The 9th European Congress on Computational Methods in Applied Sciences and Engineering
ECCOMAS Congress 2024
3–7 June 2024, Lisboa, Portugal

LATENT REGRESSION OF TRANSONIC WING PRESSURE DISTRIBUTIONS

Víctor Francés-Belda[1], Alberto Solera-Rico[2,3], Javier Nieto-Centenero[1,4], Esther Andrés[1], Carlos Sanmiguel Vila[2,3], and Rodrigo Castellanos[3,2,1]

[1] Theoretical and Computational Aerodynamics Branch, Flight Physics Department
Spanish National Institute for Aerospace Technology (INTA)
Torrejón de Ardoz, Spain
e-mail: vfranbel@inta.es, jniecen@inta.es, eandper@inta.es; www.inta.es

[2] Subdirectorate General of Terrestrial Systems
Spanish National Institute for Aerospace Technology (INTA)
San Martín de la Vega, Spain
e-mail: asolric@inta.es, csanvil@inta.es; www.inta.es

[3] Department of Aerospace Engineering
Universidad Carlos III de Madrid
Avda. Universidad, 30. Leganés, 28911, Spain.
e-mail: rcastell@ing.uc3m.es, www.uc3m.es

[4] Escuela Técnica Superior de Ingeniería Aeronáutica y del Espacio (ETSIAE)
Universidad Politécnica de Madrid (UPM)
Madrid 28223, Spain
www.upm.es

Key words: Autoencoder, Surrogate Modeling, Machine Learning, Aerodynamics

Summary. Accurate prediction of aerodynamic pressure distributions is essential for efficient aircraft design, but the computational demands of high-fidelity computational fluid dynamics (CFD) simulations are a significant challenge. Recent advances in machine learning have introduced surrogate models that combine dimensionality reduction with regression techniques to address such a challenge. We propose a surrogate model that integrates β-Variational Autoencoders (β-VAEs) with Gaussian Process Regression (GPR) to predict pressure distributions on a wing under transonic flight conditions. The β-VAE model effectively reduces the dimensionality of complex aerodynamic data while preserving critical flow features. We investigate the impact of the latent space dimension and the hyperparameter β on the performance and interpretability of the surrogate model. The resulting latent space is used to train a GPR model that maps flight conditions, such as Mach number and angle of attack, to the latent space coordinates. This approach enables accurate predictions of aerodynamic pressure distributions over a wide range of flight conditions. Our findings demonstrate that β-VAEs offer a robust and efficient solution for aerodynamic surrogate modeling, reducing the number of required CFD simulations.

1 INTRODUCTION

Accurate prediction of aerodynamic coefficients remains a critical challenge in aeronautics, particularly due to the computational expense associated with high-fidelity Computational Fluid Dynamics (CFD) simulations [1]. To address this matter, there is growing interest in surrogate modeling approaches that offer a computationally efficient alternative [2].

Traditional aerodynamic surrogate models often employ dimensionality reduction techniques such as Principal Component Analysis (PCA) or Proper Orthogonal Decomposition (POD) to distill high-dimensional data into a manageable number of dimensions. Although effective in many cases, these linear methods struggle to capture the inherent nonlinear complexities of aerodynamic flows, particularly in transonic regimes where sharp discontinuities such as shock waves occur [3, 4].

To overcome these limitations, manifold learning techniques such as Isomap and Local Linear Embedding (LLE) have been explored, aiming to identify the manifold on which the input data reside. These methods have shown the potential to improve surrogate models by capturing nonlinear structures, though they require additional regression algorithms to decode the low-dimensional representation to the original space [5, 6, 7]. Despite their promise, manifold learning remains underutilized in aeronautics. More recently, the advances in neural networks provide a promising framework for developing surrogate models that address the limitations of previous approaches [8, 9, e.g.]. In particular, auto-encoder (AE) architectures have emerged as powerful tools for managing large-scale aerodynamic data, offering nonlinear transformations that reduce dimensionality while preserving critical flow features [10, 11]. Among these, Variational Autoencoders (VAEs) have gained attention for their ability to encode fluid dynamics into low-dimensional latent spaces with a probabilistic approach [12, 13]. VAEs are a Machine Learning (ML) architecture that learns a probabilistic representation of the input data by encoding them into a lower-dimensional latent space and then reconstructing it [14]. VAEs offer a probabilistic approach to reduce dimensionality, providing a more refined latent space representation than PCA or POD. However, a key challenge with VAEs is the lack of orthogonality in the latent space, which can lead to entangled features, complicating interpretation.

A notable variant of VAEs, the β-Variational Autoencoder (β-VAE), introduces a regularization term that balances reconstruction accuracy with latent space disentanglement, leading to more interpretable and physically meaningful representations [15, 16]. Recent studies have demonstrated the potential of β-VAEs in capturing the most energetic modes of turbulent flows, indicating their suitability for developing robust and interpretable surrogate models [17, 18].

