

INVERSE AIRFOIL DESIGN USING AERODYNAMICALLY CONDITIONED VARIATIONAL AUTOENCODERS

K. Swannet[1], N.A.K. Doan[2] & C. Varriale[1]

[1]Delft University of Technology
[2]Imperial College London

Abstract

Inverse design can accelerate aerodynamic optimisation by generating candidate geometries directly from prescribed performance requirements, eliminating the need for an iterative approach. This work investigates a Conditional Variational Autoencoder (CVAE) framework for inverse airfoil design, where airfoils are represented by Kulfan coefficients and conditioned on aerodynamic parameters obtained with NeuralFoil, an XFOIL surrogate. The conditions include Reynolds number, lift and moment coefficients at maximum lift-to-drag ratio, and maximum lift coefficient, providing information related to both cruise-like and low-speed performance. In addition to the standard CVAE with a normal latent prior, a conditioned-prior variant is tested in which the prior distribution parameters are predicted from the aerodynamic condition vector. The results show that CVAEs can generate smooth and realistic airfoil shapes whose aerodynamic behaviour follows the imposed conditioning trends. However, the generated geometries do not consistently match the prescribed condition values exactly, especially when the requested aerodynamic quantities are strongly coupled or correspond to sparsely represented regions of the training distribution. The conditioned prior proves better for reconstruction quality, and marginally improves condition adherence, albeit not consistently. These findings suggest that conditional generative models are useful for exploring aerodynamic design trends, while generated shapes should still be interpreted as samples from a learned distribution rather than exact solutions to a deterministic inverse design problem. Nonetheless, the presented approach constitutes a first step towards interactive airfoil design tools that combine fast generative sampling with aerodynamic performance guidance over multiple operating conditions.

Keywords: Generative modelling, CVAE, Inverse design, Airfoils, Deep learning

1 Introduction

Airfoil geometry optimisation remains a key aspect of wing design, traditionally achieved through iterative computational fluid dynamics (CFD) analysis and manual refinement. While effective, these procedures can be prohibitively time-consuming and computationally expensive. Surrogate models can reduce the cost of evaluating candidate airfoils, but they do not remove the need to search through a large and often unintuitive geometry space. Inverse design accelerates this aerodynamic optimisation process by directly generating geometries matching specified performance characteristics, but this is not an easy problem to solve due to the complex relations between the shape and its aerodynamic characteristics.

Recent advances in deep generative modelling offer a promising solution. Variational Autoencoders (VAEs) [1] map high-dimensional data into compact latent spaces that can be interpreted and manipulated [2, 3]. Conditional VAEs (CVAEs) extend this by incorporating auxiliary labels directly into the generative process [4], enabling more targeted output generation. Their application to aerodynamic shapes has shown success in several controlled inverse-design settings. For example, CVAE-based models have been used to generate airfoils satisfying prescribed lift coefficients [5–7], while other

https://doi.org/10.71945/icas2026_0894

studies have used geometric descriptors, such as thickness and camber-related quantities [3, 8], or modal coefficients [9] to guide the generated geometries. These works demonstrate the potential of conditional generative models for airfoil design, but the conditioning is often either geometric in nature or defined for a single fixed operating condition.

This study considers a more direct aerodynamic conditioning problem. Instead of conditioning the geometry generation only on geometric quantities, or on a single aerodynamic coefficient at one fixed Reynolds number and angle of attack, the model is trained on airfoils evaluated over multiple Reynolds numbers and aerodynamic properties at moving operating points. The conditions therefore describe design-relevant aerodynamic behaviour rather than only shape descriptors. This makes the problem more challenging, since the aerodynamic parameters are strongly coupled and may correspond to different locations on the lift and drag polars, but it also makes the formulation closer to how an airfoil is selected in practice.

Two variants of the CVAE architecture are considered. The first uses the standard normal Gaussian prior, while the second uses a prior conditioned on the target aerodynamic properties. The second variant is intended to make the latent space more useful as a geometry parameterisation, by forcing the model to capture the aerodynamic conditions in the prior and leaving the latent variables to represent the remaining geometric information.

The aim of this work is therefore not only to generate airfoils from target conditions, but also to investigate whether a conditional generative model can retain an interpretable and useful latent geometry representation when the conditions are aerodynamic, coupled, and obtained from polar data over a range of Reynolds numbers.

Specifically, we train CVAE models with and without a condition-dependent latent prior, using the UIUC airfoil dataset and a differentiable PyTorch implementation of NeuralFoil to provide aerodynamic conditioning and training losses. The remainder of the paper is organised around the data construction, model formulation, and resulting model behaviour. Section 2 presents the airfoil representation, aerodynamic conditions, and surrogate analysis used to construct the training data. Section 3 introduces the CVAE formulation and the conditioned-prior extension, and training procedure. Finally, section 4 compares the trained models in terms of reconstruction quality, latent-space structure, condition adherence, and response to changes in the target conditions.

2 Training Data

The training dataset is the UIUC airfoil database [10]. The data is first preprocessed by fitting an interpolating B-spline through the coordinates, then transforming and normalising the obtained curves to the domain $(0 \leq x \leq 1)$, aligning the chord line with the x-axis, and applying a Gaussian-based thickness correction to repair entries with non-physical negative thicknesses. This results in $N = 1619$ distinct airfoils.

2.1 Airfoil Geometry Data

Airfoil coordinate data is converted to a compact representation using Class-Shape Transformation (CST) coefficients [11]. This curve parameterisation is commonly used because it can describe a wide range of airfoil shapes with a relatively small number of parameters. In this work, the modified CST formulation by Kulfan [12] is used, which includes additional terms for the leading edge and trailing-edge thickness. The exact parameter vector is chosen to match the input format required by NeuralFoil [13], the aerodynamic surrogate model used to evaluate the airfoil performance. For every airfoil n $(1 \leq n \leq N)$, the geometry vector $\mathbf{x}^{(n)} \in \mathbb{R}^X$, with $X = 18$, contains eight upper-surface CST coefficients, eight lower-surface CST coefficients, one leading-edge weight, and one trailing-edge thickness value.

2.2 Aerodynamic Conditionals

The conditional model is trained to generate airfoil geometries for prescribed aerodynamic requirements. These requirements are chosen to describe more than one operating point of the airfoil. The lift and moment coefficients at the maximum lift-to-drag ratio are extracted from the polar. These are

denoted by $C_{l_{\text{design}}}$ and $C_{m_{\text{design}}}$, and present a cruise-like design condition. The maximum lift coefficient $C_{l_{\max}}$ is included as a low-speed performance condition, which is relevant for take-off and landing. Finally, the logarithm of the Reynolds number is included because the aerodynamic behaviour of an airfoil depends on the flow regime, and because the same geometry may be used under different combinations of size, velocity, and fluid properties. Per airfoil geometry, $R = 10$ Reynolds numbers are drawn independently from the uniform distribution defined in equation (1).

$$\log Re \sim \mathscr{U}\left(\log(5 \times 10^5), \log(2 \times 10^7)\right) \tag{1}$$

Together, these quantities define the condition vector $\mathbf{c} \in \mathbb{R}^C$, with $C = 4$, as

$$\mathbf{c} = \left[\log Re,\, C_{l_{\text{design}}},\, C_{m_{\text{design}}},\, C_{l_{\max}}\right]. \tag{2}$$

The vector therefore combines information related to the flow condition, efficient cruise operation, pitching moment, and maximum-lift capability, rather than describing only a single point on the polar.

The choice of aerodynamic conditions is also motivated by the latent structure observed in the reference VAE, obtained in a previous study [2] and presented in appendix A figure 12. The latent traversals of the geometry-only model showed that several latent directions correspond to familiar airfoil features, such as camber, thickness location, leading-edge radius, and changes in the front- and aft-loaded character of the camber line. These geometric features are not independent from the selected aerodynamic quantities. Camber and thickness location are closely related to the achievable lift and maximum lift, while changes in the loading distribution and camber curvature influence the pitching moment. The selected condition vector therefore attempts to isolate part of these geometric variations through aerodynamic requirements, rather than prescribing the geometric features directly. It then remains to be seen which geometric features are still captured by the inferred latent space once these aerodynamic conditions are provided to the model.

Note that new Reynolds numbers are sampled for each airfoil every time a batch is passed through the model, so the same geometry is repeatedly presented with different operating conditions throughout training. This makes the model see the same geometry under a range of aerodynamic evaluations and reduces the chance that it simply memorises a fixed condition-geometry pairing. Moreover, conditional values are normalized with respect to their respective approximate mean and standard deviation over the dataset. By sampling 1500 random Reynolds numbers for all $N = 1619$ airfoils in the dataset, and evaluating their conditions values with NeuralFoil, an empirical distribution of condition values is obtained. Results are presented in table 1. This normalization is performed to ensure that the different components of the condition vector, which vary strongly in magnitude, are on a similar scale. The normalized condition vector is then used as input to the decoder during both training and generation. The complete distributions of the sampled conditions are shown in figure 1, where three reference airfoil geometries are marked for chosen Reynolds numbers, which are later used for the case studies in section 4.

Table 1 – Predetermined statistics used for normalising the aerodynamic condition variables.

Parameter	μ	σ	Min.	Max.
$\log Re$	14.9772	1.0744	13.1220	16.8110
$C_{l_{\text{design}}}$	1.0241	0.2810	0.1066	2.3170
$C_{m_{\text{design}}}$	-0.0739	0.0604	-0.3113	0.1142
$C_{l_{\max}}$	1.6876	0.2808	0.6555	2.9034

2.3 Aerodynamic Analysis

Aerodynamic analysis of the airfoils is performed using NeuralFoil [13], a neural network-based surrogate model trained to emulate XFOIL [14]. Specifically, the parameters of the largest available pretrained model[1] are loaded into a PyTorch implementation. NeuralFoil provides the aerodynamic

[1] `https://github.com/peterdsharpe/NeuralFoil`

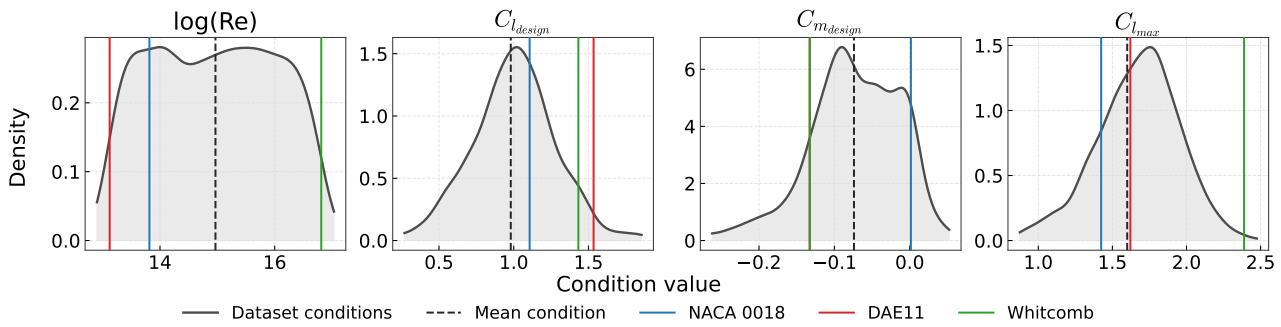


Figure 1 – Distribution of the aerodynamic condition values over the dataset used during training.

quantities normally obtained from XFOIL, including lift, drag, and moment coefficients, but at a much lower computational cost. It also avoids the convergence failures that can occur when XFOIL is applied to unconventional airfoil shapes or extreme operating conditions.

A further advantage of using a PyTorch-based model is that NeuralFoil can be included directly in the training loop. This allows the aerodynamic condition losses to remain differentiable with respect to the generated geometry and, therefore, with respect to the parameters of the generative model.

3 CVAE Model Architecture

The Conditional Variational Autoencoder (CVAE) builds upon the structure of a standard Variational Autoencoder (VAE), which itself is a generative model consisting of an encoder-decoder architecture with a probabilistic latent representation. The main novelty of the CVAE lies in its ability to condition the generative process on auxiliary variables, allowing the model to learn distributions of outputs associated with user-defined properties or characteristics. The model architecture is presented in figure 2.

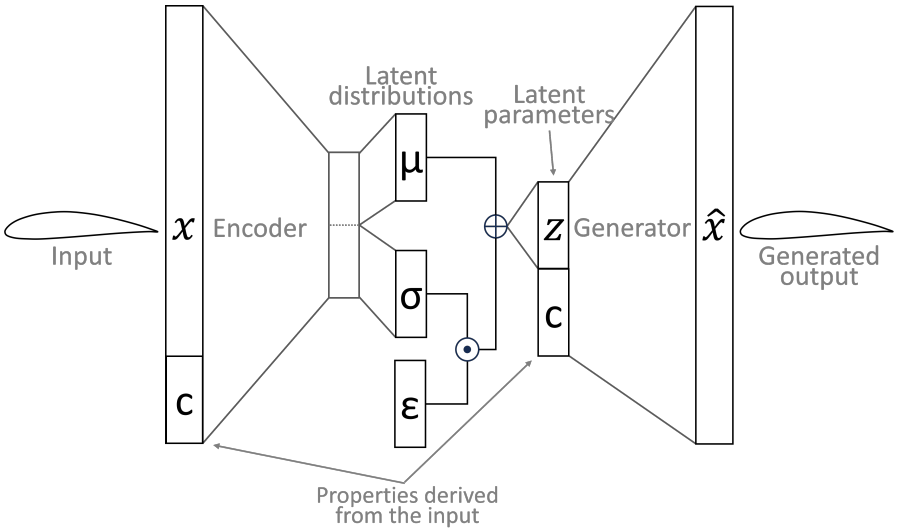


Figure 2 – The conditional variational autoencoder architecture used

The model is trained by maximising the conditional Evidence Lower Bound (ELBO) [1, 15]. For a given geometry $\mathbf{x}$ and condition vector $\mathbf{c}$, the ELBO forms a lower bound on the conditional log-likelihood $\log p(\mathbf{x} \mid \mathbf{c})$:

$$\mathscr{L}_{\mathrm{ELBO}}(\theta, \phi; \mathbf{x}, \mathbf{c}) = \underbrace{\mathbb{E}_{\mathbf{z} \sim q_\phi(\mathbf{z}|\mathbf{x},\mathbf{c})} \left[\log p_\theta(\mathbf{x} \mid \mathbf{z}, \mathbf{c})\right]}_{\text{reconstruction likelihood}} - \beta \underbrace{D_{\mathrm{KL}}\left(q_\phi(\mathbf{z} \mid \mathbf{x}, \mathbf{c}) \parallel p(\mathbf{z})\right)}_{\text{latent-space regularisation}}, \tag{3}$$

where θ and ϕ denote the network parameters, further discussed in section 3.1. The first term encourages the model to reconstruct the input geometry from the sampled latent vector and the condition vector by maximising the likelihood of the observed data. The second term regularises the

approximate posterior towards a chosen latent prior, preventing each geometry from being assigned to an arbitrary isolated region of latent space. It uses the Kullback-Leibler (KL) divergence to measure the difference between the approximate posterior and the chosen prior, which represents the true distribution of the latent variables. The weight β is a hyperparameter that controls the strength of this latent-space regularisation. If β is too low, the model overfits and fails to learn a useful latent space. If β is too high, the model overgeneralises and fails to reconstruct the geometries well. The optimal β value is determined through hyperparameter tuning. A proper balance allows the model to retain good reconstruction quality while keeping the latent space smooth and compact. In this way, the latent dimensions can capture clear and interpretable features [2, 16]. The full training objective used in the experiments, including the additional aerodynamic condition loss, is described in section 3.2.

3.1 Model Components

The three distributions in equation (3) are represented by the encoder, decoder, and prior. The approximate posterior $q(\mathbf{z} \mid \mathbf{x}, \mathbf{c})$ is modelled by the encoder network and is therefore written as $q_\phi(\mathbf{z} \mid \mathbf{x}, \mathbf{c})$, where ϕ denotes the encoder network parameters. The reconstruction distribution $p(\mathbf{x} \mid \mathbf{z}, \mathbf{c})$ is modelled by the decoder network and is written as $p_\theta(\mathbf{x} \mid \mathbf{z}, \mathbf{c})$, where θ denotes the decoder network parameters. The final distribution is the latent prior, which is used to regularize the latent space and defines the distribution from which latent variables are sampled during generation.

The encoder receives both the airfoil geometry $\mathbf{x}$ and the corresponding condition vector $\mathbf{c}$. It does not return a single deterministic latent vector, but instead predicts the parameters of a Gaussian distribution over the latent variables. This distribution is the approximate posterior

$$q_\phi(\mathbf{z} \mid \mathbf{x}, \mathbf{c}) = \mathcal{N}\left(\mu_\phi(\mathbf{x}, \mathbf{c}), \operatorname{diag}\left(\sigma_\phi^2(\mathbf{x}, \mathbf{c})\right)\right), \tag{4}$$

where μ_ϕ and σ_ϕ are outputs of the encoder network, while ϕ denotes the network weights. A latent vector $\mathbf{z}$ is sampled from this distribution and passed to the decoder together with the same condition vector $\mathbf{c}$.

The decoder models $p_\theta(\mathbf{x} \mid \mathbf{z}, \mathbf{c})$ and reconstructs, or generates, an airfoil geometry from a latent vector and a prescribed set of conditions. The decoder again outputs the parameters of a Gaussian distribution over the geometry,

$$p_\theta(\mathbf{x} \mid \mathbf{z}, \mathbf{c}) = \mathcal{N}\left(\mu_\theta(\mathbf{z}, \mathbf{c}), \operatorname{diag}\left(\sigma_\theta^2(\mathbf{z}, \mathbf{c})\right)\right). \tag{5}$$

The final output is normally a sample from this distribution, but to get the best output geometries here the mean is used as the final output with highest likelihood instead. The reconstruction loss is the negative log-likelihood of the input geometry under this distribution, seen in equation (3). In practice, the decoder can also be made deterministic, omitting the variance, and the reconstruction loss term can be implemented as a reconstruction error between the input geometry and the decoder output.

The latent vector accounts for the remaining geometric variation that is not fixed by the conditions. To make the latent space continuous and sampleable, the approximate posterior is regularised towards a chosen prior distribution. This regularization is introduced through the KL-divergence term in the ELBO, seen in equation (3). Since the true distribution of the latent variables is unknown, a simple prior is chosen to encourage the model to learn a smooth latent space. In the standard CVAE used here, the latent prior is chosen as the standard normal distribution,

$$p(\mathbf{z}) = \mathcal{N}(\mathbf{0}, \mathbf{I}). \tag{6}$$

During generation, the true geometry is not available, so a latent vector is sampled from the prior and decoded together with a user-defined condition vector.

