

Synthesizing Designs with Inter-Part Dependencies using Hierarchical Generative Adversarial Networks

Wei Chen*

Dept. of Mechanical Engineering
University of Maryland
College Park, Maryland 20742
Email: wchen459@umd.edu

Mark Fuge

Dept. of Mechanical Engineering
University of Maryland
College Park, Maryland 20742
Email: fuge@umd.edu

Real-world designs usually consist of parts with inter-part dependencies, i.e., the geometry of one part is dependent on one or multiple other parts. We can represent such dependency in a part dependency graph. This paper presents a method for synthesizing these types of hierarchical designs using generative models learned from examples. It decomposes the problem of synthesizing the whole design into synthesizing each part separately but keeping the inter-part dependencies satisfied. Specifically, this method constructs multiple generative models, the interaction of which is based on the part dependency graph. We then use the trained generative models to synthesize or explore each part design separately via a low-dimensional latent representation, conditioned on the corresponding parent part(s). We verify our model on multiple design examples with different inter-part dependencies. We evaluate our model by analyzing the constraint satisfaction performance, the synthesis quality, the latent space quality, and the effects of part dependency depth and branching factor. This paper's techniques for capturing dependencies among parts lay the foundation for learned generative models to extend to more realistic engineering systems where such relationships are widespread.

1 Introduction

Representing a high-dimensional design space with a lower-dimensional *latent space* makes it easier to explore, visualize, or optimize complex designs. This often means finding a *latent representation*, or a *manifold*, along which valid designs, such as geometries, lie [1, 2].

While this works well for single parts, designs usually

have multiple parts with inter-part dependencies. For example, the size and position of a conduit, lightening, or alignment hole in an airframe structure depend on the shape of the airfoil. We can describe inter-part dependencies in a design by using a *directed acyclic graph* (DAG). This DAG captures whether the geometry of a part depends on the geometry of its parent part(s). In this case, one may want to identify first the *parent manifold* that captures major variation of parent shapes, and then the *child manifold* that captures major variation of feasible child shapes conditioned on any parent shape (Fig. 1). Because, for example, we may first optimize the airfoil shape on the airfoil manifold (parent) to obtain the optimal lift and drag; and then given the optimal airfoil, we may optimize the hole's size and position on the hole manifold (child) for other considerations like light-weighting, *etc.*

However, finding individual part manifolds that both represent the design space well, while also satisfying part configuration, is non-trivial. Traditionally, to learn the inter-part dependency, one has to either define explicit constraints [3, 4] or learn implicit constraints via adaptive sampling [2, 5]. The former uses hard-coded (often application-specific) constraints and hence lacks flexibility; whereas the latter queries external sources by human annotation, experiment, or simulation, and thus is expensive. In this paper, we solve these problems by instead learning these constraints given examples. We assume that we only have the prior knowledge on inter-part dependencies, but not the specific types of constraints (*e.g.*, concentric, alignment, tangent, *etc.*) that confine the geometry of each part. We do this by identifying different levels of manifolds, where the higher-level (parent) manifold imposes implicit constraints on the lower-level (child) manifolds. We propose a deep generative model that synthesizes designs in a hierarchical manner ac-

*Address all correspondence to this author.




Fig. 1: Manifolds of parent and child shapes

ording to those inter-part dependencies: it first synthesizes parent parts, and then synthesizes parts conditioned on those parent parts and so on. At each level, the model simultaneously captures a parent manifold and an infinite number of child manifolds which are conditioned on parent parts. This results in latent spaces that can synthesize and search each part individually as well as their assemblies. Importantly, our method is fully data-driven and requires no hard-coded rules or querying external sources, except for providing an example dataset of designs with part correspondence and known inter-part dependencies. For designs without part correspondence, we can apply unsupervised co-segmentation [6] as a pre-processing step.

The paper’s key contributions are as follows:

1. A novel deep generative model architecture that simultaneously learns a design’s inter-part dependencies and each part’s geometry variation conditioned on the corresponding parent part(s). It decouples each part’s latent space so that we can perform design space exploration in separate low-dimensional latent spaces.
2. New benchmark datasets—both real-world and synthetic—that can be used for studying different kinds of inter-part dependencies including: type of geometric constraints, depth of hierarchy, and branching factor of parent/child relationships. This dataset can aid in future evaluation of generative models of hierarchical parts.
3. Characterizing the effects of sample size and part dependencies’ complexity (depth and branching factor) on the synthesis performance of our generative model.
4. A new evaluation metric for generative models that measures the consistency of shape variation in the latent space.

2 Related Work

Our work produces generative models that synthesize designs from latent representations. There are primarily two streams of related research from the fields of engineering de-

sign and computer graphics—Specifically, design space dimensionality reduction and design synthesis. We also review generative adversarial networks (GANs) [7], which we use to build our model.

2.1 Design Space Dimensionality Reduction

While designs can be parametrized by various techniques [8], the number of design variables (*i.e.*, the dimensionality of a design space) increases with the geometric variability of designs. In tasks like design optimization, to find better designs we usually need a design space with higher variability, *i.e.*, higher dimensionality. This demand creates the problem of exploring a high-dimensional design space. Based on the curse of dimensionality [9], the cost of exploring the design space grows exponentially with its dimensionality. Thus, researchers have studied approaches for reducing the design space dimensionality. Normally, dimensionality reduction methods identify a lower-dimensional latent space that captures most of the design space’s variability. This can be grouped into linear and non-linear methods.

Linear dimensionality reduction methods select a set of optimal directions or basis functions where the variance of shape geometry or certain simulation output is maximized. Such methods include the Karhunen-Loève expansion (KLE) [10, 11], principal component analysis (PCA) [12], and the active subspaces approach [13].

In practice, it is more reasonable to assume that design variables lie on a non-linear manifold, rather than a hyperplane. Thus, researchers also apply non-linear methods to reduce the dimensionality of design spaces. This non-linearity can be achieved by (1) applying linear reduction techniques locally to construct a non-linear global manifold [14, 15, 16, 17, 12]; (2) using kernel methods with linear reduction techniques (*i.e.*, using linear methods in a Reproducing Kernel Hilbert Space that then induces non-linearity in the original design space) [1, 12]; (3) latent variable models like Gaussian process latent variable model (GPLVM) and generative topographic mapping (GTM) [18]; and 4) neural networks based approaches such as self-organizing maps [19] and autoencoders [20, 1, 12, 21].

This work differs from these past approaches in that we aim at identifying two-level latent spaces with the lower-level encodes inter-part dependencies, rather than learning only one latent space for the complete design.

2.2 Data-Driven Design Synthesis

Design synthesis methods can be divided into two categories: rule-based and data-driven design synthesis. The former (*e.g.*, grammars-based design synthesis [22, 4, 3]) requires labeling of the reference points or surfaces and defining rule sets, so that new designs are synthesized according to this hard-coded prior knowledge; while the latter learns rules/constraints from a database and generates plausible new designs with similar structure/function to exemplars in the database.

Usually dimensionality reduction techniques allow inverse transformations from the latent space back to the de-

sign space, thus can synthesize new designs from latent variables [11, 1, 12, 21]. For example, under the PCA model, the latent variables define a linear combination of principal components to synthesize a new design [11]; for local manifold based approaches, a new design can be synthesized via interpolation between neighboring points on the local manifold [17]; and under the autoencoder model, the trained decoder maps any given point in the latent space to a new design [20, 21]. Researchers have also employed generative models such as kernel density estimation [23], Boltzmann machines [24], variational autoencoders (VAEs) [25], and generative adversarial nets (GANs) [26, 27] to learn the distribution of samples in the design space, and synthesize new designs by drawing samples from the learned distribution. Discriminative models like deep residual networks [28] are also used to generate 3D shapes.

These aforementioned models synthesize a design or a shape as a whole. There are methods that synthesize new shapes by assembling or reorganizing parts from an existing shape database, while preserving the desired structures [29, 30, 31, 32, 33]. The shapes are usually parametrized by high-level abstract representations, such as hand-crafted feature vectors [30] or shape grammars [31]. While these methods edit shapes at a high-level, they do not control the local geometry of each synthesized component.

Previously the inter-part dependencies of shapes have been modeled by grammar induction [31], kernel density estimation [34], probabilistic graph models [29, 30, 24], and recursive autoencoders [27]. Those methods handle part relations and design synthesis separately. In contrast, our method encodes part relations through the model architecture, so that it simultaneously learns the inter-part dependencies and single part geometry variation. The model can also be used for inferring the generative distribution of each part conditioned on any parent part.