This paper focuses on leveraging β-VAEs for surrogate modeling of aerodynamic pressure distributions. By using β-VAEs, we aim to provide a more effective and interpretable method for reducing the dimensionality of aerodynamic data. Then, through the mapping of physical labels with the latent space coordinates using a Gaussian process regression (GPR) model, we build a complete pipeline for predicting pressure fields, thus enhancing the accuracy and efficiency of aerodynamic predictions while accommodating the nonlinear nature of flow phenomena.

2 DATABASE

This study utilizes a dataset from the GARTEUR AD/AG60 research project, featuring Reynolds-Averaged Navier-Stokes (RANS) simulations of the XRF1 aircraft model, developed by Airbus™ to demonstrate advanced technologies for long-range, wide-body aircraft. The data set

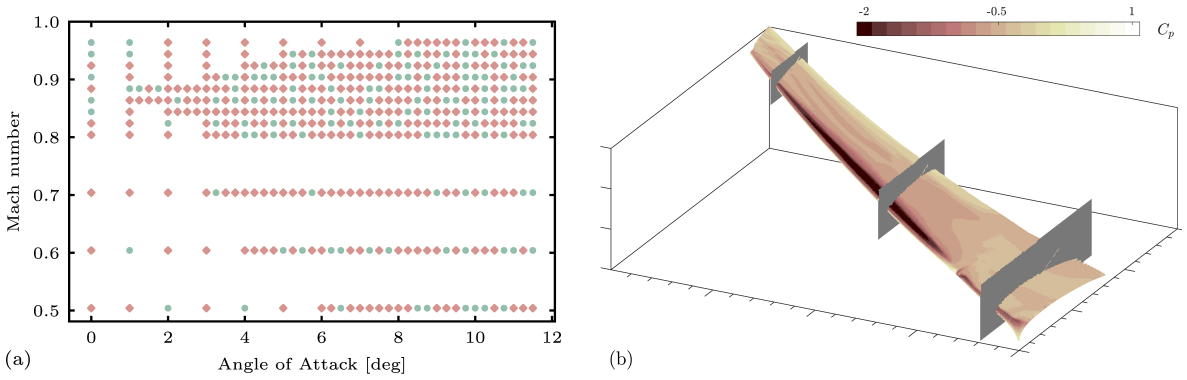

Figure 1: (a) Flight envelope with the tested flight conditions split into *training* set (◆) and *test* set (●). (b) Example of C_p distribution on the wing for $(\alpha, M) = (9.25°, 0.7)$. Three planes are highlighted at wingspan percentages $\eta = 0.1, 0.5, 0.9$ (▬).

consists of 435 simulations computed with the DLR TAU solver[19], focused on the distributions of the pressure coefficient (C_p) on the wing under various flight conditions. The simulations maintain a fixed Reynolds number ($Re = 2.5 \times 10^7$), with Mach numbers ranging from 0.5 to 0.96 and angles of attack (α) from 0° to 11.5° (see Figure 1a). The focus is on the upper surface of the wing, where significant aerodynamic effects occur under a fixed load factor of $n_z = 2.5g$, simulating extreme conditions that lead to notable wing deformation.

The dataset splits into a *training* set (305 cases, 70%) and a *test* set (130 cases, 30%). The training set is used to develop the β-VAE and GPR models, while the test set evaluates their predictive performance. The test cases cover a range of complexities, including scenarios of low $(\alpha_1, M_1 = 2.0°, 0.504)$, intermediate $(\alpha_2, M_2 = 5.75°, 0.704)$, and high $(\alpha_3, M_3 = 9.75°, 0.904)$ aerodynamic challenges, such as shockwaves on the upper wing surface.

Each C_p distribution is represented on an unstructured mesh with $q = 49,574$ points, making it a high-dimensional data point in $\mathbb{R}^q$. The encoder network reduces this dimensionality by encoding the data into a low-dimensional latent space, which the decoder then reconstructs to its original dimensionality. Once trained, the GPR model is used to map the flight condition parameters (the couples $[\alpha, M]$) with the latent space, enabling the prediction of untested flight conditions. Hereinafter, we use C_p, $\tilde{C}_p$, and $\hat{C}_p$ to denote individual samples, β-VAE reconstructions, and surrogate predictions, respectively.

3 FRAMEWORK

This section outlines the methodologies used to develop the proposed surrogate model. We begin by detailing the concept and architecture of the β -VAE, followed by a description of the GPR model used to map flight conditions to the latent space coordinates. Finally, we present the overall structure of the surrogate model, integrating these techniques to predict pressure distributions based on given flight conditions.

3.1 β-variational autoencoder

VAEs [14] are designed to derive meaningful, low-dimensional representations from high-dimensional data. They achieve this by compressing the input data into a lower-dimensional latent space through an encoder network and then reconstructing the input from this latent space using a decoder network. Let z_i represent the i-th component of the latent vector $\mathbf{Z} \in \mathcal{Z}$, where $\mathcal{Z} \subset \mathbb{R}^d$ is the latent space with dimensionality d. The encoder maps input data $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^q$ to this latent space, whereas the decoder reconstructs the data from the latent space.