A second CVAE variant is also considered, where the fixed standard normal prior is replaced by a condition-dependent prior. In the standard formulation, latent samples are drawn from the same distribution for every condition vector. This assumes that the remaining geometric variation around each aerodynamic target can be represented by the same latent distribution. Since this assumption may

be too restrictive for coupled aerodynamic conditions, a small prior network with trainable parameters ψ is used to predict the parameters of the latent prior from the condition vector, resulting in the conditional prior

$$p_\psi(\mathbf{z} \mid \mathbf{c}) = \mathcal{N}\left(\mu_\psi(\mathbf{c}), \mathrm{diag}\left(\sigma^2_\psi(\mathbf{c})\right)\right), \tag{7}$$

The encoder remains unchanged and still models $q_\phi(\mathbf{z} \mid \mathbf{x}, \mathbf{c})$, but the KL-divergence term now regularises this posterior towards $p_\psi(\mathbf{z} \mid \mathbf{c})$ instead of the fixed standard normal prior. The corresponding ELBO becomes

$$\mathcal{L}^{\mathrm{cond}}_{\mathrm{ELBO}}(\theta, \phi, \psi; \mathbf{x}, \mathbf{c}) = \mathbb{E}_{\mathbf{z}\sim q_\phi(\mathbf{z}|\mathbf{x},\mathbf{c})}\left[\log p_\theta(\mathbf{x} \mid \mathbf{z}, \mathbf{c})\right] - \beta D_{\mathrm{KL}}\left(q_\phi(\mathbf{z} \mid \mathbf{x}, \mathbf{c}) \parallel p_\psi(\mathbf{z} \mid \mathbf{c})\right) \tag{8}$$

During generation with the conditioned-prior CVAE, latent vectors are sampled from $p_\psi(\mathbf{z} \mid \mathbf{c})$ and decoded together with the same condition vector.

3.2 Training Procedure

During training, the encoder predicts the parameters of the approximate posterior distribution over the latent variables. A latent sample is obtained using the reparameterisation trick,

$$\mathbf{z} = \mu_\phi + \sigma_\phi \odot \varepsilon, \qquad \varepsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \tag{9}$$

which allows gradients to propagate through the sampling operation. Here, μ_ϕ and σ_ϕ are predicted by the encoder, and $\odot$ denotes element-wise multiplication.

For each mini-batch, the aerodynamic conditions are sampled dynamically. Each airfoil geometry $\mathbf{x}^{(i)}$ is paired with $R = 10$ Reynolds-number cases, resulting in condition vectors $\mathbf{c}^{(i,1)}, \ldots, \mathbf{c}^{(i,R)}$. Each geometry-condition pair defines its own approximate posterior, $q_\phi(\mathbf{z} \mid \mathbf{x}^{(i)}, \mathbf{c}^{(i,r)})$. A latent sample $\mathbf{z}^{(i,r)}$ is therefore drawn for each Reynolds-number case. The decoder therefore reconstructs the same input geometry for several sampled operating conditions,

$$\hat{\mathbf{x}}^{(i,r)} \sim p_\theta\left(\mathbf{x} \mid \mathbf{z}^{(i,r)}, \mathbf{c}^{(i,r)}\right), \qquad r = 1, \ldots, R. \tag{10}$$

The losses are first evaluated for each Reynolds-number case and then averaged over the R cases, such that every airfoil contributes one loss value to the mini-batch objective.

3.3 Base Loss

The ELBO in equation (3) defines the two base loss terms used during training. Following the dynamic sampling procedure described in section 3.2, both terms are evaluated for each geometry-condition pair $(\mathbf{x}^{(i)}, \mathbf{c}^{(i,r)})$ and then averaged over the R Reynolds-number cases.

The reconstruction term is implemented as the negative log-likelihood (NLL) of the Kulfan parameter vector under the distribution predicted by the decoder. For each geometry-condition pair, the likelihood is evaluated over all X Kulfan parameters, so the NLL is summed over the data dimension. The resulting scalar loss is then averaged over the R Reynolds-number cases, as defined in equation (11). Although generated airfoils are obtained from the decoder mean, the predicted decoder variance is retained during training through this likelihood term.

$$L^{(i)}_{\mathrm{NLL}} = \frac{1}{R}\sum_{r=1}^{R}\left[-\log p_\theta\left(\mathbf{x}^{(i)} \mid \mathbf{z}^{(i,r)}, \mathbf{c}^{(i,r)}\right)\right]. \tag{11}$$

For each geometry-condition pair, the KL divergence similarly corresponds to the sum of the one-dimensional KL contributions over the Z latent dimensions. The resulting scalar loss is then averaged over the R Reynolds-number cases. For the standard-prior CVAE this gives the loss in equation (12).

$$L^{(i)}_{\mathrm{KL}} = \frac{1}{R}\sum_{r=1}^{R} D_{\mathrm{KL}}\left(q_\phi\left(\mathbf{z} \mid \mathbf{x}^{(i)}, \mathbf{c}^{(i,r)}\right) \parallel p(\mathbf{z})\right), \tag{12}$$

where $p(\mathbf{z}) = \mathcal{N}(\mathbf{0}, \mathbf{I})$. For the conditioned-prior CVAE, the fixed standard prior is replaced by $p_\psi(\mathbf{z} \mid \mathbf{c}^{(i,r)})$, giving the loss in equation (13).

$$L^{(i)}_{\mathrm{KL}} = \frac{1}{R}\sum_{r=1}^{R} D_{\mathrm{KL}}\left(q_\phi\left(\mathbf{z} \mid \mathbf{x}^{(i)}, \mathbf{c}^{(i,r)}\right) \parallel p_\psi\left(\mathbf{z} \mid \mathbf{c}^{(i,r)}\right)\right). \tag{13}$$

3.4 Auxiliary Loss Terms

To get the model to learn properly, additional loss terms were required. These terms are not part of the ELBO, but are added to the training objective to provide a more informative loss signal and ensure training remains in a feasible domain.

More Informative Reconstruction Loss

In practice, determining the reconstruction loss from the Kulfan coefficients alone proved insufficient to guarantee accurate reconstruction of the airfoil geometries. Coefficient errors do not always translate uniformly to surface errors. Depending on the associated basis function, a small error in one coefficient can cause a noticeable deviation in the airfoil surface, while a similar error in another coefficient may have only a minor effect. Therefore, an additional coordinate-space loss is included to provide a more informative reconstruction signal. Both the curves defined by the input Kulfan coefficients and the reconstructed Kulfan coefficients are evaluated at $M = 199$ points along the airfoil surface. Let $\mathbf{s}(\mathbf{x}) \in \mathbb{R}^{M\times 2}$ denote the resulting surface coordinates obtained from a Kulfan vector. The coordinate reconstruction loss is then defined as the sum of squared coordinate errors, averaged over the R Reynolds-number cases,

$$L_{\text{coord}}^{(i)} = \frac{1}{R}\sum_{r=1}^{R}\sum_{m=1}^{M}\left\|\mathbf{s}_m\left(\mathbf{x}^{(i)}\right) - \mathbf{s}_m\left(\hat{\mathbf{x}}^{(i,r)}\right)\right\|_2^2. \tag{14}$$

The total reconstruction loss is defined as

$$L_{\text{rec}}^{(i)} = L_{\text{NLL}}^{(i)} + \lambda_{\text{coord}} L_{\text{coord}}^{(i)}, \tag{15}$$

where λ_{coord} controls the relative contribution of the coordinate-space reconstruction term.

Condition Losses

To encourage the generated geometries to satisfy the prescribed aerodynamic requirements, an additional condition loss is included. This loss term is not part of the ELBO, but is added to the training objective to explicitly encourage the generated geometries to satisfy the prescribed aerodynamic requirements, and is made possible by the differentiable nature and speed of the NeuralFoil predictions. For each reconstruction $\hat{\mathbf{x}}^{(i,r)}$, the aerodynamic quantities are evaluated again using NeuralFoil, producing the predicted condition vector $\hat{\mathbf{c}}^{(i,r)}$. The condition loss is computed as the mean-squared error between the normalised target and predicted conditions,

$$L_{\text{cond}}^{(i)} = \frac{1}{R}\sum_{r=1}^{R}\left\|\tilde{\mathbf{c}}^{(i,r)} - \tilde{\hat{\mathbf{c}}}^{(i,r)}\right\|_2^2, \tag{16}$$

where the tilde denotes normalised condition values. This normalisation prevents conditions with larger numerical scales from dominating the loss.

An additional condition loss is evaluated using in-batch cross-sampling. For each batch, the condition vectors already present in the batch were shuffled between airfoils and slightly perturbed. These cross-sampled condition vectors were then combined with latent vectors sampled from the prior and passed through the decoder. This creates a generation-like training case, since the latent vectors are sampled from the prior rather than inferred from an input geometry. At the same time, because the condition values originate from the current batch, the sampled condition combinations remain close to realistic regions of the condition space. This reduces the risk of training on impossible aerodynamic requests, such as combinations where $C_{l,\max}$ is lower than $C_{l,\text{design}}$, or where the moment coefficient is inconsistent with the requested lift values.

$$L_{\text{cross}}^{(i)} = \frac{1}{R}\sum_{r=1}^{R}\left\|\tilde{\mathbf{c}}_{\text{cross}}^{(i,r)} - \tilde{\hat{\mathbf{c}}}_{\text{cross}}^{(i,r)}\right\|_2^2. \tag{17}$$

Here, $\mathbf{c}_{\text{cross}}^{(i,r)}$ denotes the shuffled and perturbed condition vector, while $\hat{\mathbf{c}}_{\text{cross}}^{(i,r)}$ denotes the conditions predicted by NeuralFoil from the generated geometry. The tilde again indicates normalised condition values.

NeuralFoil Constraint Losses

Two additional penalty terms were included to discourage geometries for which NeuralFoil is expected to be unreliable. The first is based on the airfoil wiggliness metric defined by the authors of NeuralFoil. This penalty is inactive for sufficiently smooth airfoils and activates only when the wiggliness value exceeds a threshold of 2, the median wiggliness value observed in the training dataset:

$$w^{(i,r)} \approx \sum_{i=2}^{N-1} \left(c_{\text{lower},i+1} - 2c_{\text{lower},i} + c_{\text{lower},i-1}\right)^2 + \left(c_{\text{upper},i+1} - 2c_{\text{upper},i} + c_{\text{upper},i-1}\right)^2. \tag{18}$$

$$L_{\text{wiggle}}^{(i)} = \frac{1}{R} \sum_{r=1}^{R} \left[\max\left(0,\, w^{(i,r)} - 2\right)\right]^2, \tag{19}$$

where $w^{(i,r)}$ is the wiggliness of the reconstructed airfoil for the r-th Reynolds-number case, a function of the upper and lower CST coefficients $c_{\text{lower}}^{(i,r)}$ and $c_{\text{upper}}^{(i,r)}$.

The second term is based on the NeuralFoil confidence value. Low confidence indicates that the geometry or operating condition is outside the region where NeuralFoil is expected to be reliable. A penalty is therefore applied whenever the confidence drops below 0.9,

$$L_{\text{conf}}^{(i)} = \frac{1}{R} \sum_{r=1}^{R} \left[\max\left(0,\, 0.9 - \eta^{(i,r)}\right)\right]^2, \tag{20}$$

where $\eta^{(i,r)}$ denotes the NeuralFoil confidence value. Both these losses discourage generated geometries outside the reliable NeuralFoil evaluation regime, and therefore also work as a feasibility constraint on the generated geometries.

Total Training Loss

The full training objective is then written as

$$L_{\text{train}} = \frac{1}{B} \sum_{i=1}^{B} \left(L_{\text{rec}}^{(i)} + \beta L_{\text{KL}}^{(i)} + \xi L_{\text{coord}}^{(i)} + \gamma L_{\text{cond}}^{(i)} + \gamma_{\text{cross}} L_{\text{cross}}^{(i)} + \lambda_{\text{wiggle}} L_{\text{wiggle}}^{(i)} + \lambda_{\text{conf}} L_{\text{conf}}^{(i)}\right). \tag{21}$$

The coefficients β, ξ, γ, λ. in equation (21) control the relative importance of the different objectives. The parameter β weights the KL-divergence term and therefore controls the strength of the latent-space regularisation. The parameter ξ weights the coordinate-space reconstruction loss, which encourages the airfoil curves defined by the reconstructed Kulfan coefficients to match the input geometry more closely. The parameter γ weights the direct aerodynamic condition loss, while γ_{cross} weights the cross-sampled condition loss used to test generation from prior samples. Increasing either condition-loss weight places more emphasis on matching the prescribed aerodynamic quantities, but can also compete with accurate reconstruction when the requested conditions are difficult to satisfy simultaneously. Finally, λ_{wiggle} and λ_{conf} control the penalties associated with excessive airfoil wiggliness and low NeuralFoil confidence, respectively.

3.5 Network Architecture and Hyperparameters

The network architecture was chosen based on the unconditional VAE developed in previous work [2]. In that study, it was shown that the main geometric variation in the airfoil dataset can be represented with a small number of well-regularised latent variables. Starting from a latent space of dimension $Z = 8$, the KL weight β was selected such that approximately six latent dimensions remained active, while the remaining dimensions stayed close to the prior and became inactive. This provided a compact latent representation with good reconstruction quality and interpretable geometric trends. Plots showing the capabilities of the unconditional VAE are provided in appendix A.

This trained VAE is used as the initial checkpoint for the conditional models. It provides a useful starting point because the initial latent space is already regularised, reconstructs the airfoil geometries well, and contains disentangled geometric features. The CVAE experiments can therefore be

interpreted relative to a known geometry-only baseline, both in terms of reconstruction quality and in terms of the usefulness of the latent space for later design optimisation.

The conditional models retain the same basic encoder and decoder architecture as the reference VAE, with identical network sizes of 4 linear layers with 512 neurons each. For the standard CVAE, the condition vector is concatenated to the geometry before being passed to the encoder, and also concatenated to the latent vector before being passed to the decoder. For the conditioned-prior CVAE, an additional prior network maps the condition vector to the mean and variance of the conditional latent prior. This prior network is kept small compared with the encoder and decoder, with only 2 layers of 128 neurons each, since its role is only to shift and scale the latent distribution according to the requested aerodynamic conditions.

The entire implementation is in PyTorch. Models are trained using the Adam optimizer with an initial learning rate of 1×10^{-4}, and a reduction-on-plateau learning rate scheduler with a decay factor of 0.5 and a patience of 125 epochs. The dataset is split into three subsets of 80% training, 15% validation, and 5% testing data, further divided into batches of size 64 during each epoch. Training data is what the optimizer sees and uses to compute the gradients. The validation set is used to track the model's performance on unseen data during training after each epoch, and to apply early stopping when the validation loss stops improving. The test set is only used after training to evaluate the final performance of the model on completely unseen data. No maximum number of epochs is set, but instead training is stopped when the validation loss does not improve for 250 consecutive epochs. The more complete set of hyperparameters is provided in appendix E.

Remaining hyperparameters, such as the loss weights, are found by running a randomised parameter sweep, implemented using Weights & Biases [17] and selecting the best-performing model. For the results discussed in section 4, the best-performing model was defined by the lowest, unweighted cross-sampled condition loss L_{cross} on the test set.

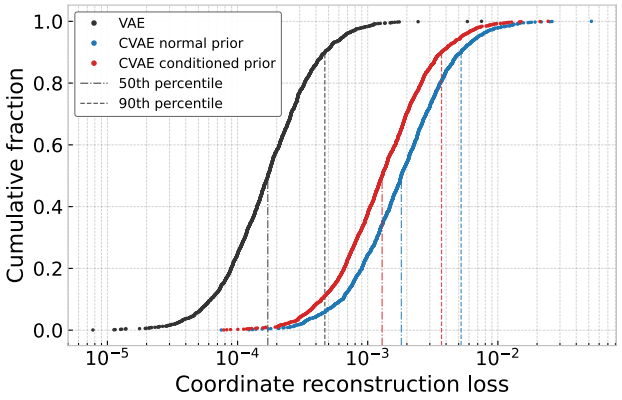
4 Results and Discussion

The normal-prior CVAE adds the aerodynamic condition vector to the decoder while retaining a fixed standard-normal latent prior. The conditioned-prior CVAE specialises this formulation by replacing the fixed prior with a condition-dependent prior distribution. The discussion below compares the results of both, and uses the geometry-only VAE as baseline.

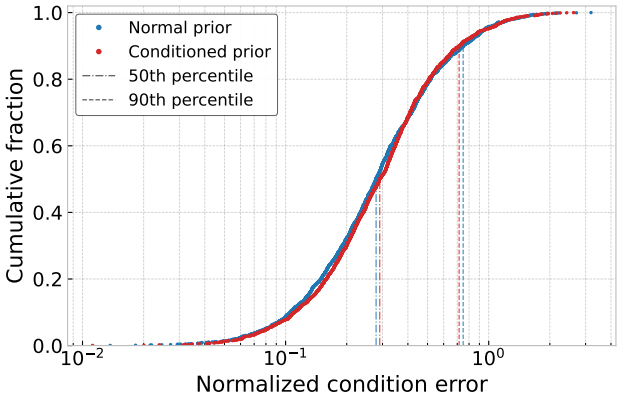
The comparison is organised to separate the different roles that the models are expected to play. First, reconstruction accuracy is considered, since this shows how much geometric fidelity is lost when using a compact latent space with aerodynamic conditioning compared to the geometry-only VAE. Next, latent and condition traversals are used to inspect which geometric features remain in the latent space and which are captured by the aerodynamic conditions. Randomly generated samples are then evaluated in terms of condition adherence, to test whether prior samples match complete target condition vectors drawn from realistic regions of the data. Finally, condition response and condition-space sensitivity are analysed to determine whether each conditioning variable behaves as an independent design control, or whether the achieved aerodynamic quantities remain coupled. These analyses are performed both around the mean condition and latent vectors, and around three reference airfoil conditions, to show how the models behave away from the mean in different regions of condition space.

