Proceedings of the ASME 2016 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference
IDETC/CIE 2016
August 21-24, 2016, Charlotte, North Carolina

DETC2016-60091

ESTIMATING AND EXPLORING THE PRODUCT FORM DESIGN SPACE USING DEEP GENERATIVE MODELS

Alexander Burnap*
Mechanical Engineering
University of Michigan
Ann Arbor, MI 48109
Email: aburnap@umich.edu

Ye Liu
Computer Science and Engineering
University of Michigan
Ann Arbor, MI 48109
Email: yeliu@umich.edu

Yanxin Pan
Design Science Program
University of Michigan
Ann Arbor, MI 48109
Email: yanxinp@umich.edu

Honglak Lee
Computer Science and Engineering
University of Michigan
Ann Arbor, MI 48109
Email: honglak@eecs.umich.edu

Richard Gonzalez
Psychology
University of Michigan
Ann Arbor, MI 48109
Email: gonzo@umich.edu

Panos Y. Papalambros
Mechanical Engineering
University of Michigan
Ann Arbor, MI 48109
Email: pyp@umich.edu

ABSTRACT

Product forms in quantitative design methods are typically expressed with a mathematical representation such as vectors, trees, graphs, and grammars. Such formal representations are restrictive in terms of realism or flexibility, and this limits their utility for human designers who typically create product forms in a design space that is restricted by the medium (e.g., free-hand sketching) and by their cognitive skills (e.g., creativity and experience). To increase the value of formal representations to human designers, this paper proposes to represent the design space as designs sampled from a statistical distribution of form and estimate a generative model of this distribution using a large set of images and design attributes of previous designs. This statistical representation approach is both flexible and realistic, and is estimated using a deep (multi-layer) generative model. The value of the representation is demonstrated in a study of two-dimensional automobile body forms. Using 180,000 form data of automobile designs over the past decade, we can morph a vehicle form into different body types and brands, thus offering human designers potential insights on realistic new design possibilities.

*Address all correspondence to this author.

1 INTRODUCTION

When developing product form for a new design concept, human designers use a mental representation of possible concept designs that implicitly defines the "true" conceptual design space [1–4] and is restricted only by the designer's cognitive skills. This true design space is searched using human creativity and experience [5–7], with a search process that is both *flexible*, i.e., moving from one design to another happens naturally and fluidly, and *realistic*, i.e., product form representations mirror their eventual embodiment or, if the representation is an abstraction such as a sketch, convey sufficient information to capture the eventual design embodiment [8].

Quantitative design methods use explicit mathematical representations of the design space, constituted of formalized elements such as vectors [9–11], trees [12], graphs [13–15], control points or handles for 2D pixels [16,17] or 3D voxels [18–22], and 2D [23–27] and 3D shape grammars [28]. Each unique combination of elements and their values represents a design. These explicit formal representations tend to be either flexible but of limited realism, due to being low dimensional (e.g., a silhouette), or realistic but of limited flexibility, due to being high dimensional

but only flexible in the local design space (e.g., 3D polygon mesh morphed by a few control points). A high-level positioning chart of different representations in terms of flexibility and realism is shown in Figure 1.

This paper describes a new approach by using a representation that is both more realistic and flexible than previous efforts. The design space is represented as designs $\mathbf{x}$ sampled from a statistical distribution $p(\mathbf{x})$, and a generative model of this distribution is estimated using a large set of images and associated design attributes of previous designs. The key design contribution in this work is the approach of changing the product form design representation to a statistical distribution, and estimating the product form design space using large-scale data of previous designs. A methodological contribution is the use of a crowd to act as an optimization algorithm for deep generative models; namely, we crowdsource opinions on whether generated designs look realistic from varying generative models. This step is important because validation of generative models is not objective, and significant differences may exist between numerical validation metrics such as "reconstruction error" versus visual quality [29].

To demonstrate these ideas, we conduct an experiment within the product form area of automobile styling. The design space distribution $p(\mathbf{x})$ is approximated using a deep generative model, in this case a variational autoencoder [30], over a data set of 2D images and design attributes of 179,702 automobile designs from the last decade. Preliminary results show that we are able to estimate a mathematical representation of the design space that is both realistic and flexible. We then explore this estimated design space by morphing vehicles via manupulation of design attributes such as body type, brand, and viewpoint.

This rest of this paper is structured as follows: Section 2 discusses related work. Section 3 introduces and interprets the approach of using a statistical distribution as a mathematical representation for product form, and the deep generative model used to approximate it. Section 4 details the numerical and crowdsourced experiment used to estimate the design space for automobile styling. Section 5 explores the design space and crowdsourcing results, and discusses the implications of this design representation, limitations, and opportunities for future work. We conclude in Section 6.

2 RELATED WORK

We discuss product form design representations in behavioral science research conducted on novice and expert designers during the conceptual design process, and in the mathematical formalization of design representations by design researchers. Next, we discuss generative models used in design and machine learning research, including the difficulty in establishing objective validation metrics for such models.

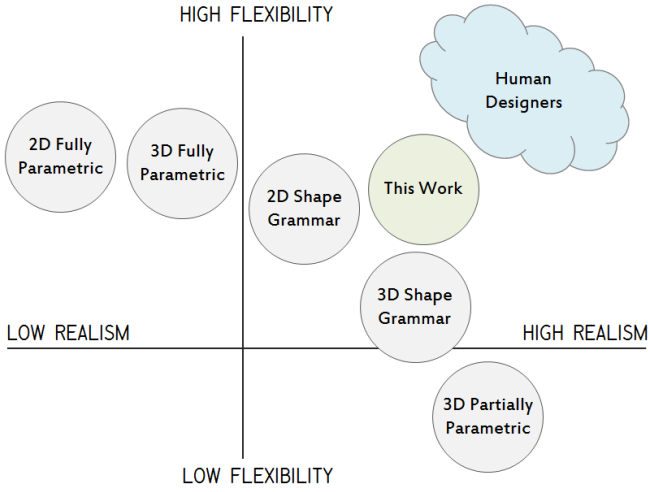

FIGURE 1. POSITIONING CHART OF PRODUCT FORM DESIGN REPRESENTATIONS ACCORDING TO LEVELS OF REALISM AND FLEXIBILITY OF REPRESENTATION. SEE APPENDIX A FOR VISUAL EXAMPLES OF EACH THESE REPRESENTATIONS.

2.1 Human Designer Mental Representation

The human designer's mental representation has been studied by design researchers extensively, with much focus on behavioral differences between novice and expert designers [6,31], and their mental representation of design knowledge [32]. Experts have been found to be better at representation *realism*, where realism is defined as the degree of design detail [33], and ability to connect design knowledge through sketches [8] and design analogies [34, 35].

Expert designers have also been found to be significantly different with regards to *flexibility* during the conceptual design search process. Expert designers make smaller "leaps" between design analogies when traversing their mental design representation [36], and are more likely to work backwards from the design solution [37], using a design problem decomposition strategy that enables "efficient" traversal of the design space.

2.2 Mathematical Design Representations

Mathematical representations of the product design space are more straightforward to compare as they are defined explicitly, thus constructing the design space according to all possible states of the representation. As noted in Section 1, these mathematical representations use a variety of formal elements that can be placed into six major categories as shown in Figure 2. While these mathematical representations have found numerous successes, including use by real designers [10, 24, 40, 41], each is limited by the tradeoff between *flexibility* and *realism* as illustrated in the positioning chart of Figure 1.

