

GENERATIVE AERODYNAMIC DESIGN WITH DIFFUSION PROBABILISTIC MODELS

T. Wagenaar*, S. Mancini*

* Airbus Defence and Space GmbH, Flight Physics, Rechliner Straße, 85077 Manching, Germany

Abstract

The optimization of geometries for aerodynamic design often relies on a large number of expensive simulations to evaluate and iteratively improve the geometries. It is possible to reduce the number of simulations by providing a starting geometry that has properties close to the desired requirements, often in terms of lift and drag, aerodynamic moments and surface areas. We show that generative models have the potential to provide such starting geometries by generalizing geometries over a large dataset of simulations. In particular, we leverage diffusion probabilistic models trained on XFOIL simulations to synthesize two-dimensional airfoil geometries conditioned on given aerodynamic features and constraints. The airfoils are parameterized with Bernstein polynomials, ensuring smoothness of the generated designs. We show that the models are able to generate diverse candidate designs for identical requirements and constraints, effectively exploring the design space to provide multiple starting points to optimization procedures. However, the quality of the candidate designs depends on the distribution of simulated designs in the dataset. Importantly, the geometries in the dataset must satisfy other requirements and constraints that are not used in conditioning of the diffusion model, to assure that the generated designs are physical.

Keywords

Aerodynamic shape design; Generative modeling; Deep learning

NOMENCLATURE

Symbols

α	Angle of attack
B	Bernstein polynomial
b	Bernstein basis polynomial
C	Coefficient
\mathbf{f}	Features
\mathbf{I}	Identity matrix
T	Maximum time step
N_1	First shape coefficient
N_2	Second shape coefficient
ϵ	Noise
\mathcal{N}	Normal distribution
\mathbf{z}	Normal sample
θ	Neural network parameters
q	Dataset sampling function
Re	Reynolds number
\mathbf{x}	Sample

t	Time step
β	Variance
A	Bernstein polynomial coefficient
y	Height

Subscripts

D	drag
l	lower
max	Maximum
M	moment
u	upper
L	lift

Abbreviations

CFD	Computational Fluid Dynamics
DPM	Diffusion Probabilistic Model
GAN	Generative Adversarial Network
MAPE	Mean Absolute Percentage Error
RMdSE	Root Median Square Error
RMSE	Root Mean Square Error

1. INTRODUCTION

Aerodynamic shape optimization often involves a large design space, which allows for more detailed and thus more efficient final designs. At the same time, larger design spaces are more tedious to explore because evaluations of designs rely on computationally expensive Computational Fluid Dynamics (CFD) simulations. As exploring the entire design space is practically infeasible, both gradient-based and gradient-free optimization methods rely on repetitive simulations of intermediate designs to iteratively improve aerodynamic shapes [1]. The number of simulations is only increased when additional disciplines besides aerodynamics are considered in a multi-disciplinary optimization procedure, as this might lead to evolving aerodynamic requirements and constraints due to its multi-objective nature.

Fortunately, recent advances in machine learning methods promise to accelerate aerodynamic shape optimization methods [2]. One class of methods aims to create a reduced representation of the design space, also termed a latent space. The rationale is that most combinations of parameters lead to non-physical designs, so that only a portion of the design space is useful. For example, it is possible to deduce a compressed design space from a set of existing feasible designs [3, 4]. Another class of methods rather focuses on reducing the computational cost of shape evaluations, for example by building surrogate models of high-fidelity simulations [5]. Both these methods have successfully been combined, leading to rapid aerodynamic shape optimization [6].

One common caveat of these methods is that they do not modify the iterative nature of the optimization procedure. This not only means that optimization procedures must be repeated whenever requirements change, but iterative optimization procedures are often sensitive to starting conditions and the definition of the objective function so that convergence is not necessarily guaranteed. An approach that promises to alleviate these issues is deep generative design, which uses machine learning methods that distill information from previous designs to directly infer new designs [7]. Depending on the accuracy of the design, it can be used as a final design or as the starting point for the traditional optimization procedure.

Deep generative design has already been applied to aerodynamic shape optimization, in particular to two-dimensional airfoils. For example, Achour et al. [8] and Tan et al. [9] use conditional Generative Adversarial Networks (GANs) to generate airfoils based on desired performance classes. However, the generated airfoils are not always smooth. Due to the sensitivity of aerodynamics to sharp features, this is problematic and Chatteraj et al. [10] reduce this issue by adding an additional smoothness objective. Since only a handful of performance classes are considered to start the generation, other authors

extend the methods to work with full polars [11] or with ranges [12].

While using GANs for deep generative aerodynamic shape design is promising, their training procedure is unstable and therefore tedious [13]. In the field of image generation, from which GANs originate, they have already been superseded by Diffusion Probabilistic Models (DPMs) [14]. Recently, the advantages of DPMs compared to GANs have also already been established for topology optimization [15]. Therefore, it seems beneficial to investigate how well DPMs perform on deep generative aerodynamic design.

