
MO-PADGAN: REPARAMETERIZING ENGINEERING DESIGNS FOR AUGMENTED MULTI-OBJECTIVE OPTIMIZATION

Wei Chen

Siemens Corporate Technology
Princeton, NJ 08540
chen.wei@siemens.com

Faez Ahmed

Massachusetts Institute of Technology
Cambridge, MA 02139
faez@mit.edu

ABSTRACT

Multi-objective optimization is key to solving many Engineering Design problems, where design parameters are optimized for several performance indicators. However, optimization results are highly dependent on how the designs are parameterized. Researchers have shown that deep generative models can learn compact design representations, providing a new way of parameterizing designs to achieve faster convergence and improved optimization performance. Despite their success in capturing complex distributions, existing generative models face three challenges when used for design problems: 1) generated designs have limited design space coverage, 2) the generator ignores design performance, and 3) the new parameterization is unable to represent designs beyond training data. To address these challenges, we propose MO-PaDGAN, which adds a Determinantal Point Processes based loss function to the generative adversarial network to simultaneously model diversity and (multi-variate) performance. MO-PaDGAN can thus improve the performances and coverage of generated designs, and even generate designs with performances exceeding those from training data. When using MO-PaDGAN as a new parameterization in multi-objective optimization, we can discover much better Pareto fronts even though the training data do not cover those Pareto fronts. In a real-world multi-objective airfoil design example, we demonstrate that MO-PaDGAN achieves, on average, a 186% improvement in the hypervolume indicator when compared to the vanilla GAN or other state-of-the-art parameterization methods.

1 Introduction

The goal of many Engineering Design applications is to find designs which perform well on multiple performance goals. For example, bike engineers want to design the geometry of road bikes which are lightweight, have less drag and exhibits good ride quality. Such design applications often have problems with huge collections of data (CAD models, images, microstructures, *etc.*) with hundreds of features and multiple functionalities. These designs are commonly represented using parametric modeling, which defines design geometry using hard-coded features and constraints [1]. Once a parametric model is defined for a design, it can be used for design synthesis (generating new designs by changing design parameters) or for design optimization (finding optimal parameters which maximize the desirable performance goals). However, there are two issues with traditional hard-coded parametric modeling: 1) the parameterization is often not flexible to cover all possible variations of the chosen type of designs, and 2) the parametric representation may contain a large number of parameters, which often leads to infeasible designs and increased optimization cost. To alleviate these issues, researchers in Engineering Design have used machine learning techniques (deep generative models in particular) to replace hard-coded parametric models with compact representations learned from data [2, 3, 4]. Those learned representations can then be treated as parameterizations that are fed into a design optimization method to find one or more designs with optimal performance measures. Variational autoencoders (VAEs) [5] and Generative Adversarial Networks (GANs) [6] are the two most commonly used deep generative models for these tasks. These models are applied in design optimization over domains including microstructural design [7], 3D modeling [8], and aerodynamic shape design [9]. However, existing generative models, whose goal is learning the distribution of existing designs, face three challenges when being used for parameterizing designs: 1) generated designs may have limited design space coverage (*e.g.*, when mode collapse happens with GANs), 2) the generator ignores design performance, and 3) the new parameterization is unable to represent designs outside the training data.

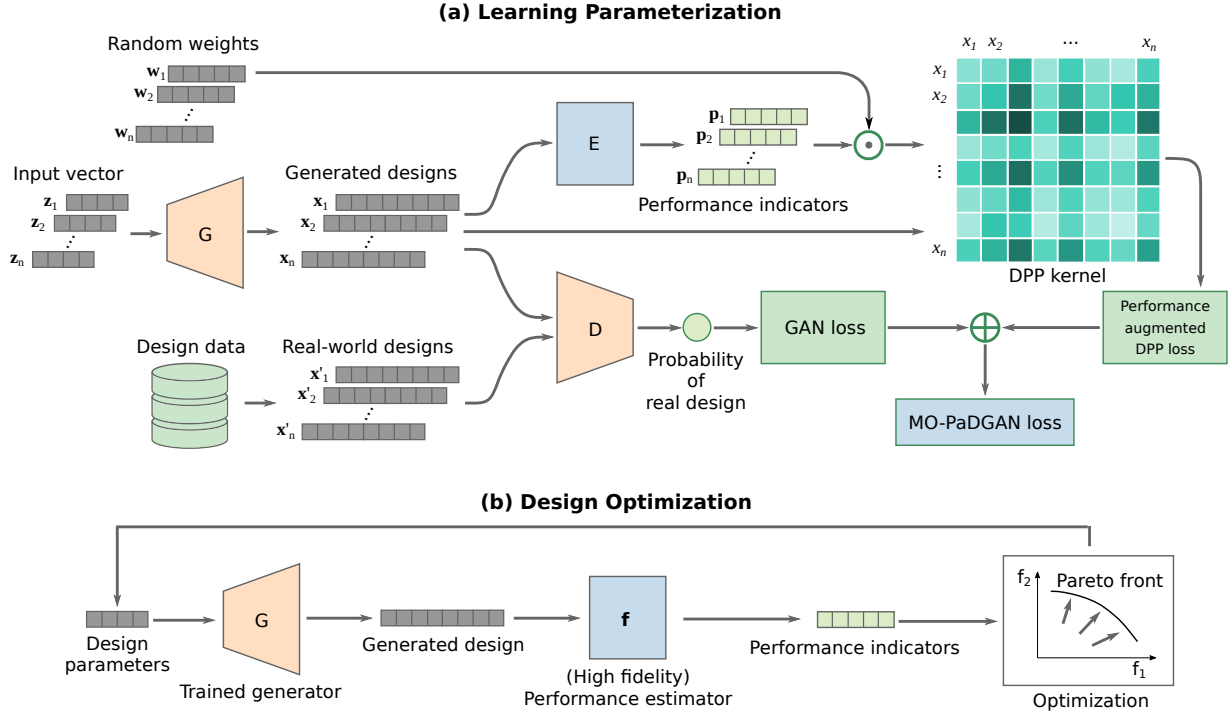


Figure 1: Overview of our method: (a) the MO-PaDGAN architecture at the training stage (The operator \odot denotes performance aggregation); (b) the multi-objective optimization loop at the design optimization stage.