Fully parametric 2D or 3D vector representations have a high degree of flexibility since they are generally capable of mor-

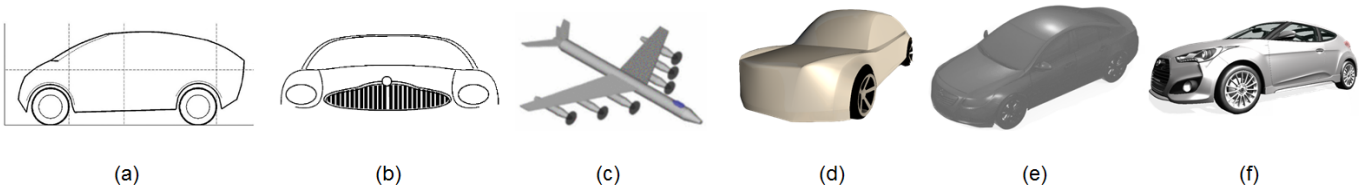

FIGURE 2. EXAMPLE DESIGNS FROM VARIOUS MATHEMATICAL PRODUCT FORM REPRESENTATIONS: (A) 2D FULLY PARAMETRIC [10], (B) 2D SHAPE GRAMMAR [24], (C) 3D SHAPE GRAMMAR [38], (D) 3D FULLY PARAMETRIC [21], (E) 3D PARTIALLY PARAMETRIC WITH ESTIMATED HANDLES [18], (F) 3D PARTIALLY PARAMETRIC WITH HAND-ENGINEERED HANDLES [39].

phing between all designs in the design space. This characteristic is important for the validity of results drawn from experiments using these representations. For example, assessing customer preferences using these representations enables full coverage of the space. The drawback of these representations, as shown in Figures 2(a) and 2(d), is that the resulting representations are often of limited realism due to their relatively low dimensionality.

Partially parametric 2D and 3D vector representations are manipulated with a lower dimensional set of “handles" or “control points" that affect the design representation’s pixels or voxels using an attachment function. This function can be defined through a functional form, such as a kernel [18, 42], or statistically estimated model [43]. The attachment functions often work on the entire design representation via manipulating all pixels or vertices, or by deforming the area or volume itself as shown in Figure 2(e). Another category shown in Figure 2(f) works on more fine details (e.g., headlight form and LEDs), but requires the use of design experts to hand-engineer various parametric handles.

Opposite of the fully parametric vector representations, partially parametric representations typically are very high-dimensional (i.e., 10,000’s of pixels or 100,000’s of voxels), and are subsequently very realistic. This comes at the cost of limited flexibility–the extent of possible manipulations is restricted to local perturbations. For larger changes, constraints must be placed between existing designs, typically through correspondence points with very strong and perhaps unrealistic assumptions on the interpolation function (e.g., linear or quadratic) [40].

Shape grammars, both 2D and 3D as shown in Figure 2(b) and Figure 2(c), occupy a middle ground in terms of flexibility and realism. These representations are powerful in that the design space they define is much larger than fully or partially parametric vector representation due to being combinatorial in in their composition of designs. Accordingly, while they offer flexibility across various designs constructed from constituent grammars, such flexibility between designs is limited. This comes with the advantage of being able to extrapolate much more reasonably as compared with vector representations. We discuss possible directions in combining random vector representations in the current work and grammar representations in Section 5.1.

2.3 Design Generative Models

Methods to generate design concepts have received attention by the design research community [9–11, 21, 26] and practicing designers [40]. These methods employ the mathematical representations noted in Section 2.2 and a variety of design generation schemes.

Human-guided design selection use queries with multiple generated designs in response or as an iterative communication between human and machine. A single query for a set of generated responses is often the focus of knowledge representation tools geared towards product form representations [32] following early ideas from Simon [44]. Iterative communication tools include interactive genetic algorithms [26, 45–47] and more recently proposed online crowdsourcing methods [21]. We note that while these previous approaches have used neural networks similar to this work, these previous uses have modeled preference elicitation rather than design representation and design generation.

2.4 Deep Generative Models

Deep generative models refer to a class of hierarchical statistical models (referred to as “deep learning" in the computer science community; see [48, 49] for survey) characterized by being composed of multiple layers of nonlinear functions, with each layer connected to its adjacent layers via a set of connecting “weights." These deep generative models work by modeling the data distribution, using assumptions on the data space, rather than the locally connective assumptions used in graph methods. Such models have recently received renewed attention due to their successes on benchmark tasks such as 2D image object recognition [50]. Here we discuss three related models that form the state-of-the-art with regards to 2D image generation.

The generative adversarial network (GAN) is a generative model that has a unique parameter estimation approach [51]. The model is divided into two parts, a generator and a discriminator; the generator is trained to generate images so that the discriminator cannot distinguish them from the ground truth images, while the discriminator is trained to discriminate generated images from the known “ground truth" images. The two parts are trained simultaneously to force the generator to produce images

as similar to the ground truth images as possible, where similarity is defined by the discriminator. Experiments have shown that this model is capable of generating highly realistic images with some exceptions. In particular, since the discriminator makes decisions based on pixel-level distance metrics, the generator can make unrealistic mistakes both obvious and important to humans, e.g., a face with a displaced nose.

The deconvolutional neural network is a multi-layer model composed of fully connected layers followed by two sets of convolutional layers–one tasked to generate design images and the other to generate segmentation masks of the design [52]. This model takes multiple design attribute to be generated (e.g., types of chairs). The model assembles a deterministic function that maps a set of input attributes to one output; however, this modelling assumption does not align well with our goal of capturing uncertainty from design attributes–we discuss this in detail in Section 3.

The variational autoencoder (VAE) [30] used in this work is an advanced version of the deconvolutional neural network, with major differences in the method of statistically estimating the model parameters of the model in its parametrization to introduce randomness to the generating process. A detailed introduction of the VAE model is in Section 3.

2.5 Validation of Generative Models

One challenging issue inherent to generative models is their the lack of straightforward validation. The requirements of this validation are twofold: The model needs to generate realistic 2D images that can be recognized by humans as a particular category of objects (e.g., cars), while at the same time, these images must be different from any image the model has seen in the data set, otherwise overfitting on the data set would model the simple solution of memorizing known training images. While the former requirement leads the model to produce similar images to the ones used in training, the later one forces the model in the opposite direction (i.e., generalization via interpolation and extrapolation).

Consequently, it is nontrivial to establish a quantitative validation that reflects model performance on both requirements. Research has been done using measurements such as pixel-level Euclidean distance, image retrieval from the known data set, and structured similarity indices from nearest neighbor images in the training set. However, none of these methods can give direct validation regarding the two requirements. In many cases, a generated 2D image that results in a favorable score under numerical measures scores very low on visual quality as assessed by a human [29]. To address this issue, we propose to utilize a crowd as a direct validation of the model's capability to generate realistic images.

