FoilGen2: Learning Coupled Latent Spaces for Hybrid and Performance-Driven Airfoil Generation

Avneh Singh Bhatia
avnehb@gmail.com

Abstract

Airfoil design plays a key role in aerospace and renewable energy applications. However, it remains bottlenecked by slow, iterative trial-and-error evaluation using CFD or experiments, which limits rapid exploration of the design space. FoilGen2 learns separate latent spaces for geometry and performance, linked via a neural mapper, enabling hybrid design, smooth interpolation, and performance-driven synthesis. Experiments on 6,359 airfoils demonstrate geometric reconstruction errors below 0.8% of chord, lift coefficient errors under 3%, drag coefficient errors below 0.3%, and L/D errors under 3%. These results highlight the effectiveness of coupled latent spaces for flexible, aerodynamically coherent airfoil generation.

1 Introduction

Airfoil design is a critical component of aerodynamic engineering, where high-performing shapes are essential for achieving favorable lift-to-drag ratios, maintaining stall margins, and satisfying lift constraints. Traditional design approaches rely on iterative evaluation using computational fluid dynamics (CFD) or wind-tunnel testing, which is both time-consuming and costly. As a result, data-driven generative models have emerged as a promising alternative for exploring the airfoil design space more efficiently.

Existing approaches typically focus on either geometry or performance alone. Models that solely focus on geometric construction offer limited aerodynamic control. Coupling geometry and performance in a single framework is challenging due to the nonlinear relationship between airfoil shape and aerodynamic characteristics, as well as the sparsity of labeled performance data.

In this work, we propose FoilGen2, a framework that learns separate latent representations for airfoil geometry and aerodynamic performance and translates between them through a neural mapping. This coupled latent formulation enables smooth interpolation between designs, hybrid airfoil generation, and performance-driven synthesis, allowing designers to explore the space of airfoil shapes while respecting aerodynamic constraints. Our contributions are threefold:

- a coupled latent space formulation that captures both geometry and performance,
- hybrid generation through interpolation that produces geometrically valid and aerodynamically coherent designs, and
- performance-driven synthesis that enables conditioning on target aerodynamic metrics while maintaining design diversity.

Experiments demonstrate accurate reconstruction of airfoil shapes, smooth latent-space interpolation, and effective generation of performance-conditioned hybrids.

The remainder of the paper is organized as follows. Section 2 reviews related work, Section 3 presents the proposed method, Section 4 details the dataset and preprocessing, Section 5 describes the training setup, Section 6 evaluates reconstruction, interpolation, and hybridization results, and Section 7 discusses limitations and concludes.

2 Related Work

2.1 Data-Driven Airfoil Modeling

Wang et al. propose a generative framework based on a VAEGAN architecture that encodes airfoil geometries into a low-dimensional latent space and synthesizes new shapes through sampling and interpolation. The learned latent representations enable smooth generation of novel airfoil geometries, while a genetic algorithm is applied to the latent variables to evolve designs toward improved aerodynamic performance. However, performance information is not explicitly incorporated into the generative latent space during training; instead, aerodynamic objectives are enforced through post hoc optimization of latent vectors, limiting direct control over performance during generation [6].

2.2 Latent Representation Learning for Design

Latent representation learning has been widely used to compress high-dimensional design spaces into low-dimensional embeddings that capture essential geometric features. For example, Chen and Ramamurthy [2] employ a deep generative model to parameterize three-dimensional aerodynamic shapes, enabling the efficient generation and interpolation of novel designs. Such latent spaces facilitate smooth exploration of design variations and support optimization in the compressed space.

2.3 Geometry–Performance Coupling

A small number of works attempt to incorporate performance objectives directly into the generative modeling of airfoils. Physics-guided generative approaches introduce aerodynamic constraints into the loss function to encourage physically plausible outputs [5].

3 Method

3.1 Geometry Latent Encoder–Decoder

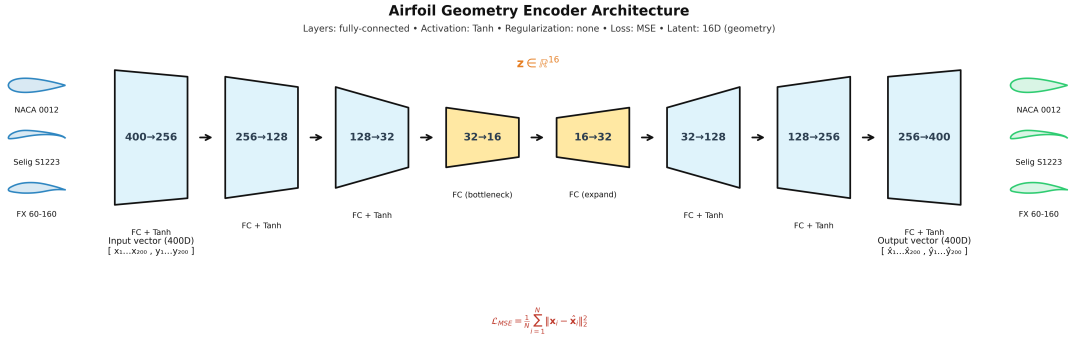

Figure 1: Geometry encoder-decoder architecture. The encoder compresses airfoil coordinates into a latent vector $z_{\text{geo}} \in \mathbb{R}^{16}$, which the decoder reconstructs back to the airfoil geometry. The latent space enables interpolation and hybrid generation of new designs.

