

MIT Open Access Articles

Design subspace learning: Structural design space exploration using performance-conditioned generative modeling

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Danhaive, Renaud and Mueller, Caitlin T. 2021. "Design subspace learning: Structural design space exploration using performance-conditioned generative modeling." *Automation in Construction*, 127.

As Published: 10.1016/J.AUTCON.2021.103664

Publisher: Elsevier BV

Persistent URL: <https://hdl.handle.net/1721.1/145568>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-NonCommercial-NoDerivs License



1 **Design subspace learning: structural design space exploration using performance-** 2 **conditioned generative modeling**

3
4 Renaud Danhaive^{a,*}, Caitlin T. Mueller^a

5
6 ^a Massachusetts Institute of Technology, Building Technology Program, Department of Architecture, Cambridge, MA 02139, USA

7 ^{*} Corresponding author. *Address:* 77 Massachusetts Avenue, Room 5-418, Cambridge, MA 02139, USA. *Tel.* +1 732 693 7447. *Email address:*
8 danhaive@mit.edu (R. Danhaive)

9 **Abstract**

10 Designers increasingly rely on parametric design studies to explore and improve structural concepts based
11 on quantifiable metrics, generally either by generating design variations manually or using optimization
12 methods. Unfortunately, both of these approaches have important shortcomings: effectively searching a
13 large design space manually is infeasible, and design optimization overlooks qualitative aspects important
14 in architectural and structural design. There is a need for methods that take advantage of computing
15 intelligence to augment a designer’s creativity while guiding—not forcing—their search for better-
16 performing solutions. This research addresses this need by integrating conditional variational autoencoders
17 in a performance-driven design exploration framework. First, a sampling algorithm generates a dataset of
18 meaningful design options from an unwieldy design space. Second, a performance-conditioned variational
19 autoencoder with a low-dimensional latent space is trained using the collected data. This latent space is
20 intuitive to explore by designers even as it offers a diversity of high-performing design options.

21
22 **Keywords:** deep generative modeling, latent space, latent variable, variational autoencoder, design space,
23 computational design

24 **Highlights**

- 25 • There is a need to balance performance and diversity in design space exploration.
- 26 • High-dimensional spaces are hard to explore without resorting to optimization.
- 27 • Deep latent learning (VAE) can usefully compress high-dimensional design spaces.
- 28 • Performance-driven sampling yields better latent spaces at less computational cost.
- 29 • A two-dimensional latent space is a natural interface for design exploration.

30 **1 Introduction**

31 Computational design offers structural designers ways to explore large arrays of options parametrically
32 generated from formalized design spaces. Design spaces with a large number of parameters are appealing
33 because they have the potential to yield solutions that are unexpected yet high-performing. Unfortunately,
34 they are also tedious to effectively explore through manual means, and human cognition is not effective at
35 processing high-dimensional information. A natural solution is thus to use automated optimization
36 procedures. While optimization certainly has a role to play as part of the arsenal of methods available for
37 design exploration, its applicability is limited because it does not effectively account for human input and
38 leaves no room for intuition. In addition, design spaces may be ill-defined and objective functions are
39 sometimes one but many aspects for a human designer to factor in, rendering optimization results close to
40 useless.

41 *1.1 Research objectives and scope*

42 There is a need for methods that allow designers to intuitively explore chaotic design landscapes without
43 resorting to automated procedures, given the importance of human factors in design. This research seeks to
44 address this need by capitalizing on advances in generative modeling and artificial intelligence to help

45 human designers explore large design spaces in more intuitive, performance-driven ways. To do so, it
46 introduces a method that generates and uses design performance data to build high-performance low-
47 dimensional design subspaces which may be explored directly and easily visualized in their entirety.
48

49 First, a newly proposed performance-driven sampling algorithm is used to generate a dataset of
50 meaningful—i.e. biased toward high-performance design regions—options from a large, unwieldy design
51 space. Second, these datasets are used to train low-dimensional deep generative models that are intuitive to
52 explore by human designers and offer a diversity of high-performing design options. The models allow
53 designers to freely and flexibly explore design options that attain performance levels prescribed by the
54 designer, without the need to rely on optimization. Instead of replacing human intuition with deterministic,
55 quantitative rules, the computer here acts as a design collaborator that augments the human intellect.
56

57 Because of the large amount of data required to train deep generative models, the proposed method is best
58 suited for applications where simulating thousands of solutions is feasible, either thanks to reasonable
59 simulation times or through massive parallel computing. In terms of simulation time, oft-used structural
60 and architectural analyses—such as linear finite element analysis or energy analysis—range between
61 seconds and minutes, meaning that thousands of solutions may be computed in hours or days at the most.
62 For more computationally expensive analyses, the tenets of the proposed method still hold but it is best to
63 substitute exact analyses for simpler ones—for example, ones that use coarser analysis resolution—or use
64 fast data-driven models. Other than constraints on simulation time, this research applies to any design
65 problem where performance must be considered alongside non-quantifiable characteristics.
66

67 *1.2 Design space transformation and reduction*

68 Some of the most recent advances in design space exploration powered by a surge of interest in data-driven
69 algorithms focus on the transformation of high-dimensional design spaces into lower-dimensional
70 representations. The established goal of such techniques is to reduce the dimensionality of the original space
71 through visualization or variable transformation or both. When used for visualization, dimensionality
72 reduction helps designers make sense of the design space in formats, 2D maps for example, that clearly
73 emphasize patterns in the design space. In that regard, self-organizing maps [1], [2] have become popular
74 as a means to organize design samples on a two-dimensional plane [3]–[5].
75

76 Variable transformation explicitly looks to build meaningful mappings from a new reduced set of new
77 variables to the original variables, such that designers can control a small number of super-variables. The
78 common denominator to these methods is that they are unsupervised, i.e. they do not directly rely on the
79 objective function values. Rather, the objective function is used indirectly to gather the data, as a means to
80 bias collected samples towards well-performing regions of the design space. The hope is that the structure
81 of these regions can be uncovered later on through unsupervised learning. Brown and Mueller [6], for
82 example, use optimization run histories as datasets for which they compute the principal components. The
83 extracted principal directions may then be used as new, composite variables—the first of which is likely to
84 represent the direction of steepest change of the objective function. A different, yet related area of research
85 seeks to objective function as a design variable of sorts through the use of inference. Conti, for example,
86 employs Bayesian networks to predict the probability that a given variable value will yield a desired
87 objective function value [7].
88

89 This research similarly seeks to build embeddings in which design exploration is facilitated, but, compared
90 to previous work, seeks to capture the nonlinear manifold structure of the collected dataset and learn a
91 reduced and continuous latent representation of the high-performance regions that can be conditioned by a
92 designer to meet prescribed performance levels.