While the usage of DPMs for geometry generation has already been investigated in several fields, for example to make a general framework [16], to generate molecules [17] and to generate various objects [18], their application to aerodynamics remains limited. To the authors' best knowledge, their only application to aerodynamic shape optimization so far still relies on an iterative optimization procedure to generate airfoils [19].

Therefore, the objective of this paper is to show that DPMs can be used for deep generative aerodynamic shape design. In addition, we will further formalize the deep generative aerodynamic shape design approach. Firstly, we will use a smooth parameterization so that the shapes are inherently smooth. Secondly, we will show that diverse candidate shapes can be generated for the same operating conditions. Thirdly, we will show that the generated shapes are generally different than those in the training dataset. Lastly, we provide a systematic way of conditioning the models to take into account both design constraints and requirements.

2. METHODOLOGY

As an example, the generation of two-dimensional airfoils is considered. In particular, the task is to generate airfoil geometries that have a certain lift coefficient C_L , drag coefficient C_D and moment coefficient C_M at a specified angle of attack. While this example is rather simple, the methodology can be extended to more complex cases such as those considering the maximum lift $C_{L,\max}$, the lift-to-drag ratio L/D or even an entire lift polar similar to Nobari et al. [12].

2.1. Diffusion probabilistic models

Diffusion models are originally inspired by nonequilibrium thermodynamics [20], but were later more successfully applied to image generation [21]. They operate through two processes, of which the first is the *forward process*. In the forward process, samples \mathbf{x} of the dataset are taken and Gaussian noise is added in steps. If at each time step t the variance of the noise added is β_t , the noised samples \mathbf{x}_t can directly be determined in the closed form given in Eq. (1):

$$(1) \quad \mathbf{x}_t = \sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon.$$

Here $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ is Gaussian noise, $\alpha_t = 1 - \beta_t$ and $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$. At the end of the forward process, at some time step $t = T$, the samples \mathbf{x}_T are diffused so much that they are essentially Gaussian noise. When the original samples are parameters describing feasible airfoils, the final samples represent highly deformed airfoils that do not resemble the original airfoil anymore.

Algorithm 1 Sampling procedure [21]

- 1: $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$;
 - 2: **for** $t = T, \dots, 1$ **do**
 - 3: $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ if $t > 1$, else $\mathbf{z} = \mathbf{0}$
 - 4: $\mathbf{x}_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left(\mathbf{x}_t - \frac{1-\alpha_t}{\sqrt{1-\bar{\alpha}_t}} \epsilon_\theta(\mathbf{x}_t, t) \right) + \sqrt{\beta_t} \mathbf{z}$
 - 5: **end for**
 - 6: **return** \mathbf{x}_0
-

The fundamental aspect of DPMs is however in the *reverse process*, where a neural network is used to estimate the noise ϵ based on a noisy sample \mathbf{x}_t . If this estimation is accurate, it allows to generate samples similar to those in the training dataset according to Algorithm 1. In other words, this allows to generate airfoils that are similarly distributed but not necessarily identical to the airfoils in the dataset. Since this process is stochastic, multiple generated airfoils will not necessarily be the same even if the same starting noisy sample \mathbf{x}_T is used. This is an advantage compared to existing non-probabilistic generative methods, which are only able to generate a single airfoil for any given features.

Algorithm 2 Training procedure [21]

- 1: **repeat**
 - 2: $\mathbf{x}_0 \sim q(\mathbf{x}_0)$
 - 3: $t \sim \text{Uniform}(1, \dots, T)$
 - 4: $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
 - 5: Take gradient descent step on $\|\epsilon - \epsilon_\theta(\sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1-\bar{\alpha}_t} \epsilon, t)\|^2$
 - 6: **until** converged
-

To obtain an accurate noise estimation ϵ_θ , a neural network with parameters θ is trained using the procedure outlined in Algorithm 2. It involves taking samples from the dataset q , calculating the noisy sample at some time step t , and then comparing the estimated noise for this sample to the actual noise. Gradient descent is then applied to the neural network parameters based on the difference between the two, slowly improving the estimation.

While generating new airfoils can be useful on its own, the current procedure does not yet allow for the generation of airfoils with certain features. One way to achieve this is to train multiple models, each trained on a certain subset of the dataset. However, a more general method to achieve *conditional diffusion* is to add (a representation of) the features \mathbf{f} of the airfoil as additional inputs to the neural network, i.e. $\epsilon_\theta(\mathbf{x}_t, t, \mathbf{f})$. In the considered example, $\mathbf{f} = (C_L, C_D, C_M)$ so that

the sampling procedure in Algorithm 1 will generate airfoils with the given lift, drag and moment coefficients.