Our geometry encoder-decoder network compresses airfoil shapes, represented as a vector $x_{\text{geo}} \in \mathbb{R}^{400}$, into a low-dimensional latent vector $z_{\text{geo}} \in \mathbb{R}^{16}$ and reconstructs them back to a coordinate vector $\hat{x}_{\text{geo}} \in \mathbb{R}^{400}$. The encoder consists of fully connected layers with Tanh activations, mapping the normalized airfoil coordinates to a 16-dimensional latent space. The decoder mirrors this structure to reconstruct the original geometry. The network is trained using a mean squared error loss between input and reconstructed shapes.

3.2 Performance Latent Encoder–Decoder

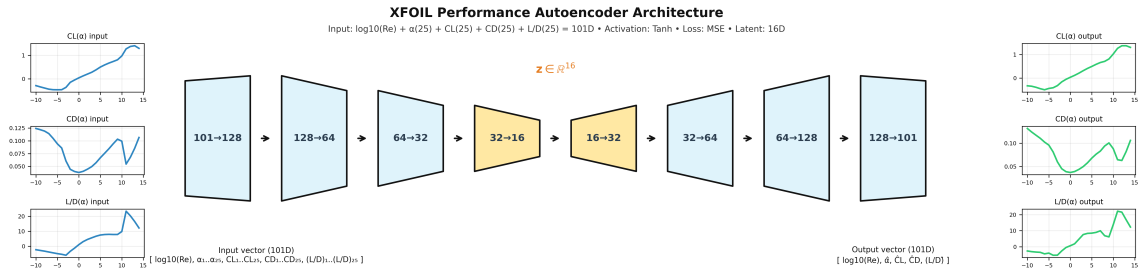

Figure 2: XFoil encoder-decoder architecture. The encoder compresses XFoil polars and the Reynolds number into a latent vector $z_{\text{xfoil}} \in \mathbb{R}^{16}$, which the decoder reconstructs back to the original polars. The latent space allows for interpretation of new polars.

Our XFoil encoder-decoder network encodes XFoil polars, along with the corresponding Reynolds number, represented as a vector $x_{\text{xfoil}} \in \mathbb{R}^{101}$, into a low-dimensional latent vector $z_{\text{xfoil}} \in \mathbb{R}^{16}$ and

reconstructs them back to the original vector $\hat{x}_{\text{xfoil}} \in \mathbb{R}^{101}$. The encoder consists of fully connected layers with Tanh activations, mapping the input vector to the 16-dimensional latent space, while the decoder mirrors this structure to reconstruct the original polars. The network is trained using a mean squared error loss between the input and reconstructed vectors. Each input vector x_{xfoil} consists of: the Reynolds number, 25 angles of attack α, ranging from $\alpha = -10°$ to $\alpha = 15°$, the corresponding lift and drag coefficients (C_L and C_D), and the lift-to-drag ratio C_L/C_D for each α.

3.3 Latent Space Mapping and Translation

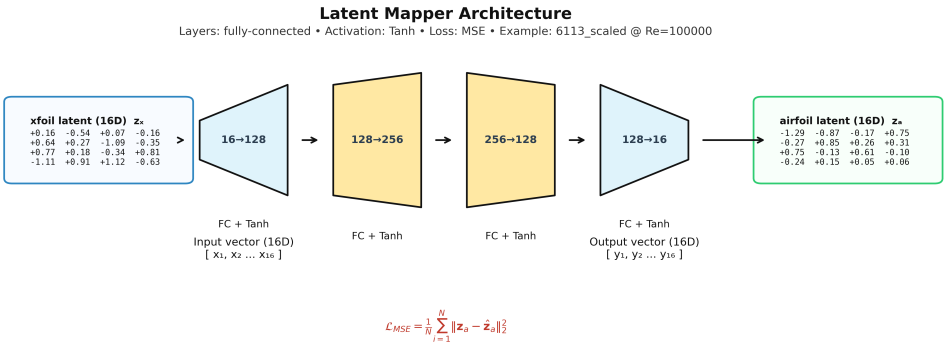

Figure 3: Learned latent space mapping architecture.