In this work, we address all the three challenges by modifying the architecture and loss of deep generative models, which allows simultaneous maximization of generated designs’ diversity and (possibly multivariate) performance. With such a setting, we develop a new variant of GAN, named MO-PaDGAN (Multi-Objective Performance Augmented Diverse Generative Adversarial Network). We use MO-PaDGAN as a parameterization in multi-objective optimization tasks and demonstrate that it contributes to large improvements in the final Pareto optimal solutions when compared to a baseline GAN model and other state-of-art parameterizations. The contributions of this work are as follows:

1. We propose MO-PaDGAN which models both diversity and (multivariate) performance using a DPP kernel-based loss function in GANs, such that the generated samples have higher quality and better coverage of the design space.
2. We show that the latent representation learned by MO-PaDGAN can be used to parameterize designs for multi-objective optimization.
3. We show that MO-PaDGAN parameterization leads to solutions that are closer to the true Pareto optimal front. For a real-world airfoil design example, the method leads to *186% average improvement* of the hypervolume indicator compared to three state-of-art methods.
4. We find that our method allows the discovery of novel high-performance designs that it had not seen from existing data.

2 Background

Below we provide background on Generative Adversarial Networks and Determinantal Point Processes, which are the two key ingredients of our method.

2.1 Generative Adversarial Nets

Generative Adversarial Networks [6] model a game between a generative model (*generator*) and a discriminative model (*discriminator*). The generator G maps an arbitrary noise distribution to the data distribution (*i.e.*, the distribution of designs in our scenario), thus can generate new data; while the discriminator D tries to perform classification, *i.e.*, to distinguish between real and generated data. Both G and D are usually built with deep neural networks. As D improves

its classification ability, G also improves its ability to generate data that fools D . Thus, a GAN has the following objective function:

$$\min_G \max_D V(D, G) = \mathbb{E}_{\mathbf{x} \sim P_{data}} [\log D(\mathbf{x})] + \mathbb{E}_{\mathbf{z} \sim P_{\mathbf{z}}} [\log(1 - D(G(\mathbf{z})))] \quad (1)$$

where \mathbf{x} is sampled from the data distribution P_{data} , \mathbf{z} is sampled from the noise distribution $P_{\mathbf{z}}$, and $G(\mathbf{z})$ is the generator distribution. A trained generator thus can map from a predefined noise distribution to the distribution of designs. The noise input \mathbf{z} is considered as the latent representation of the data, which can be used for design synthesis and exploration. Note that GANs often suffer from *mode collapse* [10], where the generator fails to capture all modes of the data distribution. In this work, by maximizing the diversity objective, mode collapse is discouraged as it leads to less diverse samples.

2.2 Determinantal Point Processes

Determinantal Point Processes (DPPs), which arise in quantum physics, are probabilistic models that model the likelihood of selecting a subset of diverse items as the determinant of a kernel matrix. Viewed as joint distributions over the binary variables corresponding to item selection, DPPs essentially capture negative correlations and provide a way to elegantly model the trade-off between often competing notions of quality and diversity. The intuition behind DPPs is that the determinant of a kernel matrix roughly corresponds to the volume spanned by the vectors representing the items. Points that “cover” the space well should capture a larger volume of the overall space, and thus have a higher probability.

DPP kernels can be decomposed into quality and diversity parts [11]. The (i, j) -th entry of a positive semi-definite DPP kernel L can be expressed as:

$$L_{ij} = q_i \phi(i)^T \phi(j) q_j. \quad (2)$$

We can think of $q_i \in R^+$ as a scalar value measuring the quality of an item i , and $\phi(i)^T \phi(j)$ as a signed measure of similarity between items i and j . The decomposition enforces L to be positive semidefinite. Suppose we select a subset S of samples, then this decomposition allows us to write the probability of this subset S as the square of the volume spanned by $q_i \phi_i$ for $i \in S$ using the equation below:

$$\mathbb{P}_L(S) \propto \prod_{i \in S} (q_i^2) \det(K_S), \quad (3)$$

where K_S is the similarity matrix of S . As item i ’s quality q_i increases, so do the probabilities of sets containing item i . As two items i and j become more similar, $\phi_i^T \phi_j$ increases and the probabilities of sets containing both i and j decrease. The key intuition of MO-PaDGAN is that if we can integrate the probability of set selection from Eq. (3) to the loss function of any generative model, then while training it will be encouraged to generate high probability subsets, which will be both diverse and high-performance.

3 Methodology

MO-PaDGAN adds a *performance augmented DPP loss* to a vanilla GAN architecture which measures the diversity and performance of a batch of generated designs during training. In this section, we first explain how to construct a DPP kernel. Then we show how to bake the DPP kernel into MO-PaDGAN’s loss function. Finally, we formulate the multi-objective design optimization problem where we use MO-PaDGAN’s generator as a parametric model and optimize designs over the learned latent space.

3.1 Creating a DPP kernel

We create the kernel L for a sample of points generated by MO-PaDGAN from known inter-sample similarity values and performance vector.