2.2. Airfoil parameterization

An important aspect of the generation process is how the airfoils are represented by the parameterization \mathbf{x} . Although the process results in airfoils similar to those in the dataset, the parameterization can still affect how well the model can be trained and how physical the resulting airfoils are. For instance, previous works on generative models for geometry generation have used point clouds [22], voxel grids [23], meshes [24] and signed distance functions [25] to represent geometries. A visualization of these parameterizations on two-dimensional airfoils is shown in Fig. 1. The first three are not continuous, so that the resulting geometries only have a finite resolution. In addition, all these methods are not inherently smooth, so that small differences in the shape can lead to drastically different aerodynamic characteristics.

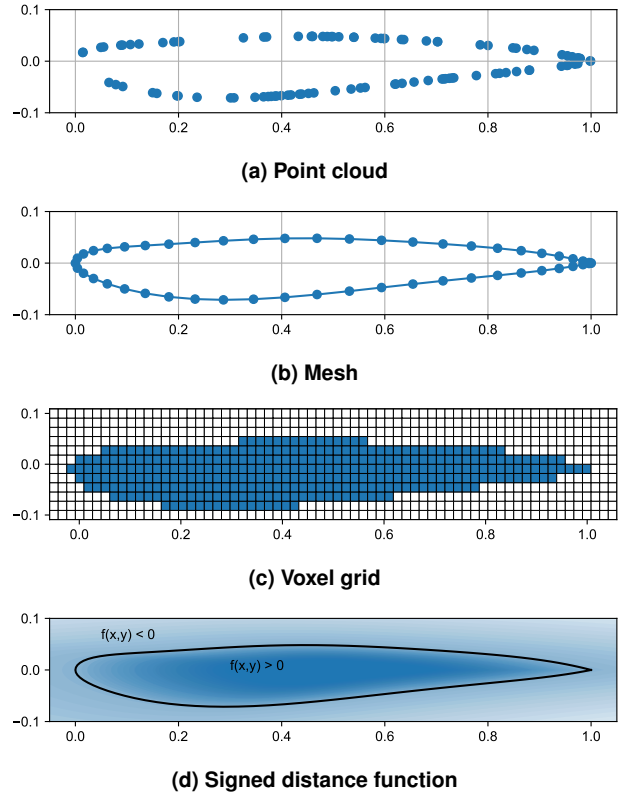


FIG 1. Different geometry parameterizations applied to two-dimensional airfoils.

Therefore, an inherently smooth and continuous parameterization is used based on Bernstein polynomials [26]. Apart from these two properties, this parameterization also allows to enforce "airfoil-like" geometries that for example have rounded leading edges and sharp trailing edges. Both the upper and lower surface of the airfoil are modelled according to Eq. (2), where N_1 and N_2 are general airfoil shape parameters and B_n is a Bernstein polynomial of order n .

$$(2) \quad \frac{y}{c} = \left(\frac{x}{c}\right)^{N_1} \left(1 - \frac{x}{c}\right)^{N_2} \cdot B_n\left(\frac{x}{c}\right)$$

The Bernstein polynomials are defined using equation Eq. (3), where A_i are the polynomial coefficients and thus the geometry parameters of the airfoil.

$$(3) \quad B_n\left(\frac{x}{c}\right) = \sum_{i=0}^n A_i \binom{n}{i} \left(\frac{x}{c}\right)^i \left(1 - \frac{x}{c}\right)^{n-i}$$

In other words, the samples considered in Section 2.1 take the form of Eq. (4) where A_l and A_u represent the parameters for the lower side and upper side of the airfoil respectively. Examples of generated airfoils using this parameterization are shown in Fig. 2. While the geometries are certainly airfoil-like, some geometries are self-intersecting, have sharp leading edges or round trailing edges.

$$(4) \quad \mathbf{x} = \{A_{l,0}, A_{l,1}, \dots, A_{u,0}, A_{u,1}, \dots\}$$




FIG 2. Randomly generated airfoils using Bernstein polynomials with $n = 6$ coefficients and random values for the shape parameters N_1 and N_2 .

Nevertheless, using this representation provides smooth and continuous airfoils, but it also significantly reduces the number of parameters that represent the geometry compared to the earlier mentioned methods. As mentioned in the introduction, design spaces are often unnecessarily large and can benefit from reduced representations, termed latent spaces. This is also the fundamental thought behind latent DPMs [27], which first convert the parameter space \mathbf{x} to a reduced latent space $\hat{\mathbf{x}}$. As a result, both the forward and reverse process occur in much lower-dimensional spaces which can significantly improve the computational cost of the method.

While the Bernstein parameterization in Eq. (2) can also be seen as a form of latent space, typically the mapping of the original space to the latent space is learned by means of an autoencoder [28]. Similar techniques have already been applied to generative

methods for airfoil geometries [3], but it is important to note that using a latent space does not necessarily solve the issues of continuity and smoothness. In the present work, the additional use of a latent space is not considered since the Bernstein design space is already significantly reduced compared to standard discretization. In fact, compressing the design space even further might actually pose problems because DPMs rely on some concept of "noise". When the design space is too compressed, even randomly sampled sets of parameters lead to a physical airfoil. As a result, it may become hard to distinguish the original airfoils at $t = 0$ and the noisy airfoils at $t = T$.