4.1 Reconstruction Accuracy

First, the models' capacity to reconstruct the training data is evaluated. The VAE provides a strong baseline for reconstruction quality, since all model capacity can be used to encode and decode geometry, without having to satisfy externally imposed aerodynamic targets.



(a) Cumulative error plot of the coordinate-space reconstruction loss.



(b) Cumulative error plot of the normalized condition error.

Figure 3 – Cumulative loss plots for the reconstruction errors (a) and the condition errors (b) over the training dataset for the VAE, normal-prior CVAE, and conditioned-prior CVAE.

The CVAEs both sacrifice reconstruction accuracy to gain condition-based control, as they must balance the geometry reconstruction task with the task of moving the decoded geometry toward the requested aerodynamic targets. The coordinate-loss distributions in the left panel of figure 3 quantify this trade-off. This also shows that the conditioned-prior CVAE is marginally better at reconstructing the training data than its normal-prior counterpart, which is likely due to the more flexible, condition-dependent latent space. However, the condition-error distributions in the right panel of figure 3 show that the slightly improved geometry reconstruction of the conditioned-prior CVAE does not translate to significant improvements in condition adherence, with negligible difference between both model variants for the training data.

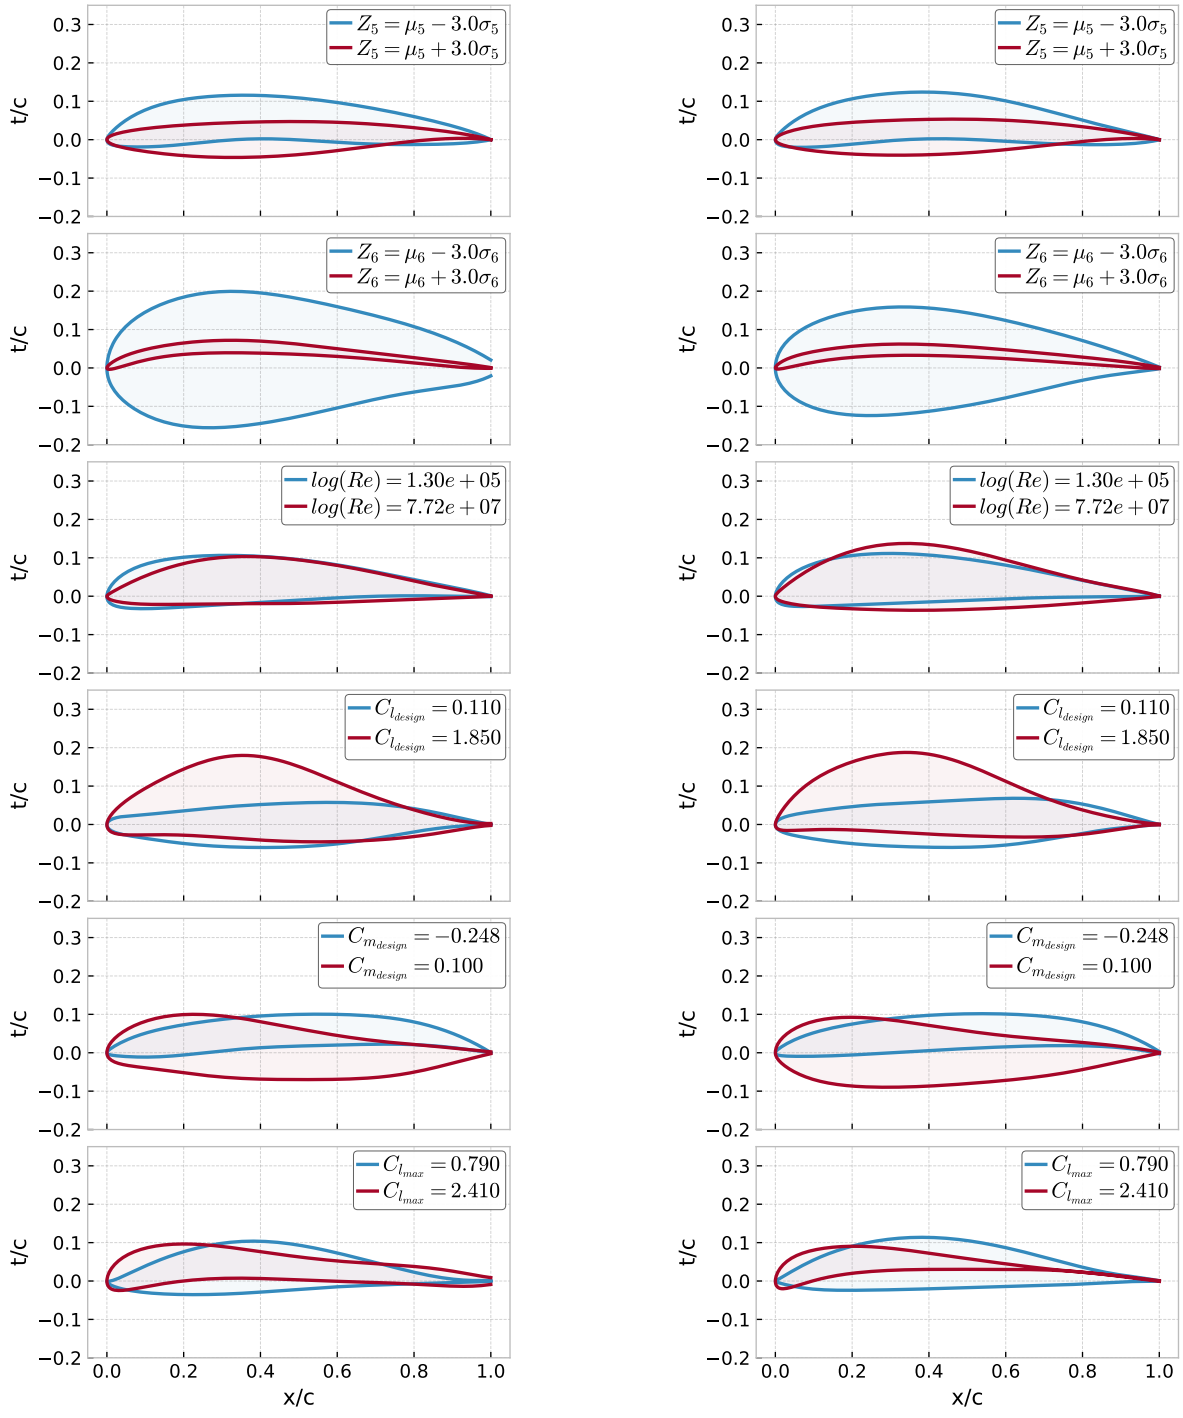
Plots showing the worst reconstructed geometries from the dataset, used as a visual indication of the lower bound of its reconstruction quality, are included in appendices A and B (specifically appendix A figure 11, and appendix B figures 13 and 14 for the VAE, normal-prior CVAE, and conditioned-prior CVAE, respectively).

The VAE has the lowest median coordinate loss (1.71×10^{-4}), while the normal-prior and conditioned-prior CVAEs have higher medians of 1.79×10^{-3} and 1.29×10^{-3}, respectively. The same trend is visible at the 90th percentile: 4.68×10^{-4} for the VAE, 5.18×10^{-3} for the normal-prior CVAE, and 3.69×10^{-3} for the conditioned-prior CVAE. These show that the conditional models are clearly solving a harder task. They must keep the geometry realistic while also allowing the decoder output to move toward prescribed aerodynamic conditions. The fact that a trade-off exists between geometry reconstruction and condition adherence also implies that fully disentangled features captured in each dimension, clearly separating the geometry and condition information, is not possible.

The inferior reconstruction accuracy of the conditioned models does not immediately diminish their value. Their value should instead be assessed from the quality of conditional generation and from how well prior samples match the requested aerodynamic targets.

4.2 Feature Extraction

The latent variables capture the geometric degrees of freedom that are not directly controlled by the condition vector. To get an idea of which features have been captured by the different latent dimensions and conditions, figures 4a and 4b show representative latent traversals for the normal-prior and conditioned-prior CVAE, respectively. Each plot shows the effect of varying one latent dimension or one condition dimension to three standard deviations from the mean in either direction, while keeping the remaining latent variables at their prior mean, and the remaining conditions at their mean baseline value (see table 1).



(a) Latent traversal for the normal-prior CVAE.

(b) Latent traversal for the conditioned-prior CVAE.

Figure 4 – Latent traversals for the normal-prior CVAE (a) and the conditioned-prior CVAE (b), with a limited range of ± 3 standard deviations. Inactive dimensions are omitted.

Both models have very similar latent traversals, with similar geometric features in the latent space or being represented through the aerodynamic condition space. Of the eight latent dimensions, just two remain active. Inactive dimensions which do not impact the output are omitted here, but the complete traversal plots are provided in appendix C. The features captured in those two dimensions are thickness, and a feature relating to camber which varies between cambered airfoil and a more neutral camber at the start of the chordline with a stronger camber curve toward the trailing edge. Just as for the base VAE the camber curvature is captured by z_5, and thickness is captured in z_6.

Both models show similar geometric variations associated with the Reynolds number condition, with the largest change occurring in the leading-edge radius. This trend is physically plausible, since

lower-Reynolds-number flows are more sensitive to leading-edge separation, whereas sharper leading edges can remain viable at higher Reynolds numbers due to the reduced relative influence of viscosity. In this case, the conditioned-prior CVAE also assigns slightly more thickness variation to this condition, which may explain why the thickness variation associated with z_6 is slightly reduced for this model variant. Reynolds number is an extrinsic flow parameter and does not prescribe the geometry directly. However, because the remaining aerodynamic conditions are evaluated at the sampled Reynolds number, the model can still learn correlations between Reynolds number and the geometric features required to satisfy those conditions. The fact that a similar trend appears in both models suggests that this correlation is present in the training data and is learned in a consistent way.

The $C_{l_{\text{design}}}$ condition controls a combination of camber and thickness, with the strongest effect on the upper surface of the airfoil. Both CVAE variants show similar trends: increasing $C_{l_{\text{design}}}$ increases the upper-surface camber and thickness, while the lower-surface shape remains comparatively less affected. The conditioned-prior model produces a slightly sharper and more angular leading edge for the high-lift case. The trailing-edge reflex visible in this case is likely related to the simultaneous requirement to maintain the prescribed $C_{m_{\text{design}}}$.

The $C_{m_{\text{design}}}$ condition captures mainly camber, most pronounced near the trailing edge for low values. This is combined with a low thickness and sharp leading edge for low moment values, and increased thickness and leading edge radius on the high end. This time it is the normal-prior CVAE which produces a sharper leading edge for higher-moment cases, combined with some lower surface reflex near the leading edge, where the conditioned-prior CVAE produces a more rounded leading edge with stronger curvature of the lower surface.

Increasing $C_{l_{\text{max}}}$ results in the point of maximum thickness shifting forward, and increased camber, while lower values shift the maximum thickness point backward, forming a sharp leading edge, and reduced camber. It is likely that the sharp leading edge and mid-chord maximum thickness are associated with early flow separation to achieve such low $C_{l_{\text{max}}}$ values.

Not all generated shapes shown are representative of practical airfoil geometries. However, the condition values used in these plots are close to the extremes of the training data and, in some cases, approach outlier or extrapolation territory. This is especially true for the $C_{l_{\text{max}}}$ sweeps, where the most extreme targets lead to geometries with very sharp leading edges, wavy surfaces, and thickness values close to or even exceeding infeasible limits. Note that for the low end of the $C_{l_{\text{max}}}$ sweep, this is partially caused due to the condition vector itself being infeasible: the value used is lower than the mean $C_{l_{\text{design}}}$ value, which is not physically realistic. Nevertheless, the geometric trends captured by the model remain broadly consistent with the expected aerodynamic behaviour.

4.3 Condition Adherence

The primary purpose of the CVAE is to generate geometries that match requested aerodynamic conditions. To verify how well the models achieve that, 2000 samples were generated using both CVAE variants. To ensure that the requested conditions are within the data-supported, feasible region of condition space, the target condition vectors are obtained using cross-data sampling: random entries are drawn from the training data and evaluated to obtain the condition vector for random Reynolds number. These condition vectors are then shuffled, perturbed with a small amount of noise, and combined with random latent samples from the prior. The respective decoded geometries are then evaluated and condition values are compared. When unrealistic or infeasible condition vectors are used, the model will still output a geometry, but the aerodynamic properties of that geometry will not match the requested conditions. This is demonstrated in appendix D.

Figure 5 compares the target and achieved conditions for both CVAE variants. The coefficient of determination (R^2) marked in each plot show that all three conditions are well-correlated with the targets. In general, differences between the two CVAE variants are negligible when considering the condition adherence of the generated samples. The conditioned-prior CVAE matches the imposed $C_{l_{\text{design}}}$ targets slightly better than the standard-prior CVAE, with an R^2 value of 0.828 compared with 0.818, and a mean absolute percentage error (MAPE) of 9.3% compared with 9.4%. Similarly, both

models perform almost identically when considering $C_{l_{\max}}$, with only a 0.1% difference in MAPE, and a 0.04 difference in R^2. For $C_{m_{\text{design}}}$, the conditioned-prior model even performs slightly worse, with a MAPE of 104.7% compared with 90.4% for the standard-prior model.

Overall, $C_{l_{\text{design}}}$ remains the hardest condition to match in terms of mean squared error (MSE), with MSE values of 1.50×10^{-2} and 1.42×10^{-2} for the standard-prior and conditioned-prior CVAEs, respectively. The largest MAPE is observed for $C_{m_{\text{design}}}$, with an average value of approximately 98% across the two model variants. This large percentage error is partly caused by the small magnitude of the moment coefficients, for which even modest absolute errors can translate into large relative errors.

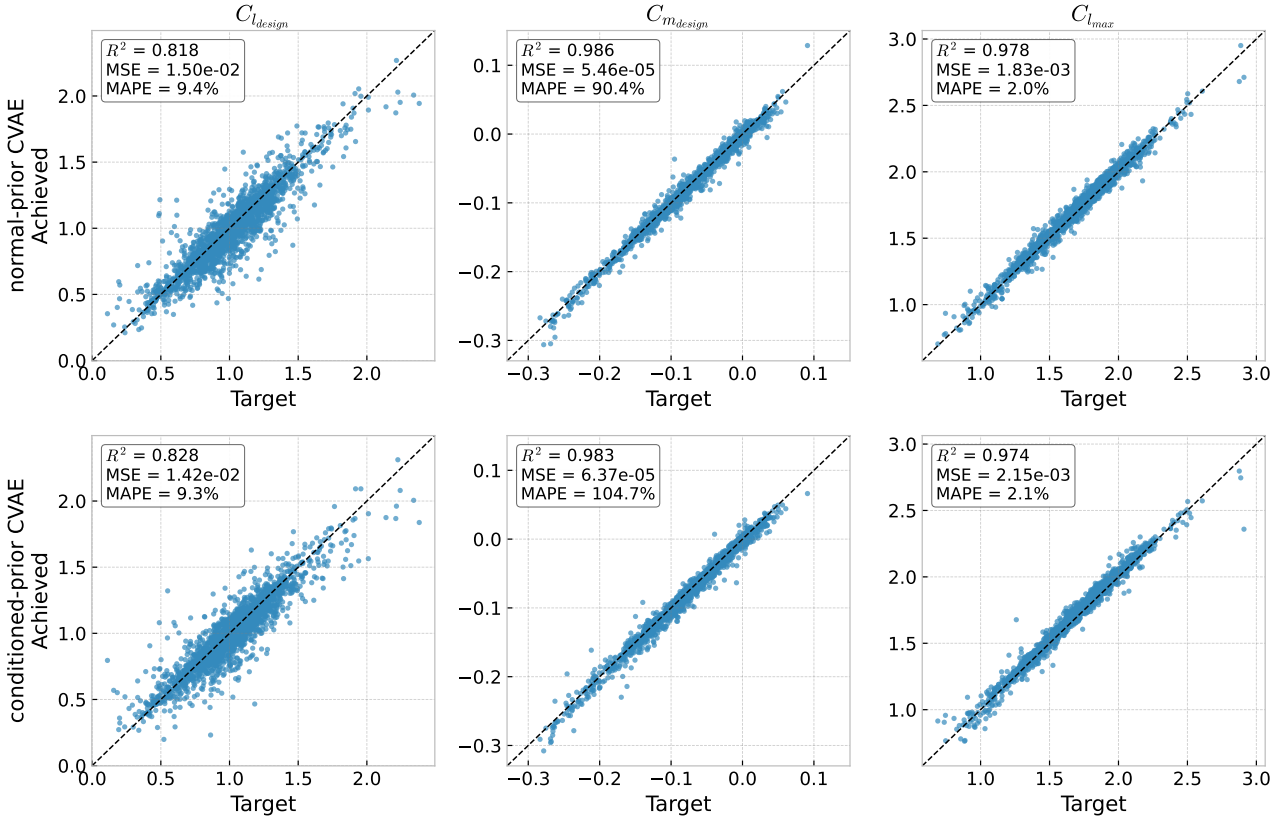


Figure 5 – Target and achieved conditions for 2000 generated samples obtained using both CVAE variants. R^2 is the coefficient of determination between the target and achieved conditions, and MSE and MAPE are the mean squared error and mean absolute percentage error, respectively.

4.4 Condition Response

While the scatter plots in figure 5 show how closely the generated airfoils match complete target condition vectors drawn from realistic operating points, they do not by themselves show how directly each conditioning variable acts as an independent design control. To get a better idea of that behavior, figure 6 introduces condition-response plots in which one target condition is varied systematically while the remaining conditions are held fixed at a representative baseline value. The resulting curves show how the achieved aerodynamic response changes as a single design target is adjusted, while the shaded spread reflects the remaining variation caused by latent sampling. In this way, the scatter plots show overall adherence for realistic requests, whereas the response plots show how clearly and how independently each condition can be used as a design knob. As expected, the spread across latent samples is largest for $C_{l_{\text{design}}}$, indicating that the latent variables still affect the achieved lift when the target lift is fixed. This means that the model can produce multiple geometries for a similar target, but it also limits precision when exact aerodynamic matching is required.

Again, both models perform similarly. The conditioned-prior CVAE is closer to the target line in some regions, while the standard-prior CVAE performs slightly better in others. Note that this time, a correction is put in place to keep the lift values in the condition vector feasible, ensuring the maximum lift coefficient is above the design lift coefficient. When needed, the $C_{l_{\max}}$ target is increased to be at least 0.1 above the $C_{l_{\text{design}}}$ target.