The similarity terms $\phi(i)^T \phi(j)$ can be derived using any similarity kernel, which we represent using $k(\mathbf{x}_i, \mathbf{x}_j) = \phi(i)^T \phi(j)$ and $\|\phi(i)\| = \|\phi(j)\| = 1$. Here \mathbf{x}_i is a vector representation of a design. Note that in a DPP model, the quality of an item is a scalar value representing design performance such as compliance, displacement, drag-coefficient. For multivariate performance, we use a *performance aggregator* to obtain a scalar quality $q(\mathbf{x}) = \mathbf{w}^T \mathbf{p}$, where $\mathbf{p} = (p_1, \dots, p_M)^T$ are M -dimensional performances and the corresponding weights $\mathbf{w} = (w_1, \dots, w_M)^T$ are positive numbers sampled uniformly at random and $\mathbf{1}^T \mathbf{w} = 1$. The performance vector \mathbf{p} can be estimated using a *performance estimator* E , i.e., $\mathbf{p} = E(\mathbf{x})$. The performance estimator can be a physics simulator or a surrogate model. Maximizing

$q(\mathbf{x})$ pushes the non-dominated Pareto set of generated samples in the performance space to have higher values. While more complex performance aggregators (*e.g.*, the Chebyshev distance from an ideal performance vector) are also applicable, we use the common linear scalarization to have fewer assumptions about the performance space.

3.2 Performance Augmented DPP Loss

The overall model architecture of MO-PaDGAN at the training stage is shown in Fig. 1(a). Our performance augmented DPP loss models diversity and performance simultaneously and gives a lower loss to sets of designs which are both high-performance and diverse. Specifically, we construct a kernel matrix L_B for a generated batch B based on Eq. (2). For each entry of L_B , we have

$$L_B(i, j) = k(\mathbf{x}_i, \mathbf{x}_j) (q(\mathbf{x}_i)q(\mathbf{x}_j))^{\gamma_0}, \quad (4)$$

where $\mathbf{x}_i, \mathbf{x}_j \in B$, $q(\mathbf{x})$ is the quality function at \mathbf{x} , and $k(\mathbf{x}_i, \mathbf{x}_j)$ is the similarity kernel between \mathbf{x}_i and \mathbf{x}_j . To allow trade-off between quality and diversity, we adjust the dynamic range of the quality scores by using an exponent (γ_0) as a parameter to change the distribution of quality. A larger γ_0 increases the relative importance of quality as compared to diversity, which provides the flexibility to a user of MO-PaDGAN in deciding emphasis on quality vs diversity.

The performance augmented DPP loss is expressed as

$$\mathcal{L}_{\text{PaD}}(G) = -\frac{1}{|B|} \log \det(L_B) = -\frac{1}{|B|} \sum_{i=1}^{|B|} \log \lambda_i, \quad (5)$$

where λ_i is the i -th eigenvalue of L_B . We add this loss to the vanilla GAN's objective in Eq. (1) and form a new objective:

$$\min_G \max_D V(D, G) + \gamma_1 \mathcal{L}_{\text{PaD}}(G), \quad (6)$$

where γ_1 controls the weight of $\mathcal{L}_{\text{PaD}}(G)$. For the backpropagation step, to update the weight θ_G^i in the generator in terms of $\mathcal{L}_{\text{PaD}}(G)$, we descend its gradient based on the chain rule:

$$\frac{\partial \mathcal{L}_{\text{PaD}}(G)}{\partial \theta_G^i} = \sum_{j=1}^{|B|} \left(\frac{\partial \mathcal{L}_{\text{PaD}}(G)}{\partial q(\mathbf{x}_j)} \frac{dq(\mathbf{x}_j)}{d\mathbf{x}_j} + \frac{\partial \mathcal{L}_{\text{PaD}}(G)}{\partial \mathbf{x}_j} \right) \frac{\partial \mathbf{x}_j}{\partial \theta_G^i}, \quad (7)$$

where $\mathbf{x}_j = G(\mathbf{z}_j)$. Equation (7) indicates a need for $dq(\mathbf{x})/d\mathbf{x}$, which is the gradient of the quality function. In practice, this gradient is accessible when the quality is evaluated through a performance estimator that is differentiable (*e.g.*, using adjoint-based solver). If the gradient of a performance estimator is not available, one can either use numerical differentiation or approximate the quality function using a differentiable surrogate model (*e.g.*, a neural network-based surrogate model, as used in our experiments).

3.3 Optimizing Designs over Latent Representation

After training with the performance augmented DPP loss, the generator forms a parametric model where the parameters are its input \mathbf{z} , as shown in Fig. 1(b). The multi-objective optimization problem can then be formulated as (when maximizing performances):

$$\begin{aligned} \max_{\mathbf{z}} \quad & f_m(G(\mathbf{z})), m = 1, \dots, M \\ \text{s.t.} \quad & g_j(G(\mathbf{z})) \leq 0, j = 1, \dots, J \\ & h_k(G(\mathbf{z})) = 0, k = 1, \dots, K \\ & z_d^{\text{inf}} \leq z_d \leq z_d^{\text{sup}}, d = 1, \dots, D, \end{aligned} \quad (8)$$

where f_m is the estimator of the m -th performance indicator for a given design. In practice, it can be, for example, a Computational Fluid Dynamics (CFD) or Finite Element Method (FEM) solver, where we can evaluate the physics properties of any given design. Note that $\mathbf{f} = (f_1, \dots, f_M)$ is not necessarily the same as the performance estimator E . To balance cost and accuracy, we can use a cheap low-fidelity model for E and an expensive high-fidelity model for \mathbf{f} , since MO-PaDGAN training requires much more frequent performance evaluations whereas design optimization needs higher accuracy as it leads to the final product. Functions g_j and h_k define the j -th inequality and the k -th equality constraints, respectively. The d -th design parameter z_d is bounded by z_d^{inf} and z_d^{sup} . There is no limitation of using any specific type of optimizer. Note that in the case where gradients are required (for gradient-based optimization), the functions f_m , g_j , and h_k need to be differentiable (*e.g.*, by using adjoint solvers). This allows access to the gradients of the objective and constraints with respect to the design $\mathbf{x} = G(\mathbf{z})$. Since we can already obtain $d\mathbf{x}/d\mathbf{z}$ using automatic differentiation, the gradients of the objective and constraints at \mathbf{z} can be computed using the chain rule.