2.3. Dataset description

Before going into detail on how the dataset itself is generated, it is important to discuss the choices in filtering or conditioning features. In the considered case, conditioning only occurs for the lift, drag and moment coefficients. Any other features that are not included as conditioning variables and may thus vary freely, such as the thickness or camber of the airfoil. If airfoils with unreasonable values for these features are included in the dataset, unfeasible airfoils may be generated. Therefore, it is important to filter these airfoils out of the dataset before training. Another option is to leave the dataset untouched and to simply condition the model on all relevant features, although this must be combined with optional or range-based conditioning to allow features to vary when they do not need to be constrained. However, this option is not investigated in the present work.




FIG 3. Subset of airfoils in the dataset along with their features (C_L, C_D, C_M).

In general, filtering should not be done based on features that are also used to condition the model. For example, removing airfoils with negative lift coefficients from the dataset when conditioning on the lift coefficient may seem logical. However, due to the strong generalization capabilities of neural networks even "bad" geometries may provide information to the model on how to generate "good" geometries. For this same reason, the choice is made to not use the airfoils in the UIUC airfoil dataset [29] as it only

includes optimized airfoils, so that it may not include geometries matching more rarely desired features. Instead, the dataset of airfoils is generated by sampling the parameters in Eq. (4) uniformly. In particular, the airfoil shape is set to $N_1 = 0.5$ and $N_2 = 1.0$ for NACA-like airfoils with rounded leading edges and sharp trailing edges. Furthermore, Bernstein polynomials of order 6 are chosen for both the upper and lower side, resulting in a total of 11 parameters since the first lower and upper coefficient are set equal for continuous curvature at the nose. All parameters are sampled in the range $[-0.5, 1.5]$ while $A_{l,0}$ and $A_{u,0}$ are sampled in the range $[0.3, 1.0]$ to force a minimum leading edge radius. After generation, airfoils with self-intersecting surfaces are removed. Examples of resulting airfoils are shown in Fig. 3.




FIG 4. Distribution of the features of the airfoils in the dataset.

Afterward, XFOIL [30] is used to simulate the aerodynamics of the airfoils. Although this is a relatively low-fidelity method, the approach can be extended to higher-fidelity methods without any changes to the methodology. The airfoils are discretized to 100 cosine-spaced elements and simulated at $Re = 1 \times 10^6$ and $\alpha = 5^\circ$ using otherwise default settings. To remove outliers, the Mahalanobis distance is calculated for each point in the feature space and any points above the 99.5th distance percentile are removed. The resulting distribution of features is shown in Fig. 4. In total, 1,000 airfoils are entered into the dataset along with their corresponding features.

2.4. Model and training setup

The dataset is split into 600 samples for training, 200 samples for validation and 200 samples for testing. Furthermore, both the samples and the features are z-score normalized based on their values in the training set. For the forward and reverse process, the vari-

ance β_t is linearly scheduled from 1×10^{-3} to 0.2 in a total of $T = 1000$ diffusion steps. A visualization of the forward process is given in Fig. 5 in both the latent and physical space, confirming the assumption in Section 2.2 that the Bernstein parameterization alone compresses the design space to an extent that the airfoil at $t = T$ is not very noisy (or out-of-distribution).

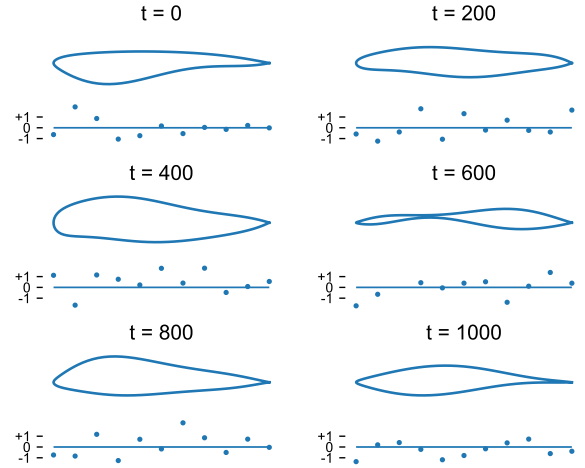


FIG 5. Forward process in the physical and latent space. The physical space is shown in terms of the geometry, while the latent space is shown with a scatter plot of the Bernstein coefficients.

The neural network used for the noise estimator $\epsilon_\theta = (\tilde{x}_t, \tilde{t}, \tilde{f})$ consists of 4 layers with 32 neurons each. The \tanh activation function is used for each layer except the output layer. Here \tilde{x} are the normalized parameters, $\tilde{t} = t/T$ is the normalized time and \tilde{f} are the normalized features. Note that the noise estimator has the same output size as the number of airfoil parameters, which is 11. The neural network architecture is visualized in Fig. 6.




FIG 6. Architecture of the neural network used for the noise estimator ϵ_θ , with N the batch size.

