Wang, Y.; Solera-Rico, A.; Sanmiguel Vila, C.; Vinuesa, R. Towards optimal beta-variational autoencoders combined with transformers for reduced-order modelling of turbulent flows. *Int. J. Heat Fluid Flow* **2024**, *105*, 109254. [\[CrossRef\]](#)
35. Murata, T.; Fukami, K.; Fukagata, K. Nonlinear mode decomposition with convolutional neural networks for fluid dynamics. *J. Fluid Mech.* **2019**, *882*, A13. [\[CrossRef\]](#)
36. Fukami, K.; Nakamura, T.; Fukagata, K. Convolutional neural network based hierarchical autoencoder for nonlinear mode decomposition of fluid field data. *Phys. Fluids* **2020**, *32*, 095110. [\[CrossRef\]](#)
37. Zhang, B. Nonlinear mode decomposition via physics-assimilated convolutional autoencoder for unsteady flows over an airfoil. *Phys. Fluids* **2023**, *5*, 095115-1–095115-17. [\[CrossRef\]](#)
38. Akkari, N.; Casenave, F.; Hachem, E.; Ryckelynck, D. A Bayesian Nonlinear Reduced Order Modeling Using Variational AutoEncoders. *Fluids* **2022**, *7*, 334. [\[CrossRef\]](#)
39. Hines, D.; Bekemeyer, P. Graph neural networks for the prediction of aircraft surface pressure distributions. *Aerosp. Sci. Technol.* **2023**, *137*, 108268. [\[CrossRef\]](#)

40. Eivazi, H.; Martínez, S.L.C.; Hoyas, S.; Vinuesa, R. Towards extraction of orthogonal and parsimonious non-linear modes from turbulent flows. *Expert Syst. Appl.* **2022**, *202*, 117038. [\[CrossRef\]](#)
41. Varney, M.; Pavia, G.; Passmore, M.; Crickmore, C. Windsor Model Experimental Aerodynamic Dataset. 2020. [\[CrossRef\]](#)
42. Lehmkuhl, O.; Piomelli, U.; Houzeaux, G. On the extension of the integral length-scale approximation model to complex geometries. *Int. J. Heat Fluid Flow* **2019**, *78*, 108422. [\[CrossRef\]](#)
43. Gasparino, L.; Spiga, F.; Lehmkuhl, O. SOD2D: A GPU-enabled Spectral Finite Elements Method for compressible scale-resolving simulations. *Comput. Phys. Commun.* **2024**, *297*, 109067. [\[CrossRef\]](#)
44. Guermond, J.L.; Pasquetti, R.; Popov, B. Entropy viscosity method for nonlinear conservation laws. *J. Comput. Phys.* **2011**, *230*, 4248–4267. [\[CrossRef\]](#)
45. Kennedy, C.A.; Gruber, A. Reduced aliasing formulations of the convective terms within the Navier-Stokes equations for a compressible fluid. *J. Comput. Phys.* **2008**, *227*, 1676–1700. [\[CrossRef\]](#)
46. Codina, R. Pressure Stability in Fractional Step Finite Element Methods for Incompressible Flows. *J. Comput. Phys.* **2001**, *170*, 112–140. [\[CrossRef\]](#)
47. Owen, H.; Chrysokentis, G.; Avila, M.; Mira, D.; Houzeaux, G.; Borrell, R.; Cajas, J.C.; Lehmkuhl, O. Wall-modeled large-eddy simulation in a finite element framework. *Int. J. Numer. Methods Fluids* **2020**, *92*, 20–37. [\[CrossRef\]](#)
48. Reichardt, H. Vollständige Darstellung der turbulenten Geschwindigkeitsverteilung in glatten Leitungen. *ZAMM-J. Appl. Math. Mech.* **1951**, *31*, 208–219. [\[CrossRef\]](#)
49. Lehmkuhl, O.; Park, G.; Bose, S.; Moin, P. Large-eddy simulation of practical aeronautical flows at stall conditions. In *2018 Summer Program, Center for Turbulence Research*; Stanford University: Stanford, CA, USA, 2018; pp. 87–96.