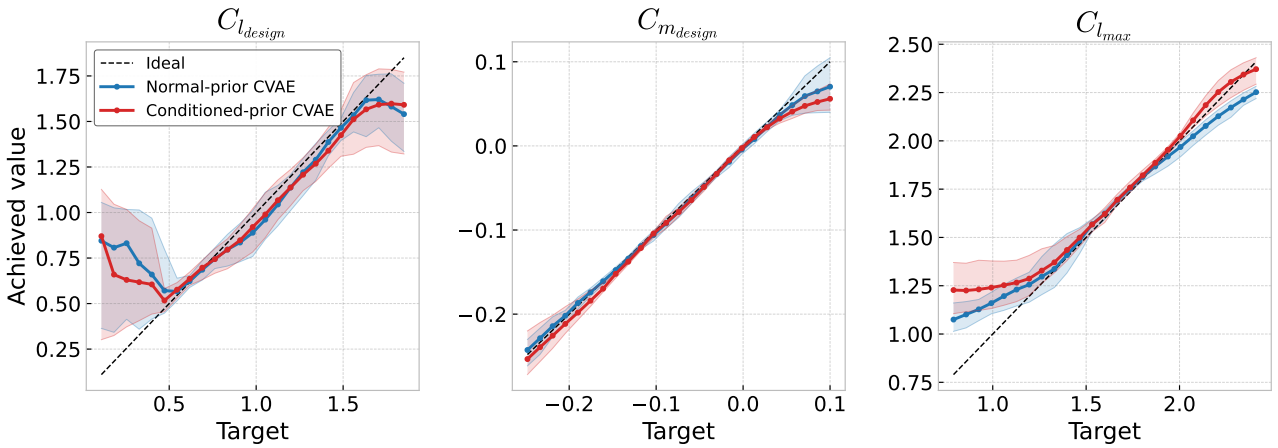


Figure 6 – Condition sweep response for both CVAE variants. A constraint is applied which ensures $C_{l_{max}} \geq C_{l_{design}} + 0.1$.

Tables 2 and 3 show the generated airfoils for representative points along each sweep, again for the normal-prior and conditioned-prior CVAE variants respectively. They confirm that the numerical trends correspond to plausible geometric changes. For the extreme ends of the sweep, we see the geometries matching those previously seen in the traversals in figures 4a and 4b.

Table 2 – Condition sweep geometries for the Normal-prior CVAE. The complementary lift condition is adjusted if needed to ensure $C_{l_{max}} \geq C_{l_{design}} + 0.1$

$C_{l_{design}}$			$C_{m_{design}}$			$C_{l_{max}}$		
Target	True	Geometry	Target	True	Geometry	Target	True	Geometry
0.11	0.998		−0.248	−0.245		0.79	1.12	
0.545	0.509		−0.161	−0.163		1.2	1.28	
0.98	0.896		−0.074	−0.0798		1.6	1.64	
1.42	1.44		0.013	0.00301		2.01	1.99	
1.85	1.53		0.1	0.0655		2.41	2.27	

The target and true columns show that the generated geometries do not always exactly realise the requested condition, especially near the extremes of the sweeps. This is a consequence of demanding extremes from these conditions, while keeping the remaining ones at their mean. This forces the model to extrapolate, where the decoder has not been trained to satisfy such targets, or the solution might simply not exist. These examples are useful because they show the main practical limitation of the method: the decoder can guide the design toward the requested aerodynamic region, but the final geometry should still be checked with an aerodynamic model before being accepted. It is also important to consider that for those non-realistic geometries, the output of Neuralfoil is likely inaccurate and using XFOIL would likely simply fail outright. Therefore user discretion remains important. Tables including all condition values are available in appendix F.

Table 3 – Condition sweep geometries for the Conditioned-prior CVAE. The complementary lift condition is adjusted if needed to ensure $C_{l_{max}} \geq C_{l_{design}} + 0.1$

$C_{l_{design}}$			$C_{m_{design}}$			$C_{l_{max}}$		
Target	True	Geometry	Target	True	Geometry	Target	True	Geometry
0.11	1.03		−0.248	−0.267		0.79	1.22	
0.545	0.53		−0.161	−0.174		1.2	1.26	
0.98	0.903		−0.074	−0.0799		1.6	1.65	
1.42	1.34		0.013	0.0104		2.01	2.03	
1.85	1.74		0.1	0.0457		2.41	2.4	

4.5 Condition-space Structure

So far, only variations around the means of single conditions have been considered. This does not give a complete picture of how the CVAE variants perform across the full condition space. In fact, the conditioned-prior CVAE is designed to have a condition-dependent latent distribution, so it may perform better in some regions of condition space than in others. Moreover, aerodynamic properties are almost always coupled, so changing one target condition may affect the achieved values of the other conditions.

The sensitivity heatmaps in figure 7 show how the conditions respond to changes in the target conditions. The normalized sweep slope s and Pearson correlation r between the swept target and the achieved response are marked in each cell. The normalized slope measures the strength of the response and calibration, with values near unity on the diagonal indicating that the achieved condition changes by approximately the same normalized amount as the swept target, while off-diagonal slopes quantify the magnitude of cross-talk between conditions. The Pearson correlation measures the consistency of the response over the sweep, independent of its scale. A high correlation therefore indicates that the achieved condition varies monotonically with the swept target, even if the slope shows that the response is too weak or too strong. When considering the diagonal terms, only the aerodynamic conditions are taking into account, as the Reynolds number is only a flow condition which remains constant. The models did however associate it with geometric features as seen in section 4.2, so it will have non-zero slopes.

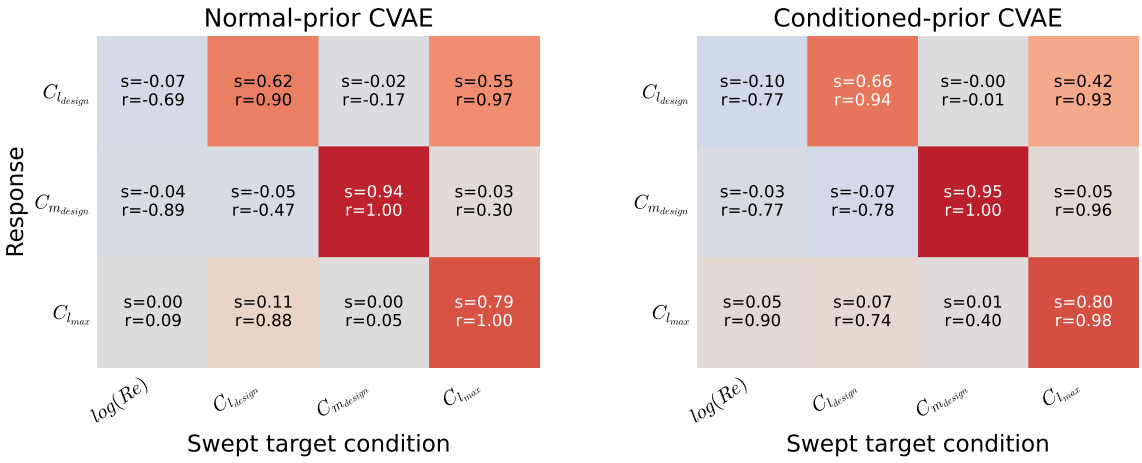


Figure 7 – Condition sensitivity heatmaps for the normal-prior and conditioned-prior CVAEs. The normalized sweep slope s, and Pearson correlation r are marked in each cell.

The largest correlations and slopes are observed for the diagonal terms, confirming that the conditions respond well to changes in their own targets. Most off-diagonal terms have smaller local slopes, indicating little effect in the observed point, but also significant correlations over the full sweep, confirming that conditions will affect each other in certain regions. This is a consequence of the coupled nature of aerodynamic properties, and it means that the CVAE is not able to treat each condition as a fully independent control knob. The significant sensitivity and correlation between $C_{l_{\max}}$ and $C_{l_{\text{design}}}$, while not surprising, is partially induced by the requirement that $C_{l_{\max}}$ must be at least 0.1 higher than $C_{l_{\text{design}}}$ during sampling. This is not applied as a bound, but rather the respective value is modified from its default value when the swept parameter gets too close, artificially introducing a slope.

The conditioned-prior CVAE manages to slightly reduce the off-diagonal correlations for most cases, while slightly increasing the diagonal slopes. Using a condition-dependent prior therefore does give some improvement in the independence of the condition controls, but it does not eliminate the inherent coupling between aerodynamic properties.

To evaluate how the model performs away from the mean values, the same sensitivity analysis is repeated for target conditions varied around chosen reference airfoil geometries instead, which is a representative sample from the training data. The geometries chosen are the NACA0018, DAE11, and Whitcomb, airfoils, each with distinct features and practical purpose, and respective reference Reynolds numbers $Re_{NACA0018} = 1 \times 10^6$, $Re_{DAE11} = 5 \times 10^5$, and $Re_{Whitcomb} = 2 \times 10^7$. These airfoils are plotted in figure 8. Note that these reference geometries are not exact copies of their occurrence in the training data, but rather their reconstructed form, with corresponding reconstruction inaccuracies induced by the models.



Figure 8 – Reference airfoil geometries used for the condition-space analyses away from the mean.

The sensitivity heatmaps obtained from sweeping over the values of these reconstructed reference airfoils are shown in figure 9. Not only are the off-diagonal sensitivities and correlations now much more present, but also the diagonal sensitivities are now lower for some. For example, changing the $C_{l_{\max}}$ target for the Whitcomb reference geometry without affecting the other conditions is clearly more difficult than for the mean conditions. It is now also common for the $C_{l_{\text{design}}}$ to be more sensitive to changes in the $C_{l_{\max}}$ target than to changes in its own target. Given the better adherence to the $C_{l_{\max}}$, it appears the model prioritises this target over the $C_{l_{\text{design}}}$ target. These examples show that the sensitivity of the achieved conditions to changes in the target conditions varies depending on the reference airfoil, and is thus not uniform across condition space. The coupling between aerodynamic properties is context-dependent and can be influenced by the specific shape being considered.

The response plots in figure 10 show the effect of modifying one target condition at a time for each reference airfoil. These confirm that the model struggles to achieve the target $C_{l_{\text{design}}}$ near the encoded reference values. This also shows that, compared to the mean condition sweeps, the conditioned prior now causes the model to be slightly worse at matching exact values for $C_{l_{\text{design}}}$ compared to the normal prior model, but overall does appear to have learned a smoother response with more consistent slopes across the sweeps.

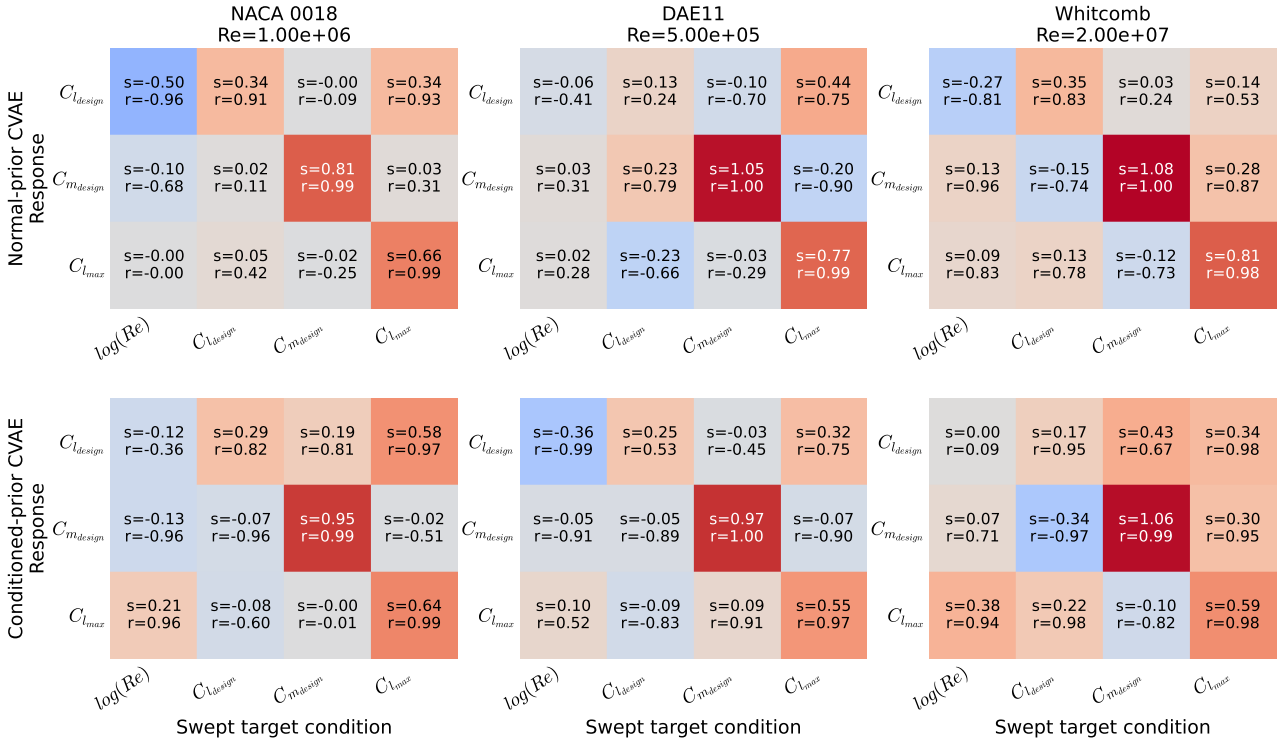


Figure 9 – Condition sensitivity heatmaps with values varied around reference airfoil conditions.

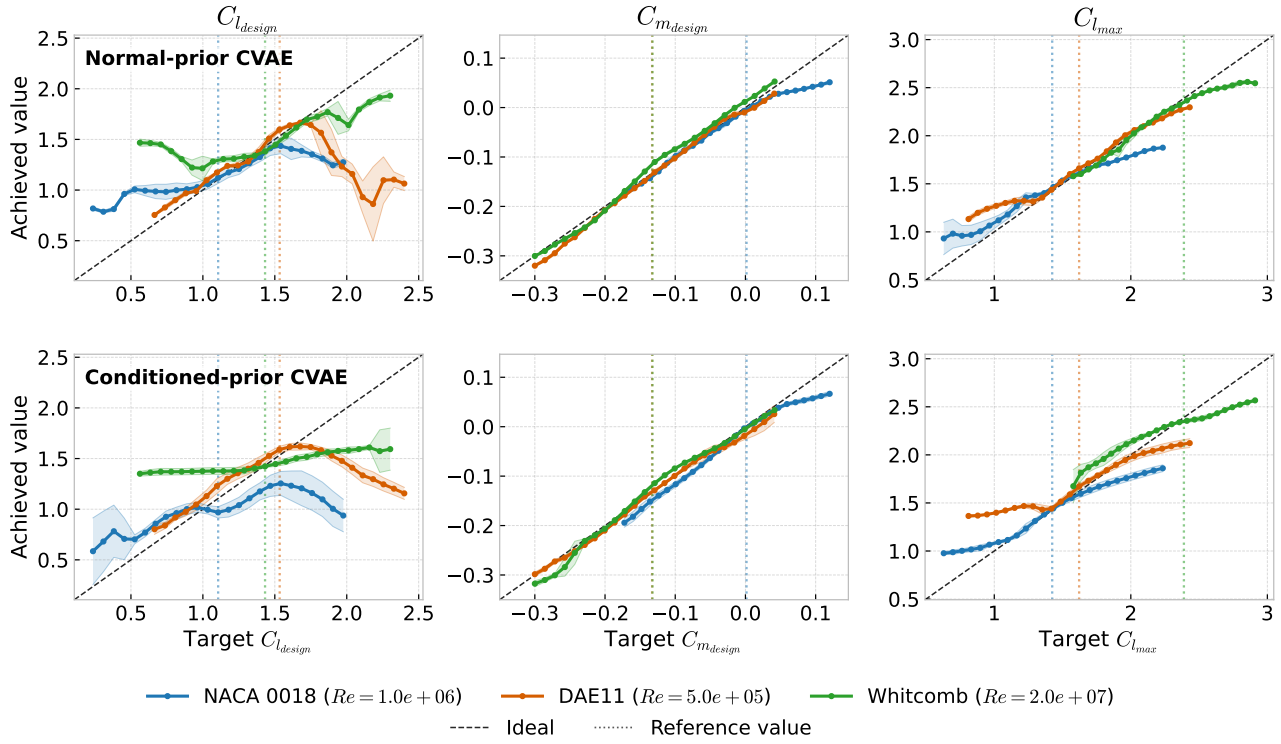


Figure 10 – Condition response for target sweeps around encoded reference airfoil baselines. The dotted vertical lines mark the encoded reference value for each airfoil.

The target conditions, corresponding achieved values, and generated geometries associated with these sweeps are presented in the geometry tables in tables 4 and 5 for the normal-prior and conditioned-prior CVAE variants, respectively. Detailed sweep data including all condition values is provided in appendix F tables 12 to 14.

Table 4 – Reference airfoil condition sweep geometries using the normal-prior CVAE. Reynolds numbers are fixed during sweeps. When sweeping $C_{l_{\text{design}}}$ or $C_{l_{\text{max}}}$, the complementary lift condition is adjusted if needed so $C_{l_{\text{max}}} \geq C_{l_{\text{design}}} + 0.1$. Dashed lines indicate the reference geometry.