For training, the Adam optimizer [31] is used with a learning rate of 1×10^{-4} . The batch size is set to 64 and training is performed until the loss on the validation set does not improve anymore. In a nutshell, the model is trained to effectively predict the noise in the parameters given the time step and desired aerodynamic features, implicitly learning aerodynamics in the process. By iteratively subtracting this estimated noise, the corresponding geometry slowly evolves to one with the matching features.

3. RESULTS

After training, the model can be used in the reverse process as described in Algorithm 1. An example of this process is shown in Fig. 7 and some resulting airfoils are shown in Fig. 8. The generated airfoils are smooth and satisfy the requirements and constraints imposed in Section 2.3, covering two of the four main objectives in the introduction. Nevertheless, some of the airfoils are nontypical because certain structural constraints have not been considered.

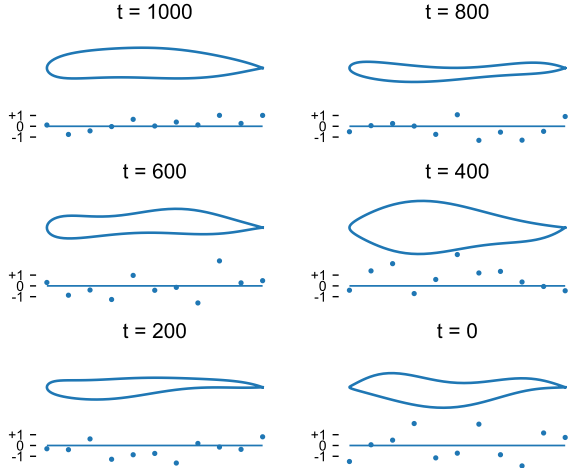


FIG 7. Reverse process in the physical and latent space.

The two objectives that remain are showing that (a) diverse candidate designs can be generated and (b) that the generated airfoils are unique and thus not identical to airfoils samples in the training dataset. But before doing so, it is important to quantify how well the aerodynamic features of the generated airfoils match the desired features.

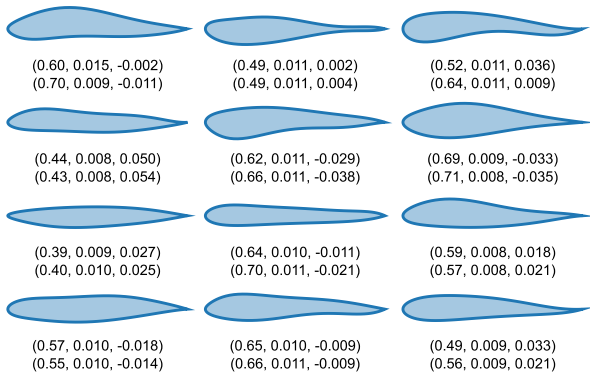


FIG 8. Examples of generated airfoils with their desired features and actual features (C_L , C_D , C_M).

3.1. Airfoil accuracy

To quantify this accuracy of the model, the aerodynamic features of the 200 airfoils in the test set are taken and used to condition the generation of new

airfoils. The resulting average Root Mean Squared Error (RMSE), Root Median Squared Error (RMdSE) and the Mean Absolute Percentage Error (MAPE) of the features are shown in Table 1. While the aerodynamic features of the generated airfoils are fairly accurate, there is certainly room for improvement.

TAB 1. Errors for the different features quantified using varying metrics. MAPE for C_M is not shown because its range is zero-centered.

Feature	C_L	C_D	C_M
RMSE	9.3e-2	7.8e-3	1.7e-2
RMdSE	3.0e-2	8.4e-4	6.4e-3
MAPE	13%	11%	-

Note that despite the smoothness of the airfoils, the present approach still suffers from inherent aerodynamic nonlinearities and modeling errors of the chosen solver. To illustrate this, consider the two airfoils in Fig. 9. While the airfoils are nearly identical, their aerodynamic features are significantly different. This may explain the large difference between the mean and median errors in Table 1; the model may sometimes produce a seemingly correct geometry, but due to the aerodynamic nonlinearities and modeling errors the corresponding features might deviate significantly leading to a skewed error distribution.




FIG 9. Two airfoils with similar geometries but vastly different aerodynamic features. Airfoil (a) is taken from the dataset and airfoil (b) is a slight modification of this airfoil.

One important aspect to take into account is that the error in the features may depend on the values of the features themselves, since the model is generally worse at generating airfoils for features that are less well represented in the training set. For example, the distribution in Fig. 4 shows that the density of the samples is not uniform in the feature space. In particular, the sample density for higher C_D values decreases. If the errors are separated by a C_D threshold as done in Table 2, the error in the drag coefficient almost increases tenfold. Therefore, it is recommended to make the dataset more uniform or to

weigh the samples by their local feature density in future work. Interestingly, note that the lift and momentum errors are less affected due to the strong generalization properties of neural networks.

TAB 2. Feature errors for different regions of drag coefficient values.