3 PROBLEM FORMULATION

The problem formulation begins with assuming a fictitious conceptual design scenario involving three ingredients: (1) A "true" product form design space $\mathscr{X}$ containing the product forms of all 2D images capturing what it means to be the particular design (e.g., a passenger vehicle); (2) a "complete" (possibly infinite) list of design attributes, denoted $\mathbf{a}^*$, and obtained by being the exact set of design attributes most preferred by the targeted customer; and (3) a "perfect" design tool f^*, able to map deterministically a single complete design attribute list $\mathbf{a}^*$ to a single design $\mathbf{x} \in \mathscr{X}$:

$$\mathbf{x} = f^*(\mathbf{a}^*) \tag{1}$$

In reality, we do not have access to this complete set of design attributes $\mathbf{a}^*$ (e.g., the customer would most prefer exactly these bodyline curves, taillight shape and illumination, etc.), and must instead settle for a dramatically smaller finite set of design attributes $\mathbf{a}$ (e.g., the customer would prefer a 'Cadillac' 'Coupe'). We now have a massive number of unknown latent variables called "design features" denoted $\mathbf{h}$ as introduced in [53]. In other words, the previous complete list of design attributes is now partitioned into known design attributes and unknown design features $\mathbf{a}^* = \{\mathbf{a}, \mathbf{h}\}$. This introduces uncertainty into our originally deterministic function, which may now be represented according to a statistical distribution p_* with unknown functional form:

$$\mathbf{h} \sim p_*(\mathbf{h}). \tag{2}$$

Since we do not know this functional form, we instead assume the uncertainty from the random vector $\mathbf{h}$ may be captured by a distribution parametrized by θ, giving us the conditional distribution we aim to estimate:

$$\mathbf{x} \sim p_\theta(\mathbf{x}|\mathbf{a}) \tag{3}$$

Practically estimating the parameters of this distribution is challenging due to the high dimensionality of $\mathscr{X}$, which itself is a subset of a universal domain [3] of all 256-bit RGB values of the number of pixels forming the 2D image (i.e., $N = 256^{(3 \cdot \text{pixels})}$). Accordingly, we turn to a variational approximation [54] of the conditional likelihood of the data, one that will be particularly suited to online mode-seeking optimization methods.

This requires introduction of a latent random vector $\mathbf{z}$ as a tool to make the variational approximation using a tractable dis-

tribution q_ϕ:

$$\begin{aligned}\log p_\theta(\mathbf{x}|\mathbf{a}) &= \sum_{\mathbf{z}} q_\phi(\mathbf{z}|\mathbf{x},\mathbf{a}) \log p_\theta(\mathbf{x}|\mathbf{a}) \\ &= -\sum_{\mathbf{z}} q_\phi(\mathbf{z}|\mathbf{x},\mathbf{a}) \log p_\theta(\mathbf{z}|\mathbf{x},\mathbf{a})) \\ &+ \sum_{\mathbf{z}} q_\phi(\mathbf{z}|\mathbf{x},\mathbf{a})) \log p_\theta(\mathbf{x},\mathbf{z}|\mathbf{a})) \\ &= KL\left(q_{\mathbf{z}|\mathbf{x},\mathbf{a}} || p_{\mathbf{z}|\mathbf{x},\mathbf{a}}\right) + \mathscr{L}(\theta,\phi;\mathbf{x})\end{aligned} \tag{4}$$

where $KL\left(q_{\mathbf{z}|\mathbf{x},\mathbf{a}} || p_{\mathbf{z}|\mathbf{x},\mathbf{a}}\right)$ is the KL divergence, which is always non-negative. Therefore, the second term $\mathscr{L}(\theta,\phi;\mathbf{x})$ becomes a lower bound of the conditional likelihood given in Equation (3), and becomes the objective function we seek to maximize.

We expand $\mathscr{L}(\theta,\phi;\mathbf{x})$ into three terms, which are then amenable to the deep generative model in Section 3.1:

$$\begin{aligned}\mathscr{L}(\theta,\phi;\mathbf{x}) &= \sum_{\mathbf{z}} q_\phi(\mathbf{z}|\mathbf{x},\mathbf{a}) \left(\log p_\theta(\mathbf{x},\mathbf{z}|\mathbf{a}) - \log q_\phi(\mathbf{z}|\mathbf{x},\mathbf{a})\right) \\ &= \sum_{\mathbf{z}} q_\phi(\mathbf{z}|\mathbf{x},\mathbf{a}) \left(\log p_\theta(\mathbf{x}|\mathbf{z},\mathbf{a}) + \log p_\theta(\mathbf{z}|\mathbf{a}) - \log q_\phi(\mathbf{z}|\mathbf{x},\mathbf{a})\right).\end{aligned} \tag{5}$$

3.1 Deep Generative Model

We estimate the parameters θ and ϕ for the conditional likelihood given in Equation (5) using a hierarchical parametric model (i.e., "deep learning") that exploits invariance in 2D images, as well as optimization techniques to obtain point estimates to the values of these parameters. In particular, we use a variational autoencoder (VAE), a variational Bayesian approach introduced by Kingma and Welling [30] that learns a directed probabilistic model by approximating the posterior expectation with a reparametrization trick.

The VAE is made up of two networks: an "encoder" that transforms the 2D images within the data space to a latent representation (i.e., the last term in Equation (5), and a generative "decoder" model that transforms the latent representation back to a 2D image reconstruction in the data space, i.e., the first term in Equation (5). We use an extension to the VAE that includes conditioning on additional data [55–57], in our case known design attributes $\mathbf{a}$ associated with a given design $\mathbf{x}$. These conditional terms allow the latent representation to instead focus on encoding the uncertainty from features $\mathbf{h}$ not contained in the known design attributes $\mathbf{a}$:

$$\mathbf{z} \sim \text{Encoder}(\mathbf{x}) = q_\phi(\mathbf{z}|\mathbf{x},\mathbf{a}) \tag{6}$$

$$\hat{\mathbf{x}} \sim \text{Decoder}(\mathbf{x}) = p_\theta(\mathbf{x}|\mathbf{z},\mathbf{a}) \tag{7}$$

The reparametrization trick discussed [30] expresses our introduced latent random variable $\mathbf{z} \sim q_\phi(\mathbf{z}|\mathbf{x},\mathbf{a})$ with a deterministic variable $\mathbf{z} = g_\phi(\varepsilon,\mathbf{x},\mathbf{a})$, where ε is an independent "auxiliary" random variable, and $g_\phi(\cdot)$ is some vector-valued function parametrized by ϕ. Further, we approximate this axillary variable using Monte Carlo sampling:

$$\begin{aligned}\mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x},\mathbf{a})}[f(\mathbf{z})] &= \mathbb{E}_{p(\varepsilon)}\left[f\left(g_\phi(\varepsilon,\mathbf{x},\mathbf{a})\right)\right] \\ &\approx \frac{1}{L}\sum_{l=1}^{L} f\left(g_\phi\left(\varepsilon^{(l)},\mathbf{x},\mathbf{a}\right)\right) \text{ with } \varepsilon^{(l)} \sim p(\varepsilon)\end{aligned} \tag{8}$$

in which l denotes Monte Carlo draws and L denotes the total number of draws. Using Equation (8), we reparametrize the lower bound of the conditional likelihood we are after in Equation (5):

$$\begin{aligned}\mathscr{L}(\theta,\phi;\mathbf{x}) \quad \approx \quad & \frac{1}{L}\sum_{l=1}^{L} \log p_\theta\left(\mathbf{x}|\mathbf{z}^{(l)},\mathbf{a}\right) \\ & + \frac{1}{L}\sum_{l=1}^{L} \log p_\theta(\mathbf{z}^{(l)}|\mathbf{a}) \\ & - \frac{1}{L}\sum_{l=1}^{L} q_\phi\left(\mathbf{z}^{(l)}|\mathbf{x},\mathbf{a}\right) \\ \text{where } & \mathbf{z}^{(l)} = g_\phi\left(\varepsilon^{(l)},\mathbf{x}\right), \varepsilon^{(l)} \sim p(\varepsilon)\end{aligned} \tag{9}$$