VAEs include probabilistic modeling by forcing the latent variables to follow a predefined prior distribution, usually a standard Gaussian distribution. This feature provides some control over the properties of the latent space. Thus, each input observation is modeled in terms of a mean $\boldsymbol{\mu}$ and a standard deviation $\boldsymbol{\sigma}$. Considering this, the VAE loss function not only accounts for the reconstruction loss but also the divergence between the learned and prior distributions using the Kullback–Leibler divergence loss $\mathcal{L}_{KL}$. Being $\tilde{\mathbf{x}}$ denotes the VAE reconstruction, the loss expressions reads

$$\mathcal{L}(\mathbf{x}) = \underbrace{\|\mathbf{x} - \tilde{\mathbf{x}}\|_2^2}_{\mathcal{L}_{rec}} \underbrace{- \frac{1}{2} \sum_{i=1}^{d} \left(1 + \log(\sigma_i^2) - \mu_i^2 - \sigma_i^2\right)}_{\mathcal{L}_{KL}}. \tag{1}$$

β-VAEs, as proposed by [15], introduce a hyperparameter β to control the trade-off between reconstruction accuracy and the disentanglement of latent variables. Higher β values promote more disentangled representations at the cost of increased reconstruction error:

$$\mathcal{L}(\mathbf{x}) = \mathcal{L}_{rec} - \beta \mathcal{L}_{KL}. \tag{2}$$

In this study, β-VAEs are employed to map pressure distribution fields, leveraging their enhanced nonlinear dimensionality reduction capabilities compared to classical linear methods.

3.2 Gaussian Process Regression

The surrogate developed in this work uses GPR to model the relationship between flight conditions and the latent variables derived from the β-VAE. GPR is well-suited for this task due to its ability to perform nonlinear regression with high accuracy and flexibility. For a detailed mathematical analysis, readers are referred to comprehensive literature on the topic [20, 21, 22].

In the context of this work, flight conditions are used as input data to predict latent space variables. The kernel choice for GPR combines a linear kernel and a Matérn 3/2 kernel. The former captures direct linear relationships, while the latter addresses nonlinearities, providing a comprehensive model for complex dependencies in the data. This combination enhances the model's predictive performance, allowing it to effectively map flight conditions to the latent space and ensure robust predictions.

3.3 Surrogate modeling for efficient aerodynamic data prediction

The proposed surrogate model combines the specified ML techniques to establish a solid pipeline to predict pressure fields based on flight conditions. The schematic of the data flow through the β-VAE+GPR model is depicted in Figure 2.

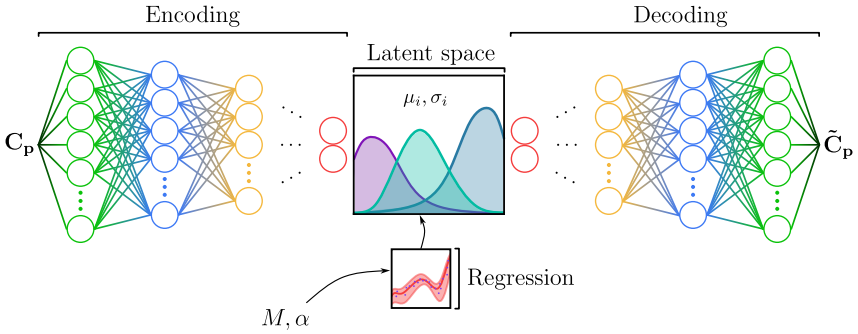

Figure 2: Schematic of the model workflow. Pressure fields are encoded into a latent space, computing the associated mean and standard deviation. The latent vectors are then reconstructed using the decoder, yielding the β-VAE reconstruction. Once trained, the regression model links the flight condition parameters with the latent space coordinates.

Table 1: $\beta-$VAE characteristics.

Optimizer	Adam
Layers type	Fully connected
Activation type	ELU
Hidden layers	10
Encoder	1024-512-256-128-64
Decoder	64-128-256-512-1024
Input size	49,574
No. of parameters	102,974,122

During the training stage, pressure distributions are fed to the encoder network to obtain their low-dimensional representation. This yields a mean and a standard deviation per each latent dimension. Then, the decoder takes these parameters and transforms them into the original pressure fields, aiming to minimize Equation 2. When this phase is completed, the parameters of the networks are frozen and the GPR model is trained to discover a relationship between the flight conditions and the means, resulting in a direct pipeline for C_p prediction.

The β-VAE architecture, detailed in Table 1, uses a multilayer perceptron (MLP) for both the encoder and decoder due to the absence of mesh connectivity in the dataset. This implies a large amount of trainable parameters, considering the vast input size of the data.