93 *1.3 Deep generative modeling*

94 Generative modeling—not to be confused with generative design—is a branch of unsupervised machine
95 learning that seeks to understand data by learning to recreate it. While they apply to diverse fields of study,
96 generative models, such as generative adversarial networks (GAN), have made headlines and captured the
97 popular imagination in particular for their ability to generate images that are nearly indistinguishable from
98 real-world pictures [8].
99

100 Given a dataset of observations, generative models are trained to retrieve the probability distribution from
101 which the dataset was drawn. Real-world data often lies on complex manifolds in high-dimensional space.
102 The interest of generative modeling in design is that it may be used to generate previously unseen yet
103 probable designs by learning the structure of that manifold. Unfortunately, sampling new designs from
104 generative models is not necessarily intuitive or controllable, and recent years have seen a sharper focus on
105 latent variable models, which overcome this issue by generating data distributions based on a fixed number
106 of variables whose mapping to the original data space is learned through data.
107

108 This research focuses on one class of deep latent generative models: variational autoencoders [9], [10].
109 Variational autoencoders (VAE) assume that high-dimensional observations in a dataset are drawn from
110 probability distributions defined over latent variables, which may be used to draw new samples by
111 navigating the learned latent subspace. This subspace thus offers a controlled way to generate new data and,
112 in the context of this research, explore new design options. This work uses the conditional variant of VAE
113 [11] to include performance as an explicit input to the design generation.
114

115 VAEs consist of two differentiable (neural) networks, one encoding high-dimensional input data into a
116 reduced latent representation and the other decoding back the low-dimensional code into its high-
117 dimensional representation (Figure 1). The two networks are chained and trained to minimize a loss
118 function with a term representing the reconstruction error (MSE) between input data fed to the encoder and
119 the output data decoded by the decoder and a regularization term (Kullback-Leibler divergence). VAEs
120 lead to continuous and smooth latent spaces, which are especially advantageous for design space
121 exploration. By continuous and smooth latent space, we mean that the latent space is decoded onto original
122 space by continuous mappings (no local jumps) that are not noisy and mostly exhibit low-frequency
123 (smooth) variations across the latent space. From the standpoint of mathematical terminology, the mappings
124 are indeed continuous, but they are not necessarily smooth—i.e. they are not necessarily differentiable
125 everywhere—since they may use non-smooth activation functions like ReLU. Here, the concept of
126 smoothness should be understood qualitatively and does not refer to the differentiability concept of
127 smoothness.
128




Figure 1: General architecture of variational autoencoders. e_{θ} projects a high-dimensional input \mathbf{z} to a compressed representation \mathbf{x} , which the decoder d_{Φ} projects back to its original representation as best as possible. θ and Φ denote their respective network's weights which are optimized during training to minimize a composite loss function combining MSE (reconstruction term) and Kullback-Leibler divergence (which can be seen as a regularization term).

129 1.4 Deep generative modeling in design

130 The ability of VAEs to pack complex data distributions into continuous and low-dimensional latent spaces
 131 makes them particularly applicable to design applications, mostly thanks to the properties of their latent
 132 spaces. For example, Umetani [12] demonstrated that a 10-dimension latent representation may be found
 133 for cars based on a dataset of three-dimensional car models. The resulting latent space encompasses large
 134 variations and allows novice users to interactively design a car body. In a similar vein, but using image
 135 data, Burnap, Liu, Pan, Lee, Gonzalez, *et al.* [13] build a latent design space of two-dimensional automobile
 136 bodies. Further away from product design, Carter and Nielsen [14] use a variational autoencoder to build
 137 an intuitive design interface for fonts. This previous work demonstrates the usefulness of VAEs to build a
 138 low-dimensional design space based on observed real-world data.

139
 140 Generative adversarial networks (GANs) have similarly been shown to be good candidates for generative
 141 design applications, though their latent spaces are often harder to explore. Some of the most impressive
 142 results produced by GANs have been images, of human faces in particular [8], and there have been similarly
 143 striking advances for design applications. For example, generative adversarial networks have been used to
 144 generate voxel-based or point-based three-dimensional models of furniture, cars, and other objects [15]–
 145 [17]. GANs and their conditional variations like pix2pix [18] have also been used to generate building floor
 146 plans [19] and indoor furniture layouts [20], [21].

147
 148 This research differs from this previous work in that it establishes workflows, which may be applied
 149 systematically for design space exploration for which previously explored data is unavailable. While
 150 thousands of car designs may have been observed in the past, a dataset of structures designed for a particular
 151 site with specific loads is unlikely to exist. Perhaps closer to this research, Yumer, Asente, Mech, and Kara
 152 [22] showed how an autoencoder network may be used to ease the burden of exploring procedural models
 153 for non-expert users. However, their method does not incorporate performance in any way but focuses on
 154 shape features as a means to differentiate data. Conversely, the proposed method is geometry-agnostic and
 155 solely focuses on design parameters and objective function values, which ensures that it can be applied
 156 systematically for performance-driven design exploration.