Feature	C_L	C_D	C_M
RMSE ($C_D < 0.01$)	7.7e-2	1.1e-3	1.7e-2
RMSE ($C_D > 0.01$)	9.8e-2	9.1e-3	1.8e-2

3.2. Airfoil uniqueness

To show that the generated airfoils from the features in the test set are unique, they are compared to the closest airfoils in the training dataset. Fig. 10 shows that while the generated airfoils might resemble the types of airfoils in the training dataset, they are generally not identical.

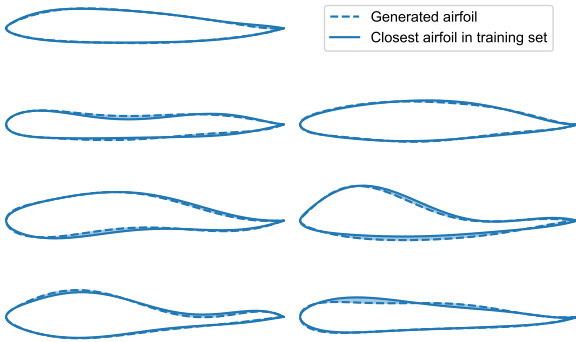


FIG 10. Examples of generated airfoils and their closest airfoils in the training set.

Note that naturally, the distance to the nearest airfoils depends on the number of airfoils and resulting feature density of the samples in the dataset. If more airfoils are included, the resemblance of generated airfoils to airfoils in the training set will increase. While this does not mean that the generated airfoils are incorrect, as is shown in Table 3 for a model trained on a larger airfoil dataset, it does make less sense to use DPMs as the geometries might as well be interpolated using simpler methods. However, when using higher-fidelity CFD techniques to quantify the features of airfoils it is unlikely that the datasets are this large.

TAB 3. Errors for the different features when training a model with 64 neurons per layer on a training dataset consisting of 6,000 airfoils.

Feature	C_L	C_D	C_M
RMSE	5.8e-2	7.3e-3	8.0e-3
RMdSE	9.5e-3	3.7e-4	2.0e-3
MAPE	7%	8%	-

Regardless of the solver used, it becomes harder to achieve high feature density when the number of features increases due to the "curse of dimensionality". In the present case only three features are considered, but other features such as thickness, camber, area as well as full polars may be considered. Neural networks in general are known to be robust to the curse of dimensionality [32] and DPMs in particular are good at generalizing high-dimensional features to produce novel samples. Therefore, increasing the number of features remains an important line of future work.

3.3. Airfoil diversity

An advantage of generative methods is that they can nearly instantly generate new geometries if requirements change. As mentioned in the introduction, this is especially useful in the context of multi-disciplinary optimization procedures, where requirements and constraints might evolve during due to their multi-objective nature. For this purpose, it is beneficial if the model can generate diverse geometries for the same desired features because it allows to have multiple starting points for the optimization procedure.

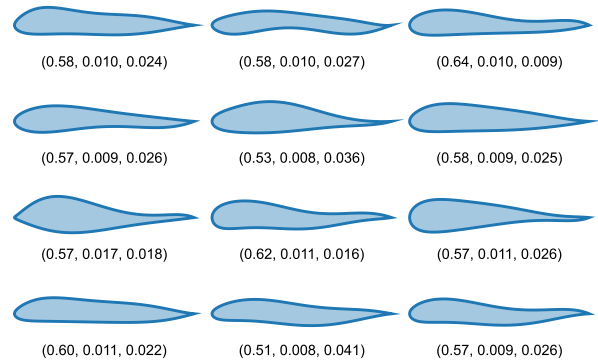


FIG 11. Generated airfoils for $C_L = 0.6$, $C_D = 0.01$ and $C_M = 0.02$.

Fig. 11 shows the geometries that are generated for the same set of desired features, showing a diverse set of airfoils. Note that apart from small derivations of the same airfoil, the samples also contain different "classes" of airfoils altogether. This is an advantage compared to non-probabilistic methods, which can only output one airfoil that is often of just one class. Nevertheless, it is important to keep in mind that that these airfoils do not necessarily have the same features as desired due to the inaccuracies described in Section 3.1.

4. CONCLUSION

All in all, the results show that diffusion probabilistic models provide an effective method for direct aerodynamic shape design. By combining them with a field-specific latent space, they allow for the generation of smooth, diverse and unique airfoils. Further-

more, they allow to impose requirements and constraints easily through conditioning and filtering of the dataset. However, we acknowledge that there are certain gaps to be further explored. For example, the options for conditioning and filtering should be expanded, and in particular future work should also consider structural requirements. In addition, the choice of model architecture and latent space representations should be further explored. Lastly, the effectivity of the method should also be tested with high-fidelity methods, as they can be even more sensitive to small changes in the design space. Nevertheless, the present work provides a starting point for future work on using DPMs for aerodynamic shape design.

Contact address:

thomas.wagenaar@airbus.com

Acknowledgments

This project has received funding from the Clean Aviation Joint Undertaking under the European Union's Horizon Europe research and innovation programme under Grant Agreement HERA (Hybrid-Electric Regional Architecture) n° 101102007. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or CAJU. Neither the European Union nor the granting authority can be held responsible.