4 MODEL ANALYSIS

Although β-VAEs are effective for dimensionality reduction into latent spaces, optimizing their performance involves several challenges. These include selecting the ideal architecture for the encoder and decoder, determining the appropriate hyperparameter β, and choosing the latent space dimension d. Given that our dataset is solely based on flight conditions represented by a tuple of two design variables ($[\alpha, M]$), setting the latent space dimension to $d = 2$ simplifies the regression from parameter space to latent space. This choice is in line with the findings of [11], who noted that a β-VAE with $d = 16$ revealed that only two latent variables were

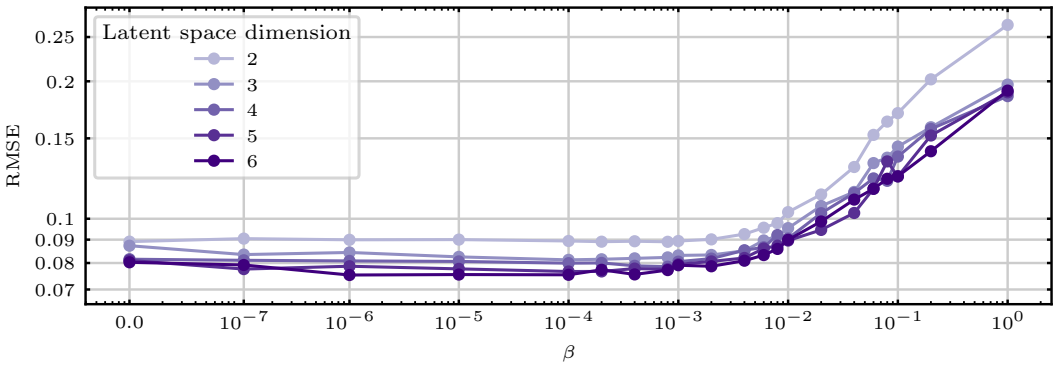

Figure 3: Reconstruction loss for β-VAE model depending on latent space dimension and β hyperparameter for the testing dataset. Root mean square error is computed between original C_p and reconstructed $\tilde{C}_p$ through the β-VAE.

significantly correlated with M and α. Hence, we explore latent space dimensions in the range $d \in [2, 6]$ to determine the appropriate dimensionality, and β values from 0 to 1 to assess model performance. Small values of β are crucial for maintaining acceptable reconstruction accuracy, achieving poor regularization. The upper limit, $\beta = 1$, corresponds to a classic VAE.

Figure 3 shows the Root Mean Square Error (RMSE) as a function of β for different latent space dimensions. The RMSE generally increases with larger β, indicating reduced reconstruction accuracy as expected in Equation 2.

The latent space dimension impacts the reconstruction error. For lower β values, the RMSE remains stable in all dimensions but increases significantly with larger β. This indicates that lower values of β help stabilize the model's performance, leading to similar results.

Further analysis of the latent space, visualized in Figure 4, reveals that the latent variables correlate well with flight conditions. For example, a value of $\beta = 0.008$ shows a clear correlation between M and α with the latent variables (increasing directions from bottom to top, and from left to right, respectively), demonstrating that the β-VAE learns to capture relevant aerodynamic parameters since flight conditions are not used during training. This correlation strongly depends on the latent space dimension, suggesting that lower values tend to yield clearer relationships.

The choice of d does not significantly affect RMSE for lower values of β. Although higher dimensions capture more information, $d = 2$ remains sufficient, aligning with the relevant features captured. This encourages us to set the latent space to $d = 2$ from now, seizing the simple subsequent regression that results as well. For β, small values yield better performance, with $\beta \in [0, 10^{-3}]$ providing the lowest RMSE and a reasonable latent space distribution. As β approaches 1, the reconstruction accuracy decreases due to excessive regularization.

4.1 Latent space prediction

To predict the latent variables from flight conditions accurately, the regression model must effectively map the parameter space to the latent space. The GPR model, using the kernels defined in subsection 3.2, captures the relationship between the parameter space and the latent space (see Figure 4). For clarity, we will refer to the latent space provided by the β-VAE and visualized in Figure 4 as the *reference* latent space, which the GPR model aims to predict.