Our latent mapper network learns a continuous transformation between the aerodynamic and geometric latent spaces. Specifically, it maps the XFoil latent vector $z_{\text{xfoil}} \in \mathbb{R}^{16}$, produced by the XFoil encoder, to a corresponding airfoil geometry latent vector $z_{\text{airfoil}} \in \mathbb{R}^{16}$. The model consists of fully connected layers with Tanh activations arranged in a $16 \to 128 \to 256 \to 128 \to 16$ architecture, enabling nonlinear expansion and compression between the two latent representations. The network is trained using a mean squared error loss between the predicted and target airfoil latent vectors, $\|z_{\text{airfoil}} - \hat{z}_{\text{airfoil}}\|^2$. This learned mapping enables bidirectional translation between aerodynamic performance features and airfoil geometry within a shared low-dimensional latent space.

4 Dataset and Preprocessing

Our study utilizes a dataset of 6,359 airfoil geometries and corresponding aerodynamic polars obtained from the Bigfoil.com airfoil database [3]. The original airfoil coordinates varied in resolution and distribution, so we first interpolated each airfoil to a uniform grid of 200 points along both the x and y axes using cubic spline interpolation. The coordinates were then normalized to the $[0, 1]$ range to facilitate stable training of the geometry encoder-decoder.

For aerodynamic performance data, we employed NeuralFoil [4] to generate lift and drag polars for each airfoil across eight Reynolds numbers: 50k, 100k, 200k, 300k, 400k, 500k, 750k, and 1M. Each polar was sampled at 25 α values ranging from $\alpha = -10°$ to $\alpha = 15°$. The Reynolds numbers were transformed using a base-10 log scale, and all polar features (C_L, C_D, C_L/C_D) were standardized using a `StandardScaler` to ensure zero mean and unit variance across the dataset. These polars, combined with the transformed Reynolds number for each case, formed the input vectors $x_{\text{xfoil}} \in \mathbb{R}^{101}$ for the XFoil encoder-decoder.

The dataset was split into training and validation sets with an 80/20 ratio. The data preparation proceeded in stages. First, the geometry encoder-decoder was trained on the normalized airfoil coordinates to learn a compact latent representation of shape. Next, the generated aerodynamic polars were used to train the XFoil encoder-decoder, producing a corresponding performance latent space. Finally, a neural mapping was trained between the geometry latent space and the performance latent space, enabling bidirectional translation between airfoil shape and aerodynamic performance. This preparation ensures that both latent spaces are aligned and suitable for hybrid airfoil generation and performance-driven synthesis.

5 Training Setup

All models were trained using the Adam optimizer with mean squared error (MSE) reconstruction loss. Geometry and performance autoencoders used an 80/20 train/validation split, batch size 32, and were trained for 200 epochs. The latent mapper network used a batch size of 64 and trained for 300 epochs.

A ReduceLROnPlateau scheduler was applied to all networks (factor 0.5, patience 10–20 depending on model). Input data were scaled using `StandardScaler`, and Reynolds numbers were transformed using $\log_{10}$. Training was performed on Apple MPS when available, otherwise CPU.