References

- [1] Joaquim R.R.A. Martins. Aerodynamic design optimization: Challenges and perspectives. *Computers & Fluids*, 239:105391, May 2022. ISSN:0045-7930. DOI: [10.1016/j.compfluid.2022.105391](https://doi.org/10.1016/j.compfluid.2022.105391).
- [2] Jichao Li, Xiaosong Du, and Joaquim R.R.A. Martins. Machine learning in aerodynamic shape optimization. *Progress in Aerospace Sciences*, 134:100849, Oct. 2022. ISSN:0376-0421. DOI: [10.1016/j.paerosci.2022.100849](https://doi.org/10.1016/j.paerosci.2022.100849).
- [3] Wei Chen, Kevin Chiu, and Mark D. Fuge. Airfoil design parameterization and optimization using bézier generative adversarial networks. *AIAA Journal*, 58(11):4723–4735, Nov. 2020. ISSN: 1533-385X. DOI: [10.2514/1.j059317](https://doi.org/10.2514/1.j059317).
- [4] Yuyang Wang, Kenji Shimada, and Amir Barati Farimani. Airfoil gan: encoding and synthesizing airfoils for aerodynamic shape optimization. *Journal of Computational Design and Engineering*, 10(4):1350–1362, June 2023. ISSN: 2288-5048. DOI: [10.1093/jcde/qwad046](https://doi.org/10.1093/jcde/qwad046).
- [5] Gang Sun and Shuyue Wang. A review of the artificial neural network surrogate modeling in aerodynamic design. *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, 233(16):5863–5872, July 2019. ISSN:2041-3025. DOI: [10.1177/0954410019864485](https://doi.org/10.1177/0954410019864485).
- [6] Xiaosong Du, Ping He, and Joaquim R.R.A. Martins. Rapid airfoil design optimization via neural networks-based parameterization and surrogate modeling. *Aerospace Science and Technology*, 113:106701, June 2021. ISSN: 1270-9638. DOI: [10.1016/j.ast.2021.106701](https://doi.org/10.1016/j.ast.2021.106701).
- [7] Sangeun Oh, Yongsu Jung, Seongsin Kim, Ikjin Lee, and Namwoo Kang. Deep generative design: Integration of topology optimization and generative models. *Journal of Mechanical Design*, 141(11), Sept. 2019. ISSN: 1528-9001. DOI: [10.1115/1.4044229](https://doi.org/10.1115/1.4044229).
- [8] Gabriel Achour, Woong Je Sung, Olivia J. Pinon-Fischer, and Dimitri N. Mavris. Development of a Conditional Generative Adversarial Network for Airfoil Shape Optimization. In *AIAA Scitech 2020 Forum*, Orlando, FL, Jan. 2020. American Institute of Aeronautics and Astronautics. ISBN: 9781624105951. DOI: [10.2514/6.2020-2261](https://doi.org/10.2514/6.2020-2261).
- [9] Xavier Tan, Dai Manna, Joyjit Chattoraj, Liu Yong, Xu Xinxing, Dao My Ha, and Yang Feng. Airfoil Inverse Design using Conditional Generative Adversarial Networks. In *2022 17th International Conference on Control, Automation, Robotics and Vision (ICARCV)*, pages 143–148, Dec. 2022. DOI: [10.1109/ICARCV57592.2022.10004343](https://doi.org/10.1109/ICARCV57592.2022.10004343).
- [10] Joyjit Chattoraj, Jian Cheng Wong, Zhang Zexuan, Manna Dai, Xia Yingzhi, Li Jichao, Xu Xinxing, Ooi Chin Chun, Yang Feng, Dao My Ha, and Liu Yong. Tailoring Generative Adversarial Networks for Smooth Airfoil Design, 2024. DOI: [10.48550/ARXIV.2404.11816](https://doi.org/10.48550/ARXIV.2404.11816), <https://arxiv.org/abs/2404.11816>.
- [11] Emre Yilmaz and Brian German. Conditional Generative Adversarial Network Framework for Airfoil Inverse Design. In *AIAA AVIATION 2020 FORUM, VIRTUAL EVENT*, June 2020. American Institute of Aeronautics and Astronautics. ISBN: 9781624105982. DOI: [10.2514/6.2020-3185](https://doi.org/10.2514/6.2020-3185).
- [12] Amin Heyrani Nobari, Wei (Wayne) Chen, and Faez Ahmed. Range-gan: Design synthesis under constraints using conditional generative adversarial networks. *Journal of Mechanical Design*, page 1–16, Sept. 2021. ISSN: 1528-9001. DOI: [10.1115/1.4052442](https://doi.org/10.1115/1.4052442).