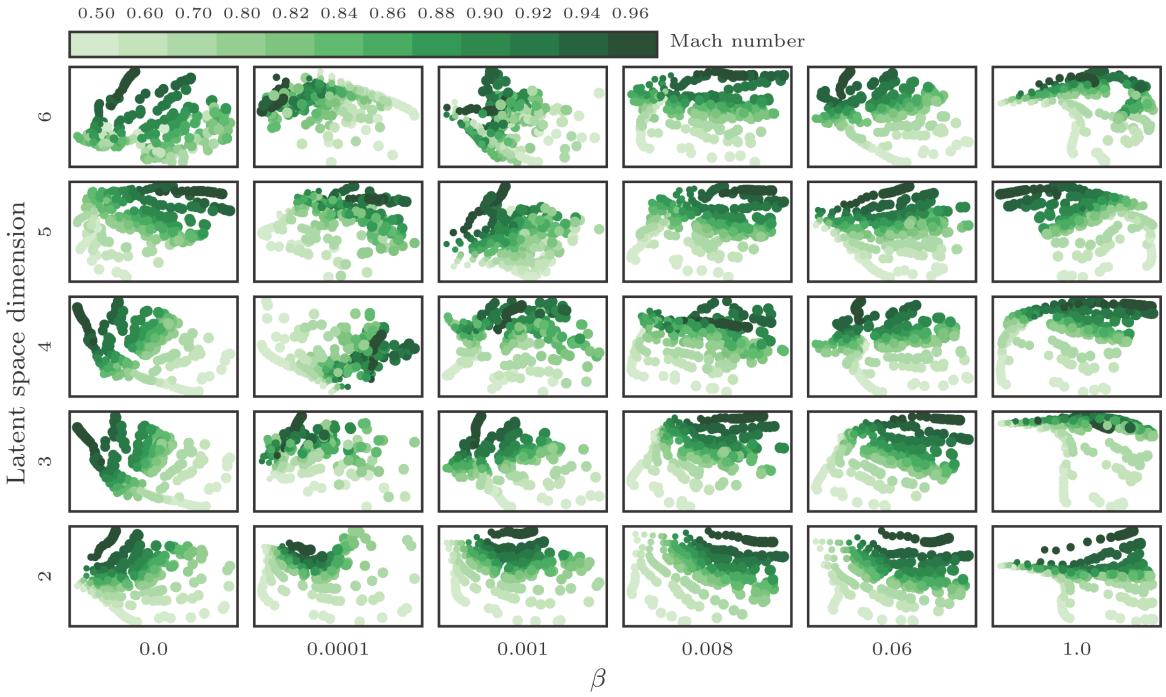

Figure 4: First two variables of the latent space depending on its dimension and β hyperparameter for β-VAE model. Value ranges are adjusted to be within the interval from -1 to 1. Latent spaces with more than two dimensions show the relationship between the first two modes that accumulate more information. Flight conditions are characterized by Mach number (point color) and angle of attack (point size).

Figure 5 illustrates the regression performance on the test set, comparing the reference latent space with the predicted one. Arrows indicate the displacement of the predicted latent variables for the 10 cases with the highest prediction errors. Overall, the regressor performs well, with most predictions aligning closely with the reference latent space points, demonstrating its ability to effectively map the parameter space to the latent space.

The regressor tends to perform less accurately at points near or on the boundaries of the parameter space, especially at high angles of attack (the right side of the latent space). This behavior suggests that the regressor struggles more in these regions, potentially due to the increased complexity and nonlinearity of aerodynamic responses, such as flow separation or shockwave displacement. Additionally, these boundary data points in the latent space have fewer neighbors, reducing the local information available to the regressor and making accurate predictions more challenging.

The regressor performance is quantified with the Normalized Root Mean Square Error (NRMSE) between the reference latent space and the regression predictions for both training and testing datasets ($\text{NRMSE}(\mu, \hat{\mu}) = \|\mu - \hat{\mu}\|_2 / \|\mu - \bar{\mu}\|_2$). The NRMSE metric quantifies the prediction error relative to the overall size of the latent space, which allows for comparison across different latent spaces regardless of β. For a data sample in the latent space μ, its mean value $\bar{\mu}$, and its prediction $\hat{\mu}$, the NRMSE is defined as:

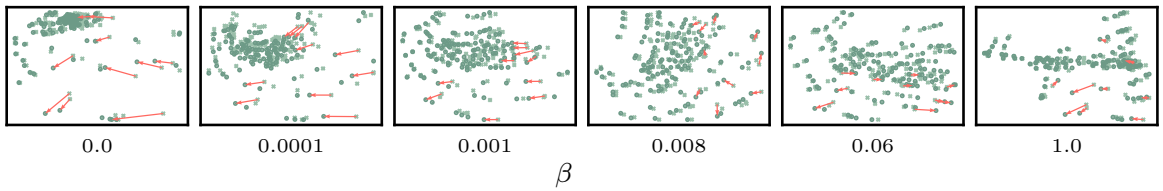

Figure 5: Regression performance in the latent space on the test set, comparing the reference latent space (×) and the prediction (●). The arrows (→) highlight the displacement of the predictions for the 10 cases with the highest error.

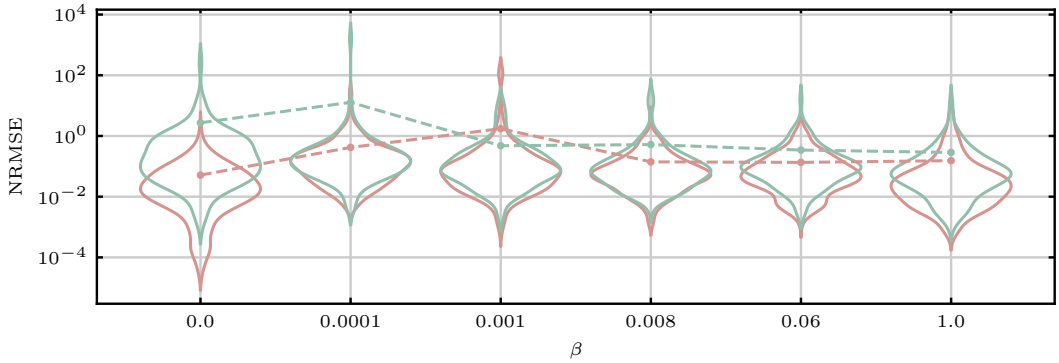

Figure 6: Normalized root mean squared error between the reference latent space and the regression prediction from flight conditions. Both training (▬) and testing (▬) datasets are considered. Dashed lines (- - -) indicate mean values.