6 Experiments

6.1 Geometry Reconstruction

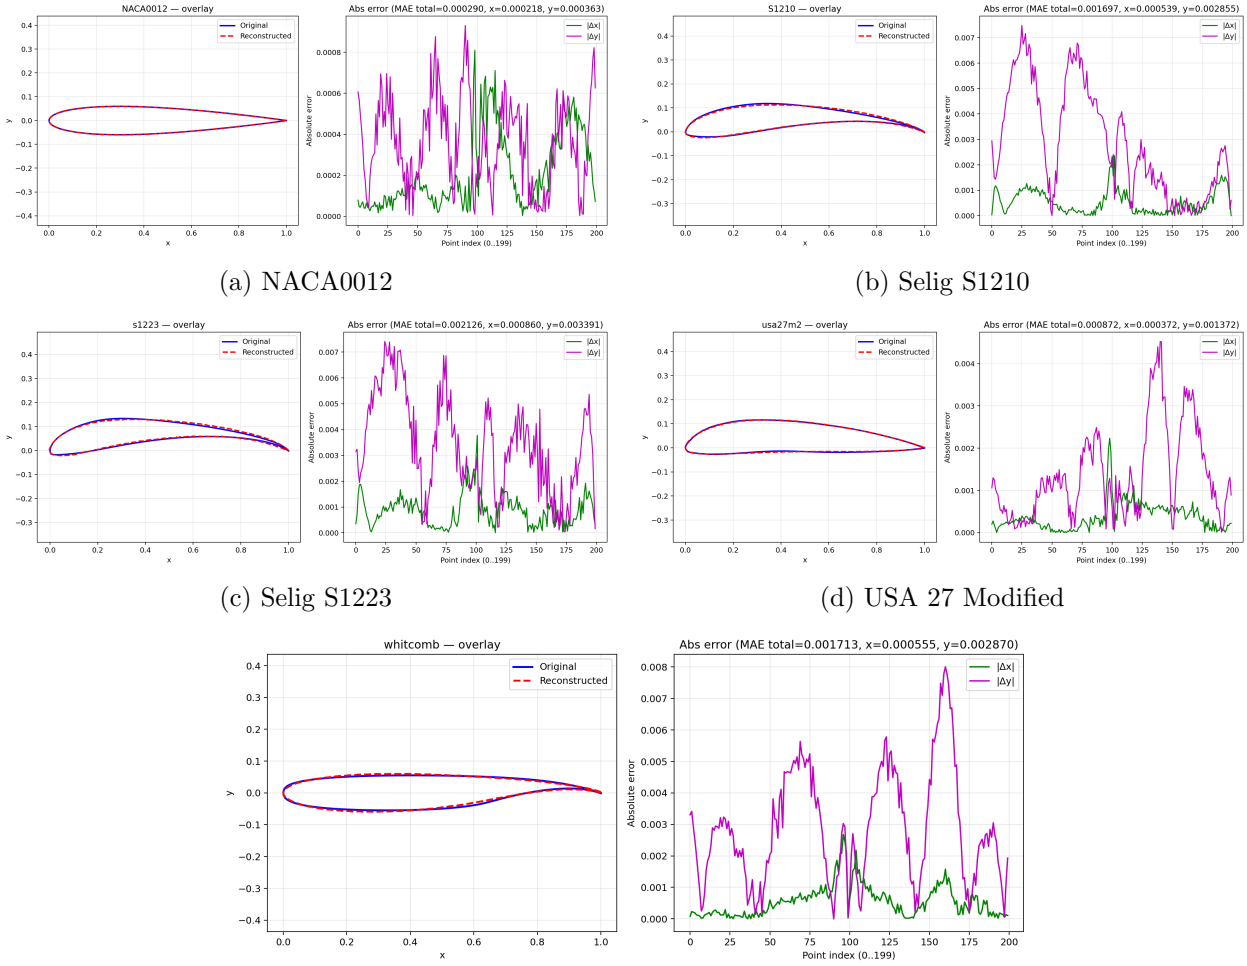

(a) NACA0012

(b) Selig S1210

(c) Selig S1223

(d) USA 27 Modified

(e) NASA Whitcomb Integral Supercritical Airfoil

Figure 4: Geometry reconstruction overlays for representative airfoils.

We evaluate the encoder-decoder across different airfoil types. Standard airfoils like NACA 0012 are reconstructed with maximum MAE below 0.001. Heavily cambered airfoils, such as S1223 and S1210, show slightly higher errors around 0.008. Non-standard airfoils, like the NASA Whitcomb Supercritical Integral and USA 27 Modified, have MAEs of approximately 0.008 and 0.005, respectively.

Overall, reconstruction errors remain below 1% of chord, with highest accuracy on standard shapes and slightly reduced accuracy on extreme camber or unconventional airfoils. This demonstrates that the latent geometry space reliably captures essential features across diverse airfoil families.

6.2 Polar Reconstruction

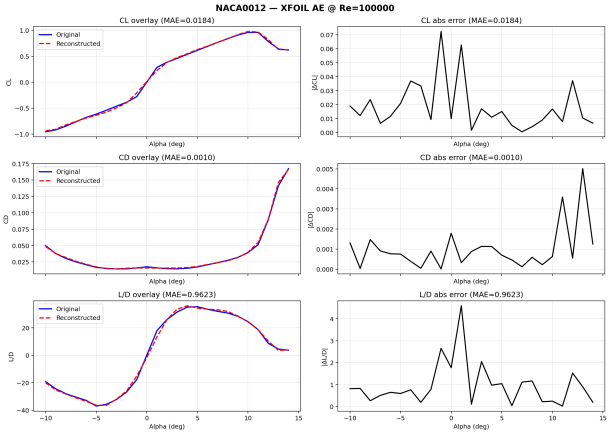

(a) NACA0012 ($Re = 100\,000$)

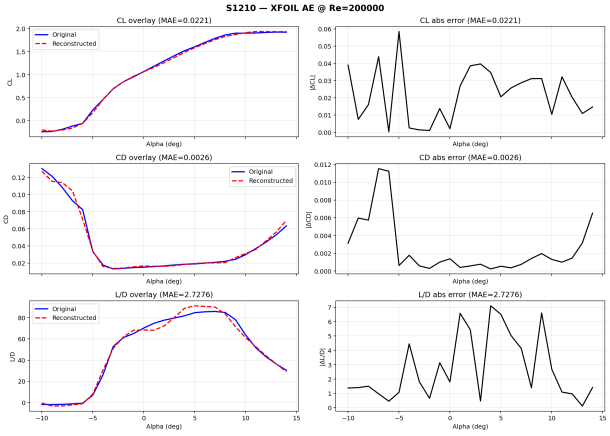

(b) S1210 ($Re = 200\,000$)