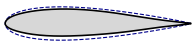
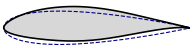
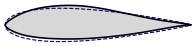
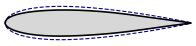
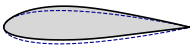
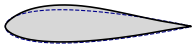
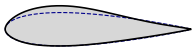
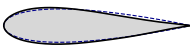
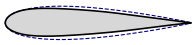
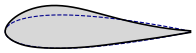
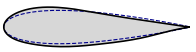
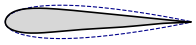
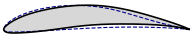
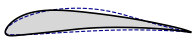
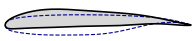
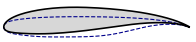
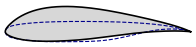
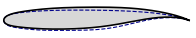
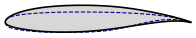
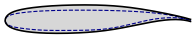
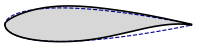
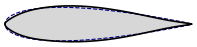
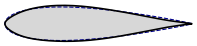
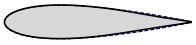
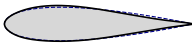
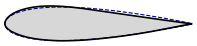
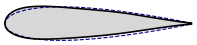
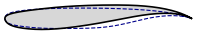
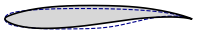
$C_{l_{design}}$			$C_{m_{design}}$			$C_{l_{max}}$		
Target	True	Geometry	Target	True	Geometry	Target	True	Geometry
NACA 0018 $Re = 1 \times 10^6$								
0.671	0.977		−0.0855	−0.0899		1.02	1.19	
0.889	1		−0.042	−0.0471		1.22	1.36	
1.11	1.1		0.00154	−0.00484		1.42	1.47	
1.32	1.27		0.045	0.0257		1.63	1.62	
1.54	1.33		0.0885	0.04		1.83	1.72	
DAE11 $Re = 5 \times 10^5$								
1.09	1.17		−0.22	−0.236		1.22	1.29	
1.3	1.26		−0.176	−0.183		1.42	1.42	
1.52	1.61		−0.133	−0.136		1.62	1.67	
1.74	1.84		−0.0894	−0.091		1.82	1.83	
1.96	1.83		−0.0459	−0.0414		2.03	2.05	
Whitcomb $Re = 2 \times 10^7$								
0.997	1.21		−0.219	−0.233		1.98	1.99	
1.21	1.33		−0.176	−0.173		2.19	2.19	
1.43	1.35		−0.132	−0.113		2.39	2.34	
1.65	1.61		−0.0889	−0.0717		2.59	2.47	
1.87	1.79		−0.0454	−0.0326		2.79	2.55	

Table 5 – Reference airfoil condition sweep geometries using the conditioned-prior CVAE. Reynolds numbers are fixed during sweeps. When sweeping $C_{l_{\text{design}}}$ or $C_{l_{\text{max}}}$, the complementary lift condition is adjusted if needed so $C_{l_{\text{max}}} \geq C_{l_{\text{design}}} + 0.1$. Dashed lines indicate the reference geometry.

$C_{l_{design}}$			$C_{m_{design}}$			$C_{l_{max}}$		
Target	True	Geometry	Target	True	Geometry	Target	True	Geometry
NACA 0018 $Re = 1 \times 10^6$								
0.671	0.887		−0.0855	−0.104		1.02	1.1	
0.889	1.04		−0.042	−0.0541		1.22	1.22	
1.11	0.952		0.00154	0.000934		1.42	1.44	
1.32	1.06		0.045	0.0388		1.63	1.59	
1.54	1.19		0.0885	0.0547		1.83	1.69	
DAE11 $Re = 5 \times 10^5$								
1.09	1.23		−0.22	−0.23		1.22	1.42	
1.3	1.44		−0.176	−0.181		1.42	1.42	
1.52	1.61		−0.133	−0.131		1.62	1.68	
1.74	1.62		−0.0894	−0.0848		1.82	1.83	
1.96	1.65		−0.0459	−0.0453		2.03	1.99	
Whitcomb $Re = 2 \times 10^7$								
0.997	1.36		−0.219	−0.223		1.98	2.13	
1.21	1.36		−0.176	−0.176		2.19	2.25	
1.43	1.42		−0.132	−0.117		2.39	2.34	
1.65	1.52		−0.0889	−0.0737		2.59	2.41	
1.87	1.56		−0.0454	−0.0454		2.79	2.52	

5 Conclusion

In this study, we investigated the use of Conditional Variational Autoencoders (CVAEs) for inverse airfoil design, where Kulfan-parameterized airfoil geometries are generated from prescribed aerodynamic condition vectors evaluated with NeuralFoil. Two different CVAE variants were considered: one based on a fixed standard-normal latent prior, and one based on a condition-dependent prior whose distribution parameters are predicted from the target aerodynamic conditions. Both CVAE variants were compared against a geometry-only VAE baseline to assess the reconstruction accuracy lost when aerodynamic conditioning is introduced, and were then compared to each other in terms of learned latent and condition-space structure, condition adherence, and sensitivity to target-condition sweeps.

The results show that the CVAE framework can condition airfoil geometry generation on aerodynamic quantities and produce meaningful, controllable shape variation within data-supported regions of condition space. The generated airfoils respond to changes in the requested lift, moment, maximum-lift, and Reynolds-number conditions, and the learned latent variables retain geometric features such as thickness and camber variation. The method is therefore useful as an inverse-design aid, but the conditioning should be interpreted as probabilistic guidance rather than exact constraint satisfaction. The generated shapes remain samples from a learned distribution, and their final aerodynamic properties should be verified by a surrogate or higher-fidelity solver before further design use.

The condition-response and sensitivity analyses show that the aerodynamic conditions cannot be treated as fully independent decision variables for inverse design problems. The diagonal responses confirm that the models generally change the intended aerodynamic quantity when a target condition is swept, but the off-diagonal sensitivities show that changes in one target can also move the other achieved conditions. This coupling is strongest away from the mean condition values and depends on the region of the condition space being considered. Condition parameters should therefore be selected so that they are physically meaningful and not too strongly coupled, although the latter might not be possible when dealing with aerodynamic properties.

For this airfoil-generation task, the training behaviour also suggests that a probabilistic decoder with a negative log-likelihood reconstruction loss is not necessarily the most effective choice when reconstruction accuracy is the main objective. Removing the stochastic output variance may reduce regularization, but a deterministic decoder with a direct coordinate-space loss can provide a more informative training signal for accurate geometry recovery. This is especially relevant when the geometry is represented by a relatively low-dimensional curve parameterization, where small parameter errors can still lead to visible coordinate differences. Therefore, even if the model is trained to generate shapes through curve parameters, the reconstruction loss should be calculated in coordinate space.

The condition-dependent prior provides a small improvement in reconstruction quality, suggesting that allowing the latent prior to depend on the aerodynamic conditions helps the model represent the conditional geometry distribution more effectively. However, the achieved aerodynamic values are not consistently closer to the prescribed targets than for the normal-prior CVAE. In several cases, the conditioned-prior model shows a more continuous geometric response to changes in the condition vector, with fewer visible jumps or fluctuations. However, this does not always translate into better agreement with the target aerodynamic values. The conditioned prior should therefore be viewed as a useful extension, rather than as a generally better replacement for the fixed normal prior. Future work could investigate whether a more expressive conditional prior, an alternative prior formulation, or different conditioning variables can improve condition adherence.

Overall, the results support the use of CVAEs as practical generative tools for airfoil inverse design, provided that their outputs are interpreted correctly. They can guide geometry generation toward requested aerodynamic performance and produce plausible shape variation, but they do not replace aerodynamic verification. Focussed regularization does result in an interpretable latent and condition space, but cannot remove the inherent coupling between performance quantities. With carefully defined conditioning variables and realistic target ranges, CVAEs can provide a compact and controllable starting point for airfoil design exploration.

6 Contact Author Email Address

k.swannet@tudelft.nl

7 Copyright Statement

The authors confirm that they, and/or their company or organization, hold copyright on all of the original material included in this paper. The authors also confirm that they have obtained permission, from the copyright holder of any third party material included in this paper, to publish it as part of their paper. The authors confirm that they give permission, or have obtained permission from the copyright holder of this paper, for the publication and distribution of this paper as part of the ICAS proceedings or as individual off-prints from the proceedings.

References

[1] Kingma D P and Welling M. Auto-Encoding Variational Bayes. 2013. doi: 10.48550/ARXIV.1312.6114.

[2] Swannet K, Varriale C, and Doan N. Latent space correlation for interpretable airfoil parameterization using variational autoencoders. In *34th Congress of the International Council of the Aeronautical Sciences (ICAS)*, Florence, Italy, sept 2024. URL https://www.icas.org/icas_archive/icas2024/data/papers/icas2024_0484_paper.pdf.

[3] Kang Y E, Lee D, and Yee K. Physically interpretable airfoil parameterization using variational autoencoder-based generative modeling. In *AIAA SCITECH 2024 Forum*, Orlando, Florida, Jan. . doi: 10.2514/6.2024-0685.

[4] Sohn K, Lee H, and Yan X. Learning structured output representation using deep conditional generative models. In Cortes C, Lawrence N, Lee D, Sugiyama M, and Garnett R, editors, *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc., 2015. URL https://proceedings.neurips.cc/paper_files/paper/2015/file/8d55a249e6baa5c06772297520da2051-Paper.pdf.

[5] Yonekura K and Suzuki K. Data-driven design exploration method using conditional variational autoencoder for airfoil design. *Structural and Multidisciplinary Optimization*, 64(2):613–624, Aug. 2021. ISSN 1615-147X, 1615-1488. doi: 10.1007/s00158-021-02851-0.

[6] Yonekura K, Wada K, and Suzuki K. Generating various airfoils with required lift coefficients by combining NACA and Joukowski airfoils using conditional variational autoencoders. *Engineering Applications of Artificial Intelligence*, 108:104560, Feb. 2022. ISSN 09521976. doi: 10.1016/j.engappai.2021.104560.

[7] Yonekura K, Tomori Y, and Suzuki K. Airfoil shape generation and feature extraction using the conditional VAE-WGAN-gp. *AI*, 5(4):2092–2103, 2024. doi: 10.3390/ai5040102.

[8] Xie H, Wang J, and Zhang M. Parametric generative schemes with geometric constraints for encoding and synthesizing airfoils, 2023. URL https://arxiv.org/abs/2205.02458.

[9] Masood Z, Khan S, Kostas K, and Kaklis P. Physics-informed geometric operators in generative airfoil design with hybrid VAEs. *Journal of Computational Design and Engineering*, 13(4):51–66, 2026. doi: 10.1093/jcde/qwag025.

[10] Selig M. Uiuc airfoil data site. https://m-selig.ae.illinois.edu/ads/coord_database.html, 1996.

[11] Kulfan B M. Universal Parametric Geometry Representation Method. *Journal of Aircraft*, 45(1):142–158, Jan. 2008. ISSN 0021-8669, 1533-3868. doi: 10.2514/1.29958.

[12] Kulfan B M. Modification of cst airfoil representation methodology, 08 2020. URL https://www.researchgate.net/publication/343615711_Modification_of_CST_Airfoil_Representation_Methodology.

[13] Sharpe P. NeuralFoil: An airfoil aerodynamics analysis tool using physics-informed machine learning. https://github.com/peterdsharpe/NeuralFoil, 2023.

[14] Drela M. Xfoil: An analysis and design system for low reynolds number airfoils. In Mueller T J, editor, *Low Reynolds Number Aerodynamics*, pages 1–12, Berlin, Heidelberg, 1989. Springer Berlin Heidelberg. ISBN 978-3-642-84010-4. doi: 10.1007/978-3-642-84010-4_1.

[15] Tomczak J M. *Deep Generative Modeling*. Springer International Publishing, 2 edition, 2022. ISBN 978-3-030-93157-5 978-3-030-93158-2. doi: 10.1007/978-3-030-93158-2.

[16] Burgess C P, Higgins I, Pal A, Matthey L, Watters N, Desjardins G, and Lerchner A. Understanding disentangling in beta-vae. *arXiv preprint arXiv:1804.03599*, 2018. doi: 10.48550/arXiv.1804.03599. URL `http://arxiv.org/abs/1804.03599`.

[17] Biewald L. Experiment tracking with weights and biases, 2020. URL `https://www.wandb.com/`. Software available from wandb.com.

A Reference VAE Results

Figure 11 shows the airfoils with the largest reconstruction loss. These are the worst reconstruction cases of the reference VAE, which was used as the initial checkpoint for the CVAE model training. The reference geometry is shown in solid lines, and the reconstructed geometry in dashed lines, and the total mean squared error is indicated in each subplot.

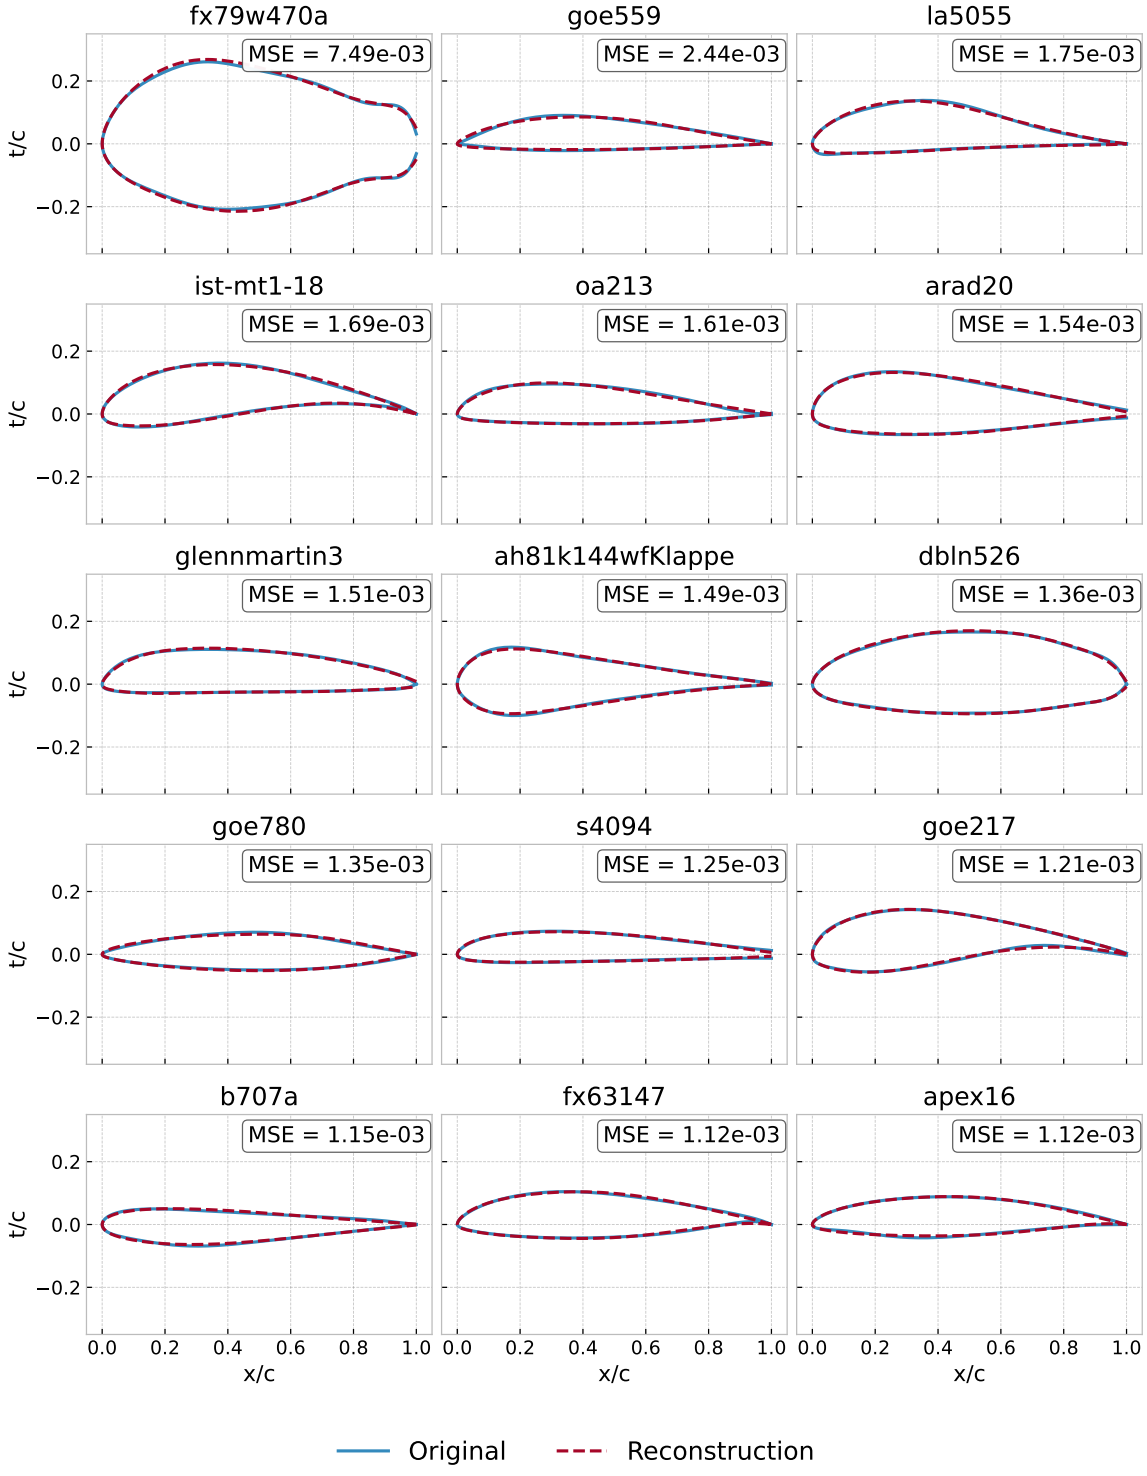


Figure 11 – Worst-case reconstruction cases for the reference VAE.

Figure 12 shows the latent traversal for the reference VAE. This gives an indication of how the initial checkpoint was already regularized, with certain latent directions corresponding to specific geometric features.