Figure 6 displays the NRMSE distribution as a function of β for the β-VAE model across training and testing datasets. The dashed lines represent the mean error evolution with β, while the violin plots show the error distribution for each β value. As β decreases, the error distribution widens, reflecting greater variability in prediction errors. This variability is consistent with the large displacements reported in Figure 5 for high angles of attack. Despite this, the mean NRMSE remains relatively constant and independent of β.

In conclusion, the GPR model and the selected kernels demonstrate reasonably accurate performance regardless of the β value. This supports the choice of GPR for this application, despite the availability of more sophisticated regressors, such as DNN-based models. The robustness and versatility of GPR make it a reliable choice in similar studies [11, 23].

4.2 Pressure prediction

The performance of the surrogate model is assessed by predicting the pressure distribution on the upper wing side of the XRF1 aircraft for various Mach numbers and angles of attack within the flight envelope.

Figure 7 presents the Mean Absolute Error (MAE) between the actual pressure distribution and the predictions for the β-VAE+GPR model. The results show that for smaller values of β (e.g., 0, 0.008), the β-VAE+GPR model generally commits a small and stable error. This

Table 2: Performance indicators of the surrogate. They include MAE, RMSE, and R^2 score, which have been computed using the testing dataset.

β	MAE	$RMSE$	R^2
0	0.062	0.114	0.812
0.0001	0.055	0.101	0.854
0.001	0.053	0.098	0.864
0.008	0.061	0.102	0.852
0.06	0.110	0.155	0.671
1	0.192	0.266	0.106

indicates that minimal regularization improves prediction accuracy, while higher β values tend to increase MAE due to over-regularization.

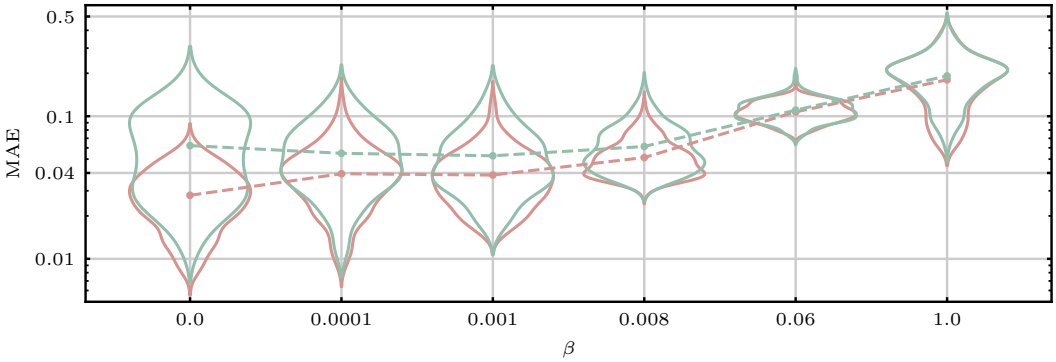

Figure 7: Mean absolute error between ground truth and C_p predictions for β-VAE+GPR model. Both training (▭) and testing (▭) datasets are showed. Dashed lines (- -) indicate mean values.

The performance metrics, displayed in Table 2 and including MAE, RMSE, and R^2 score, reveal that the β-VAE+GPR model is able to explain more information when the β is small enough, pulling down the bottom tail of the error distributions. Despite the complexity of tweaking properly the hyperparemeters of the model, a properly tuned β value significantly enhances surrogate model performance.

The global performance of the surrogate models is evaluated through detailed visualization cases using $\beta = 0.008$, which offers a good balance between prediction accuracy and latent space reconstruction. Figure 8 presents a comparison between the ground truth pressure coefficient fields from RANS simulations and the predictions from β-VAE+GPR model for various flight conditions. For the low Mach number and angle of attack case, the surrogate generally shows small prediction errors across the wing surface, though there are localized inaccuracies near the leading edge and trailing edge. In the moderate condition, prediction errors increase slightly, especially near the leading edge, with the model smoothing the pressure distribution compared to the reference. Finally, the high M and high α case shows more pronounced prediction errors, particularly around the shockwave region. The model struggles to accurately capture the pressure drop associated with shockwaves, as can be seen at the 0.9 wingspan percentage, but

it learns the general tendency of the actual distribution.