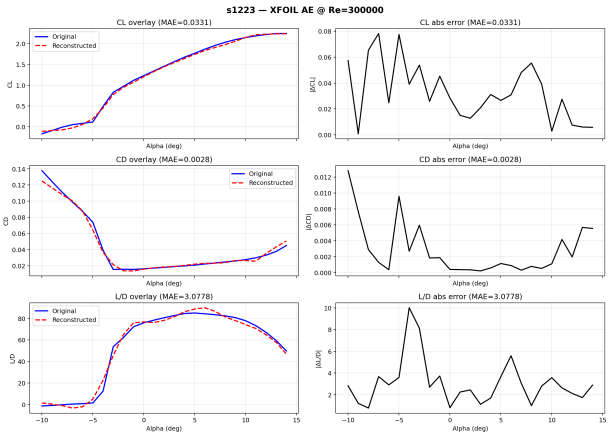

(c) S1223 ($Re = 300\,000$)

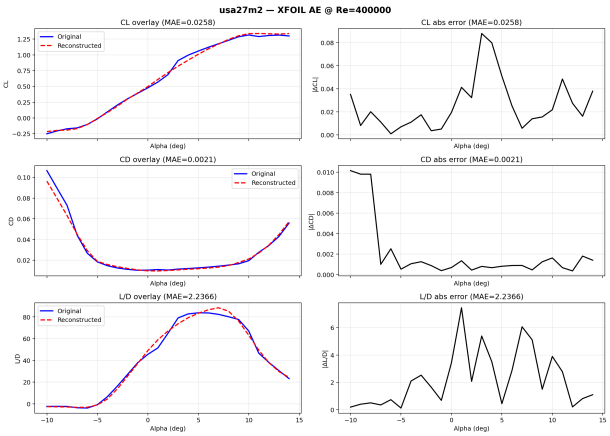

(d) USA27M2 ($Re = 400\,000$)

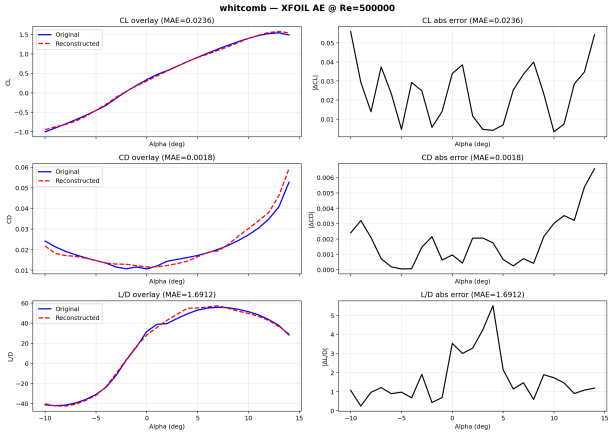

(a) Whitcomb ($Re = 500\,000$)

Figure 6: Polar reconstruction examples for representative airfoils at different Reynolds numbers.

To evaluate the performance encoder-decoder, we reconstructed XFoil polars for a representative set of airfoils across different Reynolds numbers. Table 1 summarizes the mean absolute errors (MAE) for lift coefficient (C_L), drag coefficient (C_D), and lift-to-drag ratio (C_L/C_D).

Table 1: Reconstruction errors for XFoil polars across selected airfoils and Reynolds numbers.

Airfoil	Reynolds	MAE(C_L)	MAE(C_D)	MAE(C_L/C_D)
NACA 0012	100,000	0.0184	0.0010	0.9623
Selig S1210	200,000	0.0221	0.0026	2.7276
Selig S1223	300,000	0.0331	0.0028	3.0778
USA 27 Modified	400,000	0.0258	0.0021	2.2366
Whitcomb Integral Supercritical	500,000	0.0236	0.0018	1.6912

The results indicate that the performance autoencoder accurately reconstructs C_L and C_D across a variety of airfoil shapes, with the highest MAE for C_L occurring for highly cambered airfoils such as S1223. Errors in C_L/C_D are larger due to compounding of small deviations in lift and drag, but remain within acceptable bounds for rapid design exploration. Overall, the encoder-decoder demonstrates robust generalization across standard, cambered, and unconventional airfoils.