4 Experimental Results

In this section, we demonstrate the merit of using MO-PaDGAN as a parameterization in multi-objective optimization problems via two synthetic examples and a real-world airfoil design example. Please refer to our code and data for reproducing the experimental results¹.

4.1 Synthetic Examples




Figure 2: Synthetic data and test quality (objective) functions.

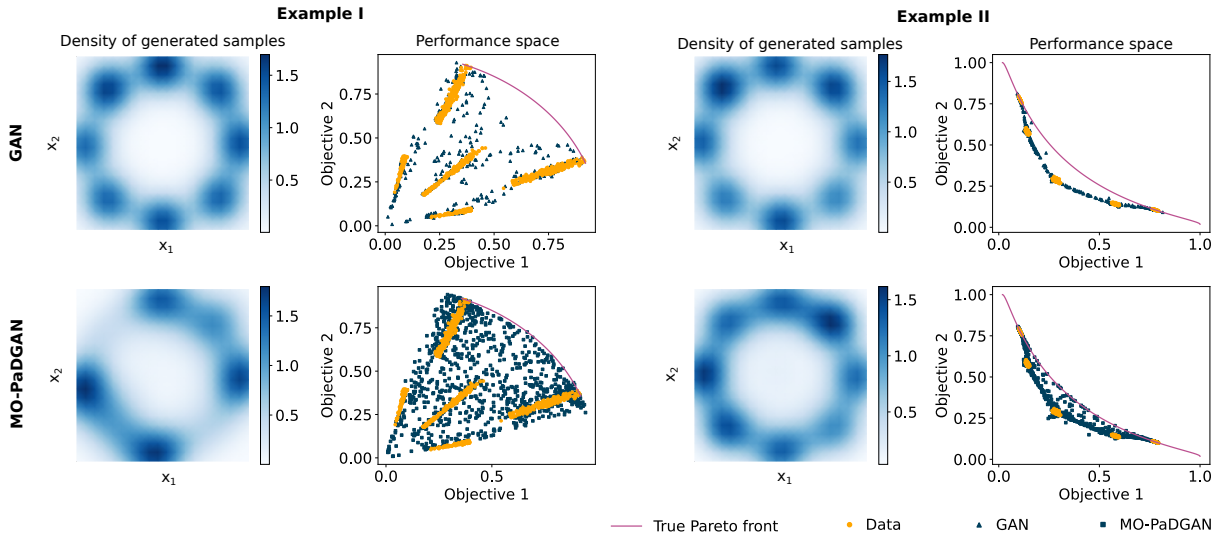


Figure 3: Distribution and performance of generated points for the synthetic examples. While the data has few points near the true Pareto fronts, many of the generated samples from MO-PaDGAN are near the Pareto front.

We create two synthetic examples, with eight clusters each, where cluster centers are evenly spaced around a circle (Fig. 2). The sample size is 10,000 points. Both the synthetic examples have two parameters and two objectives, which enables us to visualize the results and compare the performance against a known ground truth Pareto optimal front. In the example following this section, we discuss a real-world airfoil example, which has larger dimensionality.

The objective functions of both synthetic examples are plotted in Fig. 2. Please see the appendix for more details about the objective functions used.

¹Code and data is included in the appendix and will be open-sourced when the paper is accepted.

We use the same network architecture and hyperparameter settings in both examples. The generator has four fully connected layers. The input dimension is two. Each hidden layer has 128 nodes and a LeakyReLU activation. We apply a hyperbolic tangent activation at the last layer. The discriminator has three fully connected layers. Each hidden layer has 128 nodes and a LeakyReLU activation. We apply a Sigmoid activation at the last layer. For MO-PaDGAN, we set $\gamma_0 = 2$ and $\gamma_1 = 0.5$. We use a RBF kernel with a bandwidth of 1 when constructing the DPP kernel matrix L_B in Eq. (4), *i.e.*, $k(\mathbf{x}_i, \mathbf{x}_j) = \exp(-0.5\|\mathbf{x}_i - \mathbf{x}_j\|^2)$. We train both the MO-PaDGAN and the vanilla GAN for 10,000 iterations with 32 training samples drawn randomly from data at each iteration. For both generator and discriminator, we use Adam optimizer with a learning rate of 0.0001.

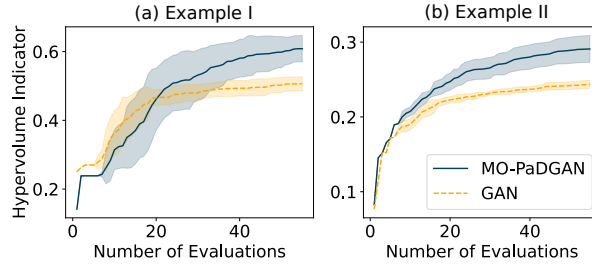


Figure 4: Multi-objective optimization history for synthetic examples. Results shown are mean and standard deviation from ten optimization runs.




Figure 5: The union of solution sets obtained by ten runs of multi-objective optimization for synthetic examples. MO-PaDGAN solutions are closer to the Pareto front than GAN.

Figure 3 visualizes the parameter space and the performance (objective) space of points generated by a vanilla GAN and MO-PaDGAN, respectively. It shows that in the performance space, data are separated into several clusters, just as in the parameter space. The vanilla GAN approximates the data distribution well, except that it generates some points in between clusters, which is reasonable since the noise input \mathbf{z} is continuous. Most points from the training data or generated by GAN are away from the (ground truth) Pareto fronts. In contrast, MO-PaDGAN generates points that fill up the gaps between clusters in the performance space and has many solutions near the Pareto fronts. This shows that guided by the performance augmented DPP loss, MO-PaDGAN pushes generated points towards higher performance in all objectives, although these new points are not seen in the dataset.