2.3 Generative Adversarial Networks

Generative adversarial nets [7] model a game between a generative model (*generator*) and a discriminative model (*discriminator*). The generative model 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 discriminative model tries to perform classification, *i.e.*, to distinguish between real and generated data (Fig. 2). The generator G and the discriminator 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, the objective of GAN is

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

where \mathbf{x} is sampled from the data distribution P_{data} , \mathbf{z} is sampled from the noise distribution P_z , and $G(\mathbf{z})$ is the generator distribution. A trained generator thus can map from the pre-defined noise distribution to the distribution of designs. The noise input \mathbf{z} is considered as the latent representation in the dimensionality reduction scenario, since \mathbf{z} captures the vari-

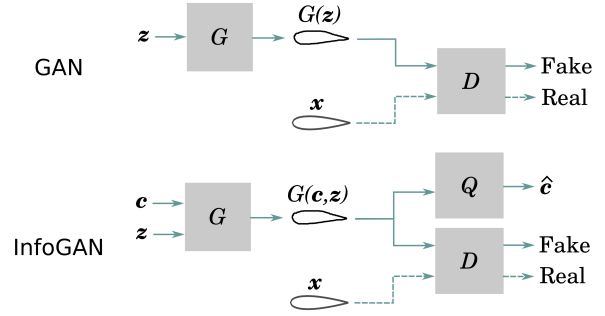


Fig. 2: Architectures of the standard GAN and the InfoGAN

ability of the data.

One advantage GANs hold over Variational Autoencoders (VAEs) [35] is that GANs tend to generate more realistic data [36]. But a disadvantage of the original GAN formulation is that it cannot learn an interpretable latent representation. Built upon these “vanilla” GANs, the InfoGAN [37] aims at regularizing the latent representation of the data space by maximizing a lower bound of the mutual information between a set of *latent codes* \mathbf{c} and the generated data. The generator is provided with both \mathbf{z} and \mathbf{c} . Thus, the generator distribution $P_G = G(\mathbf{c}, \mathbf{z})$ is conditioned on \mathbf{c} . The mutual information lower bound L_I is

$$L_I(G, Q) = \mathbb{E}_{\mathbf{c} \sim P(\mathbf{c}), \mathbf{x} \sim G(\mathbf{c}, \mathbf{z})} [\log Q(\mathbf{c}|\mathbf{x})] + H(\mathbf{c}) \quad (2)$$

where $H(\mathbf{c})$ is the entropy of the latent codes, and $Q(\mathbf{c}|\mathbf{x})$ is called the auxiliary distribution which approximates $P(\mathbf{c}|\mathbf{x})$. We direct interested readers to [37] for the derivation of L_I . The InfoGAN objective combines L_I with the standard GAN objective:

$$\min_{G, Q} \max_D V_{InfoGAN}(D, G, Q) = V(D, G) - \lambda L_I(G, Q) \quad (3)$$

where λ is a weight parameter.

In practice, $H(\mathbf{c})$ is treated as constant if the distribution of \mathbf{c} is fixed. The auxiliary distribution Q is parametrized by a neural network—here we call it the *auxiliary network*.

In our design synthesis scenario, the latent codes \mathbf{c} can represent any continuous or discrete factor that controls the geometry of the design, *e.g.*, the upper/lower surface protrusion of the airfoil.

Previously there are GAN-based models that generates specific types of data (*i.e.*, images and videos) by a two-level hierarchy [38, 39]. For example, a generator first generates the structure of an image, and then conditioned on that structure, another generator generates the texture of that image. In the case of design synthesis, it is intuitive to generate each design component given all its dependencies, which results in a hierarchical model structure with multiple levels of generators.

3 Method

In this section we introduce our proposed deep neural network architecture and its training details.

3.1 Problem Formulation

We can use a directed acyclic graph to define inter-part dependencies of a design. We call this graph a *part dependency graph*. For example, suppose we want to design an airfoil with two holes inside (top left in Fig. 3). We might first design the airfoil (Part A), and then set the position and diameter of one hole (Part B) based on the shape of the airfoil, followed by the design of the second hole (Part C) based on both the airfoil shape and the first hole. Thus, the dependencies can be expressed by the graph shown in the bottom left of Fig. 3.

This paper deals with a design synthesis and design space exploration problem, where we have a database of geometric designs, each of which has multiple parts. We have no prior knowledge on the specific types of constraints that confine the geometry of each part, but only inter-part dependencies. We propose a model that learns to synthesize designs based on those inter-part dependencies. We want to use this model to (1) correctly synthesize each part that follows both the shape feasibility and the dependency constraints; and (2) use a low-dimensional *latent spaces* C to represent the design spaces \mathcal{X} of each part. The ability to learn separate representations of each part is useful for decomposing a design space exploration problem (*e.g.*, in design optimization).

Desired latent spaces satisfy the following requirements:

1. Any child latent space should be conditioned on a parent shape (*e.g.*, for the design in Fig. 3, any latent space of the first hole should be conditioned on an airfoil).
2. Major variability across designs in the database should be captured by those latent spaces.
3. Designs should change consistently as we move along any basis of the latent space. This regularity/consistency/smoothness will improve latent space design exploration and optimization.

To meet the first requirement we construct a composite generative model—a model with multiple generators, each of which learns a (conditional) generative distribution of a part from the design. We ensure the rest of the requirements by adapting InfoGAN’s architecture and objective, *i.e.*, conditioning the generator on latent codes and maximizing a lower bound of the mutual information between the synthesized designs and their latent representations. This is known to (1) make latent codes target salient features of the designs, and (2) disentangle the latent representation of the data space [37]. We introduce the details of our model’s architecture in the following section.

3.2 Model Architecture

To learn separate distributions/representations of each part, we use a generative adversarial net with multiple generators/auxiliary networks. We call this network the Hierarchical Generative Adversarial Networks (HGAN). An example of its architecture is shown in Fig. 3. We use multiple generators to synthesize different parts in a design. The interaction between generators/auxiliary networks changes with the

part dependency graph. The latent code \mathbf{c}_j defines the latent representation of the j -th part, where $j = 1, \dots, n$, and n is the number of parts. The generator G_j learns a (conditional) shape distribution $P(\mathbf{x}_j | \mathbf{c}_j, \mathbf{x}_{\text{Par}(j)})$, where $\text{Par}(j)$ denotes the set of parent(s) of the j -th part. For the example in Fig. 3, $\text{Par}(A) = \emptyset$, $\text{Par}(B) = \{A\}$, and $\text{Par}(C) = \{A, B\}$. The output distribution $P_{G_j} = G_j(\mathbf{c}_j, \mathbf{z}_j, \hat{\mathbf{x}}_{\text{Par}(j)})$ represents the distribution of the synthesized part $\hat{\mathbf{x}}_j$, where $\hat{\mathbf{x}}_{\text{Par}(j)}$ denotes the synthesized parent part(s) of \mathbf{x}_j , and \mathbf{z}_j is the noise input to G_j .

The auxiliary network Q_j predicts the latent code distribution of the corresponding part, *i.e.*, to estimate the conditional distributions $P(\mathbf{c}_j | \mathbf{x}_j, \mathbf{x}_{\text{Par}(j)})$. The discriminator D predicts whether a full design is from the database or generated by generators. A properly trained D should distinguish designs with unrealistic or mismatched parts.

The objective of HGAN is expressed as

$$\min_{\{G_j\}, \{Q_j\}} \max_D V_{HGAN}(D, \{G_j\}, \{Q_j\}) = V(D, \{G_j\}) - \lambda L_I(\{G_j\}, \{Q_j\}) \quad (4)$$

where $\{G_j\}$ and $\{Q_j\}$ denotes the set of generators and auxiliary networks, respectively. The first term in Eqn. (4) denotes the standard GAN objective:

$$V(D, \{G_j\}) = \mathbb{E}_{\{\mathbf{x}_j\} \sim P_{data}} [\log D(\{\mathbf{x}_j\})] + \mathbb{E}_{\{\mathbf{c}_j \sim P(\mathbf{c}_j)\}, \{\mathbf{z}_j \sim P(\mathbf{z}_j)\}} [\log(1 - D(\{\hat{\mathbf{x}}_j\})] \quad (5)$$

where $\hat{\mathbf{x}}_j = G_j(\mathbf{c}_j, \mathbf{z}_j, \hat{\mathbf{x}}_{\text{Par}(j)})$. The second term in Eqn. (4) is the lower bound of the mutual information between the latent codes and the synthesized designs:

$$L_I(\{G_j\}, \{Q_j\}) = \sum_{j=1}^n \mathbb{E}_{\mathbf{c}_j \sim P(\mathbf{c}_j), \mathbf{x}_j \sim G_j(\mathbf{c}_j, \mathbf{z}_j, \hat{\mathbf{x}}_{\text{Par}(j)})} [\log Q_j(\mathbf{c}_j | \mathbf{x}_j, \hat{\mathbf{x}}_{\text{Par}(j)})] + H(\mathbf{c}_j) \quad (6)$$

4 Experiments

To demonstrate the performance of our model, we built six datasets with different ground-truth inter-part dependencies. We train the proposed network on these datasets, and evaluate the generative performance, constraint satisfaction, and the latent space property through qualitative and quantitative measures.