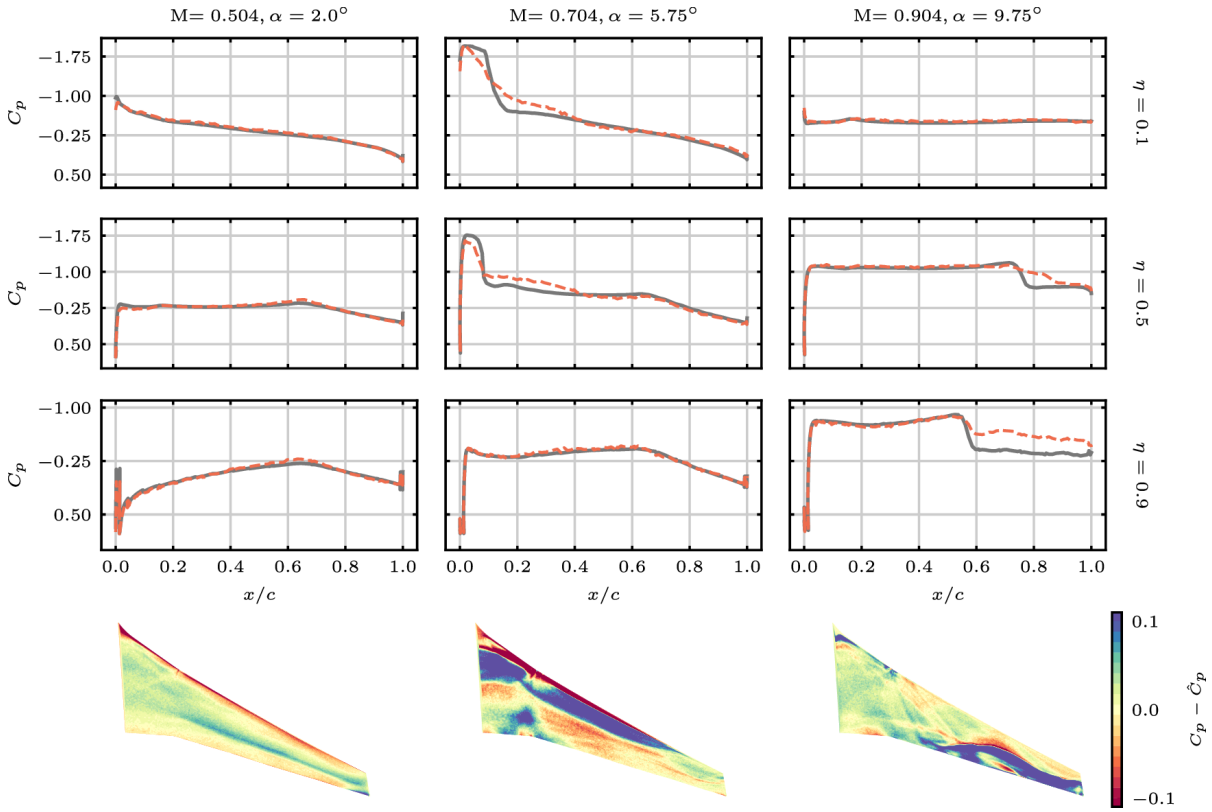

Figure 8: Difference between ground truth and predicted coefficients for the visualization test cases with $\beta = 0.008$. Chordwise pressure distributions from β-VAE+GPR (- - -) model at span percentages $\eta = 0.1, 0.5, 0.9$ are displayed and contrasted with ground truth (——).

In summary, the proposed methodology effectively produces a surrogate model capable of capturing the essential features of complete pressure distributions under various flight conditions. The architecture efficiently projects these fields into a low-dimensional space, thereby conserving computational resources, and subsequently reconstructs the pressure distributions from this latent representation while smoothing out the most nonlinear phenomena. Overall, the proposed model demonstrates superior performance compared to traditional surrogates based on linear reduction techniques [24].

5 CONCLUSIONS

This research investigates the use of β-VAEs combined with GPR for predicting pressure fields over a wing in transonic flight conditions. The study employs a MLP architecture for both the encoder and decoder in the β-VAE, effectively handling high-dimensional pressure coefficient samples to produce their latent representations. The analysis demonstrates that β-VAE models, when combined with GPR, accurately capture key aerodynamic features and maintain strong correlations with flight conditions across various latent space dimensions and β values.

The GPR model's performance in predicting latent space coordinates proves reliable, showing

minimal prediction errors and consistent results between training and testing datasets. The surrogate model efficiently predicts wing pressure distributions, exhibiting robust performance even in challenging conditions such as high angles of attack and shockwave regions. Overall, the β-VAE+GPR framework provides accurate and efficient predictions of pressure distributions, making it a valuable tool for aerodynamic analysis and optimization. The results underscore the effectiveness of this approach in balancing interpretability and prediction accuracy across different flight conditions.

ACKNOWLEDGEMENTS

This work has been supported by the TIFON project, ref. PLEC2023-010251/MCIN/AEI/10.13039/501100011033, funded by the Spanish State Research Agency, and by AIRBUS Defence & Space through the CETACEO project, ref. PTAG-20231008, funded by the CDTI. The authors would like to thank AIRBUS for providing the XRF1 database.