50. Pope, S. *Turbulent Flows*; Cambridge University Press: Cambridge, UK, 2000. [\[CrossRef\]](#)
51. Jeong, J.; Hussain, F. On the identification of a vortex. *J. Fluid Mech.* **1995**, *285*, 69–94. [\[CrossRef\]](#)
52. Lorite-Díez, M.; Jiménez-González, J.; Pastur, L.; Cadot, O.; Martínez-Bazán, C. Drag Reduction on a Three-Dimensional Blunt Body with Different Rear Cavities under Cross-Wind Conditions. *J. Wind Eng. Ind. Aerodyn.* **2020**, *200*, 104145. [\[CrossRef\]](#)
53. Roshko, A. Experiments on the flow past a circular cylinder at very high Reynolds number. *J. Fluid Mech.* **1961**, *10*, 345–356. [\[CrossRef\]](#)
54. Mercker, E. A blockage correction for automotive testing in a wind tunnel with closed test section. *J. Wind Eng. Ind. Aerodyn.* **1986**, *22*, 149–167. [\[CrossRef\]](#)
55. Varney, M.; Passmore, M.; Gaylard, A. Parametric Study of Asymmetric Side Tapering in Constant Cross Wind Conditions. *SAE Int. J. Passeng. Cars Mech. Syst.* **2018**, *11*, 213–224. [\[CrossRef\]](#)
56. Kingma, D.P.; Welling, M. Auto-Encoding Variational Bayes. *arXiv* **2013**, arXiv:1312.6114.
57. Arcucci, R.; Xiao, D.; Fang, F.; Navon, I.M.; Wu, P.; Pain, C.C.; Guo, Y.K. A reduced order with data assimilation model: Theory and practice. *Comput. Fluids* **2023**, *257*, 105862. [\[CrossRef\]](#)
58. Ioffe, S.; Szegedy, C. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. *arXiv* **2015**, arXiv:1502.03167.
59. Glorot, X.; Bengio, Y. Understanding the Difficulty of Training Deep Feedforward Neural Networks. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, PMLR, Chia Laguna Resort, Sardinia, Italy, 13–15 May 2010; Volume 9, pp. 249–256.
60. Kingma, D.P.; Ba, J. Adam: A Method for Stochastic Optimization. *arXiv* **2017**, arXiv:1412.6980.
61. Bai, Y.; Yang, E.; Han, B.; Yang, Y.; Li, J.; Mao, Y.; Niu, G.; Liu, T. Understanding and Improving Early Stopping for Learning with Noisy Labels. *arXiv* **2021**, arXiv:2106.15853. [\[CrossRef\]](#)
62. Schmid, P.J. Dynamic mode decomposition of numerical and experimental data. *J. Fluid Mech.* **2010**, *656*, 5–28. [\[CrossRef\]](#)
63. Cooper, K.R. The Effect of Front-Edge Rounding and Rear-Edge Shaping on the Aerodynamic Drag of Bluff Vehicles in Ground Proximity. *SAE Trans.* **1985**, *94*, 727–757. [\[CrossRef\]](#)
64. Eiximeno, B.; Begiashvili, B.; Miro, A.; Valero, E.; Lehmkuhl, O. pyLOM: Low order modelling in python. 2024. [\[CrossRef\]](#)
65. Miró, A.; Eiximeno, B.; Rodríguez, I.; Lehmkuhl, O. Self-Induced Large-Scale Motions in a Three-Dimensional Diffuser. In *Flow, Turbulence and Combustion*; Hemisphere: New York, NY, USA, 2023; pp. 1–18. [\[CrossRef\]](#)
66. Eiximeno, B.; Tur-Mongé, C.; Lehmkuhl, O.; Rodríguez, I. Hybrid Computation of the Aerodynamic Noise Radiated by the Wake of a Subsonic Cylinder. *Fluids* **2023**, *8*, 236. [\[CrossRef\]](#)
67. Tu, J.H.; Rowley, C.W.; Luchtenburg, D.M.; Brunton, S.L.; Kutz, J.N. On dynamic mode decomposition: Theory and applications. *J. Comput. Dyn.* **2014**, *1*, 391–421. [\[CrossRef\]](#)
68. Jovanović, M.R.; Schmid, P.J.; Nichols, J.W. Sparsity-promoting dynamic mode decomposition. *Phys. Fluids* **2014**, *26*, 024103. [\[CrossRef\]](#)

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.