4.1 Experimental Setup

4.1.1 Dataset.

We created the following datasets, as shown in Fig. 4:

1. **AH**: An 2D airfoil (Part A) with a hole (circle, Part B) *inside*.
2. **AHH**: An 2D airfoil (Part A) with two *non-intersecting* holes (Parts B and C) *inside*. The centers of the two holes lie on a *horizontal* line.
3. **SE**: A superformula [40] (Part A) with a *concentric* ellipse (Part B) *inside*.

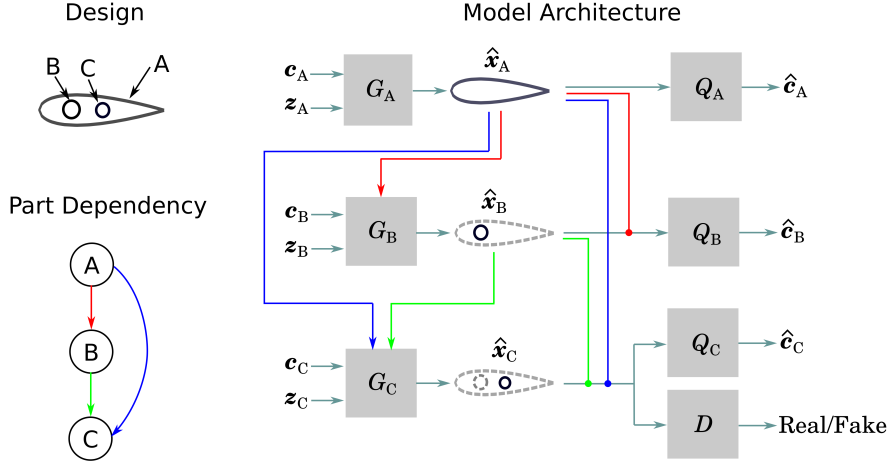


Fig. 3: An example of part dependency and the corresponding Hierarchical GAN architecture. The interaction between generators (G_A , G_B , and G_C) and auxiliary networks (Q_A , Q_B , and Q_C) is based on the connection of the part dependency graph. This figure is best viewed in color.




Fig. 4: Datasets with different inter-part dependencies

4. **SEoEi**: A superformula (Part A) with two ellipses *inside*. The second ellipse (Part C) is also *inside* the first one (Part B). All three shapes are *concentric*.
5. **SEiEo**: A superformula (Part A) with two ellipses *inside*. The second ellipse (Part C) is also *outside* the first one (Part B). All three shapes are *concentric*.¹
6. **SCC**: A superformula (Part A) with two *tangent* circles—one (Part B) inside and the other (Part C) outside.

In addition, we control the part dependency graph’s depth and branching factor by adding more circles/ellipses to SCC/SEoEi. This results in six extra datasets—**S**, **SC**, **SCCC**, **SCCCC**, **SEEE**, and **SEEEE**. Here the dataset of superformulas (**S**) is used as a baseline where no inter-part dependency is presented.

Specifically, the airfoil shapes are from the UIUC airfoil coordinates database², which provides the Cartesian coordinates for nearly 1,600 airfoils. Each airfoil is re-parametrized and represented with 64 Cartesian coordinates, resulting in a 64×2 matrix.

¹Note that the assembly of SEiEo is the same as SEoEi, but since we synthesize its parts in a different order, the inter-part dependency changes (see Fig. 4). We create this dataset for comparing different inter-part dependencies, rather than simulating practical use case.

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

Though targeted for real-world applications, the airfoils may not be a perfect experimental dataset to visualize the latent space, because the ground truth intrinsic dimension (ID)³ of the airfoil dataset is unknown. Thus, we create another synthetic dataset using superformulas and ellipses, the IDs of which are controllable. The superformula is a generalization of the ellipse [40]. We generate superformulas using the following equations:

$$\begin{aligned}
 n_1 &= 10s_1 \\
 n_2 &= n_3 = 10(s_1 + s_2) \\
 r(\theta) &= (|\cos \theta|^{n_2} + |\sin \theta|^{n_3})^{-\frac{1}{n_1}} \\
 (x, y) &= (r(\theta) \cos \theta, r(\theta) \sin \theta)
 \end{aligned} \tag{7}$$

where $s_1, s_2 \in [0, 1]$, and (x, y) is a Cartesian coordinate. For each superformula, we sample 64 evenly spaced θ from 0 to 2π , and get 64 grid-point Cartesian coordinates. Equations (7) show that we can control the deformation of the superformula shape with s_1 and s_2 . Thus, the ground truth ID of our superformula dataset is two. For further details on how to modify the superformula parameters to adjust the ID, complexity, and number of manifolds see [1].

The other shapes (*e.g.*, circles and ellipses) are also rep-

³The intrinsic dimension is the minimum number of variables required to represent the data. It indicates the degrees of freedom we have to control the shape and position.

resented using 64 Cartesian coordinates. The ground truth ID of ellipses in SE, SEoEi, and SEiEo is two, since we fix their centers and only change their semi-major axis and semi-minor axis lengths. The ground truth ID of circles in SCC is also two, since they can change their radii and move tangential to the superformulas. The Part B (circle) in AH and AHH has a ground truth ID of three, as both their centers and radii can change; while the Part C in AHH has a ground truth ID of two, since the y-coordinate of its center is fixed. Figure 5 shows samples drawn from each dataset.

For each dataset, we run experiments with sample sizes ranging from 500 to 10,000.

4.1.2 Network and Training.

We implement our airfoil/superformula generators by adopting the generator architecture of the BézierGAN [41, 42]. The circle/ellipse generators first generate shape parameters (*e.g.*, the center coordinates and the radius for a circle), and then convert them into grid-point coordinates using corresponding parametric functions. For a generator with parent parts as inputs, we use an encoder to convert each parent part into a feature vector before concatenating all inputs and feeding them into the generator. Similarly, we use an encoder to convert each generated part into a feature vector before feeding them into the auxiliary networks and the discriminator. The use of encoders reduces the model complexity by reducing the input dimension of the generators, the auxiliary networks and the discriminator. But one has to carefully choose the feature vector dimensions to avoid loss of information from dimensionality reduction.

At training, we sample the latent codes from uniform distribution $\mathcal{U}(0, 1)$, and the noise inputs from normal distribution $\mathcal{N}(0, 0.25)$. The hyper-parameter λ in Eqn. (4) was set to 0.1 in all experiments. The network was optimized using Adam [43] with the momentum terms $\beta_1 = 0.5$ and $\beta_2 = 0.999$. The learning rates were set to 0.0001. The total number of training steps was 100,000. The batch size was 32. Interested readers who wish to reproduce our exact architectures, hyper-parameters, and training procedures are directed to our code located on GitHub⁴. The training procedure is summarized in Algorithm 1.

We used TensorFlow [44] to build the networks. We trained our networks on a Nvidia Titan X GPU. For each experiment, the training process took around 2.2 hours for SE and AH, and 3 hours for other examples. The inference took less than 10 seconds.

5 Results and Discussion

We evaluated the performance of our trained generative models using both visual inspection (Figs. 6-10) and quantitative measures (Fig. 11). We analyze the effect of sample size and problem complexity on those quantitative performance metrics.