- [13] Maciej Wiatrak, Stefano V. Albrecht, and Andrew Nystrom. Stabilizing generative adversarial networks: A survey, 2019. DOI: [10.48550/ARXIV.1910.00927](https://arxiv.org/abs/1910.00927), <https://arxiv.org/abs/1910.00927>.
- [14] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 8780–8794. Curran Associates, Inc., 2021.
- [15] François Mazé and Faez Ahmed. Diffusion models beat gans on topology optimization. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(8):9108–9116, June 2023. ISSN: 2159-5399. DOI: [10.1609/aaai.v37i8.26093](https://arxiv.org/abs/2303.08133).
- [16] Yanxuan Zhao, Peng Zhang, Guopeng Sun, Zhigong Yang, Jianqiang Chen, and Yueqing Wang. Ccdpm: A continuous conditional diffusion probabilistic model for inverse design. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(15):17033–17041, Mar. 2024. ISSN: 2159-5399. DOI: [10.1609/aaai.v38i15.29647](https://arxiv.org/abs/2303.08133).
- [17] Alex Morehead and Jianlin Cheng. Geometry-complete diffusion for 3d molecule generation and optimization. *Communications Chemistry*, 7(1), July 2024. ISSN: 2399-3669. DOI: [10.1038/s42004-024-01233-z](https://arxiv.org/abs/2112.10752).
- [18] Xiaohui Zeng, Arash Vahdat, Francis Williams, Zan Gojcic, Or Litany, Sanja Fidler, and Karsten Kreis. Lion: Latent point diffusion models for 3d shape generation. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 10021–10039. Curran Associates, Inc., 2022.
- [19] Cashen Diniz and Mark Fuge. Optimizing diffusion to diffuse optimal designs. In *AIAA SCITECH 2024 Forum*. American Institute of Aeronautics and Astronautics, Jan. 2024. DOI: [10.2514/6.2024-2013](https://arxiv.org/abs/1503.03585).
- [20] Jascha Sohl-Dickstein, Eric A. Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics, 2015. DOI: [10.48550/ARXIV.1503.03585](https://arxiv.org/abs/1503.03585), <https://arxiv.org/abs/1503.03585>.
- [21] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models, 2020. DOI: [10.48550/ARXIV.2006.11239](https://arxiv.org/abs/2006.11239), <https://arxiv.org/abs/2006.11239>.
- [22] Shitong Luo and Wei Hu. Diffusion probabilistic models for 3d point cloud generation. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, June 2021. DOI: [10.1109/cvpr46437.2021.00286](https://arxiv.org/abs/2303.08133).
- [23] Edward J. Smith and David Meger. Improved adversarial systems for 3d object generation and reconstruction. In *Proceedings of the 1st Annual Conference on Robot Learning*, volume 78 of *Proceedings of Machine Learning Research*, pages 87–96. PMLR, 13–15 Nov 2017.
- [24] Zhen Liu, Yao Feng, Michael J. Black, Derek Nowrouzezahrai, Liam Paull, and Weiyang Liu. Meshdiffusion: Score-based generative 3d mesh modeling, 2023. DOI: [10.48550/ARXIV.2303.08133](https://arxiv.org/abs/2303.08133), <https://arxiv.org/abs/2303.08133>.
- [25] Gene Chou, Yuval Bahat, and Felix Heide. Diffusion-sdf: Conditional generative modeling of signed distance functions. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*. IEEE, Oct. 2023. DOI: [10.1109/iccv51070.2023.00215](https://arxiv.org/abs/2112.10752).
- [26] Brenda Kulfan. Universal parametric geometry representation method. *Journal of Aircraft - J AIRCRAFT*, 45:142–158, 01 2008. DOI: [10.2514/1.29958](https://arxiv.org/abs/2112.10752).
- [27] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models, 2021. DOI: [10.48550/ARXIV.2112.10752](https://arxiv.org/abs/2112.10752), <https://arxiv.org/abs/2112.10752>.
- [28] Yasi Wang, Hongxun Yao, and Sicheng Zhao. Auto-encoder based dimensionality reduction. *Neurocomputing*, 184:232–242, Apr. 2016. ISSN: 0925-2312. DOI: [10.1016/j.neucom.2015.08.104](https://arxiv.org/abs/1412.6980).
- [29] Michael S. Selig. Uiuic airfoil data site, 1996.
- [30] Mark Drela. Xfoil: An analysis and design system for low reynolds number airfoils. In Thomas J. Mueller, editor, *Low Reynolds Number Aerodynamics*, pages 1–12, Berlin, Heidelberg, 1989. Springer Berlin Heidelberg.
- [31] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2014. DOI: [10.48550/ARXIV.1412.6980](https://arxiv.org/abs/1412.6980), <https://arxiv.org/abs/1412.6980>.
- [32] Zheyuan Hu, Khemraj Shukla, George Em Karniadakis, and Kenji Kawaguchi. Tackling the curse of dimensionality with physics-informed neural networks. *Neural Networks*, 176:106369, Aug. 2024. ISSN: 0893-6080. DOI: [10.1016/j.neunet.2024.106369](https://arxiv.org/abs/2006.11239).