6.3 Geometric Reconstruction from Polars

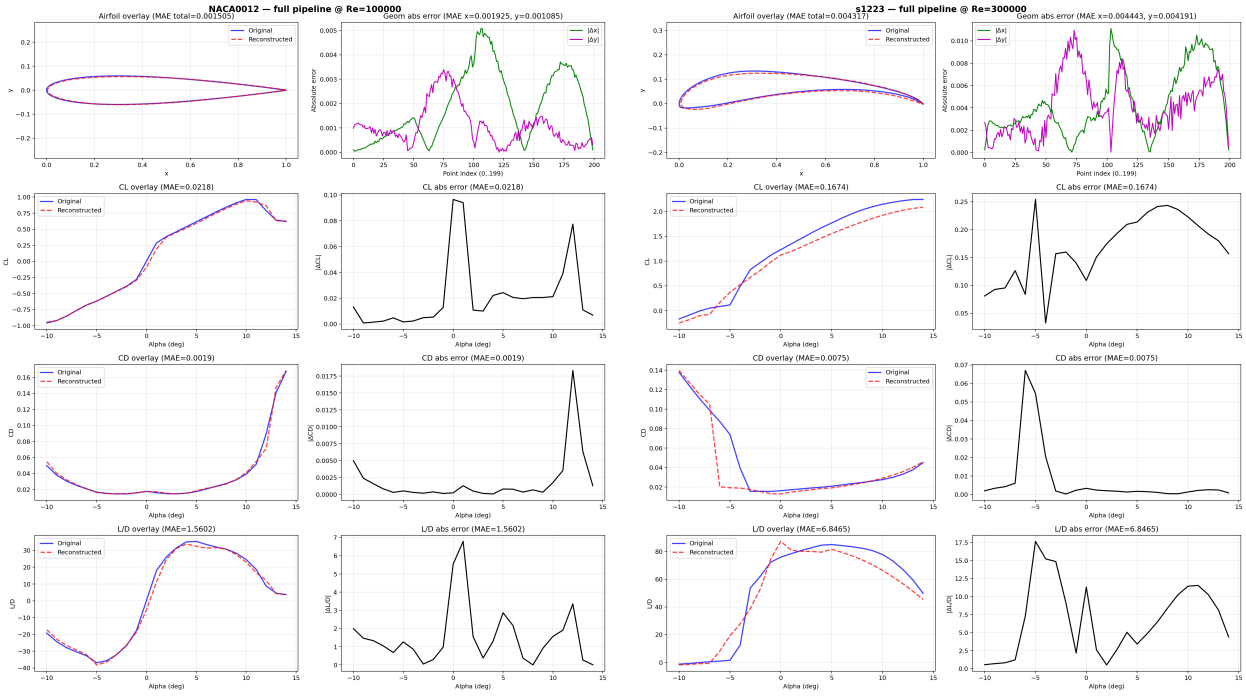

(a) NACA0012 (best case)

(b) S1223 (worst case)

Figure 7: Airfoils reconstructed purely from polars and Reynolds number. Any discrepancies between the original and reconstructed airfoils (and their corresponding polars) are most likely attributable to NeuralFoil-generated polar noise/approximation rather than the latent-space mapping itself.

Reconstruction through the full pipeline shows high accuracy on standard airfoils, shown in the reconstruction of NACA 0012 at $Re = 100{,}000$, with geometric MAEs below 0.002, C_L MAE 0.022,

C_D MAE 0.0019, and L/D MAE 1.56. Errors peak near lift-curve transitions but remain small overall. More challenging cases, such as S1223 at $Re = 300{,}000$, exhibit higher errors—particularly in C_D at negative angles of attack—leading to larger L/D deviations, with geometric MAEs around 0.0044 and C_L, C_D, L/D MAEs of 0.167, 0.0075, and 6.85, respectively. These examples demonstrate that the pipeline reliably reconstructs standard and moderately cambered airfoils, while highly cambered or extreme shapes are more prone to localized deviations. Some of the deviations may also be attributable to noise from NeuralFoil