Algorithm 1 Train HGAN for designs with n parts

```

1:  $\triangleright T$ : number of training steps
2:  $\triangleright m$ : batch size
3:  $\triangleright X$ : training set
4: procedure TRAIN( $T, m$ )
5:   for  $t = 1 : T$  do
6:     Sample  $\{\mathbf{x}^1, \dots, \mathbf{x}^m\}$  from  $X$ , where  $\mathbf{x}^i =$ 
7:        $[\mathbf{x}_1^i, \dots, \mathbf{x}_n^i], i = 1, \dots, m$ 
8:     for  $j = 1 : n$  do
9:       Sample  $\{\mathbf{c}_j^1, \dots, \mathbf{c}_j^m\}$  from a uniform distribu-
10:      tion
11:       Sample  $\{\mathbf{z}_j^1, \dots, \mathbf{z}_j^m\}$  from a normal distribu-
12:      tion
13:     end for
14:      $\triangleright$  Based on Eqn. (4-6):
15:     Train  $D$  using  $\{\mathbf{x}^1, \dots, \mathbf{x}^m\}$ , fixing  $G_1, \dots, G_n$ 
16:      $\hat{\mathbf{x}}^i := [G_1(\mathbf{c}_1^i, \mathbf{z}_1^i, \hat{\mathbf{x}}_{\text{Par}(1)}^i), \dots, G_n(\mathbf{c}_n^i, \mathbf{z}_n^i, \hat{\mathbf{x}}_{\text{Par}(n)}^i)]$ 
17:     Train  $Q_1, \dots, Q_n$  and  $D$  using  $\{\hat{\mathbf{x}}^1, \dots, \hat{\mathbf{x}}^m\}$ , fixing
18:      $G_1, \dots, G_n$ 
19:     Train  $G_1, \dots, G_n$ , fixing  $Q_1, \dots, Q_n$  and  $D$ 
20:   end for
21: end procedure

```

5.1 Visual Inspection.

The captured latent spaces for different examples are visualized in Figs. 6-9. All these plots are generated using a sample size of 10,000. The results show that each child latent space adjusts itself according to its parent part, so that the sizes/positions of child parts matches their parent parts. This indicates that the child generator figures out the implicit constraints encoded in data. The latent spaces capture major shape variations and show consistent shape change. For example, in Fig. 6, the outer ellipse (middle subplot) has a consistently decreasing width from left to right, and increasing height from top to bottom. Interestingly, Fig. 8 shows that the circles (middle and right subplots) change in the latent space according to a polar coordinate system, instead of a Cartesian coordinate system like in other examples. For example, the inner circle’s radius decreases with the radial coordinate, and its position moves with the angular coordinate. This behavior is interesting because we did not explicitly encode this polar-coordinate representation into the HGAN architecture—rather, the HGAN automatically learns that such a representation is appropriate for this constraint. Figure 10 shows that when we removed the mutual information lower bound L_I in Eqn. (4), latent spaces either failed to capture major shape variation, or became entangled and inconsistent.

In the AHH example (Fig. 9), unfeasible synthesized designs occur when the two holes intersect. The figure shows that while the second hole moves from one side of the airfoil to the other side, it has to pass through a narrow unfeasible region. This unfeasible region cuts off the latent space, but the generator ignores this fact and learns a continuous latent space by interpolating designs inside the unfeasible region. The small volume of this unfeasible region may cause the discriminator to ignore it. In other words, the generator is

⁴https://github.com/IDEALLab/hgan_jmd_2019




Fig. 5: Samples drawn from datasets




Fig. 6: Latent space visualization of the SEoEi example. **Left:** synthesized superformulas in a 2-D latent space; **middle:** synthesized outer ellipses in a 2-D latent space conditioned on a random superformula; **right:** synthesized inner ellipses in a 2-D latent space conditioned on a random outer ellipse.

willing to take the minor loss incurred by this small region of the infeasible design space in order to avoid making the latent space representation more complicated. To solve this problem, we can perform adaptive sampling in the latent space to more accurately identify the feasible region(s) [2, 5].

5.2 Constraint Satisfaction.

We measure the precision of constraint satisfaction by computing the proportion of feasible designs among all the synthesized designs. We call this metric the Constraint Satisfaction Score (CSS). Specifically, different constraints define feasibility in different examples:

1. **AH:** Each point on the hole should be inside the airfoil.
2. **AHH:** (i) Each point on both holes should be inside the airfoil; (ii) the centers of the two holes should have a vertical distance of less than 0.01 (alignment constraint);⁵ and (iii) the distance of the two centers should be larger than the sum of the two radii (non-intersection constraint).

3. **SE:** (i) Each point on the ellipse should be inside the superformula; and (ii) The distance between the origin and the center of the ellipse should not exceed 0.01 (concentric constraint).⁶
4. **SEoEi:** (i) Each point on both ellipses should be inside the superformula; (ii) both the semi-major and semi-minor axis lengths of the first ellipse should be larger than the second one; and (iii) for both ellipses, the distance between their centers and the origin should not exceed 0.01 (concentric constraint).
5. **SEiEo:** Same as the SEoEi example, except that the first and the second ellipses are swapped.
6. **SCC:** For each circle with a center C_o and a radius r , the difference between r and the distance from C_o to the superformula should be less than 0.03 (tangent constraint).

The results on CSS in Fig. 11 show that the tangent constraint in the example SCC is the hardest to learn. But the learning performance improves with a larger sample size. The SE example outperforms SEoEi and SEiEo on CSS,

⁵All designs are rescaled such that the airfoils and the superformulas have unit widths.

⁶For all the examples we assume that the superformula is centered at the origin.




Fig. 7: Latent space visualization of the SEiEo example. **Left:** synthesized superformulas in a 2-D latent space; **middle:** synthesized inner ellipses in a 2-D latent space conditioned on a random superformula; **right:** synthesized outer ellipses in a 2-D latent space conditioned on that same superformula and a random outer ellipse.




Fig. 8: Latent space visualization of the SCC example. **Left:** synthesized superformulas in a 2-D latent space; **middle:** synthesized inner circles in a 2-D latent space conditioned on a random superformula; **right:** synthesized outer circles in a 2-D latent space conditioned on that same superformula. Interestingly, the HGAN automatically learns a polar-coordinate representation for the tangent constraint.

which is expected since the task of learning constraints in SE can be considered as a sub-task in SEoEi and SEiEo. This also applies to AH and AHH. It is also expected that SEoEi outperforms SEiEo, since the two are dealing with the same design but the former has fewer dependencies.

5.3 Distance between Data and Generator Distributions.

We measure how well our generator approximates the real data distribution by computing the kernel maximum mean discrepancy (MMD) [45] between the data and the generator distribution:

$$\text{MMD}^2(P_{data}, P_G) = \mathbb{E}_{\mathbf{x}_d, \mathbf{x}'_d \sim P_{data}; \mathbf{x}_g, \mathbf{x}'_g \sim P_G} [k(\mathbf{x}_d, \mathbf{x}'_d) - 2k(\mathbf{x}_d, \mathbf{x}_g) + k(\mathbf{x}_g, \mathbf{x}'_g)]$$

where $k(\mathbf{x}, \mathbf{x}') = \exp(-\|\mathbf{x} - \mathbf{x}'\|^2 / (2\sigma^2))$. A lower kernel MMD indicates that the generator distribution is closer to the data distribution.

The top middle plot in Fig. 11 shows the results of MMD. Here the SEoEi and SEiEo examples have similar MMD values, which suggests that synthesized designs in both examples have similar perceptual quality, despite the fact that SEiEo requires more model parameters to learn the additional dependency. The MMD plot provides insight into how realistic the generated designs can get as the training sample size changes. In general, MMD first decreases steeply with the sample size, and then reaches a plateau at some point. This point indicates the smallest sample size required to reach a “perceptually good” synthesis performance.

5.4 Diversity of Generated Designs.

A common problem in GANs’ training is mode collapse, during which the generator only generates a few types of designs to fool the discriminator instead of properly learning the complete data distribution. Therefore it is important to measure the diversity of the synthesized designs. We use

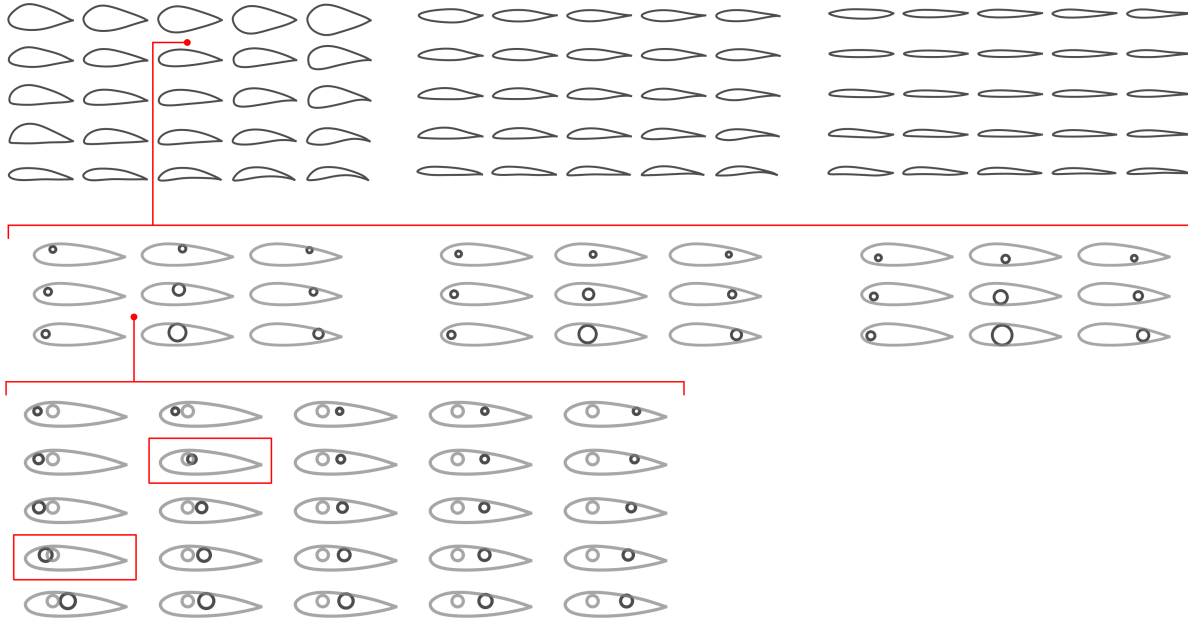


Fig. 9: Latent space visualization of the AHH example. **Top**: Synthesized airfoils in a 3-D latent space (visualized by multiple slices of 2-D latent spaces); **middle**: Synthesized holes in a 3-D latent space conditioned on a random airfoil; **Bottom**: Synthesized holes in a 2-D latent space conditioned on that same airfoil and another random hole. Unfeasible synthesized designs occur when the two holes intersect (indicated by red boxes).