157 2 Methodology

158 This section describes the workflow to build subspace representations of the high-performing regions.

159 In the following, a design option is specified by a n -dimensional vector of design variables $\mathbf{x} \in D \subset \mathbb{R}^n$.
160 The domain D formally defines the design space. In this research, D is a bounded domain
161 $\times_i [x_{i,min}, x_{i,max}] (i = 1, \dots, n)$. Each design may be evaluated by a performance metric $f(\mathbf{x})$, which is
162 computationally expensive to query and unknown except when computed for discrete samples. This
163 formalism corresponds to the definitions broadly adopted in most design space exploration research and is
164 useful for describing the algorithms below.

165 First, we introduce a performance-driven sampling algorithm used to efficiently collect a dataset
166 $\{(\mathbf{x}^{(i)}, f(\mathbf{x}^{(i)})) \mid i = 1, \dots, N\}$ of design samples. Then, we show how this dataset is used to train conditional
167 generative models and how these can be used for design exploration.

168 2.1 Performance-driven sampling

169 Training deep latent generative models requires significant amounts of data, which may be collected by
170 sampling through the design space. Because the explicit goal of such models is to uncover a reduced number
171 of latent variables that explain the distribution of the data, they are only effective when the data is non-
172 uniformly distributed. Sampled datasets produced by uniform sampling schemes are thus not adequate to
173 train a latent variable model. Instead, the datasets used for training need to be biased and mostly include
174 those design samples that present desirable attributes, which are assumed to be grouped in specific regions
175 of the design space. In practice, these attributes are measured by some objective function. For the spatial
176 truss example used to illustrate the proposed method (see Figure 5), the structural mass required to resist
177 imposed loads is the objective of interest. If multiple objectives are simultaneously considered, they may
178 be grouped in a composite objective function, an approach used successfully in previous related work [6],
179 or multi-objective sampling schemes may be considered altogether.

180
181 The performance bias may be introduced downstream: samples are then obtained through uniform sampling
182 schemes and filtered based on their objective values. If a fine enough sampling resolution may be achieved,
183 this scheme ensures uniform coverage of all high-performance regions. In practice, however, it is typically
184 computationally prohibitive to sample high-dimensional design spaces at a high resolution both because of
185 the curse of dimensionality and the slow simulations typically required in structural and architectural
186 design. An alternative is to sample using schemes that directly incorporate bias. Previous work [6], [23]
187 proposes using optimization histories to collect design samples with a performance bias. Optimization,
188 however, imposes a strong prejudice against regions in the design space that are suboptimal but still may
189 be of interest to designers. In addition, optimization algorithms oversample best-performance regions
190 before reaching convergence. This is true for stochastic optimization algorithms as well, albeit to a lesser
191 extent. Though useful, optimization methods are not designed for sampling.

192
193 Sampling is often used to build fast surrogate models substituting for complex and slow engineering
194 analyses, whether this sampling is accomplished through physical experiments as in early surrogate
195 modeling research [24] or using computational fluid dynamics [25] and thermal [26] simulations, and much
196 research has been devoted to devising sampling plans that yield surrogate models that are as accurate as
197 possible for a minimum number of samples [27]–[30]. In contrast to previous work, the proposed sampling
198 algorithm is designed to generate samples from pockets of high-performance more than it is geared toward
199 optimal model accuracy. In performance-driven design, low-performance regions are of little interest, and
200 model accuracy there is less important as a result. In other words, the performance distribution of the
201 samples matters more than the quality of the surrogate model they could be used to build because these
202 samples are generated to train unsupervised learning algorithms or building visualizations used to
203 understand, explore, or generate high-performance design options.

204
205 The proposed algorithm is a sequential, model-based sampling scheme, similar conceptually to Bayesian
206 optimization, that uses filtering to introduce performance bias. It starts by building an initial surrogate
207 model based on a limited number of samples evaluated using the true objective function. These samples are

208 used to build an initial surrogate model. This surrogate model is then used to evaluate another set of samples.
 209 These samples are filtered using an acceptance criterion which is conditioned on their performance as
 210 evaluated by the surrogate model. The filtered-in samples are in turn evaluated using the true objective
 211 function and added to the set of collected samples. Based on the samples collected so far, a new surrogate
 212 model is built, and the steps above are repeated a prescribed number of times.

213
 214 The introduced sampling scheme is non-deterministic and uses a tunable sigmoid-like gating function
 215 defined in Algorithm 1 to determine which samples to evaluate with the true objective function. The
 216 algorithm (see Algorithm 1 for the detailed pseudocode) starts with a low-resolution, unbiased sampling—
 217 such as Latin hypercube sampling—of the design space to build a dataset of n_{init} samples $D =$
 218 $\{(\mathbf{x}_i, y_i = f(\mathbf{x}_i)), i = 1, \dots, n_{init}\}$ where \mathbf{x}_i is the design vector of sample i and y_i is its corresponding
 219 score computed with f , the design performance function (calls to f are assumed to be slow). This initial
 220 dataset is used to build a surrogate model f^* that is hoped to approximate the function f as well as possible,
 221 but, given the limited sample size, it is expected to be flawed. Once the surrogate model is built, the design
 222 space is sampled again, but, this time around, the samples are evaluated using f^* instead. These evaluations
 223 are cheap and fast. Based on their predicted performance, designs are filtered using a sigmoid-like
 224 probabilistic gate with a user-specified growth rate g and a performance threshold p . Accepted designs are
 225 then evaluated using the true objective function f , and they are added to the dataset D . The augmented
 226 dataset is then used to build an updated surrogate model f^* , and this process is repeated $(n_{steps} - 1)$ times.
 227 It is worth noting that the growth rate is multiplied by a factor of $1 + \beta$ in each subsequent iteration, where
 228 $\beta \geq 0$ is usually small and allows to progressively increase the growth rate as the uncertainty of the
 229 surrogate model decreases.