6.4 Hybrid Airfoils

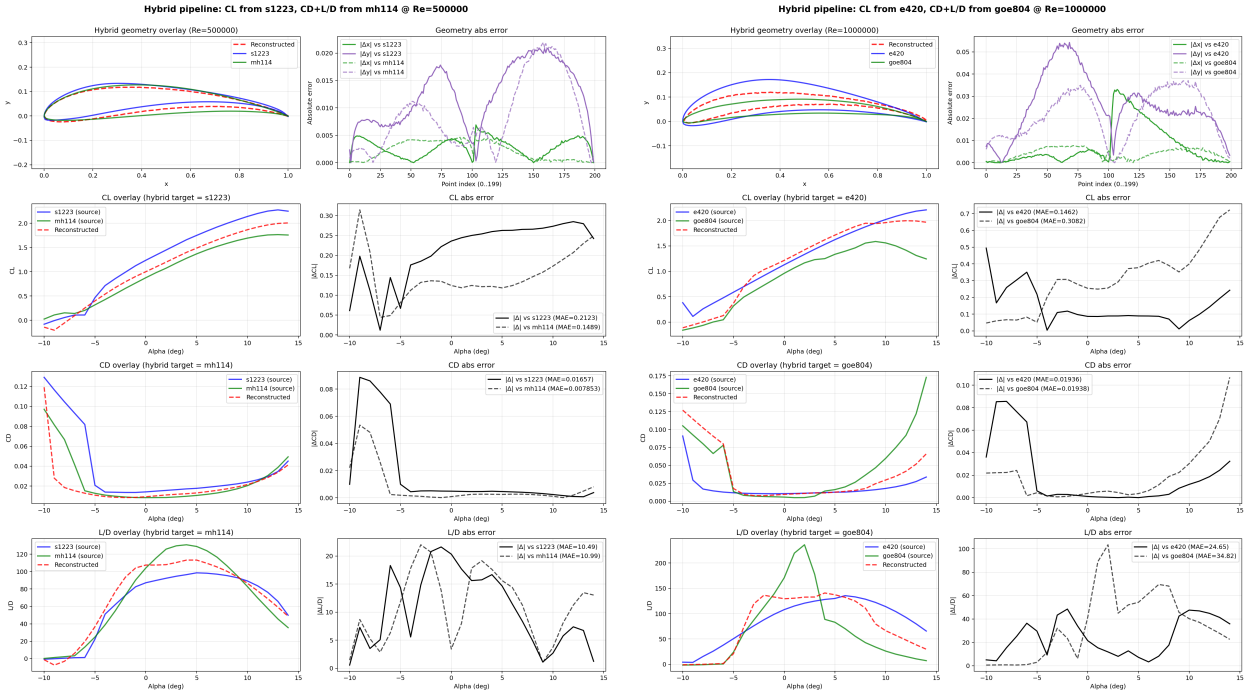

(a) S1223–MH114 (best case; parent polars are relatively close)

(b) E420–GOE804 (worst case; stark L/D differences in parent polars)

Figure 8: Hybrid airfoil generation by interpolating between parent geometries and translating the hybrid back into performance space. When parent polars are close, the hybrid polar tends to vary smoothly; when parent L/D behavior differs sharply, hybridization can produce larger deviations.

Hybridization of two parent airfoils is achieved by combining the lift polar (C_L) from one parent with the drag and lift-to-drag ratio (C_D and L/D) from the other. This produces a hybrid airfoil that effectively averages the characteristics of its parents, attempting to capture the best of both designs. This flexibility is enabled by the model's modularity: different components of the input vector can be spliced from multiple parents, allowing targeted transfer of performance traits during generation. When the parent polars are similar in shape, the resulting hybrid shows smooth, well-behaved performance. However, if the parent polars differ sharply in shape, the hybrid, while still relatively intermediate in shape and performance, can exhibit larger deviations and reduced aerodynamic quality.

6.5 Custom Airfoils from User–Defined Polars

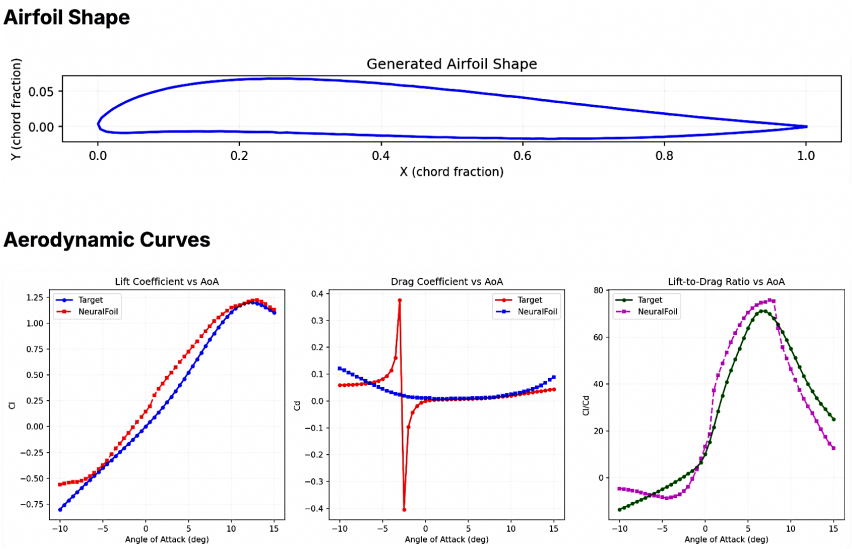

Figure 9: User-generated airfoil at $Re = 250\,000$ using user-inputted spline points for C_L and L/D; C_D was computed from these inputs. Drag spike is a numerical artifact.