We then perform multi-objective Bayesian optimization [12] using the trained generators as parameterizations. We sample 5 points via Latin Hypercube Sampling (LHS) [13] for initial evaluations. This is followed by Bayesian optimization to select 50 additional evaluations. We use the hypervolume indicator [14] to assess the quality of a solution set. The hypervolume measures the volume of the dominated portion of the performance space bounded from below by a reference point. We use $(0, 0)$ as the reference points for both examples. Figure 4 shows the history of hypervolume indicators during optimization from ten runs. Figure 5 shows the union of final solution sets from these ten runs. The results indicate that MO-PaDGAN achieves a much better performance than the vanilla GAN in discovering the Pareto front. This is expected since, as shown in Fig. 3, the parameterization learned by MO-PaDGAN can represent high-performance samples near the true Pareto front, whereas a vanilla GAN cannot.




Figure 6: Airfoil shapes randomly drawn from the training data or generated by parameterizations.

4.2 Airfoil Design Example

An airfoil is the cross-sectional shape of an airplane wing or a propeller/rotor/turbine blade. Airfoil shape optimization is crucial when designing a wing or a blade, as we usually optimize multiple 2D airfoils at different cross-sections and interpolate between sections to get the full 3D geometry. In this example, we use the UIUC airfoil database² as our data source. It provides the geometries of nearly 1,600 real-world airfoil designs. We preprocessed and augmented the dataset similar to [9], which led to a dataset of 38,802 airfoils, each of which is represented by 192 surface points (*i.e.*, $\mathbf{x}_i \in \mathbb{R}^{192 \times 2}$). Figure 6(a) shows airfoil shapes randomly drawn from the training data. We use two performance indicators for designing the airfoils — the lift coefficient (C_L) and the lift-to-drag ratio (C_L/C_D). These two are common objectives in aerodynamic design optimization problems and have been used in different multi-objective optimization studies [15]. We use XFOIL [16] for CFD simulations and compute C_L and C_D values³. We scaled the performance scores between 0 and 1. To provide the gradient of the quality function for Eq. (7), we trained a neural network-based surrogate model on all 38,802 airfoils to approximate both C_L and C_D .

We use a residual neural network (ResNet) [17] as the surrogate model to predict the performance indicators and a BézierGAN [9, 18, 19] to parameterize airfoils. Please refer to the appendix and the code for details on their network architectures. Different from the vanilla GAN’s architecture, BézierGAN’s generator has two inputs — the noise vector and the *latent code*, which follows the InfoGAN’s setting [20]. The additional latent code provides a more disentangled representation and hence will be used as our design parameters. We set the noise dimension and the latent dimension to 10 and 5, respectively. For simplicity, we refer to the BézierGAN as a vanilla GAN and the BézierGAN with the loss \mathcal{L}_{PaD} as a MO-PaDGAN in the rest of the paper. We follow the same training configuration (*i.e.*, training iterations, batch size, optimizer, and learning rate) as in the synthetic examples. For MO-PaDGAN, we set $\gamma_0 = 5$ and $\gamma_1 = 0.2$. We use the RBF kernel with a bandwidth of 1 in the DPP kernel matrix.




Figure 7: Randomly sampled airfoils embedded into a 2D space via t-SNE.

Airfoils randomly generated by MO-PaDGAN and GAN are shown in Fig. 6(b) and (c), respectively. Both MO-PaDGAN and GAN can generate realistic airfoil shapes. To compare the distribution of real and generated airfoils in the design space, we map randomly sampled airfoils into a two-dimensional space through t-SNE, as shown in Figure 7. The results indicate that compared to a vanilla GAN, MO-PaDGAN can generate airfoils that are outside the boundary of the training data, driven by the tendency to maximize diversity using the DPP loss. Figure 8 visualizes the joint distribution of C_L and C_L/C_D for randomly sampled/generated airfoils. It shows that MO-PaDGAN generates airfoils with performances exceed randomly sampled airfoils from training data and the vanilla GAN (*i.e.*, the non-dominated Pareto set of generated samples is pushed further in the performance space to have higher values). Hence, figures 7 and 8 indicate that MO-PaDGAN can expand the existing boundary of the design space towards high-performance regions

²http://m-selig.ae.illinois.edu/ads/coord_database.html

³We set $C_L = C_L/C_D = 0$ for unsuccessful simulations.




Figure 8: Performance space visualization for airfoil samples shown in Fig. 7.

outside the training data. This directed expansion is allowed since the MO-PaDGAN’s generator is updated with the quality gradients (*i.e.*, $dq(\mathbf{x})/dx$ in Eq. (7)).




Figure 9: Hypervolume indicator history for airfoil optimization using different parameterizations. Results shown are mean and standard deviation from ten optimization runs. MO-PaDGAN achieves on average a 186% improvement in the final hypervolume indicator compared to the second-best parameterization.




Figure 10: The union of solution sets obtained by ten runs of multi-objective airfoil design optimization using MO-PaDGAN, GAN, SVD, and FFD. We see that MO-PaDGAN solutions dominate solutions from all other methods.

We then perform multi-objective Bayesian optimization using the trained generators as parameterizations. We use only the latent codes as design parameters and fix the noise vector to zeros. We sample 15 points via LHS for initial evaluations and 150 subsequent evaluations for Bayesian optimization. To assess the optimization performance, we use $(0, 0)$ as the reference points in the hypervolume indicator.