Lastly, we define $q_\phi(\mathbf{z}|\mathbf{x}), p_\theta(\mathbf{z}), p_\theta(\mathbf{x}|\mathbf{z})$ as Gaussian distributions, whose parameters θ and ϕ we estimate using the VAE:

$$q_\phi(\mathbf{z}|\mathbf{x},\mathbf{a}) = \mathscr{N}\left(\mathbf{z};\mu_\phi(\mathbf{x},\mathbf{a}),\sigma^z_\phi(\mathbf{x},\mathbf{a})\right) \tag{10}$$

$$p_\theta(\mathbf{z}) = \mathscr{N}(\mathbf{z};\mathbf{0},\mathbf{I}) \tag{11}$$

$$p_\theta(\mathbf{x}|\mathbf{z},\mathbf{a}) = \mathscr{N}\left(\mathbf{x};\mu(\mathbf{z},\mathbf{a}),\sigma^2(\mathbf{z},\mathbf{a})\right) \tag{12}$$

3.2 Model Architecture

The architecture of a deep hierarchical model concerns the types of (i) "layers", i.e., vectors of functions that define function compositions between layers; (ii) "neurons" or "filters" making up the various layers, particularly their functional form; and (iii) connectivity linking layers to each other via parameters θ and ϕ.

The chosen architecture significantly influences the performance of the deep generative model, as architecture decisions constrain the flow of information throughout the model. Poor architecture choices increase the number of parameters of the model θ and ϕ or sub-optimal generative performance. For a VAE, a number of special layers is used to reduce the number of parameters while trading off information capture of the underlying data distribution. We show in Figure 3.1 the model architecture that uses four types of layers as described below:

Fully Connected Layers With a fully connected layer, the input $\mathbf{x} \in \mathbb{R}^{BxN}$ and the output $\mathbf{y} \in \mathbb{R}^{BxM}$ are associated with the function of $\mathbf{y} = f(\mathbf{x}^T\mathbf{w} + \mathbf{b})$, where $\mathbf{w} \in \mathbb{R}^{NxM}, \mathbf{b} \in \mathbb{R}^M$, and $f(\cdot)$ denotes a nonlinear function–in our case, the Rectified Linear (ReLu).

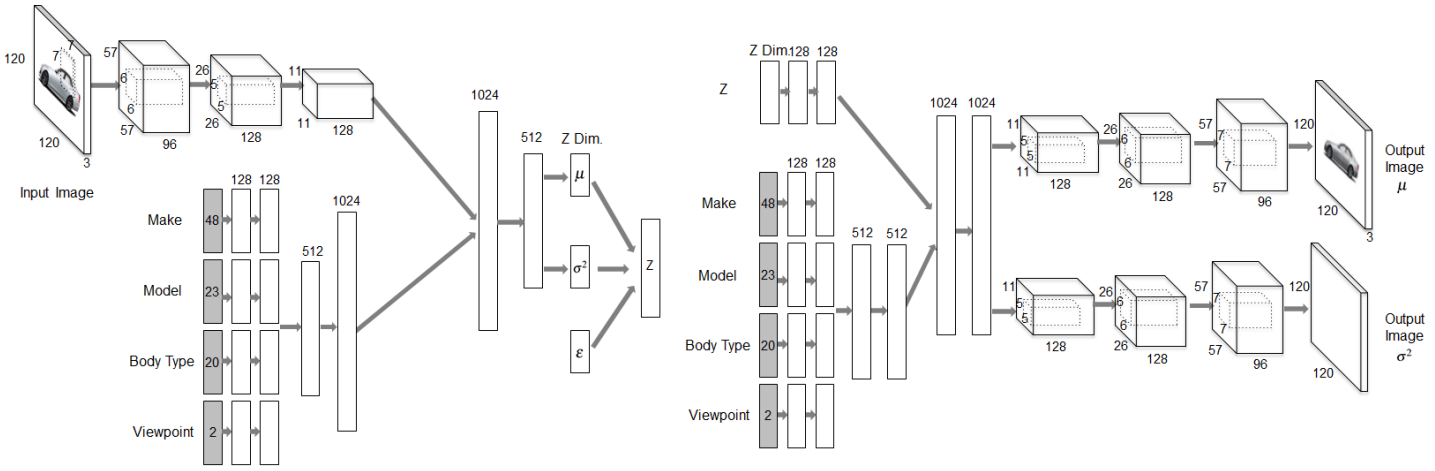

FIGURE 3. DEEP GENERATIVE MODEL ARCHITECTURE OF VARIATIONAL AUTOENCODER; ON THE LEFT IS THE ENCODER, WHILE THE RIGHT IS THE DECODER. SHADED BOXES REPRESENT INPUTS, WHITE BOXES REPRESENT FULLY CONNECTED LAYERS, AND RECTANGULAR PRISMS REPRESENT CONVOLUTIONAL AND POOLING LAYERS IN THE ENCODER, AND UPSAMPLING LAYERS IN THE DECODER.

Convolutional Layers Convolutional layers capture the notion that there are translation and rotation invariance, such as local regions forming image components that exist globally across the image. Such convolutional filters greatly reduce the number of parameters necessary in the layer relative to a fully connected layer, while still capturing a similar amount of information.

Similar to the fully connected layer, in a convolutional layer the input $\mathbf{x}$ and the output $\mathbf{y}$ are associated with the function of $\mathbf{y} = f(\mathbf{w} \bigotimes \mathbf{x} + \mathbf{b})$, where $\bigotimes$ denotes the 2D convolution operation and $f(\cdot)$ denotes a nonlinear function, in which we use Rectified Linear (ReLu) except the last layer in which no nonlinearity function is employed.

Pooling Layers Because pixel value information is highly redundant in images (i.e., neighbor pixels values are highly correlated) additional measurements are taken to reduce the number of parameters in the model. In a pooling layer, one output value will be used to replace a square area of input values. In this work, we use max pooling layers with a pooling size of 2 by 2, i.e., the maximum value of a 2 by 2 pixels grid, with the rest discarded to reduce parameters by a factor of 2 for each dimension.

Upsampling Layers Upsampling refers to the inverse of a "pooling" operation, which is used to "upsample" the coded information back to the same dimension as the input images. This operation necessarily lacks reconstruction information; however, the choices for such approximate inversion are varied (e.g., fixed location upsampling, average upsampling, and upsampling with switch units). In the current work, we use average upsampling.

4 EXPERIMENT

Our goal in the experiment was to statistically estimate the design space $p_\theta(\mathbf{x})$ to obtain a mathematical representation with realism and flexibility advantages as described in Section 1, using the model described in Section 3.1 and optimized using both numerical techniques and crowdsourcing.

4.1 Dataset

The data set consisted of 179,702 data points, with each point made up of a 2D image and four design attributes–make, model, body type, and viewpoint–with corresponding dimensionality shown in Figure 3.1. Each 2D image was downscaled to 120x120 pixels using OpenCV, an open source computer vision library [58]. We then split this data of previous designs into a "training set" and "validation set" with a 3:1 split ratio.