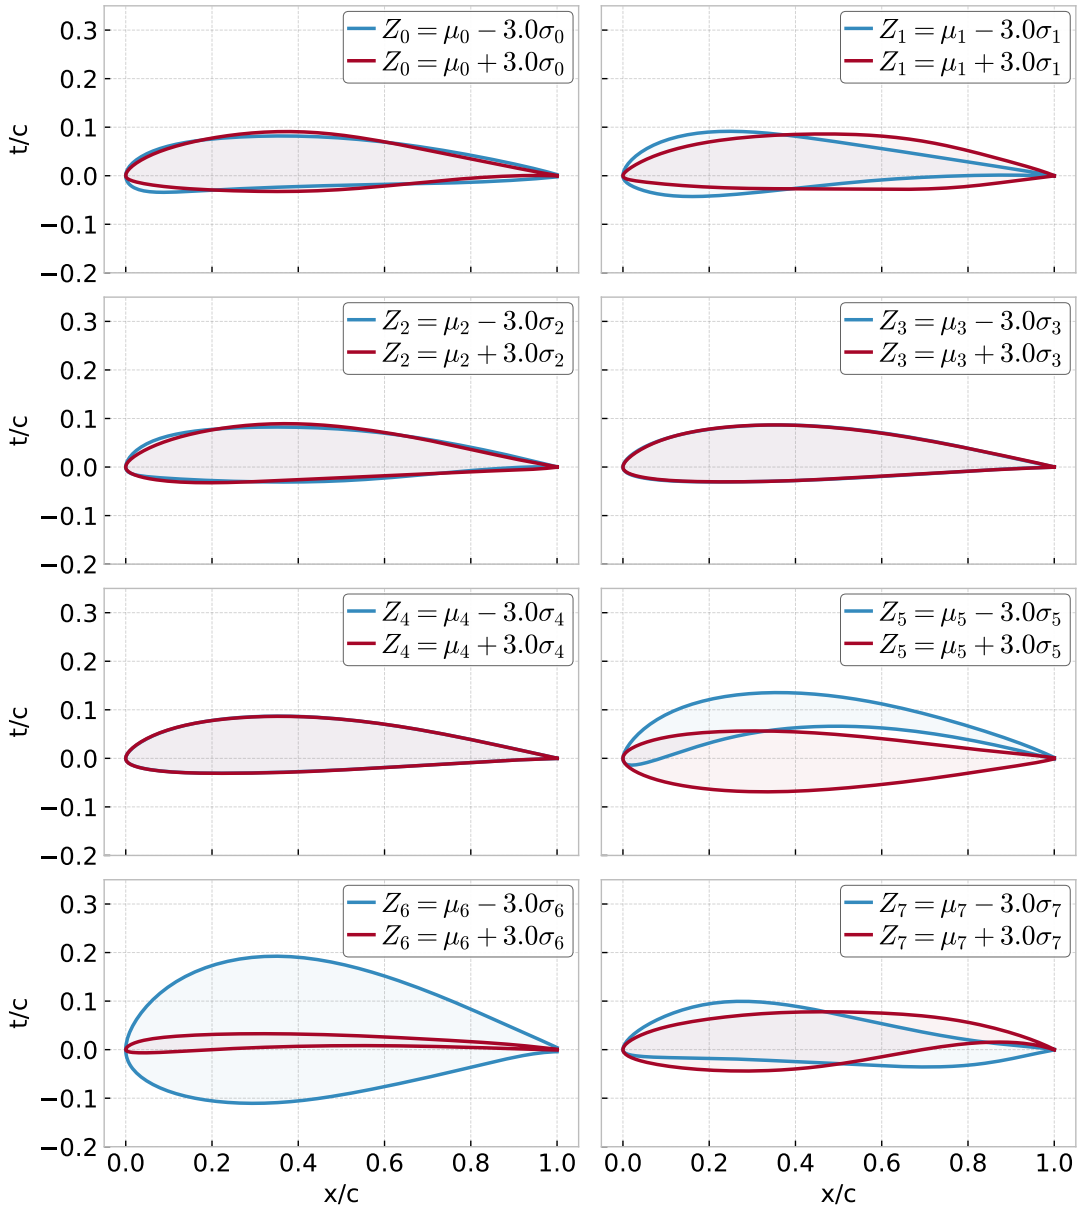


Figure 12 – Latent traversal for the base VAE.

B CVAE Reconstruction plots

Figure 13 shows the worst reconstructions for the normal-prior CVAE, and figure 14 for the conditioned-prior CVAE. They provide a visual indication of the lower bound of reconstruction quality for each model, and to show how the worst-case reconstructions differ between models. Compared to figure 11, they also highlight the loss of reconstruction quality after introducing the aerodynamic conditions.

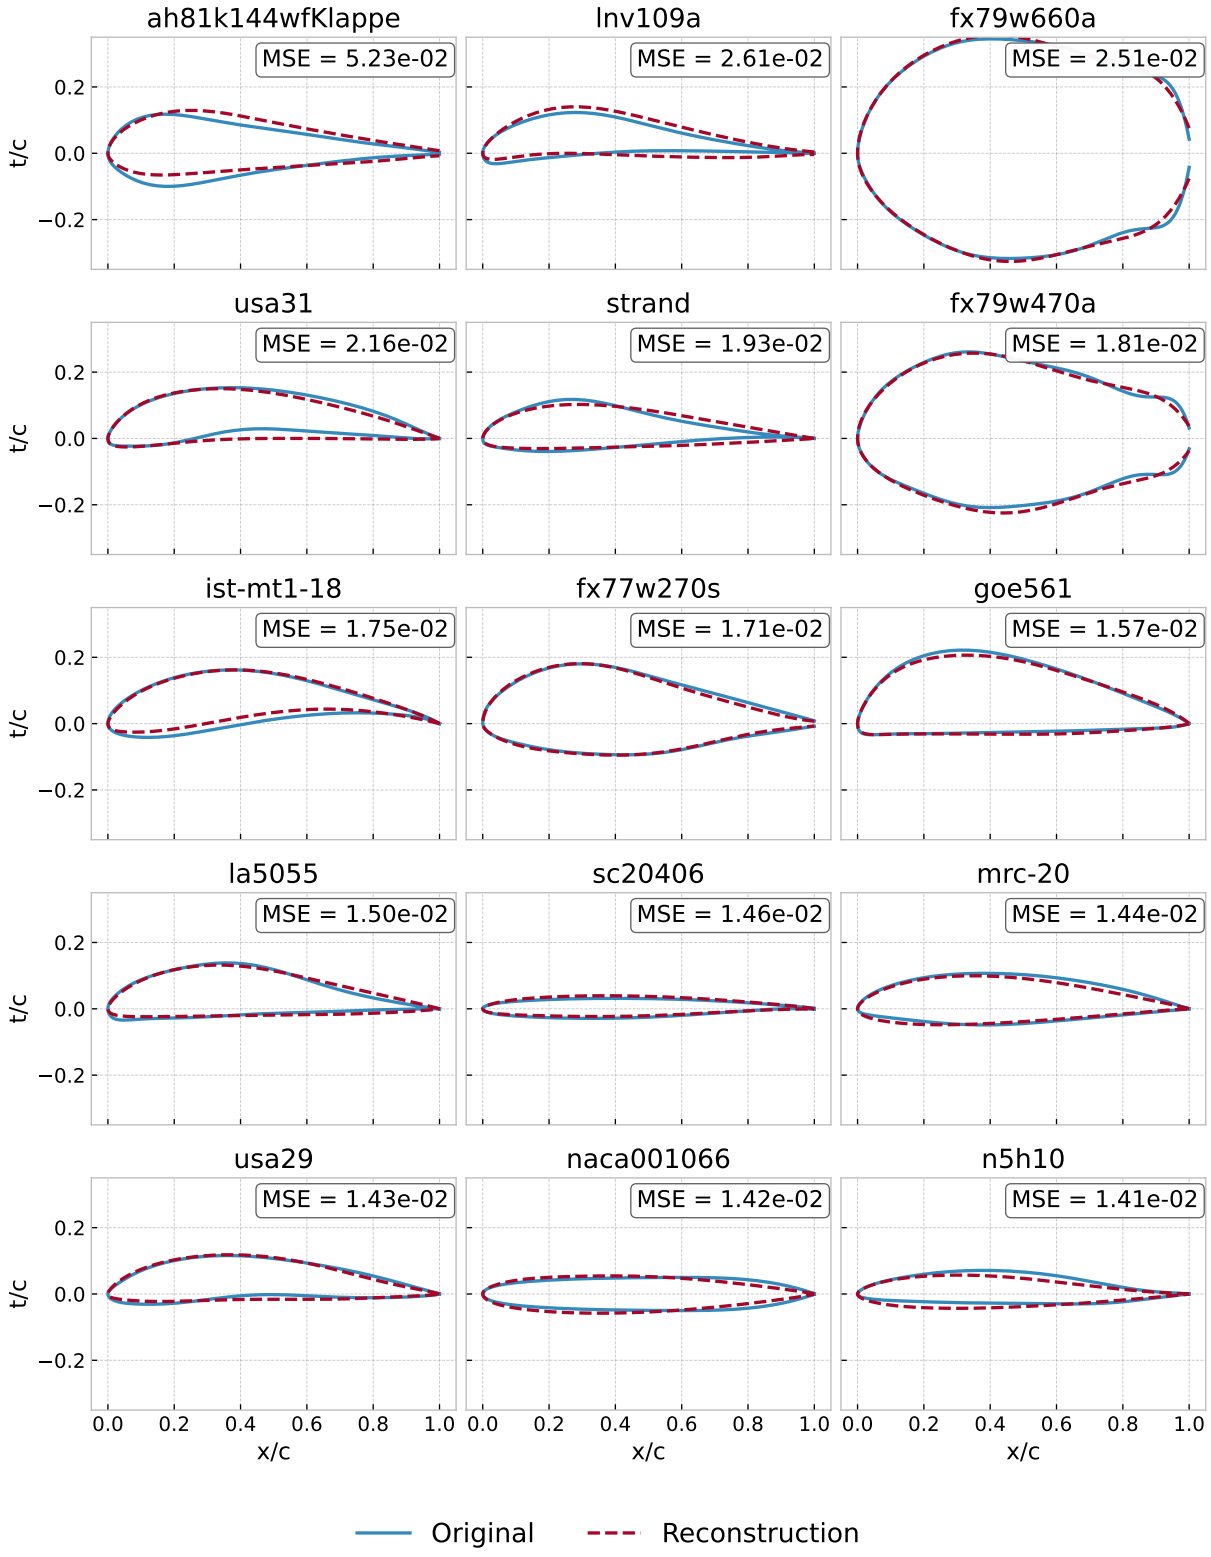


Figure 13 – Worst-case reconstruction cases for the normal-prior CVAE.

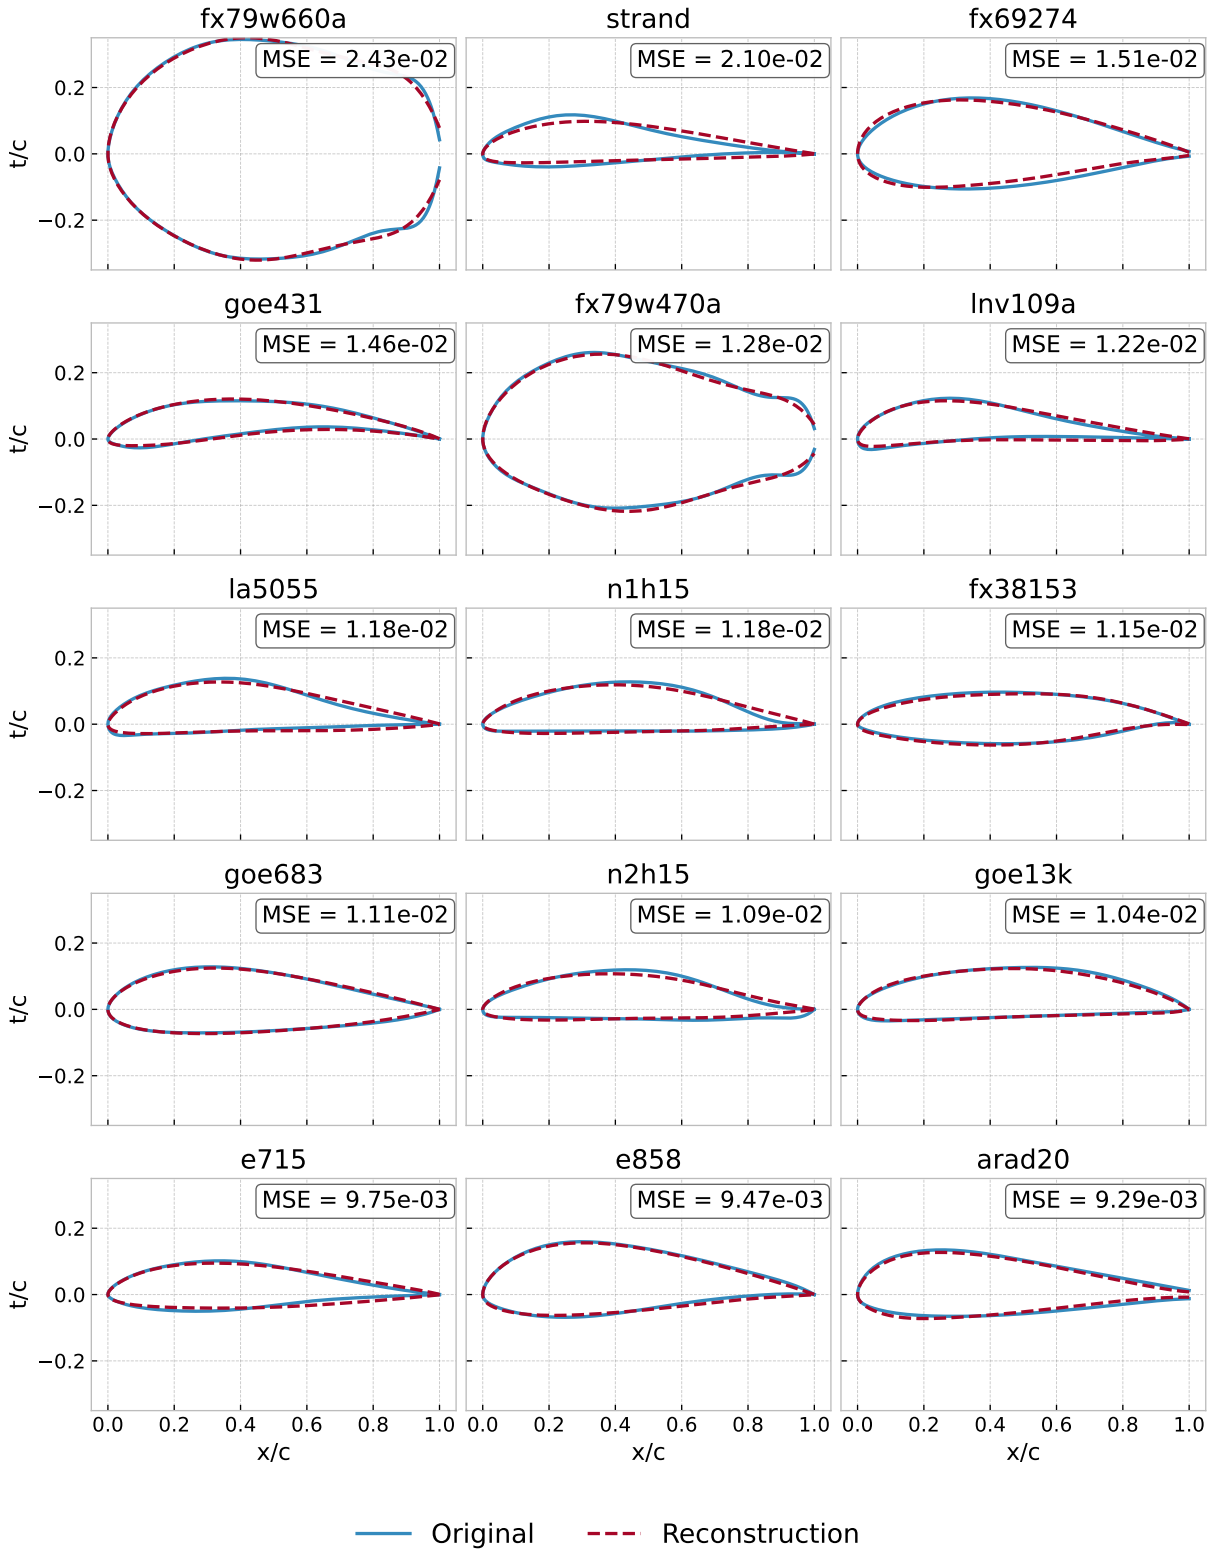


Figure 14 – Worst-case reconstruction cases for the conditioned-prior CVAE.

C CVAE Complete Latent Traversal Plots

The full traversal plots for the normal-prior CVAE and conditioned-prior CVAE are shown in figures 15 and 16 respectively. These show the effect of varying each latent variable individually, while keeping the others fixed at the mean of the prior distribution. These plots are identical to the ones shown in figures 4a and 4b, but now include all latent dimensions, not only the active ones.