Algorithm 1: Sequential performance-gated sampling

Input:

- $\Omega = \times_i [x_{i,min}, x_{i,max}] (i = 1, \dots, d) \subset \mathbb{R}^d$, the design space,
- $f: \Omega \rightarrow \mathbb{R}$, the true objective function,
- n_{init} , the initial number of samples,
- N , desired number of samples in addition to initial samples
- n_{steps} , the number of sampling steps,
- $p \in]0,1[$, the performance threshold,
- $g > 0$, the growth rate,
- $\beta \geq 0$, the increase rate for g

Output:

- Dataset $D = \{(\mathbf{x}_i, y_i = f(\mathbf{x}_i))\}$ of generated design samples

Initialization:

- Sample n_{init} initial designs in Ω using LH sampling.
- Evaluate designs using f to build dataset $D = \{(\mathbf{x}_i, y_i = f(\mathbf{x}_i)), i = 1, \dots, n_{init}\}$.
- Train surrogate model f^* based on D .

for i in $[0, \dots, n_{steps} - 1]$ **do**

- Initialize empty set $D_{step} \leftarrow \{\}$.
 - Set $g_{step} \leftarrow g * (1 + \beta)^i$.
 - **while** $|D_{step}| < \frac{N}{n_{steps}}$ **do**
 - sample $\frac{N}{n_{steps}}$ designs \mathbf{x}_j ($j = 1, \dots, \frac{N}{n_{steps}}$) in Ω using LH sampling,
 - evaluate design scores $y_j^* = f^*(\mathbf{x}_j)$ using surrogate model f^* ,
-

- compute p -values $p_j = \frac{\left| \left\{ y_j^* \leq y_k^*, k=1, \dots, \frac{N}{n_{steps}} \right\} \right| - 1}{N-1}$ for each design,
- compute acceptance probabilities $\alpha_j = \text{gate}_{g_{step}, p}(1 - p_j)$ for each design,
- accept design j with probability α_j , evaluate using true objective function f , and add to D_{step} :

$$D_{step} \leftarrow D_{step} \cup \left\{ (x_j, y_j = f(x_j)) \mid \text{Ber}(\alpha_j) = 1, j = 1, \dots, \frac{N}{n_{steps}} \right\}$$
, where Ber is the Bernoulli distribution.
- **end while**
- Add samples generated and accepted in this step to D : $D \leftarrow D \cup D_{step}$
- Train surrogate model f^* based on D .

end for

return D

Where

$$\text{gate}_{g,p}(x) = \frac{\sigma_{g,p}^*(x) - \sigma_{g,p}^*(0)}{\sigma_{g,p}^*(1) - \sigma_{g,p}^*(0)} \text{ with } \sigma_{g,p}^*(x) = \begin{cases} \sigma_g \left(\frac{x-p}{1-p} \right), & \text{if } x > p \\ \sigma_g \left(\frac{x}{p} - 1 \right), & \text{otherwise} \end{cases} \text{ and } \sigma_g(x) = \frac{1}{1+e^{-gx}}$$

230

231 The role of the probabilistic performance gate is to balance exploration and exploitation of the initial
 232 surrogate model as well as the ones built at every subsequent step. Since the initial surrogate model is built
 233 with a limited number of samples, it is not expected to be very accurate in many regions of the design space.
 234 However, it is used to evaluate the next batch of samples because it is extremely cheap to query compared
 235 to the true objective function.

236

237 If the sample filtering used a hard threshold and was completely deterministic, there would be a risk that,
 238 based on the incomplete picture provided by the surrogate model, only samples in a narrow area may be
 239 accepted. These would then be evaluated with the true objective function and added to the initial set of
 240 samples to train a new surrogate model, which, compared to the first one, will have only gained information
 241 about that narrow area. This may cause the next sampling steps to drill down (i.e. exploit the model) in that
 242 one area without exploring other wells of performance that may have been missed by the initial sampling.
 243 The probabilistic gating strategy solves this problem by allowing samples that are predicted to be
 244 performing worse than the specified performance threshold to be accepted, albeit with a lesser chance. This
 245 results in a wider exploration of the design space and smaller odds that good regions are missed. How much
 246 weight is given to exploration versus exploitation is controlled by the growth rate g , as is explained in the
 247 previous section: the larger the growth rate, the more onus is put on exploitation, and vice-versa. In fact, if
 248 $g \rightarrow \infty$, the gate is a step function and the filtering is fully deterministic.

249

250 As the algorithm progresses, more samples are collected, and the quality of the surrogate model improves
 251 and hopefully identifies regions of good performance correctly. When it does, it makes sense to put more
 252 weight on exploitation. In other words, the first sampling steps require more exploration while the later
 253 ones demand more exploitation. To achieve this, an additional parameter β introduced and is used to
 254 increase the growth rate at each new step by multiplying its previous value by $1 + \beta$. For example, for $\beta =$
 255 0.2 , a starting growth rate of 5 translates into a growth rate of about 31 after 10 iterations.

256

257 The sampling algorithm depends on 6 hyperparameters: the number of initial samples n_{init} , the minimum
 258 number of desired samples, the number of sampling steps, the percentile threshold, the growth rate g , and
 259 the parameter β . The first four hyperparameters are readily interpretable and may be chosen with reasonable
 260 engineering judgment. The number of initial samples may be increased or decreased depending on the

261 complexity of the objective function for which designers generally have a good working intuition. The
262 choice of the number of desired samples would likely be based on the requirements of the data-driven
263 method they may be used (in our case, it is reasonable to assume that a deep neural network will require
264 thousands of samples to effectively train).

265
266 Choosing the number of sampling steps essentially amounts to choosing how many surrogate models will
267 be built throughout the sampling process. Choosing as many steps as the desired number of samples would
268 mean building a new surrogate model for every new sampled design, surely an inefficient strategy since
269 any single point does not contribute so much new information as to warrant training a new model, so
270 keeping the number of steps relatively low is reasonable. This is confirmed by experimental results
271 discussed in Section 3.1 which demonstrate that the quality of collected samples quickly improves only
272 after a few steps. The number of sampling steps and the minimum number of desired samples constitute
273 rigid stopping criteria that are commonly used for sampling algorithms. Convergent behavior is clearly
274 observed in experimental results, especially at the distribution level (Figure 7), but basing a stopping
275 criterion on a relative metric like relative convergence of sample mean would not be practical because the
276 algorithm’s primary objective is to provide a desired number of high-quality samples. The percentile
277 threshold explicitly controls the targeted percentile of the sampled designs compared to the performance
278 distribution of a uniformly sampled population of designs.