Fig. 10: Latent spaces learned without maximizing the mutual information lower bound. **Left**: the latent code \mathbf{c} failed to capture major shape variation; **right**: inconsistent shape variation along latent space bases (LSC= 0.562 ± 0.008).

Relative Diversity (R-Div) to measure the relative level of variability captured by the generator.

We define the diversity of N samples $\mathbf{X} \in \mathbb{R}^{d \times N}$ as⁷

$$\text{Div}(\mathbf{X}) = \frac{1}{N} \text{trace}(\text{cov}[\mathbf{X}, \mathbf{X}])$$

The R-Div metric can then be expressed as

$$\text{R-Div} = \frac{\text{Div}(\mathbf{X}_g)}{\text{Div}(\mathbf{X}_{data})}$$

where \mathbf{X}_g and \mathbf{X}_{data} denote the set of synthesized designs and designs from the dataset, respectively. A R-Div close

⁷Here each sample is represented as a column vector containing all the coordinates of the design.

to 0.0 means that there is little variation within the synthesized designs, which could be an indicator of mode collapse. A R-Div around 1.0 indicates that the synthesized designs have a similar level of variability with the dataset. Note that high diversity does not always indicate good performance, as there could be unrealistic designs being synthesized, which also contribute to diversity. Thus, we should view this metric in concert with the kernel MMD to determine how well the generator performs.

Figure 11 shows that synthesized designs tend to have high divergence of R-Div when the sample size is small. Particularly, in the AHH example with a sample size of 500, the low R-Div combined with the high kernel MMD indicates the occurrence of mode collapse (Fig. 12). Increasing the sample size stabilized the R-Div and eventually bounded it between 0.8 and 1.1.

5.5 Latent Space Consistency.

A desirable latent space has two properties: 1) *disentanglement*: each latent variable is related to only one factor; and 2) *consistency*: shapes change consistently along any basis of the latent space. Note that this consistency is evaluated along one direction at a time, since scales along different directions may vary. To the best of our knowledge, existing quantitative measurements for the first property—latent space disentanglement—are supervised, *i.e.*, the ground-truth independent factors causing shape deformation have to be provided [46, 47, 48]. The second property is important for latent space design exploration. When searching for designs along a direction in the latent space, optimiza-

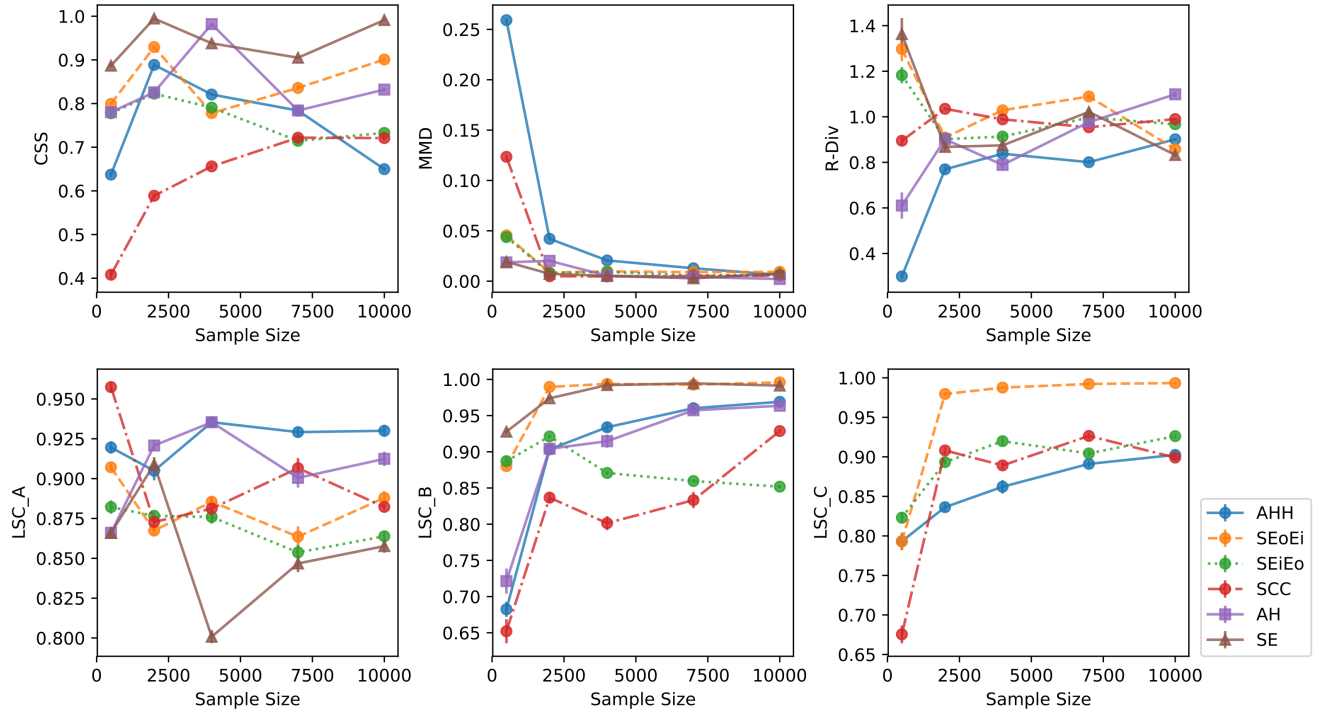


Fig. 11: Quantitative measure of synthesized design quality and latent space properties. LSC_A, LSC_B, LSC_C denote the LSC for the latent spaces of Parts A, B, C, respectively.




Fig. 12: Mode collapse occurs when the sample size is 500. The diversity is low in these randomly generated designs.

tion algorithms and humans usually prefer if shapes change consistently, such that the objective function over the latent space is less complicated (*i.e.*, has better Lipschitz continuity) and has fewer local optima.

We propose Latent Space Consistency (LSC) as a quantitative measure of how consistently shapes change along any basis of the latent space. Since change from one shape to another can be measured by their dissimilarity, distances between samples along a certain direction in the latent space should be consistent with the dissimilarity between those samples. We use Pearson correlation coefficient to measure this consistency. Algorithm 2 describes how to compute the LSC. The choice of the dissimilarity function d is not central to the overall method. In our experiments, we simply use the Euclidean distance to measure the dissimilarity of designs.

Both the middle plots in Figs. 6 and 7 show latent spaces with LSCs above 0.9. In contrast, the right plot in Fig. 10 provides a visual example of an LSC of around 0.56. The bottom plots in Fig. 11 shows that a larger sample size does

Algorithm 2 Evaluate Latent Space Consistency

```

1: procedure LATENTCONSISTENCY( $G, m, n, d$ )
2:    $\triangleright G$ : the mapping from a latent space to a design space
3:    $\triangleright m$ : the number of lines to be evaluated
4:    $\triangleright n$ : the number of points sampled on each line
5:    $\triangleright d$ : a dissimilarity function
6:    $\triangleright C$ : the latent space
7:    $\triangleright X$ : the design space
8:    $sum = 0$ 
9:   for  $i = 1 : m$  do
10:    Sample a line  $\mathcal{L}$  parallel to any basis of  $C$ 
11:    Sample  $n$  points  $\{\mathbf{c}^1, \mathbf{c}^2, \dots, \mathbf{c}^n\}$  along  $\mathcal{L}$ 
12:     $\{\mathbf{x}^1, \mathbf{x}^2, \dots, \mathbf{x}^n\} := \{G(\mathbf{c}^1), G(\mathbf{c}^2), \dots, G(\mathbf{c}^n)\}$ 
13:     $D_C := \{\|\mathbf{c}^i - \mathbf{c}^j\|\}$ ,  $D_X := \{d(\mathbf{x}^i, \mathbf{x}^j)\}$ , where
       $i, j \in \{1, \dots, n\}$ 
14:    Compute Pearson correlation coefficient:
      
$$\rho := \frac{\text{cov}(D_C, D_X)}{\sigma(D_C)\sigma(D_X)}$$

15:     $sum := sum + \rho$ 
16:   end for
17:   Return  $sum/m$ 
18: end procedure

```

not improve LSCs of Part A (at least with sample sizes in the range from 500 to 10,000), but improves the LSCs of Parts B and C in most cases.