We also compare MO-PaDGAN and GAN with other two state-of-art airfoil parameterizations — Singular Value Decomposition (SVD) [21, 22] and Free-Form Deformation (FFD) [23]. The SVD extracts salient airfoil deformation modes from a dataset using truncated SVD and uses the weights of those modes as design parameters. We set the number of extracted modes to 5 in this experiment (*i.e.*, 5 design parameters). The FFD represents a new shape by deforming some initial shape via moving a set of $m \times n$ control points. We set the number of control points to 3 and only move their y coordinates (*i.e.*, 12 design parameters). We refer interested readers to our code for detailed algorithms and settings of these two parameterizations. Figure 6(d) and (e) shows randomly synthesized airfoils using SVD and FFD. Compared to airfoils drawn from data or generated by MO-PaDGAN/GAN, airfoils synthesized by SVD/FFD either have unrealistic (*e.g.*, self-intersecting) shapes or insufficient diversity. This is expected since, as generative models, the




Figure 11: Non-dominated set among solutions shown in Fig. 10. We see that MO-PaDGAN outperforms other parameterizations on both objectives.




Figure 12: Novelty indicators of airfoils shown in Fig. 11, shows that MO-PaDGAN solutions are most novel.

objective of GAN and its variants is to learn data distribution, so it is able to generate designs with similar (realistic) appearance and diversity to the data if trained properly, whereas non-generative models do not capture the feasible boundary.

Figure 9 shows that, when using a MO-PaDGAN for parameterization, the hypervolume indicators are significantly higher than other parameterizations throughout the whole optimization process. The MO-PaDGAN achieves on average a 186% improvement in the final hypervolume indicator compared to SVD, which is the second-best parameterization. Figure 10 shows the union of solution sets obtained by ten optimization runs, which illustrates a notable gain in the performances of optimized MO-PaDGAN airfoils. We further obtain the non-dominated Pareto set for each parameterization given the points shown in Fig. 10 and plot the shapes of airfoils in the Pareto sets in Fig. 11. It shows that solutions from MO-PaDGAN’s Pareto set dominate all the solutions in other Pareto sets. To assess whether the optimized shapes are novel, *i.e.*, different from any existing shapes (from the dataset), we compute the *novelty indicator* for each optimized shape from Fig. 11. Specifically, the novelty indicator is represented by the difference between each optimized shape and the most similar shape (*i.e.*, the nearest neighbor) from the original dataset. While the difference can be measured with various metrics based on the specific application, in this experiment we use Hausdorff distance, which measures the distance between any two sets of airfoil surface points. The novelty indicators are shown in Fig. 12. It shows that while GAN, as a generative model, can generate realistic designs which exhibit similar diversity as data, it is harder to discover novel solutions in design optimization, compared to SVD or FFD. In contract, also as a generative model, MO-PaDGAN can discover optimal solutions with much higher novelty. This shows that the solutions found on the trade-off front are far from the original dataset. Because by simultaneously modeling diversity and quality, MO-PaDGAN encourages high-performance designs to be diverse and hence expands the space of high-performance designs.

5 Conclusion

We proposed MO-PaDGAN to encourage diverse and high-quality sample generation, where the quality of a sample can be specified by a multivariate metric. This is particularly useful for Engineering Design applications, which often require simultaneous improvement in several performance indicators. By using the new model as a design parameterization, we showed that it expands the design space towards high-performance regions, whereas a vanilla GAN only generates

designs within the original design space bounded by data. As a result, MO-PaDGAN can represent designs that have higher-performance but also are novel compared to existing ones. We further used MO-PaDGAN as a new parameterization in multi-objective optimization tasks. Through two synthetic examples, we found that MO-PaDGAN allows the discovery of the underlying full Pareto fronts even though data do not cover those Pareto fronts. In the airfoil design example, we compared MO-PaDGAN with three state-of-art parameterization methods — a vanilla GAN, SVD, and FFD. We showed that MO-PaDGAN achieves on average a 186% improvement in the hypervolume indicator compared to the second-best parameterization. We also discovered that the Pareto-optimal solutions by using MO-PaDGAN exhibit higher novelty than GAN.

Although we demonstrated the real-world application of MO-PaDGAN through an airfoil design example, this model shows promise in extending to other Engineering problems like material microstructural design and molecule discovery. The proposed model also generalizes to other generative models like VAEs, where we can add the performance augmented DPP loss to the Kullback-Leibler divergence and the reconstruction loss.

Future work will investigate questions focused on two key limitations of the method — a) How to extend the method to domains where gradient information of design performance is not available, and b) How to capture the entire Pareto front uniformly for many-objective optimization problems.

Appendix A Synthetic Examples: Objective Functions

In Example I, we use the modified KNO1 [24] as the objective functions:

$$\begin{aligned} \max_{x'_1, x'_2} f_1 &= (r/20) \cos(\phi) \\ \max_{x'_1, x'_2} f_2 &= (r/20) \sin(\phi) \\ \text{s.t. } r &= 9 - \left[3 \sin \left(\frac{5}{2(x_1 + x_2)^2} \right) + 3 \sin(4(x_1 + x_2)) \right. \\ &\quad \left. + 5 \sin(2(x_1 + x_2) + 2) \right] \\ \phi &= \frac{\pi(x_1 - x_2 + 3)}{12} \\ x_1 &= 3(x'_1 + 0.5) \\ x_2 &= 3(x'_2 + 0.5) \\ x'_1, x'_2 &\in [-0.5, 0.5]. \end{aligned}$$

The Pareto front for this function lies on the line defined by $x_1 + x_2 = 0.4705$ (in the parametric space). There are multiple locally optimal fronts.

In Example II, we use the modified VLMOP2 [25] as the objective functions:

$$\begin{aligned} \max_{x_1, x_2} f_1 &= \exp \left(- \left(x_1 - \frac{1}{\sqrt{2}} \right)^2 - \left(x_2 - \frac{1}{\sqrt{2}} \right)^2 \right) \\ \max_{x_1, x_2} f_2 &= \exp \left(- \left(x_1 + \frac{1}{\sqrt{2}} \right)^2 - \left(x_2 + \frac{1}{\sqrt{2}} \right)^2 \right) \\ \text{s.t. } x_1, x_2 &\in [-0.5, 0.5]. \end{aligned}$$

The Pareto front for this bi-objective problem is concave (in the objective space) and lies on the line defined by $x_1 - x_2 = 0$ (in the parametric space).