279
280 The influence of the last two hyperparameters is also significant as shown experimentally in Section 3.2.
281 The growth rate g controls how lax the performance filtering is at each step of the sampling process.
282 Because trust in the surrogate model is low at the beginning of sampling, using a relatively low growth rate
283 during the first steps of the algorithm prevents it from getting stuck in a limited region of the design space.
284 However, as the algorithm progresses, the surrogate model improves and the stringency of filtering can be
285 increased. To do so, the initial growth rate is progressively increased at each sampling step with a rate
286 controlled by the parameter β . Figure 8 shows the influence of g and β on the algorithm’s progress on the
287 long-span roof case study used to illustrate this research.

288
289 One of the objectives of the proposed sampling algorithm is to be able to control the diversity and the
290 quality (performance) of the generated samples. The trade-off between sample diversity and quality is
291 usually hard to navigate with existing sampling algorithms, but, here, it can be explicitly mediated by the
292 growth rate and the performance threshold. It is also worth noting that, while the probabilistic gating looks
293 to ensure a good balance of exploitation and exploration at each step, additional unfiltered LH samples may
294 be collected at each step to further increase the odds of widespread coverage.

295 2.2 Performance-conditioned VAE

296 This research uses a conditional VAE instead of a standard based on two observations. First, conditioning
297 the generative model on performance scores improves the model by helping it encode and decode designs
298 more easily. Instead of having to learn projections that work for all sampled designs at once, the
299 performance-conditioned VAE (Figure 2) can adapt its mappings to different performance values: this
300 facilitates disentangling performance contours that may otherwise not be discernable. Second, the
301 performance condition can be used after the model is trained to control the decoder’s output. This is
302 particularly useful when designers are looking to trade-off performance for other qualitative design
303 attributes: they can start by generating high-performance design options and can progressively relax their
304 performance condition to explore suboptimal solutions that may fulfill unformulated objectives.




Figure 2: Performance-conditioned VAE. The performance condition is fed to both the encoder and the decoder.

305 2.2.1 Performance condition: the p -value

306 The performance-conditioned VAE (PVAE) is trained on a dataset of design samples (represented by their
 307 design vectors) and their corresponding performance scores evaluated based on engineering simulation
 308 results. However, absolute performance values are not directly used. Instead, we reuse the concept of the
 309 p -value, which was introduced in the last chapter and can be seen as a normalized rank of sorts. Given a set
 310 of N design samples $\{(\mathbf{x}_i, y_i), i = 1, \dots, N\}$ with design vectors \mathbf{x}_i and scores y_i , the p -value of design i is
 311 computed as $p_i = \frac{|\{y_j \leq y_i, j=1, \dots, N\}| - 1}{N - 1}$. The definition of the p -value is slightly modified compared to the
 312 definition in the sampling algorithm such that lower p -values represent designs with lower performance
 313 scores. Again, using this strategy allows to map the scores of all designs used for training the PVAE, which
 314 can vary widely in magnitude, to the fixed interval $[0, 1]$. In addition, compared to the absolute performance
 315 scores, the design p -values are distributed evenly on the unit segment.

316 Using the p -value as the performance condition also makes exploiting the PVAE more intuitive after it is
 317 trained: the performance condition can then be seen as a simple knob or slider that can be tuned from 0 to
 318 1 to generate designs with lower or higher objective values.

319 2.3 Latent space dimensionality

320 One of the key decisions to make when building a VAE is to choose the dimensionality of the latent space
 321 onto which data will be projected. Latent spaces with fewer dimensions generally extract more meaningful
 322 directions of change from the data, but they also incur a larger compression loss, which means that it is
 323 harder to reconstruct input data from its latent representations. In the context of design space exploration,
 324 this means that lower-dimensional latent spaces are likely to generate less diverse design candidates. On
 325 the other hand, the lower the dimensionality of the latent space, the easier it is to explore.

326
 327 In practice, designers may explore the trade-off between ease of exploration by varying the dimensionality
 328 of the latent space. This research chooses to build two-dimensional latent spaces, which heavily compress
 329 input data, because it allows for the representation of complex design subspaces as two-dimensional
 330 landscapes that are easy to explore and provide designers a global view of the design candidates meeting a
 331 prescribed performance level. In the long-span roof case study used here, the latent spaces built displayed
 332 good diversity such that the benefits brought about by their low-dimensionality outweighed the diversity
 333 trade-off. With that said, the proposed method can be readily applied to latent spaces with higher dimensions
 334 that are nonetheless significantly easier to explore than the original.

335 2.4 Decoding the latent space

336 Once the PVAE is trained, we can detach the decoder and use it to generate designs by moving through the
 337 latent space, whose directions are effectively synthetic and reduced design variables. With a two-
 338 dimensional latent space, this navigation can be done through a computer screen by moving a point across
 339 a two-dimensional slider. Because the latent space dimensionality is small enough, it is also possible to
 340 precompute many design candidates and evaluate their performance to build two-dimensional performance
 341 maps (Figure 3) that can be overlaid directly on the latent space.
 342




Figure 3: Thanks to the low-dimensionality of the latent space, explorable performance maps can be precomputed by decoding a grid of designs in the latent space and computing their performance scores.

343 These performance maps can then be navigated with a cursor whose location in the latent space is decoded
 344 to generate designs (Figure 4). This allows designers to intently explore diverse, high-performing structures.
 345 The p -value constitutes an additional control that allows designers to relax performance requirements and
 346 modify the performance landscape to explore other design options.
 347




Figure 4: Navigating the latent performance map. A designer can move through the performance map—adjusting it by modifying the p -value—with a cursor, whose location is decoded to generate its corresponding design (represented here by an abstract shape with arrows suggesting the morphing that occurs as the cursor moves across the performance map).