4.2 Numerical Parameter Estimation

The variational autoencoder described in Section 3.1 requires a number of hyperparameters (i.e., user-defined values such as learning rate and batch size) during the statistical estimation of the parameter sets θ and ϕ, as well as hyperparameters of the architecture itself (e.g., number of neurons or filter in a layer). We give these architecture hyperparameters in Figure 3.1. This architecture was implemented using Theano [59], an open source symbolic differentiation library with a graph-based compiler and GPU-acceleration support.

First-order methods have shown to be often better suited to estimation of deep generative models, particularly when extended with terms that mitigate being affected by saddle points

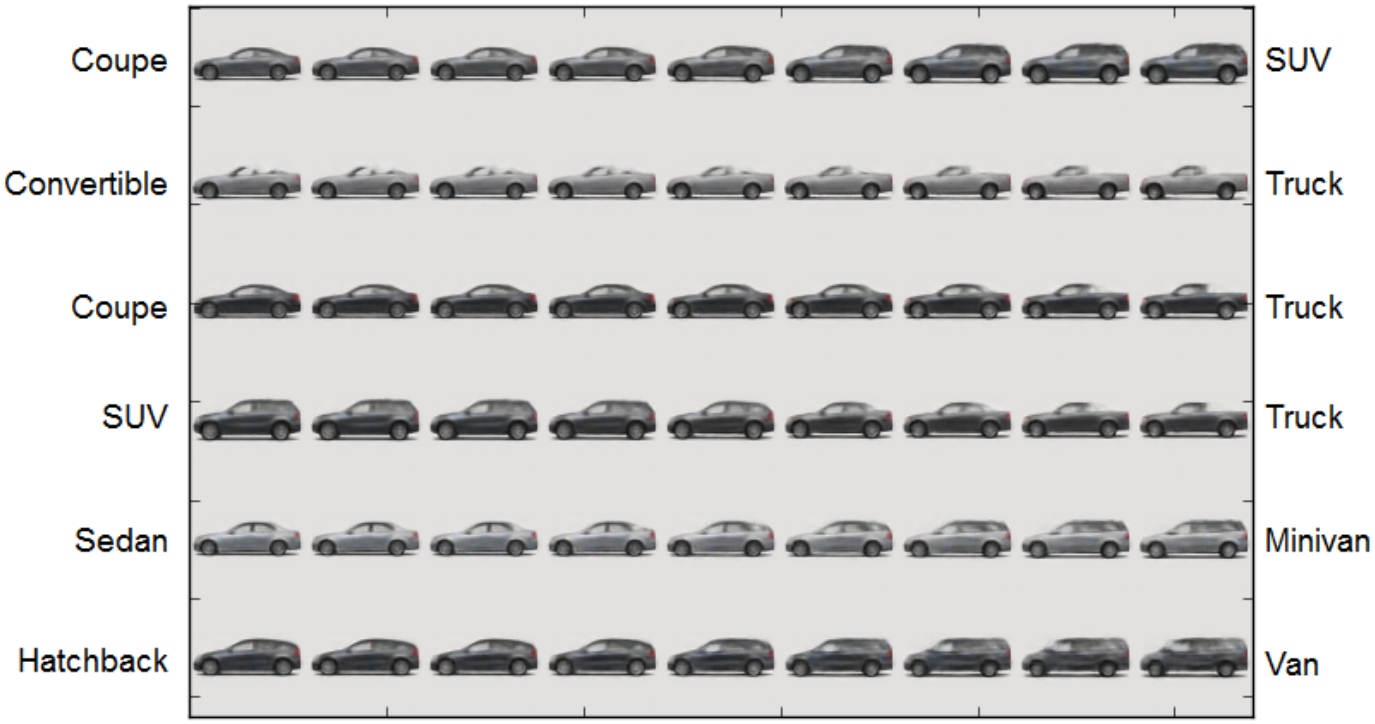

FIGURE 4. MORPHING BETWEEN VARIOUS BODY TYPES FROM THE ESTIMATED PRODUCT FORM DESIGN SPACE.

[60] and sharp discontinuities [61]. For these reasons, we use the ADAM optimizer [62], which has is particularly suited to parameter estimation of deep generative models. We use the ADAM optimization parameter of $\beta_1 = 0.1$ and $\beta_2 = 0.2$ with a learning rate of $\alpha = 0.0002$. Moreover, estimation of the parameter sets θ and ϕ in practice requires the use of "mini-batches" due to large data set sizes; accordingly, we used a mini-batch size of 100.

4.3 Crowdsourced Hyperparameter Estimation

The goal for crowdsourced hyperparameter estimation was to assess whether there were significant differences in visual quality of generated 2D images when varying the number of latent random variables Z used in the model architecture as described in Section 3.1. This assessment was performed as it has been shown that using numerical performance measures (e.g., log-likelihood) does not necessarily correspond with human perception of visual quality [29]. While certain theories (e.g., manifold hypothesis [63]) suggest that the effective dimensionality may not be best modeled as fixed, the addition of humans-in-the-loop may provide complementary advantages.

4.3.1 Participants A total of 69 participants were gathered using the crowdsourcing platform Amazon Mechanical Turk using an open call and a monetary incentive. We filtered out participants that "clicked through" the online application if their responses took less than 3 seconds per "survey question."

4.3.2 Procedure A web application with a database backend was developed to collect participant responses to generated 2D images with varying model architecture hyperparameters. Participants were first directed to a home page, which described the instructions for inputting visual quality responses to 2D images.

After clicking to proceed past the instructions, participants were presented with an ordered set of 2D images, and asked to select the 2D image that was most realistic. Each ordered set contained three 2D images, corresponding to three settings of the hyperparameters controlling the number of dimensions (i.e., 32, 128, 256 dimensions) in the latent representation Z. Each ordering was random, in order to not bias participants, while all the same design attributes were held constant (e.g., 'Cadillac,' 'coupe.'). The possible set of 2D image triplets contained all pairwise combinations of bodytypes from the sideview of the vehicle.

Participants were only allowed to click one of the three 2D images, and were able to change their selection. After participants completed 20 randomly selected 2D image triplets, they were redirected to web page thanking them for their time and presenting them with a unique code for monetary redemption.

5 RESULTS AND DISCUSSION

We explore the estimated product form design space by morphing between various pairs of design attributes. To show the

FIGURE 5. **MORPHING BETWEEN VARIOUS BRANDS FROM THE ESTIMATED PRODUCT FORM DESIGN SPACE.**

FIGURE 6. **ROTATIONS OF VARIOUS BODY TYPES FROM THE ESTIMATED PRODUCT FORM DESIGN SPACE.**

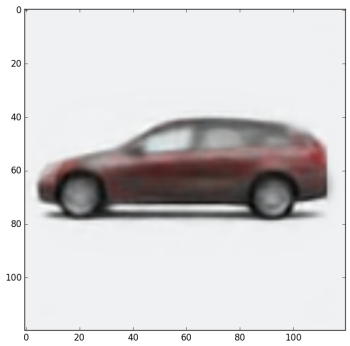

FIGURE 7. GENERATED VEHICLE BETWEEN 'SEDAN' AND 'SUV' FOR RANDOMLY SET BRAND THAT LOOKS LIKE A 'CROSSOVER.'