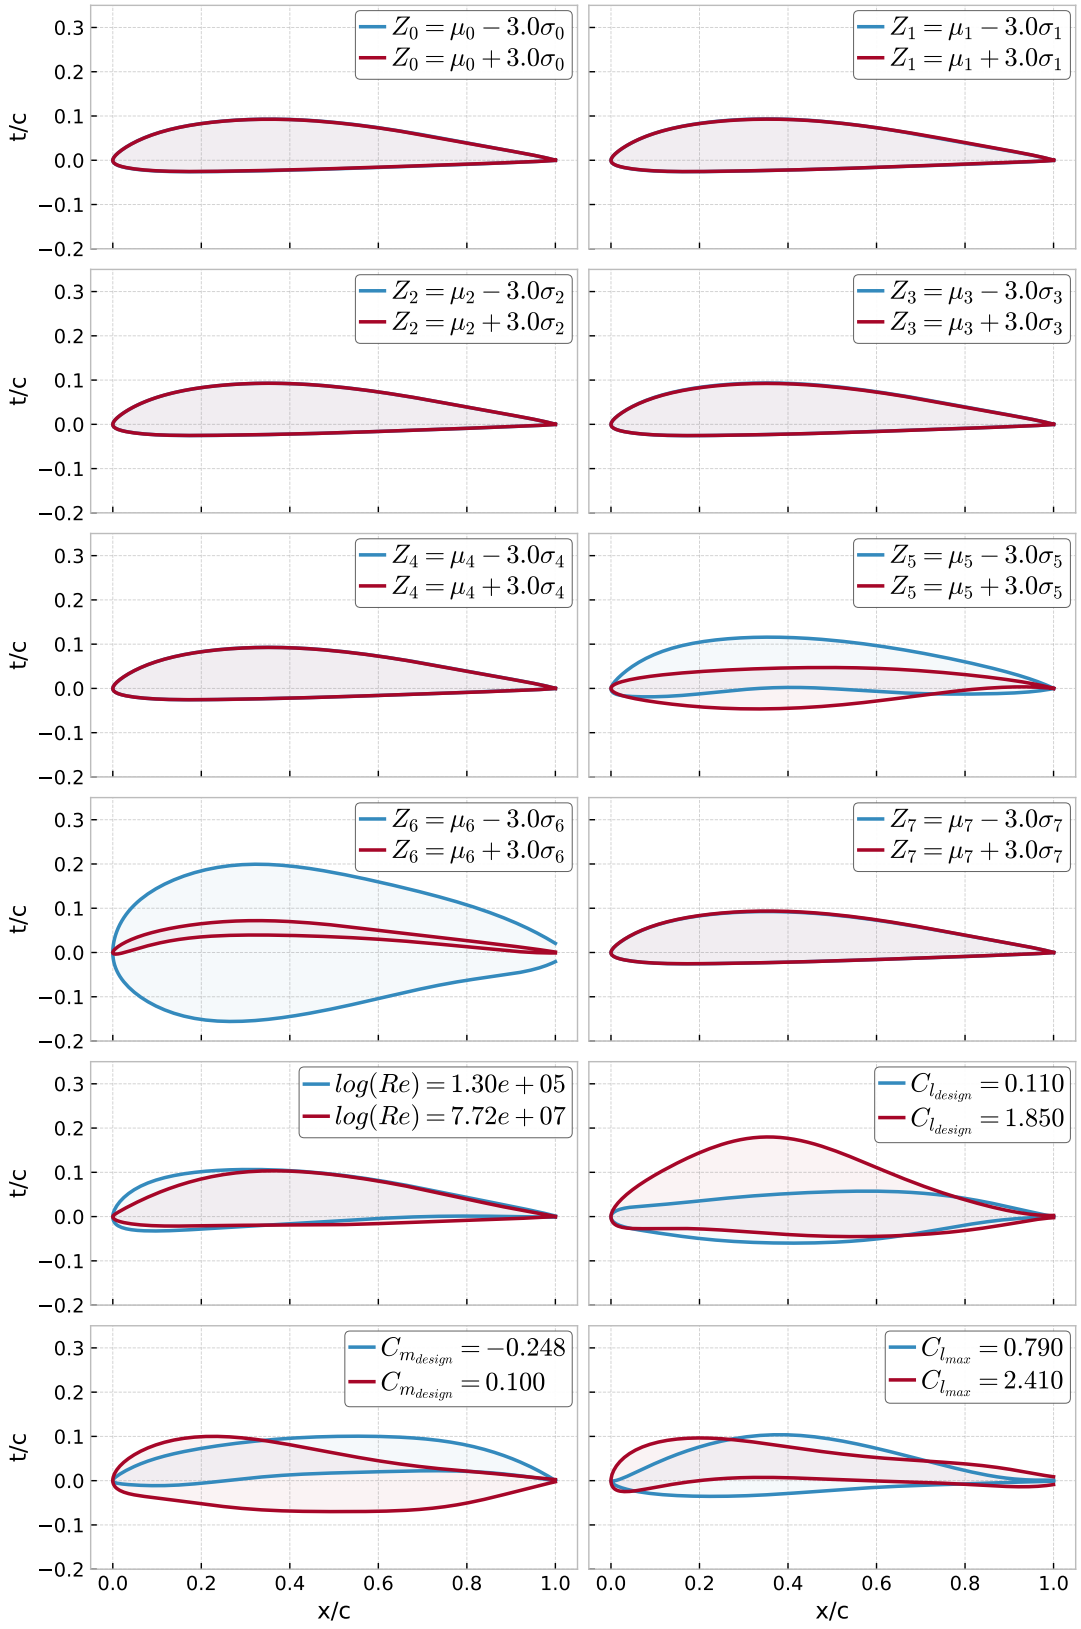


Figure 15 – Latent traversal for the normal-prior CVAE.

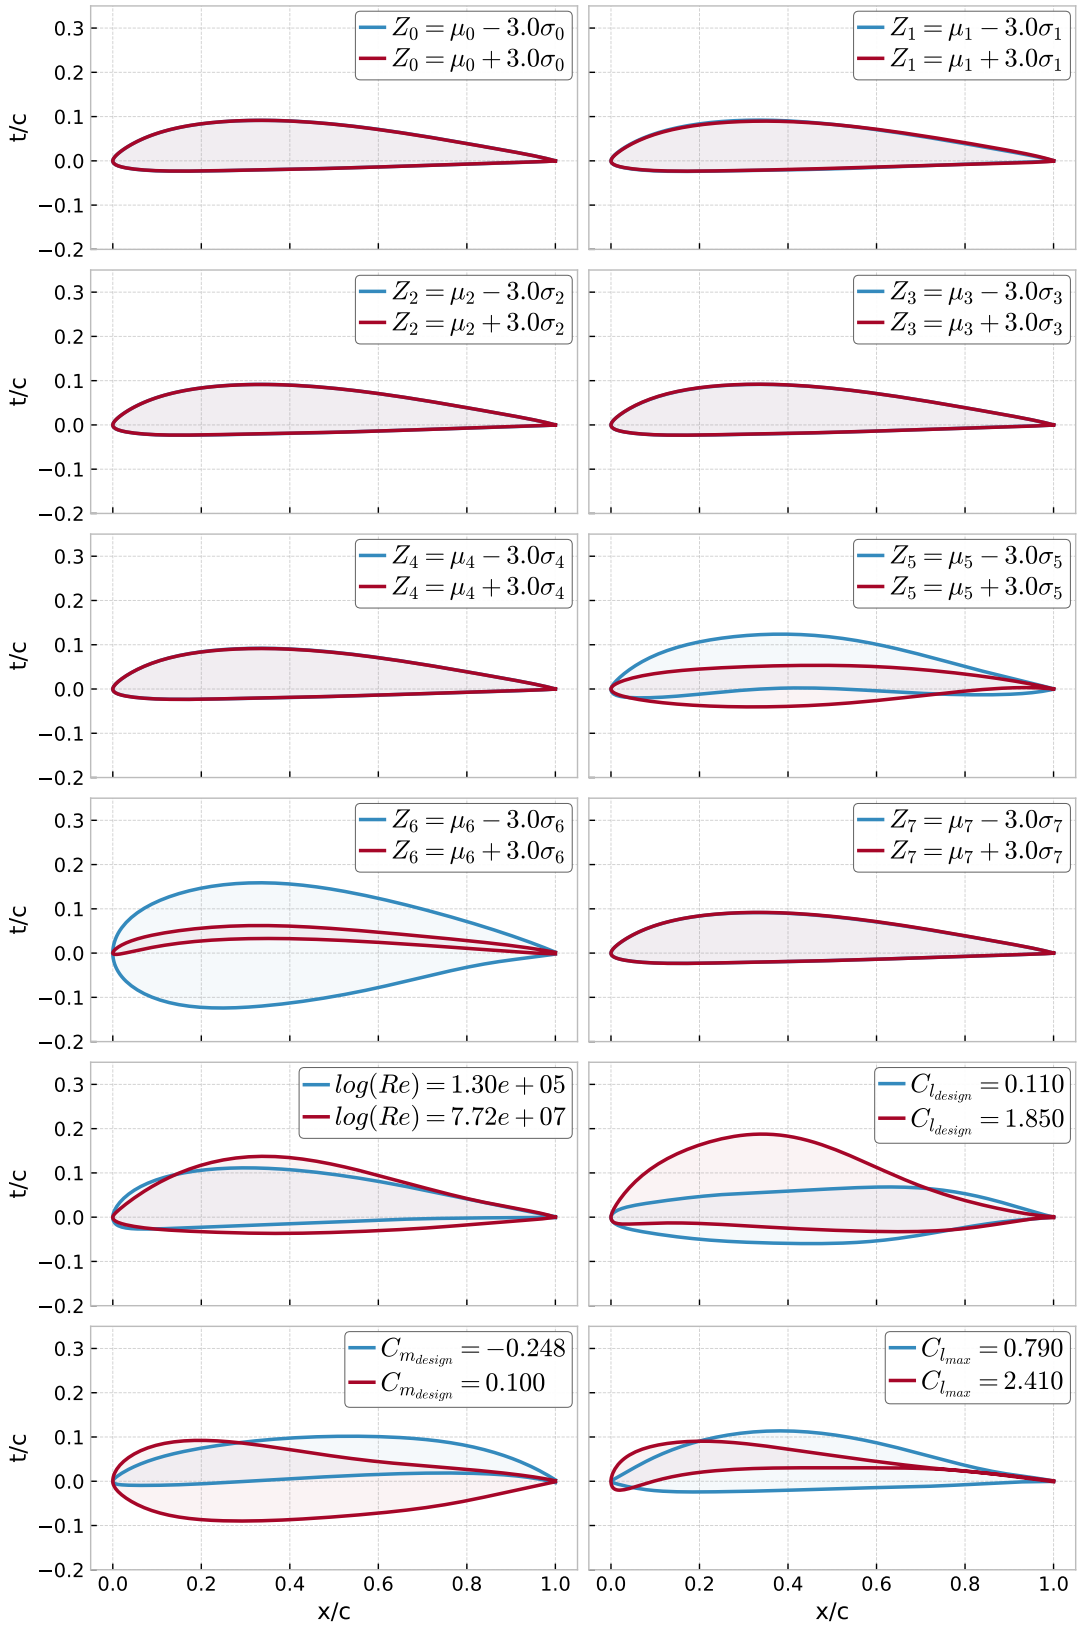


Figure 16 – Latent traversal for the conditioned-prior CVAE.

D Importance of Realistic Condition Sampling

If care is not taken to ensure that the imposed conditions are realistic, the model cannot generate meaningful samples. This leads to poor adherence to the target conditions, as shown in figure 17. The full random sampling mode results in a much wider spread of achieved conditions around the targets, and a much lower correlation between the target and achieved conditions. If conditions are infeasible or unrealistic, the model cannot generate geometries that meet those conditions. Examples of the difference in geometries generated with the full random sampling versus realistic sampling are shown in figures 19 and 20.

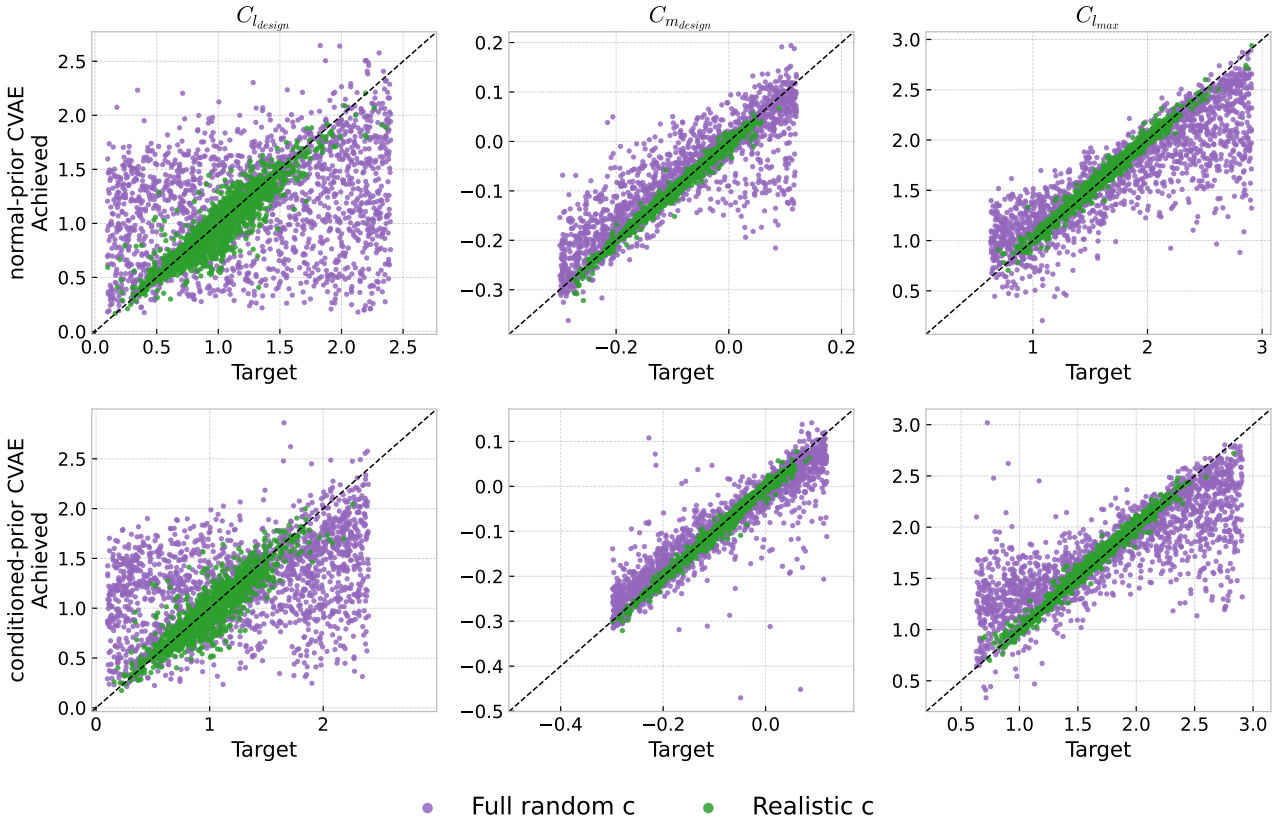


Figure 17 – Effect of full random condition sampling.

The condition-error CDFs in figure 18 show the same behaviour in a distributional form as previously presented in figure 5. For realistic condition vectors, the normal-prior CVAE has a mean total normalised condition error of 0.279 and a 90th percentile error of 0.749. The conditional-prior CVAE does slightly better, with a mean total normalised condition error of 0.370 and a 90th percentile error of 0.715. These values are small compared with the corresponding full-random condition case, where the mean total error increases to $> \pm 2.1$ for both models, and the 90th percentile error exceeds ± 4.3.

This shows that the CVAE should be used inside a physically plausible region of condition space. If the requested condition vector is sampled without respecting the correlations present in the training data, the decoder often cannot satisfy all targets simultaneously.

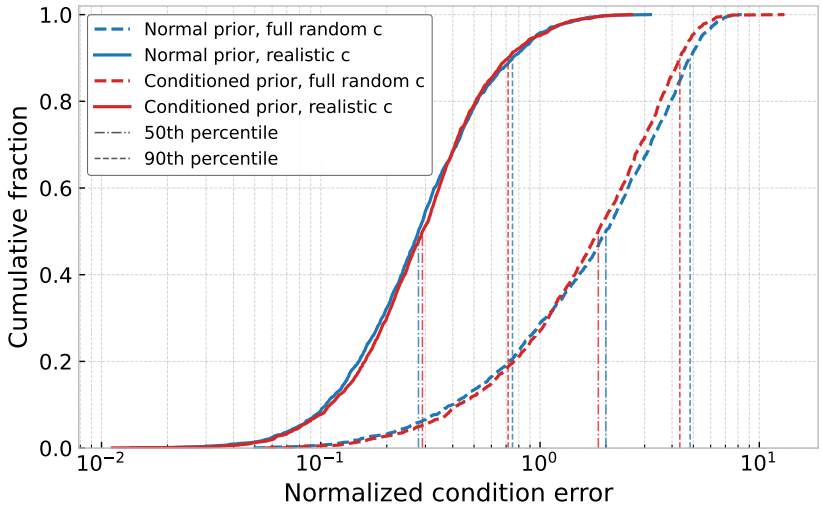


Figure 18 – Condition-adherence error distributions.

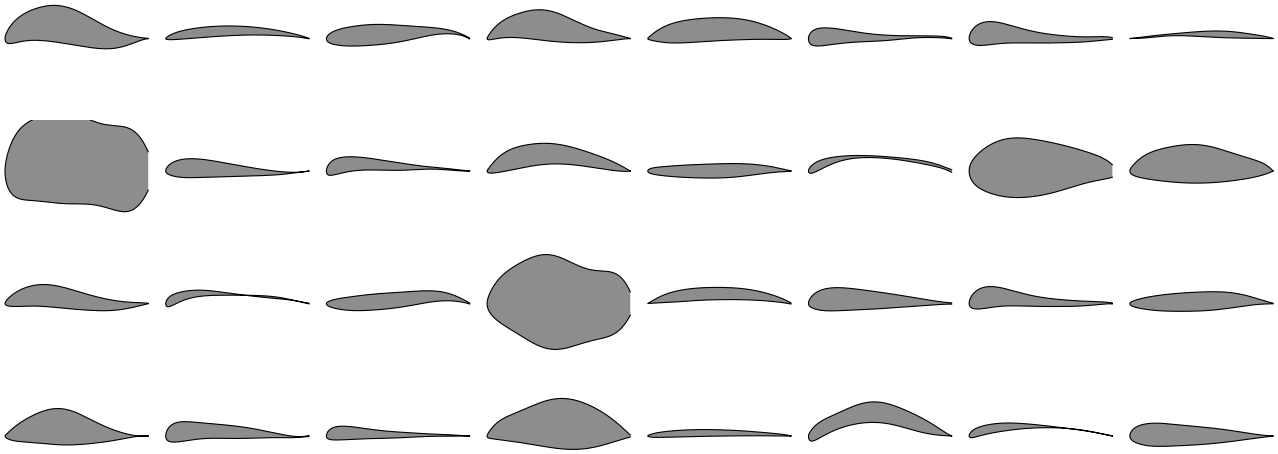


Figure 19 – Examples of random generations using the conditioned-prior CVAE.

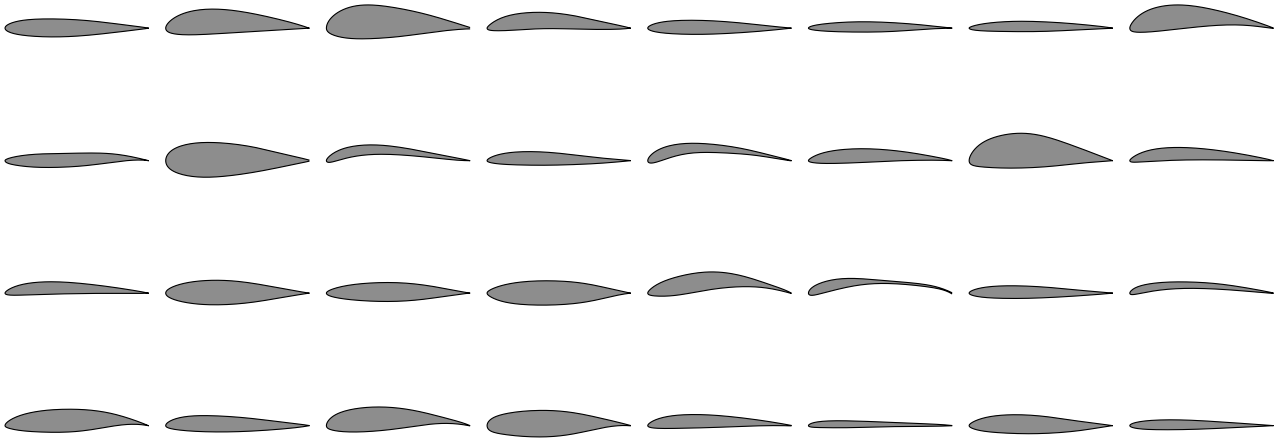


Figure 20 – Examples of realistic generations using the conditioned-prior CVAE.