348 3 Results

349 To demonstrate the effectiveness of the algorithm on a high-dimensional design space, we use a long-span
 350 roof example controlled by 36 design variables (Figure 5). The roof geometry is defined by two surfaces,

351 whose control points can be moved vertically by to modify the depth of the space truss generated between
 352 them. Space trusses or space frames are often used in structural design to span long distances: in addition
 353 to their inherent structural performance, they can be easily shaped for greater structural efficiency or to
 354 conform to an architectural designer’s vision or both.
 355 Shaping the space truss in this example is thus an exercise in both structural and architectural design since
 356 its shape will affect its structural performance, its architectural form, and the space it spans. The proposed
 357 space truss intentionally presents a substantial cantilever to exacerbate the architectural and structural
 358 impact of the design variables.




Figure 5: Long span roof example: summary of initial geometry, design variables, and performance measure. Design variables 1 through 18 control the bottom surface and variables 19 through 36 control the top surface.

359 Controlling the shape of the space truss with two NURBS surfaces (each of degree 2) allows a level of
 360 control that can be arbitrarily increased or decreased by introducing additional control points without
 361 needing to parametrize the location of every single node in the space truss. For this example, we deliberately
 362 use surfaces with a moderate number of control points (25 per surface) to generate a design space that
 363 contains diverse and potentially surprising solutions. Because the geometry is constrained to be symmetric
 364 and only the z-coordinates of the control points are in play (the roof footprint is fixed), this results in a total
 365 of 36 design variables (18 per surface). By parametric design standards, this is a high-dimensional design
 366 space, which is likely to contain good solutions both aesthetically and structurally as well as grotesque ones.

367 3.1 Structural modeling and performance metric

368 The spatial truss is connected to 4 pin supports by a total of 16 columns and is subject to a load of 50 psf
 369 or 2.39 kN/m² applied as individual downward point loads on each node of the space frame (17.3 kN per
 370 node) in addition to its self-weight. The structure is modeled as a perfect truss (elements only deform
 371 axially) built with steel (S355; $\sigma_{yield} = 355 \text{ MPa}$) circular hollow sections. Each truss member is sized
 372 automatically using the cross-section optimizer of Karamba [31], a structural analysis plug-in. Based on a

373 user-provided catalog of cross-sections, the cross-section optimizer starts by assigning each member with
 374 the smallest cross-section possible and analyzes the structure—using linear static FEA—accordingly to
 375 obtain its internal forces and displacements. These are then used to resize each member by searching the
 376 smallest cross-section possible among the specified catalog that satisfies the strength requirements of the
 377 Eurocode EN-1993 (Design of Steel Structures). Because this modifies the self-weight of the structure, the
 378 structure is analyzed again to ensure that the cross-sections can resist the updated loads and that the structure
 379 does not deform excessively. If these requirements are not met, the process is repeated until they are. For
 380 this example, the optimizer sizes the design candidates by picking sections from a catalog of 46 circular
 381 hollow tubes with diameters ranging from 10 to 100 cm—very large sections are required to be able to
 382 characterize poor-performing designs well—in increments of 2 cm and wall thicknesses equal to 5% of the
 383 tube diameters or 20 mm, whichever is smallest. Members are sized against yielding (elastic design) with
 384 a safety factor of 1.5 and against buckling with a safety factor of 3. The structure is also subject to a
 385 displacement limit of $L/300$ ($=13.3$ cm), albeit for a smaller load of 33 psf or 1.58 kN/m². The sizing process
 386 is used to automatically evaluate design candidates to the amount of material they require to resist the
 387 imposed load, which are normalized by the structure’s footprint (1600 m²) for easier interpretation.

388 3.2 Performance-driven sampling

389 The proposed performance-driven algorithm is run on the space truss design space for 10 steps with $n_{init} =$
 390 1000 , $N = 5000$, a performance threshold $p = 0.8$, a growth rate $g = 2$, and a growth increase rate $\beta =$
 391 0.3 . The surrogate model is a support vector regressor with a Gaussian kernel whose parameter γ and C are
 392 searched among 5 log-spaced values between 10^3 and 10^{-3} using 3-fold cross-validation.

393 In total, 6771 samples (including 1000 initial LH samples) are collected at the end of sampling. Figure 6
 394 shows a kernel density estimation (Gaussian kernel; bandwidth=2) of the structural mass—the performance
 395 metric of interest—of the samples resulting from the performance-driven sampling and compares it against
 396 the same estimation obtained for 6771 Latin hypercube samples, which can be seen as an estimate of the
 397 true performance distribution of the design space, and the 1000 LH samples collected to initialize the
 398 algorithm. Such comparison shows the benefits brought by the performance-driven algorithm compared to
 399 a standard sampling scheme in terms of sample quality. The proposed method leads to a much denser
 400 sampling of designs with higher performance as the shift in density toward higher-performing scores (low
 401 mass) for the performance-driven samples demonstrate.




Figure 6: Kernel density estimations of 6771 performance-driven samples scores vs. 6771 and 1000 LH samples scores with strip plot of the scores of the performance-driven samples

402 To assess the effectiveness of the algorithm on this high-dimensional example, we can also look at the
 403 evolution of the sample quality. Figure 7 shows two graphs that shine a light on the algorithm’s progress
 404 and confirm its value for sampling high-dimensional design spaces. The first (on the left) shows how the
 405 performance density of the samples collected at each step gradually moves to the left and peaks higher for
 406 lower i.e. better mass values. Particularly noteworthy is the big jump between step 0 and step 1: this shows
 407 that even a surrogate model relying on little data can significantly help sample much better design

408 candidates. The second graph (right) confirms this trend: the mean sample score progressively decreases
 409 and the samples (each represented by a single dot) cluster downward.