5.6 Effect of Encoding Inter-Part Dependencies.

To further study the effect of encoding inter-part dependencies in the GAN, we compared the results of our HGAN with a standard InfoGAN, where there is only one generator and all parts are synthesized from a single latent space. The latent space dimension was set to the sum of all latent space dimensions in the HGAN experiments. Figure 13 visualizes the latent space learned by a standard InfoGAN. It shows that in each latent dimension, three parts change simultaneously, which indicates that the latent space does not disentangle each part’s shape variation.

We also study the effects of part dependency graph’s depth and branching factor on the model’s synthesis performance. The examples S, SE, SEE, SEEE, and SEEEE simulate increasing depths, respectively; while the examples S, SC, SCC, SCCC, and SCCCC simulate increasing branch factors, respectively. The results are shown in Fig. 14. The InfoGAN model is used as a base line where we do not encode inter-part dependencies. As expected for both HGAN and InfoGAN, CSS decreases with increased number of parts, since additional parts brings extra constraints. Theoretically, encoding inter-part dependencies introduces extra constraints to the model and hence complicates the overall task. However, based on the results, the HGAN shows no significant performance drop comparing to the InfoGAN, except for the CSS when increasing the branching factor and the R-Div when increasing the depth. Note that the MMD values are below 0.05 and do not change notably. The lower CSS or R-Div indicates that HGAN compromises its synthesis precision (*i.e.*, the precision of satisfying constraints) or generator distribution’s coverage for disentangling each part’s latent space. We also included the training history of HGAN and InfoGAN in the online supplemental material.

6 Limitations and Future Work

One limitation of our approach is that it is difficult to achieve high precision in satisfying some constraints. This problem exists in all purely data-driven design methods, where there is no process of incorporating restricted constraints in the generative model or validating the constraint satisfaction of the outputs. While we can address this by encoding constraints explicitly in the generative model, this requires us to know these constraints in advance and create hand-coded rules for all types of constraints. For example, to generate concentric ellipses, one can simply fix the center of the second ellipse to have the same coordinate as the first one. However, it may be difficult to incorporate some constraints in the generative model (*e.g.*, the constraint of one part being inside another). Also, it is sometimes hard to even know the type of the exact constraint between parts (*e.g.*, aesthetic preferences when placing handles on a vase). Thus, here we assume that we have no prior knowledge on the types of constraints, and inter-part dependencies are the only knowledge we need for our model. Despite the limitation of purely data-driven design methods, they can be used in the conceptual design stage for exploring a wide range of design alternatives and inspiring novel designs. We can also

use validation based on simulation, experiment, or human annotation to exclude infeasible synthesized designs when performing latent space exploration [2, 49], which could be an interesting avenue for future work.

Another limitation is that all designs must have the same part dependency graph, which is impractical in some cases. For example, not all tables have four legs, thus for some designs their part dependency graphs might miss some nodes. Future study needs to address this situation.

Sometimes the design data is not partitioned into separate components. One possible solution to this problem is to apply unsupervised co-segmentation [6] to partition designs and establish correspondences between common components in different designs. Then we can use HGAN to learn the latent spaces and generate new designs.

7 Conclusion

We introduced a GAN-based generative model for synthesizing designs with inter-part dependencies. It decomposes the synthesis into synthesizing each part conditioned on the corresponding parent part. This also creates a conditioned low-dimensional latent representation that allows accelerated design exploration. This model is build for problems where design space exploration over the latent space has to be staged since the optimal solution of one part depends on the geometry of another. Such models can accelerate design optimization problems, which we are exploring in our future work.

An advantage of neural-network-based generative models (*e.g.*, GANs and VAEs), compared to other dimensionality reduction models (*e.g.*, PCA and GPLVM), is that one can define or regularize latent distributions. Our model adapts InfoGAN’s mutual information objective to derive a consistent latent space, where the change of shapes is consistent along any basis of the latent space. This property is desirable in latent space design exploration, as the objective function over the latent space is less complicated and has less local optima.

We also created new benchmark datasets for studying different kinds of inter-part dependencies including: type of geometric constraints, depth of hierarchy, and branching factor of parent/child relationships. By using these datasets, we characterized the effects of sample size and part dependencies’ complexity (depth and branching factor) on the synthesis performance of our generative model. We also proposed a new evaluation metric for generative models that measures the consistency of shape variation in the latent space. Compared to a standard InfoGAN, the HGAN disentangles each part’s latent space at the cost of weakened synthesis precision when the branching of the part dependency increases.

The concept of decoupling the latent space of a design is important for design space exploration in general. It allows separate exploration and synthesis of each part, and helps us understand how different constraints control shape variation. Though we use GANs to achieve this goal, the idea of encoding inter-part dependencies, learning conditioned generative distributions, and the latent space analysis is also applicable to other generative models. Thus, this paper’s techniques




Fig. 13: The latent space learned by a standard InfoGAN. This plot visualizes a 6-D latent space of the S E o E i example by showing each two dimensions while setting the latent coordinates of other dimensions to zero. In each dimension, three parts change simultaneously. This shows that the latent space learned by a standard InfoGAN does not disentangle each part's shape variation. In contrast, the results of the HGAN disentangle and separate each part's latent space.

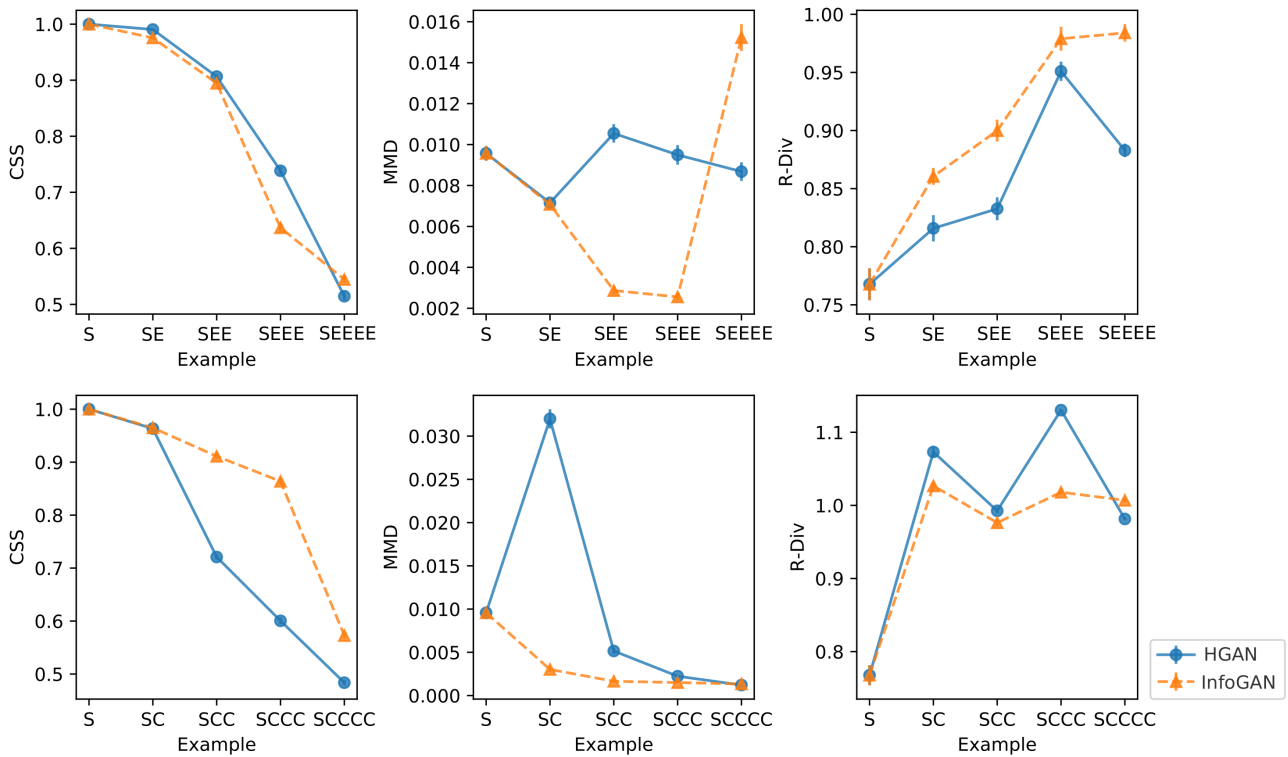


Fig. 14: Effects of depth (top) and branching factor (bottom) on synthesis performance. Here we denote the example S E o E i as SEE for simplicity.

lay the foundation for learned generative models to extend to more realistic engineering systems where inter-part dependencies are widespread.