E Model Hyperparameters

Table 6 lists the main hyperparameters for the two CVAE models used in the results. The hyperparameters resulted from a Weights & Biases parameter sweep with random search. Models with the lowest cross-condition loss were chosen for the final evaluation, as this was found to correlate well with improved condition adherence and adequate geometry reconstruction quality. This is also why both models have some different weight values.

Table 6 – Main hyperparameters for the two CVAE models used in the results.

Parameter	Normal-prior CVAE	Conditioned-prior CVAE
Condition vector	$\log(Re)$, $C_{l,\text{design}}$, $C_{m,\text{design}}$, $C_{l,\max}$	$\log(Re)$, $C_{l,\text{design}}$, $C_{m,\text{design}}$, $C_{l,\max}$
Geometry representation	$2\times 8+2$ Kulfan parameters	$2\times 8+2$ Kulfan parameters
Latent dimension	8	8
Encoder/decoder hidden layers	4×512	4×512
Activation	CELU	CELU
Latent prior	Standard normal	Condition-dependent normal
Prior Net	-	2×128
Batch size	64	64
Train/validation/test split	$0.80/0.15/0.05$	$0.80/0.15/0.05$
Split seed	42	42
Optimizer	Adam	Adam
Learning rate	1×10^{-4}	1×10^{-4}
Maximum epochs	10000	10000
Early-stopping patience	250	250
Warm-start VAE run	`2pxf5reu`	`2pxf5reu`
KL weight β	385.89	231.16
Condition loss weight γ	5.68	5.29
Cross-condition loss weight γ_{cross}	84.79	87.36
Coordinate loss weight ξ	1×10^{5}	1×10^{5}
Wiggliness loss weight ω	0.3	0.3
Confidence loss weight κ	1×10^{-2}	1×10^{-2}
Minimum NeuralFoil confidence	0.90	0.90
Cross-condition target shuffle	Enabled	Enabled

F Full Geometry Sweep Tables

Tables 7 to 14 show the detailed results from the geometry sweeps presented in section 4.4, including also the values of the non-swept conditions. This gives a more complete picture of the model capabilities, and how different conditions are affected when varying each aerodynamic target.

Table 7 – Mean-condition sweep geometries for the Normal-prior CVAE at $Re = 3.16e+06$. The first block is the unaltered mean condition. Subsequent blocks sweep each aerodynamic target by ± 1 and ± 2 standard deviations in ascending order, omitting the mean center row. The complementary lift condition is adjusted if needed to ensure $C_{l_{max}} \geq C_{l_{design}} + 0.1$.

$C_{l_{design}}$		$C_{m_{design}}$		$C_{l_{max}}$		Geometry
Target	True	Target	True	Target	True	
Unaltered mean condition $Re = 3.16e+06$						
0.98	0.896	−0.074	−0.0798	1.6	1.64	
Sweep $C_{l_{design}}$						
0.4	0.349	−0.074	−0.0646	1.6	1.58	
0.69	0.658	−0.074	−0.0731	1.6	1.6	
1.27	1.29	−0.074	−0.0919	1.6	1.62	
1.56	1.55	−0.074	−0.075	1.66	1.8	
Sweep $C_{m_{design}}$						
0.98	1	−0.19	−0.187	1.6	1.66	
0.98	1.02	−0.132	−0.135	1.6	1.64	
0.98	0.975	−0.016	−0.0179	1.6	1.62	
0.98	0.877	0.042	0.0388	1.6	1.62	
Sweep $C_{l_{max}}$						
0.96	0.437	−0.074	−0.0611	1.06	1.21	
0.98	0.733	−0.074	−0.0668	1.33	1.32	
0.98	0.951	−0.074	−0.0735	1.87	1.89	
0.98	1.25	−0.074	−0.0601	2.14	2.11	

Table 8 – Mean-condition sweep geometries for the Conditioned-prior CVAE at $Re = 3.16e+06$. The first block is the unaltered mean condition. Subsequent blocks sweep each aerodynamic target by ± 1 and ± 2 standard deviations in ascending order, omitting the mean center row. The complementary lift condition is adjusted if needed to ensure $C_{l_{max}} \geq C_{l_{design}} + 0.1$.

$C_{l_{design}}$		$C_{m_{design}}$		$C_{l_{max}}$		Geometry
Target	True	Target	True	Target	True	
Unaltered mean condition $Re = 3.16e+06$						
0.98	0.903	−0.074	−0.0799	1.6	1.65	
Sweep $C_{l_{design}}$						
0.4	0.427	−0.074	−0.07	1.6	1.55	
0.69	0.664	−0.074	−0.0737	1.6	1.62	
1.27	1.23	−0.074	−0.0834	1.6	1.62	
1.56	1.48	−0.074	−0.0863	1.66	1.74	
Sweep $C_{m_{design}}$						
0.98	1.06	−0.19	−0.204	1.6	1.64	
0.98	1.03	−0.132	−0.14	1.6	1.65	
0.98	0.821	−0.016	−0.0175	1.6	1.62	
0.98	0.975	0.042	0.0335	1.6	1.64	
Sweep $C_{l_{max}}$						
0.96	0.702	−0.074	−0.0711	1.06	1.18	
0.98	0.873	−0.074	−0.0703	1.33	1.34	
0.98	0.969	−0.074	−0.0797	1.87	1.9	
0.98	1.37	−0.074	−0.0767	2.14	2.21	

Table 9 – Reference airfoil condition sweep table for NACA 0018 using the normal-prior CVAE at fixed $Re = 1.00e+06$. The first block is the unaltered reference condition. Subsequent blocks sweep each aerodynamic target by ± 1 and ± 2 standard deviations, in ascending order. The complementary lift condition is adjusted if needed to ensure $C_{l_{max}} \geq C_{l_{design}} + 0.1$.

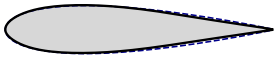
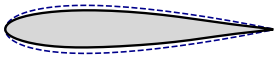
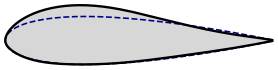
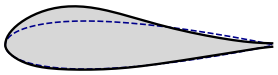
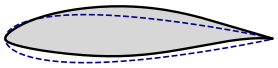
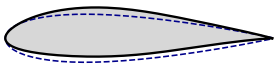
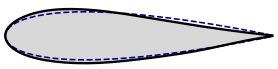
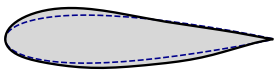
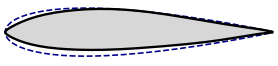
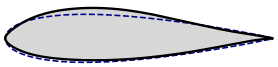
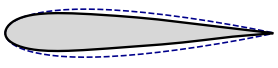
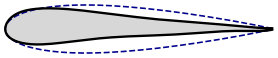
$C_{l_{design}}$		$C_{m_{design}}$		$C_{l_{max}}$		Geometry
Target	True	Target	True	Target	True	
Unaltered reference condition $Re = 1 \times 10^6$						
1.11	1.1	0.00154	−0.00484	1.42	1.47	
Sweep $C_{l_{design}}$						
0.526	0.981	0.00154	0.00524	1.42	1.38	
0.816	0.958	0.00154	−0.000541	1.42	1.45	
1.4	1.34	0.00154	−0.0101	1.5	1.5	
1.69	1.64	0.00154	−0.00443	1.79	1.77	
Sweep $C_{m_{design}}$						
1.11	1.14	−0.114	−0.12	1.42	1.43	
1.11	1.16	−0.0565	−0.0577	1.42	1.46	
1.11	1.15	0.0595	0.0304	1.42	1.45	
1.11	1.17	0.118	0.0479	1.42	1.39	
Sweep $C_{l_{max}}$						
0.784	0.555	0.00154	−0.0238	0.884	1.06	
1.05	0.977	0.00154	−0.0236	1.15	1.2	
1.11	1.11	0.00154	0.000214	1.69	1.67	
1.11	1.2	0.00154	0.011	1.96	1.77	

Table 10 – Reference airfoil condition sweep table for DAE11 using the normal-prior CVAE at fixed $Re = 5.00e+05$. The first block is the unaltered reference condition. Subsequent blocks sweep each aerodynamic target by ± 1 and ± 2 standard deviations, in ascending order. The complementary lift condition is adjusted if needed to ensure $C_{l_{max}} \geq C_{l_{design}} + 0.1$.

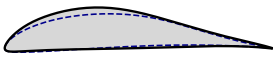
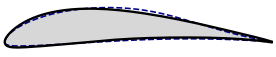
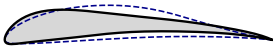
$C_{l_{design}}$		$C_{m_{design}}$		$C_{l_{max}}$		Geometry
Target	True	Target	True	Target	True	
Unaltered reference condition $Re = 5 \times 10^5$						
1.52	1.61	−0.133	−0.136	1.62	1.67	
Sweep $C_{l_{design}}$						
0.941	0.987	−0.133	−0.133	1.62	1.56	
1.23	1.25	−0.133	−0.137	1.62	1.63	
1.81	1.89	−0.133	−0.145	1.91	2	
2.1	1.78	−0.133	−0.165	2.2	2.24	
Sweep $C_{m_{design}}$						
1.52	1.66	−0.249	−0.269	1.62	1.67	
1.52	1.6	−0.191	−0.199	1.62	1.65	
1.52	1.6	−0.0749	−0.0751	1.62	1.7	
1.52	1.58	−0.0169	−0.0112	1.62	1.71	
Sweep $C_{l_{max}}$						
0.981	0.96	−0.133	−0.137	1.08	1.2	
1.25	1.25	−0.133	−0.135	1.35	1.38	
1.52	1.4	−0.133	−0.148	1.89	1.93	
1.52	1.43	−0.133	−0.145	2.16	2.13	

Table 11 – Reference airfoil condition sweep table for Whitcomb using the normal-prior CVAE at fixed $Re = 2.00e+07$. The first block is the unaltered reference condition. Subsequent blocks sweep each aerodynamic target by ± 1 and ± 2 standard deviations, in ascending order. The complementary lift condition is adjusted if needed to ensure $C_{l_{max}} \geq C_{l_{design}} + 0.1$.

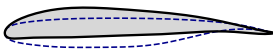
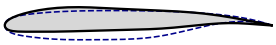
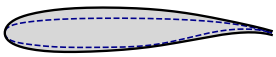
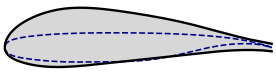
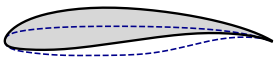
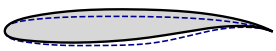
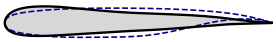
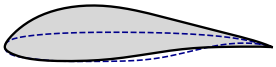
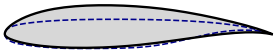
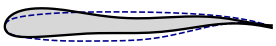
$C_{l_{design}}$		$C_{m_{design}}$		$C_{l_{max}}$		Geometry
Target	True	Target	True	Target	True	
Unaltered reference condition $Re = 2 \times 10^7$						
1.43	1.35	−0.132	−0.113	2.39	2.34	
Sweep $C_{l_{design}}$						
0.852	1.29	−0.132	−0.091	2.39	2.2	
1.14	1.28	−0.132	−0.0896	2.39	2.23	
1.72	1.71	−0.132	−0.146	2.39	2.42	
2.01	1.66	−0.132	−0.136	2.39	2.38	
Sweep $C_{m_{design}}$						
1.43	1.56	−0.248	−0.257	2.39	2.39	
1.43	1.45	−0.19	−0.195	2.39	2.39	
1.43	1.39	−0.0744	−0.0605	2.39	2.36	
1.43	1.45	−0.0164	−0.00105	2.39	2.26	
Sweep $C_{l_{max}}$						
1.43	1.45	−0.132	−0.149	1.85	1.81	
1.43	1.14	−0.132	−0.153	2.12	2.12	
1.43	1.45	−0.132	−0.0874	2.66	2.5	
1.43	1.7	−0.132	−0.0735	2.91	2.55	

Table 12 – Reference airfoil condition sweep table for NACA 0018 using the conditioned-prior CVAE at fixed $Re = 1.00e+06$. The first block is the unaltered reference condition. Subsequent blocks sweep each aerodynamic target by ± 1 and ± 2 standard deviations, in ascending order. The complementary lift condition is adjusted if needed to ensure $C_{l_{max}} \geq C_{l_{design}} + 0.1$.

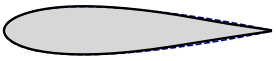
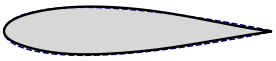
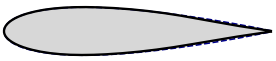
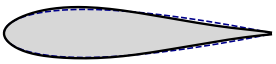
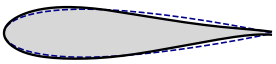
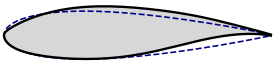
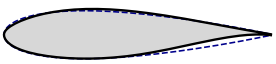
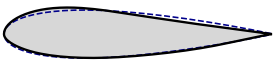
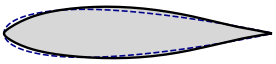
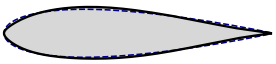
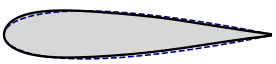
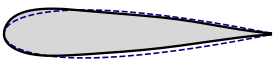
$C_{l_{design}}$		$C_{m_{design}}$		$C_{l_{max}}$		Geometry
Target	True	Target	True	Target	True	
Unaltered reference condition $Re = 1 \times 10^6$						
1.11	0.952	0.00154	0.000933	1.42	1.44	
Sweep $C_{l_{design}}$						
0.526	0.724	0.00154	0.00288	1.42	1.41	
0.816	0.944	0.00154	0.00162	1.42	1.46	
1.4	1.09	0.00154	−0.0115	1.5	1.46	
1.69	1.3	0.00154	−0.0286	1.79	1.72	
Sweep $C_{m_{design}}$						
1.11	1.01	−0.114	−0.133	1.42	1.47	
1.11	0.995	−0.0565	−0.072	1.42	1.51	
1.11	1.07	0.0595	0.0459	1.42	1.44	
1.11	1.24	0.118	0.0657	1.42	1.48	
Sweep $C_{l_{max}}$						
0.784	0.533	0.00154	−0.0113	0.884	1.02	
1.05	0.847	0.00154	−0.00535	1.15	1.16	
1.11	1.16	0.00154	0.000841	1.69	1.63	
1.11	1.22	0.00154	−0.00473	1.96	1.75	

Table 13 – Reference airfoil condition sweep table for DAE11 using the conditioned-prior CVAE at fixed $Re = 5.00e+05$. The first block is the unaltered reference condition. Subsequent blocks sweep each aerodynamic target by ± 1 and ± 2 standard deviations, in ascending order. The complementary lift condition is adjusted if needed to ensure $C_{l_{max}} \geq C_{l_{design}} + 0.1$.

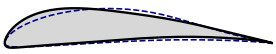
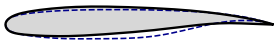
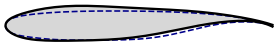
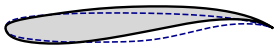
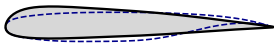
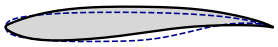
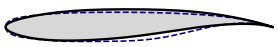
$C_{l_{design}}$		$C_{m_{design}}$		$C_{l_{max}}$		Geometry
Target	True	Target	True	Target	True	
Unaltered reference condition $Re = 5 \times 10^5$						
1.52	1.61	−0.133	−0.131	1.62	1.68	
Sweep $C_{l_{design}}$						
0.941	1.06	−0.133	−0.127	1.62	1.69	
1.23	1.39	−0.133	−0.128	1.62	1.69	
1.81	1.6	−0.133	−0.146	1.91	1.99	
2.1	1.68	−0.133	−0.153	2.2	2.21	
Sweep $C_{m_{design}}$						
1.52	1.55	−0.249	−0.259	1.62	1.58	
1.52	1.6	−0.191	−0.198	1.62	1.64	
1.52	1.47	−0.0749	−0.0693	1.62	1.69	
1.52	1.53	−0.0169	−0.0333	1.62	1.72	
Sweep $C_{l_{max}}$						
0.981	0.848	−0.133	−0.141	1.08	1.3	
1.25	1.24	−0.133	−0.14	1.35	1.43	
1.52	1.39	−0.133	−0.141	1.89	1.89	
1.52	1.36	−0.133	−0.139	2.16	2.04	

Table 14 – Reference airfoil condition sweep table for Whitcomb using the conditioned-prior CVAE at fixed $Re = 2.00e+07$. The first block is the unaltered reference condition. Subsequent blocks sweep each aerodynamic target by ± 1 and ± 2 standard deviations, in ascending order. The complementary lift condition is adjusted if needed to ensure $C_{l_{max}} \geq C_{l_{design}} + 0.1$.

$C_{l_{design}}$		$C_{m_{design}}$		$C_{l_{max}}$		Geometry
Target	True	Target	True	Target	True	
Unaltered reference condition $Re = 2 \times 10^7$						
1.43	1.42	−0.132	−0.117	2.39	2.34	
Sweep $C_{l_{design}}$						
0.852	1.36	−0.132	−0.0914	2.39	2.23	
1.14	1.36	−0.132	−0.0977	2.39	2.27	
1.72	1.49	−0.132	−0.145	2.39	2.44	
2.01	1.57	−0.132	−0.172	2.39	2.48	
Sweep $C_{m_{design}}$						
1.43	0.45	−0.248	−0.275	2.39	2.52	
1.43	1.49	−0.19	−0.194	2.39	2.42	
1.43	1.45	−0.0744	−0.0638	2.39	2.27	
1.43	1.44	−0.0164	−0.0211	2.39	2.32	
Sweep $C_{l_{max}}$						
1.43	1.33	−0.132	−0.166	1.85	2.02	
1.43	1.33	−0.132	−0.145	2.12	2.24	
1.43	1.59	−0.132	−0.0983	2.66	2.45	
1.43	1.72	−0.132	−0.105	2.91	2.57	