Figure 7: Performance of samples collected for each sampling step. Each point represents a single sample, its vertical position indicates the sample performance, and its horizontal position indicates when it was collected (jitter is introduced to improve readability). The mean performance for each step is shown by the solid line. The performance axis is bounded to $[40, 140]$, and samples with larger performance values are crowded at the top of the graph.

410 The results discussed and shown above were produced for a specific set of growth rate and β parameters,
 411 and, as was already discussed in 2.1, these two parameters can alter the result of the sampling process in
 412 ways that are worth highlighting. By running the performance-driven algorithm on the long-span roof
 413 example for 9 combinations of these two parameters (and using the values listed at the beginning of this
 414 section for the other hyperparameters) and tracking the mean of the samples collected at each step, we show
 415 that this intuition is confirmed experimentally (Figure 8). Choosing an initial high-value for the growth rate
 416 means that the in-step sample mean decreases drastically after the first step with only marginal
 417 improvements in subsequent steps. While this yields successive sampling steps with lower sample means,
 418 this may indicate that the algorithm is drilling down on a given region and ignoring other promising ones,
 419 such that one should not assume that a sudden drop followed by a plateau is necessarily best. A low growth
 420 rate and a low β also make the algorithm stall because it is simply not strict enough when filtering proposed
 421 LH samples. In summary, too high a growth rate does not yield enough exploration and low values for both
 422 parameters yield less exploitation. As Figure 8 shows, a choice of a low initial growth rate (1-5) and a
 423 moderate value of β (0.1-0.5) strikes a good balance between exploration and exploitation, but regardless
 424 of g and β , the algorithm yields samples with much better performance than the initial LH step.




Figure 8: Influence of growth rate and β parameters on the mean performance of samples collected at each step.

425 3.3 Performance-conditioned VAE




Figure 9: Architecture of performance-conditioned VAE used for the long span roof example.

426 The performance-conditioned variational autoencoder with 39,516 trainable parameters diagrammed in
 427 Figure 9 is trained based on the data collected by the performance-driven algorithm detailed in the previous

428 subsection. Before training, design variables used are normalized from $[-3,3]$ to $[0,1]$. It is worth noting
 429 that the final activation of the network is a rectified linear unit activation (ReLU) which, in contrast to a
 430 sigmoid activation, does not restrict the output of the decoder to the $[0,1]$ domain, technically allowing
 431 decoded samples to have out-of-bound design variables. In practice, this accelerates training, and the
 432 bounds of the design variables are softly integrated by the model through learning.
 433




Figure 10: Evolution of reconstruction loss and Kullback-Leibler divergence of the performance-conditioned VAE during training.

434 The PVAE is trained for 1000 epochs with 5755 samples using the RMSprop gradient descent algorithm
 435 [32] to minimize the VAE loss, itself obtained by averaging the MSE reconstruction loss and the Kullback-
 436 Leibler (KL) divergence between the encoded samples and the normal distribution. The PVAE is validated
 437 with 1016 samples. Figure 10 shows the evolution of the loss components (MSE and KL divergence) for
 438 the training and validation sets as training progresses: the MSE is expectedly minimized, but the KL
 439 divergence progressively increases to an asymptotic value. The latter result may seem a little
 440 counterintuitive given the PVAE is trained to minimize the average of both the MSE and the KL divergence.
 441 However, the KL divergence should be understood as a regularization term, which is added to the MSE to
 442 ensure that the latent space has a continuous and smooth structure. Without the KL divergence, it would be
 443 possible to reduce the reconstruction loss even more, but that would be at the expense of the latent space
 444 continuity. Conversely, minimizing the KL divergence would negatively impact the reconstruction MSE.
 445 From that perspective, the asymptotically increasing behavior of the KL divergence is only a reflection of
 446 the trade-off between the two terms of the VAE loss.

447 3.4 Latent space encoding

448 To assess the impact of the KL divergence on the structure of the latent space, it is useful to look at the
 449 projections of the high-dimensional training data points onto the latent space. Figure 11 shows how the
 450 5755 training samples are encoded onto the 2D latent space by the encoder, organized according to a slightly
 451 asymmetrical normal-like distribution. This structure is a direct consequence of the use of a VAE (with its
 452 KL divergence term in the loss function) over the use a regular autoencoder and is indicative of a latent
 453 space that is continuous and smooth, both important features for design exploration.




Figure 11: Projection of training samples onto the latent space by the encoder after training.

454 Figure 11 also shows that samples with different performance scores are superimposed in the latent space.
 455 This reflects the role that the performance condition plays: it essentially allows the PVAE encoder to use
 456 the same real estate in the latent space to pack multiple strata of the high-dimensional design space.
 457 Similarly, the performance condition allows the decoder to unfold the latent space into different manifolds
 458 in the high-dimensional design space.

459 3.5 Latent space decoding: an atlas of design subspaces

460 The PVAE does not only provide access to a single latent space but rather an atlas of performance maps
 461 that designers can sift through by adapting the performance condition. This is particularly useful because it
 462 allows designers to investigate potentially interesting trade-offs between performance and other intangible
 463 design factors. Figure 12 shows the performance maps obtained by decoding and evaluating the latent space
 464 for different p -values.

465 The maps show that the p -value (the performance condition) has the intended effect on the performance of
 466 the latent spaces: the performance scores increase—they get worse in this case—as the p -value is increased.
 467 The performance map for $p = 0$ includes designs with excellent performance scores, many of them
 468 hovering right above 40 kg/m^2 , i.e. the best objective function value in the design dataset used for training
 469 the PVAE and obtained using performance-driven sampling. The latent spaces corresponding to larger p -
 470 values unsurprisingly contain only slightly worse designs that demonstrate the potential of the proposed
 471 approach to explore design candidates that are slightly suboptimal but qualitatively better. Interestingly, the
 472 performance contours are not smooth everywhere, even though the decoder mapping is continuous, because
 473 the objective function, the structural mass required to support the imposed loads, is itself non-smooth with
 474 respect to the original design variables. This is particularly salient for $p = 1$ where the weight of some
 475 designs in the latent space skyrocket: these are designs where the space truss has areas with small structural
 476 depths, resulting in large internal axial forces and section sizes.