This figure shows the generated airfoil shape and its aerodynamic performance at $Re = 250\,000$ compared to the target data. The lift curve closely matches the target across most angles of attack, indicating that the model successfully learned the primary lift behavior. The large spike in the drag coefficient curve is not physical and is caused by a divide by zero error, since lift coefficient and lift to drag ratio are provided as inputs while drag coefficient is computed afterward. Despite this artifact, the lift to drag curve follows the target trend well, demonstrating that the model captures overall efficiency behavior even when derived quantities introduce numerical instability.

7 Discussion and Limitations

FoilGen2 shows promise for hybrid airfoil generation and performance-driven design, but several limitations remain. The performance polars are generated using NeuralFoil, which provides uniform angle-of-attack sampling but lower fidelity than high-resolution CFD or experiments. Small Reynolds number differences (e.g., $Re = 500{,}000$ vs $Re = 500{,}001$) can produce non-physical discrepancies, which could be reduced with a larger dataset and denser Reynolds sampling.

The latent spaces for geometry and performance are not fully continuous or physics-informed. While suitable for interpolation and hybrid generation, they may omit key aerodynamic effects, limiting interpretability and creativity. Important metrics, such as pitching moments, pressure distributions, or 3D effects, are not encoded.

Occasionally, the encoder-decoder produces jagged airfoils; training on smoothed coordinates (e.g., Venkataraman fits) could address this. Because the model abstracts patterns rather than physical laws, predictions may not generalize well outside the training distribution.

Finally, with 6,359 airfoils, extreme or unconventional shapes are underrepresented, potentially limiting hybrid diversity. These limitations suggest directions for future work, including physics-informed latent spaces, higher-fidelity performance data, and expanded datasets with more Reynolds numbers and airfoil classes.

8 Conclusion

We presented FoilGen2, a coupled-latent framework for airfoil generation that jointly learns low-dimensional representations of geometry and aerodynamic performance, along with a neural mapper that translates between them. This structure enables multiple design workflows within a single model, including geometry and polar reconstruction, performance-to-geometry synthesis from polars and Reynolds number alone, and hybrid airfoil generation through latent interpolation and polar splicing. Across representative cases, the model reconstructs aerodynamic coefficients with low error and produces geometries whose induced performance trends remain consistent with the target polars.

The approach is primarily limited by the fidelity and coverage of the training data. Performance labels are generated using NeuralFoil rather than high-fidelity CFD or experiments, and the learned latent spaces are not explicitly physics-informed, omitting quantities such as pitching moment and pressure distributions. These factors can reduce accuracy near nonlinear regimes and limit generalization to out-of-distribution airfoils. Minor geometric artifacts also suggest that additional regularization or smoother parameterizations could further improve surface quality.

Future work will focus on higher-fidelity aerodynamic data, denser Reynolds sampling, and the incorporation of additional physical constraints and metrics into the latent objectives. Overall, FoilGen2 demonstrates that aligning geometry and performance in coupled latent spaces is a viable and flexible approach to airfoil generation, while clearly identifying pathways toward improved physical accuracy and robustness.

Acknowledgements

The author thanks Mike Quayle for the airfoil dataset, Marco Passalacqua (University of Sapienza, Rome) for providing validation cases, and the AmadorUAVs team for their collaboration and feedback.

References

[1] Wei Chen, Kevin Chiu, and Mark Fuge. Airfoil design parameterization and optimization using bézier generative adversarial networks, 2020.

[2] Wei Chen and Arun Ramamurthy. Deep generative model for efficient 3d airfoil parameterization and generation. In *AIAA Scitech 2021 Forum*. American Institute of Aeronautics and Astronautics, January 2021.

[3] Mike Quayle. BigFoil.com airfoil database. `https://www.bigfoil.com`, 2025.

[4] Peter Sharpe and R. John Hansman. Neuralfoil: An airfoil aerodynamics analysis tool using physics-informed machine learning, 2025.

[5] Kazunari Wada, Katsuyuki Suzuki, and Kazuo Yonekura. Physics-guided training of gan to improve accuracy in airfoil design synthesis, 2023.

[6] Yuyang Wang, Kenji Shimada, and Amir Barati Farimani. Airfoil gan: encoding and synthesizing airfoils for aerodynamic shape optimization. *Journal of Computational Design and Engineering*, 10(4):1350–1362, June 2023.

[7] Jinouwen Zhang, Junjie Ren, Qianhong Ma, Jianyu Wu, Aobo Yang, Yan Lu, Lu Chen, Hairun Xie, Jing Wang, Miao Zhang, Wanli Ouyang, and Shixiang Tang. Funcgenfoil: Airfoil generation and editing model in function space, 2025.