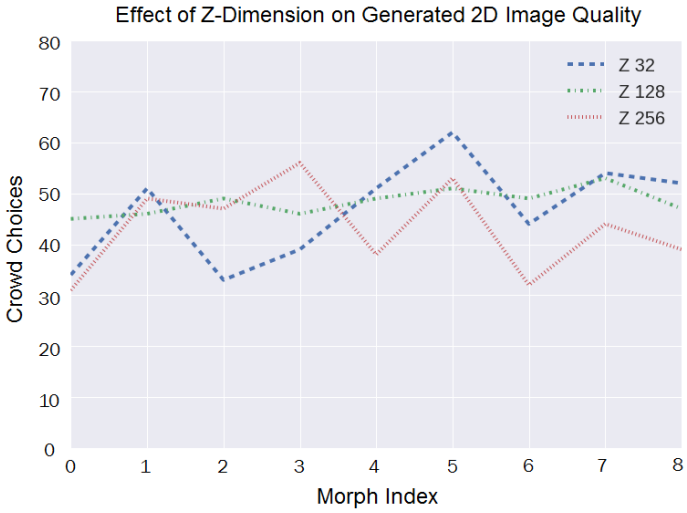

FIGURE 8. EFFECT OF VARIOUS NUMBER OF z RANDOM VARIABLES IN HIDDEN REPRESENTATION ON GENERATED 2D IMAGE QUALITY AS ASSESSED BY CROWD. MORPH INDEX REFERS TO HOW FAR BETWEEN TWO KNOWN DESIGN ATTRIBUTES A DESIGN WAS MORPHED–E.G., 0 AND 8 MAY BE 'CONVERTIBLE' AND 'TRUCK', RESPECTIVELY.

flexibility of the estimated mathematical representations $p_\theta(\mathbf{x})$, we morph between various body types in Figure 4, various brands in Figure 5, and rotational viewpoints in Figure 6.

The designs $\mathbf{x}$ sampled for these results are all artificially generated (i.e., none are in the dataset). Moreover, we show multiple steps in between each design attribute pair to indicate that we are not overfitting on the data set, as none of these generated designs exist. In particular, we observe that the visual quality of the generated designs is uniform across various morphing steps between known design attributes (e.g., from coupe to SUV); this reinforces the notion that we are not simply overfitting, and instead we are estimating the "true" product design space $p_\theta(\mathbf{x}|\mathbf{a})$.

The motivation for this work was developed in part from working with practicing designers in the automotive industry and recognizing the necessity of a design representation that can morph between various brands and body types, yet realistic enough to convey sufficient meaning to designers [39]. This representation is not limited to brand studies. A number of design questions can be explored. For example, we show in Figure 7 a generated vehicle between a 'sedan' and an 'SUV,' which is currently the fastest growing design segment in the automotive market. This type of design generation can serve as inspiration to designers working on designs for new market segments [5].

Figure 8 shows the results of the crowdsourced parameter optimization. Preliminary results indicate that we cannot conclusively state whether the crowd was able to discriminate between various hyperparameter settings during the design space estimation. Further research is required into using a crowd to fine-tune parameters affecting image quality after an initial computer-only optimization is performed.

Thus, the hypothesized value of using crowdsourced optimization requires deeper investigation. If the crowd is shown to improve the statistical estimation procedure, we may be able to build more robust crowd-powered optimization systems for these generative models. The current approach is not in real time; however, a worthwhile goal may be to build a real-time optimization loop including the crowd, particularly if incentivized as in the emerging area of gamification [64].

5.1 Limitations and Future Work

Interpretation of the latent space remains a challenging and potentially rewarding goal in this research. Recent nonlinear predictive models offer significantly improved generalization performance of design task prediction accuracy, and consequently improved capture of the underlying physics of the design task, e.g., preference prediction and market segmentation [65–67].

These nonlinear and network models are in contrast with the interpretable linear models commonly used in design task modeling, which often work on strong modeling assumptions comprised of main effects and pairwise interaction terms, and thus neglect all other statistical dependencies amongst design variables. Future work towards interpretation of these latent representation may offer much deeper insight into underlying design perceptions and preferences of customers, translated into actionable design decisions that capture how the designer can adjust design attributes to elicit desired preferences within a specific population.

We show in Figure 9 an example of the design feature "color" that we do not currently capture. While this feature may be simpler to capture using crowds, numerous other features still exist that are not as straightforward. Current work towards such "feature interpretation" has shown preliminary promise, including data-driven approaches to predict which visual features of a

FIGURE 9. **GENERATED DESIGN DISPLAYING THE DESIGN FEATURES 'COLOR,' WHICH WE DO NOT YET CONTROL.**

design most elicit attention [68]. These approaches aim to move into the causality behind features in deep convolutional neural networks [69, 70].

Interpreting such features may lead to new shape grammars. The combination of top-down statistical estimation of the design space and bottom-up definition of the space using shape grammars may be a valuable direction for future research, particularly by analyzing and subsequently merging high-level statistical features with hand-engineered shape grammars. Validating such combinations, and more generally of visual fidelity of statistical design representations, may be aided by methods that use humans-in-the-loop more advanced than in the present work, such as online crowdsourcing [71], or in-person eye tracking [72, 73].

6 CONCLUSION

Human designers use a mental design representation of product form that is both flexible and realistic, allowing efficient exploration of the design space during the conceptual design process. Mathematical representations of the product form design space impose constuctivist restrictions on the design space and trade off representation flexibility for representation realism.

We changed a statistical distribution as a mathematical representation that is more flexible and realistic than previously proposed representations. We formulated this representation by assuming a "perfect" conceptual design scenario and progressively introduced uncertainty as dictated by the real world. We approximated this true statistical distribution using a deep generative model, in particular a variational autoencoder, which is amenable to efficient computing of large-scale data sets.

We conducted an experiment in the product form domain of automotive styling, using design attributes and 2D images of automobile designs from the last decade. The results showed that we are able to find a design representation that is both flexible and realistic in exploring the design space over design attributes such as body type, brand, and viewpoint. We also examined using a crowd to improve the parameter estimation process of the deep generative model; our preliminary results showed that we are not yet able to improve our generative model results to statistically significant levels in this way.

Lastly, we discussed a number of possible improvements to this work within the emerging area of data-driven design. In particular, interpretation of design features otherwise wrapped up in uncertainty offers design researchers and practicing designers opportunities for valuable design insight. Further investigation into crowdsourcing mechanisms, real-time and gamified, may prove fruitful. Aligning design augmentation tools with practicing designers and design researchers who study expert designers remains important and can further the value of design automation tools.

ACKNOWLEDGMENTS

This work has been supported by the National Science Foundation under Grant No. CMMI-1266184. This support is gratefully acknowledged. The authors would like to thank Dr. Jeff Hartley and General Motors Corporation for industry motivation, as well as Evox Image Holdings for use of image data. We also thank Aria Collopy, Melissa Greene, and Yuqing Zhou for helpful discussions.

REFERENCES

[1] Rosenman, M. A., and Gero, J. S., 1993. "Creativity in design using a design prototype approach". *Modeling creativity and knowledge-based creative design*, pp. 111–138.

[2] Goldschmidt, G., 1997. "Capturing indeterminism: representation in the design problem space". *Design Studies,* ***18***(4), Oct., pp. 441–455.

[3] Gero, J. S., and Maher, M. L., 2013. *Modeling Creativity and Knowledge-Based Creative Design*. Psychology Press, May.