Acknowledgements

This work was supported by The Defense Advanced Research Projects Agency (DARPA-16-63-YFA-FP-059) via

the Young Faculty Award (YFA) Program. The views, opinions, and/or findings contained in this article are those of the author and should not be interpreted as representing the official views or policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the Department of Defense.

References

- [1] Chen, W., Fuge, M., and Chazan, N., 2017. “Design manifolds capture the intrinsic complexity and dimension of design spaces”. *Journal of Mechanical Design*, **139**(5), pp. 051102–051102–10.
- [2] Chen, W., and Fuge, M., 2017. “Beyond the known: Detecting novel feasible domains over an unbounded design space”. *Journal of Mechanical Design*, **139**(11), pp. 111405–111405–10.
- [3] Königseder, C., and Shea, K., 2016. “Visualizing relations between grammar rules, objectives, and search space exploration in grammar-based computational design synthesis”. *Journal of Mechanical Design*, **138**(10), p. 101101.
- [4] Königseder, C., Stanković, T., and Shea, K., 2016. “Improving design grammar development and application through network-based analysis of transition graphs”. *Design Science*, **2**, p. e5.
- [5] Chen, W., and Fuge, M., 2018. “Active expansion sampling for learning feasible domains in an unbounded input space”. *Structural and Multidisciplinary Optimization*, **57**(3), pp. 925–945.
- [6] Sidi, O., van Kaick, O., Kleiman, Y., Zhang, H., and Cohen-Or, D., 2011. “Unsupervised co-segmentation of a set of shapes via descriptor-space spectral clustering”. *ACM Trans. Graph.*, **30**(6), Dec., pp. 126:1–126:10.
- [7] Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y., 2014. “Generative adversarial nets”. In *Advances in neural information processing systems*, Curran Associates, Inc., pp. 2672–2680.
- [8] Samareh, J. A., 2001. “Survey of shape parameterization techniques for high-fidelity multidisciplinary shape optimization”. *AIAA journal*, **39**(5), pp. 877–884.
- [9] Bellman, R., 1957. *Dynamic programming*. Princeton University Press, Princeton, NY.
- [10] Diez, M., Campana, E. F., and Stern, F., 2015. “Design-space dimensionality reduction in shape optimization by karhunen–loève expansion”. *Computer Methods in Applied Mechanics and Engineering*, **283**, pp. 1525–1544.
- [11] Chen, X., Diez, M., Kandasamy, M., Zhang, Z., Campana, E. F., and Stern, F., 2015. “High-fidelity global optimization of shape design by dimensionality reduction, metamodels and deterministic particle swarm”. *Engineering Optimization*, **47**(4), pp. 473–494.
- [12] D’Agostino, D., Serani, A., Campana, E. F., and Diez, M., 2018. “Nonlinear methods for design-space dimensionality reduction in shape optimization”. In *Machine Learning, Optimization, and Big Data*, Springer International Publishing, pp. 121–132.
- [13] Tezzele, M., Salmoiraghi, F., Mola, A., and Rozza, G., 2018. “Dimension reduction in heterogeneous parametric spaces with application to naval engineering shape design problems”. *Advanced Modeling and Simulation in Engineering Sciences*, **5**(1), p. 25.
- [14] Raghavan, B., Breitkopf, P., Tourbier, Y., and Villon, P., 2013. “Towards a space reduction approach for efficient structural shape optimization”. *Structural and Multidisciplinary Optimization*, **48**(5), pp. 987–1000.
- [15] Raghavan, B., Xia, L., Breitkopf, P., Rassineux, A., and Villon, P., 2013. “Towards simultaneous reduction of both input and output spaces for interactive simulation-based structural design”. *Computer Methods in Applied Mechanics and Engineering*, **265**, pp. 174–185.
- [16] Raghavan, B., Le Quilliec, G., Breitkopf, P., Rassineux, A., Roelandt, J.-M., and Villon, P., 2014. “Numerical assessment of springback for the deep drawing process by level set interpolation using shape manifolds”. *International Journal of Material Forming*, **7**(4), pp. 487–501.
- [17] Le Quilliec, G., Raghavan, B., and Breitkopf, P., 2015. “A manifold learning-based reduced order model for springback shape characterization and optimization in sheet metal forming”. *Computer Methods in Applied Mechanics and Engineering*, **285**, pp. 621–638.
- [18] Viswanath, A., J. Forrester, A., and Keane, A., 2011. “Dimension reduction for aerodynamic design optimization”. *AIAA journal*, **49**(6), pp. 1256–1266.
- [19] Qiu, H., Xu, Y., Gao, L., Li, X., and Chi, L., 2016. “Multi-stage design space reduction and metamodeling optimization method based on self-organizing maps and fuzzy clustering”. *Expert Systems with Applications*, **46**, pp. 180–195.
- [20] Burnap, A., Liu, Y., Pan, Y., Lee, H., Gonzalez, R., and Papalambros, P. Y., 2016. “Estimating and exploring the product form design space using deep generative models”. In *ASME 2016 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, American Society of Mechanical Engineers, pp. V02AT03A013–V02AT03A013.
- [21] D’Agostino, D., Serani, A., Campana, E. F., and Diez, M., 2018. “Deep autoencoder for off-line design-space dimensionality reduction in shape optimization”. In *2018 AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, p. 1648.
- [22] Gmeiner, T., and Shea, K., 2013. “A spatial grammar for the computational design synthesis of vise jaws”. In *ASME 2013 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, American Society of Mechanical Engineers, pp. V03AT03A006–V03AT03A006.
- [23] Talton, J. O., Gibson, D., Yang, L., Hanrahan, P., and Koltun, V., 2009. “Exploratory modeling with collaborative design spaces”. *ACM Trans. Graph.*, **28**(5), Dec., pp. 167:1–167:10.
- [24] Huang, H., Kalogerakis, E., and Marlin, B., 2015. “Analysis and synthesis of 3d shape families via deep-learned generative models of surfaces”. *Computer Graphics Forum*, **34**(5), pp. 25–38.
- [25] Nash, C., and Williams, C. K., 2017. “The shape variational autoencoder: A deep generative model of part-segmented 3d objects”. *Computer Graphics Forum*,