For simplicity, both performance estimators E (used in the training stage) and f (used in the multi-objective optimization stage) are defined by the above objective functions. In practice, one can use a lower-fidelity but differentiable E during MO-PaDGAN training and a higher-fidelity f during design optimization, as shown in the airfoil design example.

Appendix B Airfoil Design Example: Neural Network Surrogate Model

To obtain a differentiable performance estimator E for airfoil designs, we train a ResNet as a surrogate model of a higher-fidelity physics solver. The model architecture is shown in Fig. 13. We use eight residual blocks, three of which

has down-sampling. The convolution is performed along the sequence of airfoil surface point coordinates. We use batch normalization and a leaky ReLU activation following each hidden layer, and a Sigmoid activation at the output layer. We use XFOIL to compute the target C_L and C_L/C_D values of the airfoil data and scale the target values between 0 and 1. We use 38,802 airfoil designs as the training data. We train the surrogate model for 10,000 iterations with a batch size of 256 at each iteration. We use the Adam optimizer with a learning rate of 0.0001, $\beta_1 = 0.5$, and $\beta_2 = 0.999$. Please refer to our code for more details on the surrogate model's network and training configurations.

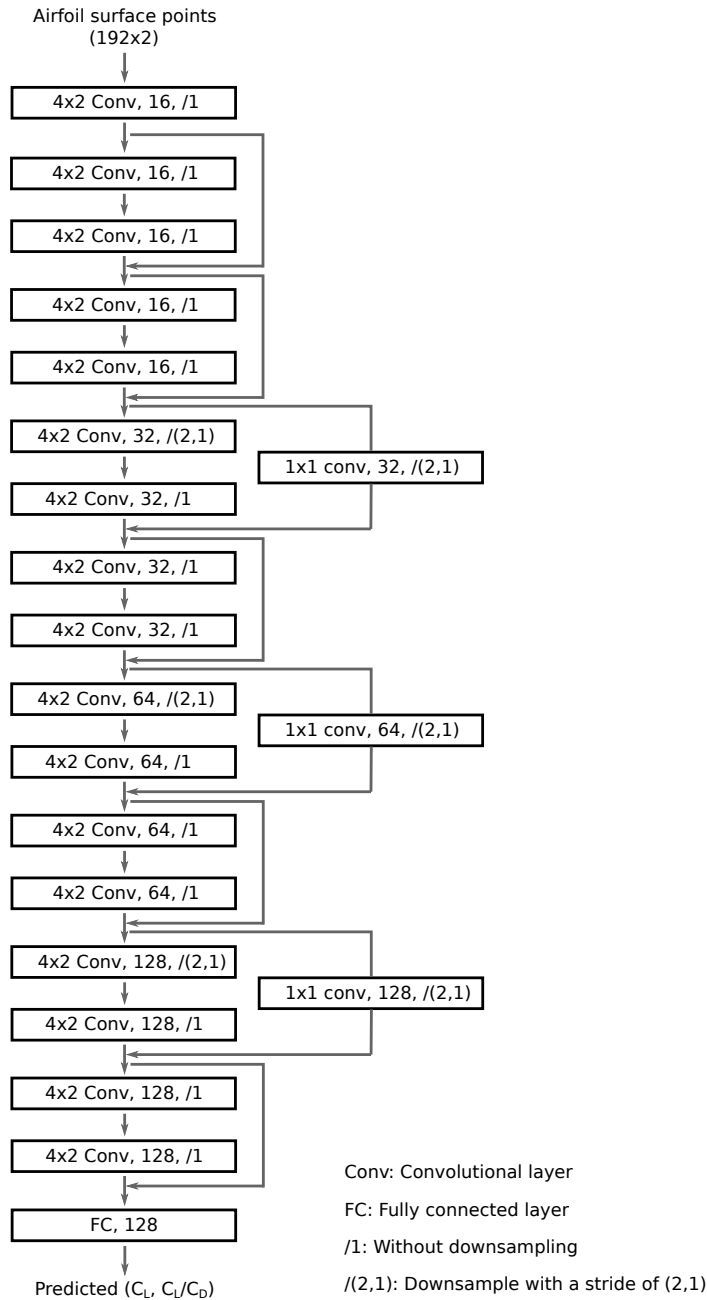


Figure 13: Architecture of the neural network surrogate model for airfoil performance prediction.

Appendix C Airfoil Design Example: BézierGAN for Smooth Airfoil Generation

Samples generated by conventional GANs can be noisy. This noisiness may not be an issue for images, but would cause airfoils to have non-smooth surfaces and greatly impair their performances. Therefore, we replace the conventional GAN with BézierGAN to guarantee the smoothness of generated airfoils. Instead of generating airfoil surface points freely in a 2D space, it uses a Bézier layer to constrain generated points to be on a rational Bézier curve. We follow [9] to setup the model architecture. We use the same training data as for the neural network surrogate model. For tests with or without the performance augmented DPP loss (*i.e.*, MO-PaDGAN or GAN), we train the model for 10,000 iterations with a batch size of 32. Both the generator and the discriminator use the Adam optimizer with a learning rate of 0.0001, $\beta_1 = 0.5$, and $\beta_2 = 0.999$. Please refer to our code for more details on BézierGAN’s network and training configurations.

Note that BézierGAN’s generator has both latent codes and the noise vector as inputs. While one can use both inputs as design parameters, we only use the latent codes as variable design parameters in our optimization experiments. How either input affects optimization performance is not central to this work (please refer to [9] for related details).