[4] Crilly, N., Moultrie, J., and Clarkson, P., 2004. "Seeing things: consumer response to the visual domain in product design". *Design Studies,* ***25***(6), Nov., pp. 547–577.

[5] Hartley, J., 1996. "Brands Through the Lens of Style". Quest and Associates.

[6] Cross, N., 2004. "Expertise in design: an overview". *Design Studies,* ***25***(5), Sept., pp. 427–441.

[7] Eckert, C., and Stacey, M., 2000. "Sources of inspiration: a language of design". *Design Studies,* ***21***(5), Sept., pp. 523–538.

[8] Kavakli, M., and Gero, J. S., 2001. "Sketching as mental imagery processing". *Design Studies,* ***22***(4), pp. 347–364.

[9] Orsborn, S., Cagan, J., and Boatwright, P., 2009. "Quantifying Aesthetic Form Preference in a Utility Function". *Journal of Mechanical Design,* ***131***(6).

[10] Reid, T. N., Gonzalez, R. D., and Papalambros, P. Y., 2010. "Quantification of Perceived Environmental Friendliness for Vehicle Silhouette Design". *Journal of Mechanical Design,* ***132***(10), p. 101010.

[11] Petiot, J.-F., Salvo, C., Hossoy, I., and Papalambros, P. Y., 2009. "A cross-cultural study of users' craftsmanship perceptions in vehicle interior design". *International Journal of Product Development,* ***7***(1), pp. 28–46.

[12] Orbay, G., Fu, L., and Kara, L. B., 2015. "Deciphering the Influence of Product Shape on Consumer Judgments Through Geometric Abstraction". *Journal of Mechanical Design,* ***137***(8), p. 081103.

[13] Bayrak, A. E., Ren, Y., and Papalambros, P. Y., 2013. "Design of Hybrid-Electric Vehicle Architectures Using Auto-Generation of Feasible Driving Modes". In ASME 2013 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference, American Society of Mechanical Engineers, pp. V001T01A005–V001T01A005.

[14] Zhang, B., and Rai, R., 2014. "Materials Follow Form and Function: Probabilistic Factor Graph Approach for Automatic Material Assignments to 3d Objects". In ASME 2014 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference, American Society of Mechanical Engineers, pp. V007T07A012–V007T07A012.

[15] Münzer, C., Helms, B., and Shea, K., 2013. "Automatically transforming object-oriented graph-based representations into boolean satisfiability problems for computational design synthesis". *Journal of Mechanical Design,* ***135***(10), p. 101001.

[16] Toh, C. A., and Miller, S. R., 2014. "The Impact of Example Modality and Physical Interactions on Design Creativity". *Journal of Mechanical Design*.

[17] Bao, Q., El Ferik, S., Shaukat, M. M., and Yang, M. C., 2014. "An Investigation on the Inconsistency of Consumer Preferences: A Case Study of Residential Solar Panels". In Proceedings of the 2014 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference.

[18] Yumer, M. E., Chaudhuri, S., Hodgins, J. K., and Kara, L. B., 2015. "Semantic shape editing using deformation handles". *ACM Transactions on Graphics (TOG),* ***34***(4), p. 86.

[19] Mukherjee, A., Zhang, Y., and Rai, R., 2014. "Probabilistic Design Miming". In ASME 2014 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference, American Society of Mechanical Engineers, pp. V007T07A011–V007T07A011.

[20] Kang, S. W., and Tucker, C. S., 2015. "Automated Concept Generation Based on Function-Form Synthesis". In ASME 2015 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference, American Society of Mechanical Engineers, pp. V02AT03A008–V02AT03A008.

[21] Ren, Y., Burnap, A., and Papalambros, P., 2013. "Quantification of Perceptual Design Attributes Using a Crowd". In Proceedings of the 19th International Conference on Engineering Design.

[22] Tovares, N., Boatwright, P., and Cagan, J., 2014. "Experiential Conjoint Analysis: An Experience-Based Method for Eliciting, Capturing, and Modeling Consumer Preference". *Journal of Mechanical Design,* ***136***(10), p. 101404.

[23] Pugliese, M. J., and Cagan, J., 2002. "Capturing a rebel: modeling the Harley-Davidson brand through a motorcycle shape grammar". *Research in Engineering Design,* ***13***(3), pp. 139–156.

[24] McCormack, J. P., Cagan, J., and Vogel, C. M., 2004. "Speaking the Buick language: capturing, understanding, and exploring brand identity with shape grammars". *Design Studies,* ***25***(1), Jan., pp. 1–29.

[25] Orsborn, S., Cagan, J., Pawlicki, R., and Smith, R. C., 2006. "Creating cross-over vehicles: Defining and combining vehicle classes using shape grammars". *AIE EDAM: Artificial Intelligence for Engineering Design, Analysis, and Manufacturing,* ***20***(03), pp. 217–246.

[26] Yannou, B., Dihlmann, M., and Awedikian, R., 2008. "Evolutive Design of Car Silhouettes". In ASME 2008 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference, American Society of Mechanical Engineers, pp. 15–24.

[27] Königseder, C., and Shea, K., 2016. "Comparing strategies for topologic and parametric rule application in automated computational design synthesis". *Journal of Mechanical Design,* ***138***(1), p. 011102.

[28] Perez Mata, M., Ahmed-Kristensen, S., and Shea, K., 2015. "Spatial Grammar for Design Synthesis Targeting Perceptions: Case Study on Beauty". p. V01AT02A013.

[29] Theis, L., Oord, A. v. d., and Bethge, M., 2015. "A note on the evaluation of generative models". *arXiv:1511.01844 [cs, stat]*, Nov.

[30] Kingma, D. P., and Welling, M., 2013. "Auto-encoding variational bayes". *arXiv preprint arXiv:1312.6114*.

[31] Dinar, M., Shah, J. J., Cagan, J., Leifer, L., Linsey, J., Smith, S. M., and Hernandez, N. V., 2015. "Empirical studies of designer thinking: Past, present, and future". *Journal of Mechanical Design,* ***137***(2), p. 021101.

[32] Chandrasegaran, S. K., Ramani, K., Sriram, R. D., Horváth, I., Bernard, A., Harik, R. F., and Gao, W., 2013. "The evolution, challenges, and future of knowledge representation

in product design systems". *Computer-aided design,* ***45***(2), pp. 204–228.

[33] Björklund, T. A., 2013. "Initial mental representations of design problems: Differences between experts and novices". *Design Studies,* ***34***(2), pp. 135–160.

[34] Fu, K., Murphy, J., Yang, M., Otto, K., Jensen, D., and Wood, K., 2015. "Design-by-analogy: experimental evaluation of a functional analogy search methodology for concept generation improvement". *Research in Engineering Design,* ***26***(1), pp. 77–95.

[35] Linsey, J. S., 2007. *Design-by-analogy and representation in innovative engineering concept generation*. ProQuest.

[36] Ozkan, O., and Dogan, F., 2013. "Cognitive strategies of analogical reasoning in design: Differences between expert and novice designers". *Design Studies,* ***34***(2), pp. 161–192.

[37] Ho, C.-H., 2001. "Some phenomena of problem decomposition strategy for design thinking: differences between novices and experts". *Design Studies,* ***22***(1), pp. 27–45.