- 36(5), pp. 1–12.
- [26] Wu, J., Zhang, C., Xue, T., Freeman, B., and Tenenbaum, J., 2016. “Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling”. In *Advances in Neural Information Processing Systems 29*, Curran Associates, Inc., pp. 82–90.
- [27] Li, J., Xu, K., Chaudhuri, S., Yumer, E., Zhang, H., and Guibas, L., 2017. “Grass: Generative recursive autoencoders for shape structures”. *ACM Trans. Graph.*, **36**(4), July, pp. 52:1–52:14.
- [28] Sinha, A., Unmesh, A., Huang, Q., and Ramani, K., 2017. “Surfnet: Generating 3d shape surfaces using deep residual networks”. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 6040–6049.
- [29] Chaudhuri, S., Kalogerakis, E., Guibas, L., and Koltun, V., 2011. “Probabilistic reasoning for assembly-based 3d modeling”. *ACM Trans. Graph.*, **30**(4), July, pp. 35:1–35:10.
- [30] Kalogerakis, E., Chaudhuri, S., Koller, D., and Koltun, V., 2012. “A probabilistic model for component-based shape synthesis”. *ACM Trans. Graph.*, **31**(4), July, pp. 55:1–55:11.
- [31] Talton, J., Yang, L., Kumar, R., Lim, M., Goodman, N., and Měch, R., 2012. “Learning design patterns with bayesian grammar induction”. In *Proceedings of the 25th Annual ACM Symposium on User Interface Software and Technology, UIST '12*, ACM, pp. 63–74.
- [32] Xu, K., Zhang, H., Cohen-Or, D., and Chen, B., 2012. “Fit and diverse: set evolution for inspiring 3d shape galleries”. *ACM Trans. Graph.*, **31**(4), July, pp. 57:1–57:10.
- [33] Zheng, Y., Cohen-Or, D., and Mitra, N. J., 2013. “Smart variations: Functional substructures for part compatibility”. *Computer Graphics Forum*, **32**(2pt2), pp. 195–204.
- [34] Fish, N., Averkiou, M., Van Kaick, O., Sorkine-Hornung, O., Cohen-Or, D., and Mitra, N. J., 2014. “Meta-representation of shape families”. *ACM Trans. Graph.*, **33**(4), July, pp. 34:1–34:11.
- [35] Kingma, D. P., and Welling, M., 2014. “Auto-encoding variational bayes”. In *Proceedings of the 2nd International Conference on Learning Representations*.
- [36] Radford, A., Metz, L., and Chintala, S., 2015. “Unsupervised representation learning with deep convolutional generative adversarial networks”. *arXiv preprint arXiv:1511.06434*.
- [37] Chen, X., Duan, Y., Houthoofd, R., Schulman, J., Sutskever, I., and Abbeel, P., 2016. “Infogan: Interpretable representation learning by information maximizing generative adversarial nets”. In *Advances in Neural Information Processing Systems*, pp. 2172–2180.
- [38] Wang, X., and Gupta, A., 2016. “Generative image modeling using style and structure adversarial networks”. In *European Conference on Computer Vision*, Springer, pp. 318–335.
- [39] Ohnishi, K., Yamamoto, S., Ushiku, Y., and Harada, T., 2018. “Hierarchical video generation from orthogonal information: Optical flow and texture”. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- [40] Gielis, J., 2003. “A generic geometric transformation that unifies a wide range of natural and abstract shapes”. *American journal of botany*, **90**(3), pp. 333–338.
- [41] Chen, W., and Fuge, M., 2018. “Béziérgan: Automatic generation of smooth curves from interpretable low-dimensional parameters”. *arXiv preprint arXiv:1808.08871*.
- [42] Chen, W., Chiu, K., and Fuge, M., 2019. “Aerodynamic design optimization and shape exploration using generative adversarial networks”. In *AIAA 2019 Science and Technology Forum and Exposition*, no. AIAA 2019-2351.
- [43] Kingma, D. P., and Ba, J., 2014. “Adam: A method for stochastic optimization”. *arXiv preprint arXiv:1412.6980*.
- [44] Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., and Zheng, X., 2015. TensorFlow: Large-scale machine learning on heterogeneous systems. Software available from tensorflow.org.
- [45] Gretton, A., Borgwardt, K. M., Rasch, M. J., Schölkopf, B., and Smola, A., 2012. “A kernel two-sample test”. *Journal of Machine Learning Research*, **13**(Mar), pp. 723–773.
- [46] Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., and Lerchner, A., 2017. “Beta-VAE: Learning basic visual concepts with a constrained variational framework”. In *Proceedings of the 5th International Conference on Learning Representations*.
- [47] Kim, H., and Mnih, A., 2018. “Disentangling by factorising”. In *Proceedings of the 35th International Conference on Machine Learning*, Vol. 80 of *Proceedings of Machine Learning Research*, PMLR, pp. 2649–2658.
- [48] Chen, T. Q., Li, X., Grosse, R. B., and Duvenaud, D. K., 2018. “Isolating sources of disentanglement in variational autoencoders”. In *Advances in Neural Information Processing Systems 31*, Curran Associates, Inc., pp. 2610–2620.
- [49] Guo, T., Herber, D. R., and Allison, J. T., 2019. “Circuit synthesis using generative adversarial networks (GANs)”. In *AIAA 2019 Science and Technology Forum and Exposition*, no. AIAA 2019-2350.

Supplemental Material: Synthesizing Designs with Inter-Part Dependencies using Hierarchical Generative Adversarial Networks

S1 Training History

Theoretically, modeling inter-part dependency requires extra model parameters (*e.g.*, in the SEoEi example, the HGAN model has 86,213,654 parameters; while the InfoGAN has 33,270,374), and hence should be harder to optimize. This effect will amplify as more inter-part dependencies are introduced. In practice, to demonstrate the difference, we compared the training history of the HGAN and the InfoGAN, as shown in Fig. S1. It shows that there is no significant difference between the training histories of HGAN and InfoGAN except for the example of SCCCC, where it is more difficult for the HGAN's generator to reduce its training loss comparing to InfoGAN. This is an interesting observation and is worth future study to identify specific cause of difference in the training difficulty.

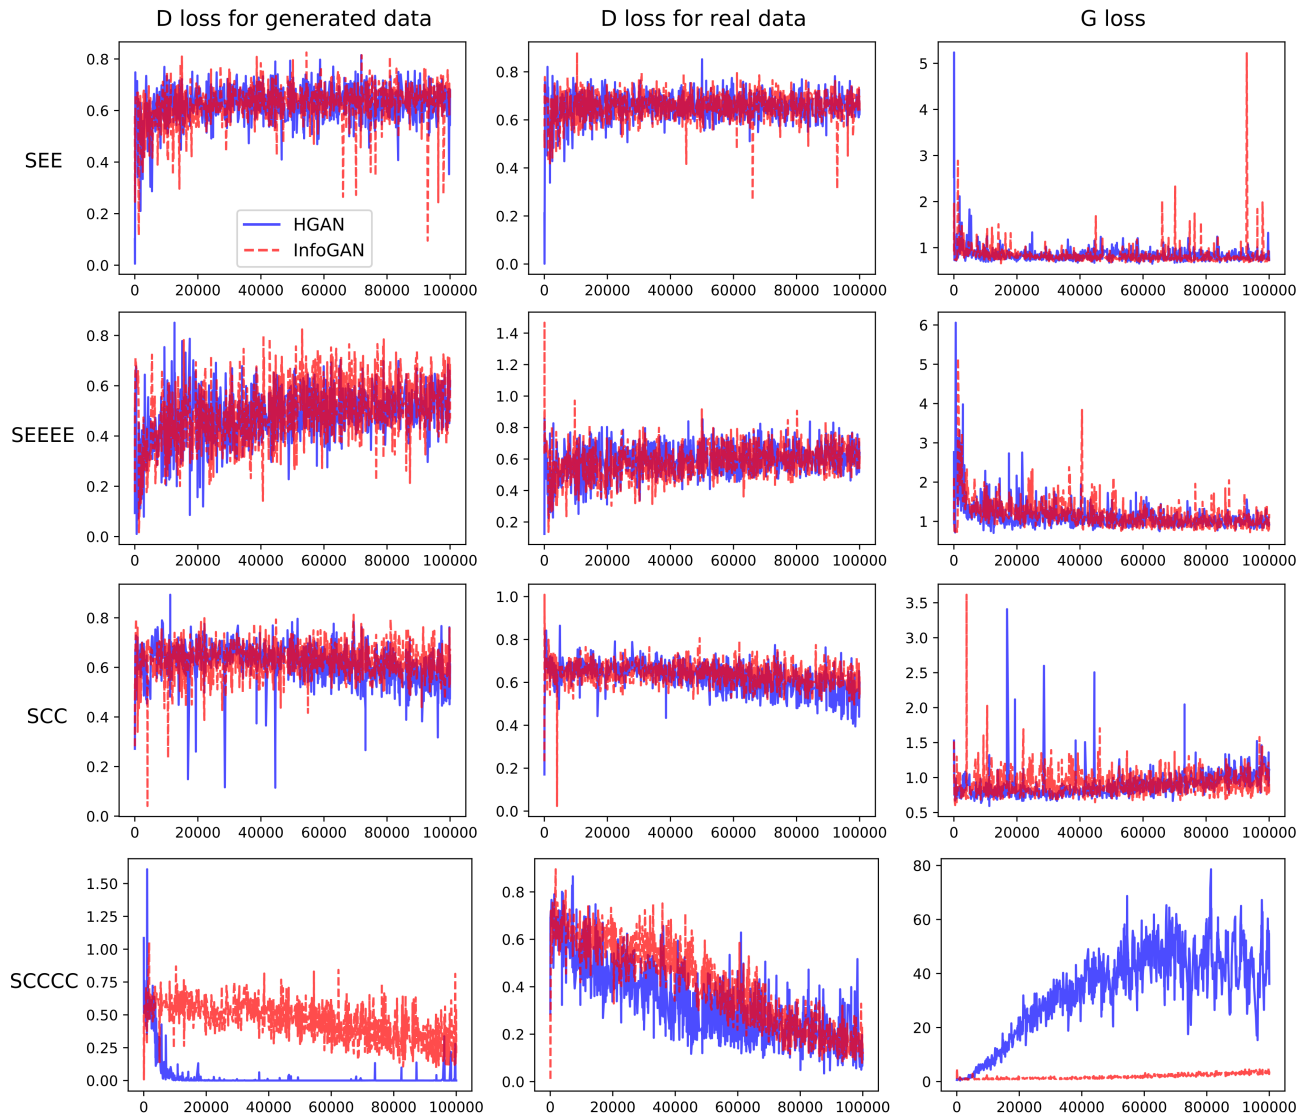


Fig. S1: Training history of HGAN and InfoGAN