Appendix D Airfoil Design Example: Heuristics for Improving MO-PaDGAN Stability

While training a MO-PaDGAN, the performances are predicted by E and used in the performance augmented DPP loss to provide feedback for updating the generator. If the quality gradients are not accurate, the generator learning can go astray. This is not a problem when the quality estimator is a simulator that can reasonably evaluate (even with low-fidelity) any design in the design space, irrespective of the designs being invalid or unrealistic. However, it creates problems when we use a data-driven surrogate model, as in our airfoil design example. A data-driven surrogate model is normally trained only on realistic designs and hence may perform unreliably on unrealistic ones. In the initial stages of training, a GAN model will not always generate realistic designs. This makes it difficult for the surrogate model to correctly guide the generator’s update and may cause stability issues. We use the follow heuristics to improve the stability:

1. *Realisticly weighted quality.* We weight the predicted quality at \mathbf{x} by the probability of \mathbf{x} being the real design (predicted by the discriminator):

$$q(\mathbf{x}) = D(\mathbf{x})q'(\mathbf{x})$$

where $q'(\mathbf{x})$ is the predicted quality (by a surrogate model for example), and $D(\mathbf{x})$ is the discriminator’s output at \mathbf{x} .

2. An *escalating schedule* for setting γ_1 (the weight of the performance augmented DPP loss). A GAN is more likely to generate unrealistic designs in its early stage of training. Thus, we initialize γ_1 at 0 and increase it during training, so that MO-PaDGAN focuses on learning to generate realistic designs at the early stage, and takes quality into consideration later when the generator can produce more realistic designs. The schedule is set as:

$$\gamma_1 = \gamma'_1 \left(\frac{t}{T} \right)^p$$

where γ'_1 is the value of γ_1 at the end of training, t is the current training step, T is the total number of training steps, and p is a factor controlling the steepness of the escalation.

References

- [1] Kuang-Hua Chang. *e-Design: computer-aided engineering design*. Academic Press, 2016.
- [2] 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), 2019.
- [3] Alex Burnap, John R Hauser, and Artem Timoshenko. Design and evaluation of product aesthetics: A human-machine hybrid approach. *Available at SSRN 3421771*, 2019.
- [4] Dule Shu, James Cunningham, Gary Stump, Simon W Miller, Michael A Yukish, Timothy W Simpson, and Conrad S Tucker. 3d design using generative adversarial networks and physics-based validation. *Journal of Mechanical Design*, 142(7), 2020.
- [5] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [6] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.

- [7] Zijiang Yang, Xiaolin Li, L Catherine Brinson, Alok N Choudhary, Wei Chen, and Ankit Agrawal. Microstructural materials design via deep adversarial learning methodology. *Journal of Mechanical Design*, 140(11), 2018.
- [8] Wentai Zhang, Zhangsihao Yang, Haoliang Jiang, Suyash Nigam, Soji Yamakawa, Tomotake Furuhashi, Kenji Shimada, and Levent Burak Kara. 3d shape synthesis for conceptual design and optimization using variational autoencoders. In *ASME 2019 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*. American Society of Mechanical Engineers Digital Collection, 2019.
- [9] Wei Chen, Kevin Chiu, and Mark D Fuge. Airfoil design parameterization and optimization using bézier generative adversarial networks. *arXiv preprint arXiv:2006.12496*, 2020.
- [10] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *Advances in neural information processing systems*, pages 2234–2242, 2016.
- [11] Alex Kulesza and Ben Taskar. Determinantal point processes for machine learning. *arXiv preprint arXiv:1207.6083*, 2012.
- [12] Ivo Couckuyt, Dirk Deschrijver, and Tom Dhaene. Fast calculation of multiobjective probability of improvement and expected improvement criteria for pareto optimization. *Journal of Global Optimization*, 60(3):575–594, 2014.
- [13] Michael D McKay, Richard J Beckman, and William J Conover. A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, 42(1):55–61, 2000.
- [14] Eckart Zitzler and Lothar Thiele. Multiobjective evolutionary algorithms: a comparative case study and the strength pareto approach. *IEEE transactions on Evolutionary Computation*, 3(4):257–271, 1999.
- [15] Kyoungwoo Park and Juhee Lee. Optimal design of two-dimensional wings in ground effect using multi-objective genetic algorithm. *Ocean Engineering*, 37(10):902–912, 2010.
- [16] Mark Drela. Xfoil: An analysis and design system for low reynolds number airfoils. In *Low Reynolds number aerodynamics*, pages 1–12. Springer, 1989.
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [18] Wei Chen, Kevin Chiu, and Mark Fuge. Aerodynamic design optimization and shape exploration using generative adversarial networks. In *AIAA SciTech Forum*, San Diego, USA, Jan 2019. AIAA.
- [19] Wei Chen and Mark Fuge. Béziergan: Automatic generation of smooth curves from interpretable low-dimensional parameters. *arXiv preprint arXiv:1808.08871*, 2018.
- [20] Xi Chen, Yan Duan, Rein Houthoofd, John Schulman, Ilya Sutskever, and Pieter Abbeel. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In *Advances in neural information processing systems*, pages 2172–2180, 2016.
- [21] Daniel J Poole, Christian B Allen, and Thomas CS Rendall. Metric-based mathematical derivation of efficient airfoil design variables. *AIAA Journal*, 53(5):1349–1361, 2015.
- [22] Daniel J Poole, Christian B Allen, and T Rendall. Efficient aero-structural wing optimization using compact aerofoil decomposition. In *AIAA Scitech 2019 Forum*, page 1701, 2019.
- [23] Dominic A Masters, Nigel J Taylor, TCS Rendall, Christian B Allen, and Daniel J Poole. Geometric comparison of aerofoil shape parameterization methods. *AIAA Journal*, 55(5):1575–1589, 2017.
- [24] Joshua Knowles. Parego: a hybrid algorithm with on-line landscape approximation for expensive multiobjective optimization problems. *IEEE Transactions on Evolutionary Computation*, 10(1):50–66, 2006.
- [25] David A Van Veldhuizen and Gary B Lamont. Multiobjective evolutionary algorithm test suites. In *Proceedings of the 1999 ACM symposium on Applied computing*, pages 351–357, 1999.