[38] Oberhauser, M., Sartorius, S., Gmeiner, T., and Shea, K., 2015. "Computational design synthesis of aircraft configurations with shape grammars". In *Design Computing and Cognition'14*. Springer, pp. 21–39.

[39] Burnap, A., Hartley, J., Pan, Y., Gonzalez, R., and Papalambros, P. Y., 2015. "Balancing Design Freedom and Brand Recognition in the Evolution of Automotive Brand Character". In Proceedings of the 2015 International Design Engineering Technical Conferences.

[40] Kókai, I., Finger, J., Smith, R. C., Pawlicki, R., and Vetter, T., 2007. "Example-based conceptual styling framework for automotive shapes". In Proceedings of the 4th Eurographics workshop on Sketch-based interfaces and modeling, ACM, pp. 37–44.

[41] Telenko, C., Sosa, R., and Wood, K. L., 2016. "Changing conversations and perceptions: The research and practice of design science". In *Impact of Design Research on Industrial Practice*. Springer, pp. 281–309.

[42] Murugappan, S., Piya, C., Ramani, K., and others, 2013. "Handy-potter: Rapid exploration of rotationally symmetric shapes through natural hand motions". *Journal of Computing and Information Science in Engineering,* ***13***(2), p. 021008.

[43] Yumer, M. E., Asente, P., Mech, R., and Kara, L. B., 2015. "Procedural modeling using autoencoder networks". In Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology, ACM, pp. 109–118.

[44] Simon, H. A., 1996. *The sciences of the artificial*. MIT press.

[45] Smith, J. R., 1991. "Designing biomorphs with an interactive genetic algorithm.". In ICGA, pp. 535–538.

[46] Poirson, E., Petiot, J.-F., Boivin, L., and Blumenthal, D., 2013. "Eliciting user perceptions using assessment tests based on an interactive genetic algorithm". *Journal of Mechanical Design,* ***135***(3), p. 031004.

[47] Chan, K. Y., Wong, Y. S., and Dillon, T. S., 2012. *Computational intelligence techniques for new product design*, Vol. 403. Springer Science & Business Media.

[48] Bengio, Y., 2009. "Learning deep architectures for AI". *Foundations and trends in Machine Learning,* ***2***(1), pp. 1–127.

[49] Schmidhuber, J., 2015. "Deep learning in neural networks: An overview". *Neural Networks,* ***61***, pp. 85–117.

[50] Krizhevsky, A., Sutskever, I., and Hinton, G. E., 2012. "Imagenet classification with deep convolutional neural networks". In Advances in neural information processing systems, pp. 1097–1105.

[51] 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, pp. 2672–2680. 00017.

[52] Dosovitskiy, A., Springenberg, J. T., Tatarchenko, M., and Brox, T., 2014. "Learning to Generate Chairs, Tables and Cars with Convolutional Networks". *arXiv:1411.5928 [cs]*, Nov.

[53] Burnap, A., Pan, Y., Liu, Y., Ren, Y., Lee, H., Gonzalez, R., and Papalambros, P. Y., 2016. "Improving design preference prediction accuracy with feature learning".

[54] Wainwright, M. J., and Jordan, M. I., 2008. "Graphical models, exponential families, and variational inference". *Foundations and Trends® in Machine Learning,* ***1***(1-2), pp. 1–305.

[55] Kingma, D. P., Mohamed, S., Rezende, D. J., and Welling, M., 2014. "Semi-supervised learning with deep generative models". In Advances in Neural Information Processing Systems, pp. 3581–3589.

[56] Yan, X., Yang, J., Sohn, K., and Lee, H., 2015. "Attribute2image: Conditional Image Generation from Visual Attributes". *arXiv preprint arXiv:1512.00570*.

[57] Louizos, C., Swersky, K., Li, Y., Zemel, R., and Welling, M., 2015. "The Variational Fair Autoencoder". *arXiv:1511.00830 [cs, stat]*, Nov.

[58] Bradski, G., and Kaehler, A., 2008. *Learning OpenCV: Computer vision with the OpenCV library*. " O'Reilly Media, Inc.".

[59] Bergstra, J., Breuleux, O., Bastien, F., Lamblin, P., Pascanu, R., Desjardins, G., Turian, J., Warde-Farley, D., and Bengio, Y., 2010. "Theano: a CPU and GPU math expression compiler". In Proceedings of the Python for Scientific Computing Conference (SciPy). Oral Presentation.

[60] Dauphin, Y. N., Pascanu, R., Gulcehre, C., Cho, K., Ganguli, S., and Bengio, Y., 2014. "Identifying and attacking the saddle point problem in high-dimensional non-convex optimization". In Advances in Neural Information Processing Systems, pp. 2933–2941.

[61] Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan,

D., Goodfellow, I., and Fergus, R., 2013. "Intriguing properties of neural networks". *arXiv preprint arXiv:1312.6199*.
[62] Kingma, D. P., and Adam, J. B., 2015. "Adam: A method for stochastic optimization". In International Conference on Learning Representation.
[63] Bengio, Y., Goodfellow, I. J., and Courville, A., 2015. "Deep learning". *An MIT Press book in preparation. Draft chapters available at http://www. iro. umontreal. ca/ bengioy/dlbook*.
[64] Ren, Y., Bayrak, E., and Papalambros, P. Y., 2015. "Ecoracer: Optimal Design and Control of Electric Vehicles Using Human Game Players". In Proceedings of the ASME 2015 International Design Engineering Technical Conferences.
[65] Haaf, C. G., Michalek, J. J., Morrow, W. R., and Liu, Y., 2014. "Sensitivity of vehicle market share predictions to discrete choice model specification". *Journal of Mechanical Design,* ***136***(12), p. 121402.
[66] Wang, M., and Chen, W., 2015. "A data-driven network analysis approach to predicting customer choice sets for choice modeling in engineering design". *Journal of Mechanical Design,* ***137***(7), p. 071410.
[67] Burnap, A., Pan, Y., Liu, Y., Ren, Y., Lee, H., Gonzalez, R., and Papalambros, P. Y., 2016. "Improving design preference prediction accuracy using feature learning". *Journal of Mechanical Design*.
[68] Pan, Y., Burnap, A., Liu, Y., Lee, H., Gonzalez, R., and Papalambros, P., 2016. "A Quantitative Model for Identifying Regions of Design Visual Attraction and Application to Automobile Styling". In Proceedings of the 2016 Internation Design Conference.
[69] Zeiler, M. D., and Fergus, R., 2014. "Visualizing and understanding convolutional networks". In *Computer Vision ECCV 2014*. Springer, pp. 818–833.
[70] Simonyan, K., Vedaldi, A., and Zisserman, A., 2013. "Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps". *arXiv:1312.6034 [cs]*, Dec.
[71] Burnap, A., Ren, Y., Gerth, R., Papazoglou, G., Gonzalez, R., and Papalambros, P. Y., 2015. "When Crowdsourcing Fails: A Study of Expertise on Crowdsourced Design Evaluation". *Journal of Mechanical Design,* ***137***(3), p. 031101.
[72] Reid, T. N., MacDonald, E. F., and Du, P., 2013. "Impact of Product Design Representation on Customer Judgment". *Journal of Mechanical Design,* ***135***(9), p. 091008.
[73] Du, P., and MacDonald, E. F., 2014. "Eye-Tracking Data Predict Importance of Product Features and Saliency of Size Change". *Journal of Mechanical Design,* ***136***(8), p. 081005.