REFERENCES

[1] Jichao Li, Xiaosong Du, and Joaquim R.R.A. Martins. Machine learning in aerodynamic shape optimization. *Progress in Aerospace Sciences*, 134:100849, October 2022.

[2] R. Yondo, E. Andrés, and E. Valero. A review on design of experiments and surrogate models in aircraft real-time and many-query aerodynamic analyses. *Progress in aerospace sciences*, 96:23–61, 2018.

[3] E. Iuliano and D. Quagliarella. Proper orthogonal decomposition, surrogate modelling and evolutionary optimization in aerodynamic design. *Computers & Fluids*, 84:327–350, 2013.

[4] V. Dolci and R. Arina. Proper orthogonal decomposition as surrogate model for aerodynamic optimization. *International Journal of Aerospace Engineering*, 2016(1):8092824, 2016.

[5] R. Castellanos, J. Nieto-Centenero, A. Gorgues, S. Discetti, A. Ianiro, and E. Andrés. Towards aerodynamic shape optimisation by manifold learning and neural networks. In *15th International Conference on Evolutionary and Deterministic Methods for Design, Optimization and Control*, EUROGEN 2023. ISAAR-NTUA, 2023.

[6] J. Nieto-Centenero, A. Martínez-Cava, and E. Andrés. Multifidelity surrogate model for efficient aerodynamic predictions. In *8th European Congress on Computational Methods in Applied Sciences and Engineering*, ECCOMAS Congress, 2024.

[7] Ruixue Li and Shufang Song. Manifold learning-based reduced-order model for full speed flow field. *Physics of Fluids*, 36(8), 2024.

[8] C. Duru, H. Alemdar, and O. U. Baran. A deep learning approach for the transonic flow field predictions around airfoils. *Computers & Fluids*, 236:105312, 2022.

[9] Z. Wang, X. Liu, J. Yu, H. Wu, and H. Lyu. A general deep transfer learning framework for predicting the flow field of airfoils with small data. *Computers & Fluids*, 251:105738, 2023.

[10] J. Wang, C. He, R. Li, H. Chen, C. Zhai, and M. Zhang. Flow field prediction of supercritical airfoils via variational autoencoder based deep learning framework. *Physics of Fluids*, 33(8), 2021.

[11] Y.-E. Kang, S. Yang, and K. Yee. Physics-aware reduced-order modeling of transonic flow via β-variational autoencoder. *Physics of Fluids*, 34(7):076103, 2022.

[12] H. Eivazi, H. Veisi, M. Hossein Naderi, and V. Esfahanian. Deep neural networks for nonlinear model order reduction of unsteady flows. *Physics of Fluids*, 32(10):105104, 2020.

[13] B. Zhang. Nonlinear mode decomposition via physics-assimilated convolutional autoencoder for unsteady flows over an airfoil. *Physics of Fluids*, 35(9), 2023.

[14] D. P. Kingma and M. Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

[15] I. Higgins, L. Matthey, A. Pal, C. Burgess, X. Glorot, M. Botvinick, S. Mohamed, and A. Lerchner. β-vae: Learning basic visual concepts with a constrained variational framework. In *International conference on learning representations*, 2017.

[16] C. P. Burgess, I. Higgins, A. Pal, L. Matthey, N. Watters, G. Desjardins, and A. Lerchner. Understanding disentangling in β-VAE. *arXiv preprint arXiv:1804.03599*, 2018.

[17] A. Solera-Rico, C. Sanmiguel Vila, M. Gómez-López, Y. Wang, A. Almashjary, S.T.M. Dawson, and R. Vinuesa. β-variational autoencoders and transformers for reduced-order modelling of fluid flows. *Nature Communications*, 15(1):1361, 2024.

[18] Y. Wang, A. Solera-Rico, C. Sanmiguel Vila, and R. Vinuesa. Towards optimal β-variational autoencoders combined with transformers for reduced-order modelling of turbulent flows. *International Journal of Heat and Fluid Flow*, 105:109254, 2024.

[19] N. Kroll, S. Langer, and A. Schwöppe. *The DLR Flow Solver TAU - Status and Recent Algorithmic Developments*. American Institute of Aeronautics and Astronautics, 2014.

[20] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. The MIT Press, 11 2005.

[21] R. B. Gramacy. *Surrogates: Gaussian process modeling, design, and optimization for the applied sciences*. Chapman and Hall/CRC, 2020.

[22] J. Melo. Gaussian processes for regression: a tutorial. *Technical Report*, 2012.

[23] J. Nieto-Centenero, R. Castellanos, A. Gorgues, and E. Andrés. Fusing aerodynamic data using multi-fidelity gaussian process regression. In *15th International Conference on Evolutionary and Deterministic Methods for Design, Optimization and Control*, EUROGEN 2023. ISAAR-NTUA, 2023.

[24] R. Castellanos, J. Bowen Varela, A. Gorgues, and E. Andrés. An assessment of reduced-order and machine learning models for steady transonic flow prediction on wings. *ICAS 2022*, 2022.