Figure 12: Objective function heat maps in the latent spaces learned by the performance-conditioned VAE for different p -values.

477 The consistent and anticipated link between the performance condition displayed in Figure 12 is even better
 478 illustrated by Figure 13, which shows how the performance of 100 random samples in the latent space
 479 changes as the p -value is increased. Figure 13 also highlights an interesting trend: for p -values smaller than
 480 0, performance continues to improve (decrease) until around $p = -0.3$. Even though the PVAE is not

481 trained with any samples associated with negative p -values, it learns to extrapolate beyond $p = 0$. These
 482 extrapolated trends broadly correspond to an increase of the space truss depth beyond the original bounds
 483 of the design space. As the depth of the space truss increases further for values lower than -0.3 , the negative
 484 impact of the increased length of the structural members starts to outweigh the benefits of the larger
 485 structural depth, and the structural mass increases slightly.




Figure 13: Evolution of the performance score of 100 random designs sampled from the latent space $([-3,3] \times [-3,3])$ (left) as the performance condition is increased. For values of p lower than 0, the VAE successfully extrapolates trends linked with an improvement (decrease) of the performance score beyond the original bounds of the design space to make the space truss deeper.

486 To understand the structure of the latent space learned by the PVAE, it is useful to look at how each point
 487 in the 2D latent space is mapped to each design variable of the original, high-dimensional design space.
 488 Figure 14 shows the values that each of the original design variables takes in different locations of the latent
 489 space for $p = 0$. It illustrates the non-linearity of the latent space, which allows to pack more complex
 490 distributions of designs than linear encoding techniques like principal component analysis. It also
 491 demonstrates that the latent space is smooth and that any exploration path in the latent space continuously
 492 morphs or interpolate between designs.
 493




Figure 14: Mapping from latent space $[-3,3]^2$ to original variables for $p = 0$. Numbers indicate the indices of each design variable as defined in Figure 5.

494 An interesting question is whether we can derive any meaning from the latent directions. Sometimes,
 495 dimensionality reduction techniques yield lower-dimensional representations with directions to which
 496 humans can ascribe meaning *a posteriori*. For example, research on lighting control has shown human-
 497 derived criteria for sensor lighting control can be extracted using PCA [33]. Interpreting the meaning of the
 498 directions of a learned latent space is typically easier for linear dimensionality reduction methods like PCA.
 499 Because the PVAE encoding and decoding are nonlinear, interpreting each latent direction globally is not
 500 as straightforward. Nevertheless, despite its apparent complexity, the latent space is intuitive to explore
 501 because it can be rendered at once on a 2D computer screen.

502 Of course, the latent space maps also change as the performance condition is modified: Figure 15 shows
 503 the evolution of the design variable maps for different p -values. The evolution of the latent space is smooth
 504 and shows that any individual design in the latent space is continuously morphed to match a prescribed

505 performance condition. Figure 15 also highlights how the PVAE extrapolates trends for p -values under 0
 506 or above 1.




Figure 15: Evolution of latent space maps of original variables for increasing (left to right) values of the p -value.

507 In addition to the performance and variable maps, it is also important to look at the actual design geometries
 508 corresponding to different points of the latent space to understand the design subspace learned by the PVAE.
 509 Figure 16 shows 36 designs decoded from a regular grid of 6-by-6 samples in the latent space (for $p = 0$)
 510 and their respective performance scores. These designs demonstrate that the latent space contains
 511 geometrically and visually diverse. All of these designs perform very well, and the low-dimensional design
 512 subspace can be exhaustively searched by designers to find high-performance options with different
 513 qualitative properties.




Figure 16: Renderings of a grid of designs in the latent space for $p = 0$ and their corresponding structural mass [kg/m^2]. The vertical and horizontal axes indicate the position of each design in the latent space.

514 While Figure 16 offers a snapshot of a grid of designs contained in the latent space for a specific value of
 515 the performance condition, the influence of the p -value on design geometry and performance is even better
 516 highlighted by looking at individual locations in the latent space with varying performance conditions.
 517 Figure 17 shows the morphological evolution of 4 designs at 4 different locations of the latent space as the
 518 performance condition is increased, and it illustrates the impact of the performance condition on both design
 519 geometry and performance: generally, designs in the latent space with $p = 1$ are essentially shallower
 520 versions of the ones in the latent space with $p = 0$. It is noteworthy that much of the latent space with $p =$
 521 1 still contains many well-performing designs (see Figure 12) with much shallower morphologies. This
 522 demonstrates once again the potential of the proposed approach to explore design options that are sub-
 523 optimal quantitatively but potentially better fits for a host of other reasons.
 524 These results confirm that the performance condition preserves most of the broad design characteristics of
 525 each design and mostly participates in making the space truss shallower. For experienced structural
 526 designers, this result is not surprising: structural depth in areas with high-bending moments, in this case
 527 over the supports, is almost always conducive to greater structural efficiency. However, the fact that the
 528 PVAE inferred it from data is non-trivial, and the PVAE does not simply operate by moving up the lower
 529 surface and moving down the upper surface. For example, some areas of the space truss become shallower
 530 more quickly than others.




Figure 17: Morphological and performance evolution of 4 designs (right) in the latent space (left) as p increases from 0 to 1. Each column corresponds to a single point in the latent space (the location of which is shown at the top). As the performance condition increases, so does the structural mass, given in kg/m^2 .

531 In addition, Figure 17 further shows the usefulness of the performance condition for exploring non-optimal
 532 designs. For example, the second and fourth columns of designs show that it can be used to dramatically
 533 change the geometry of near-optimal designs while remaining in acceptable territories from a performance
 534 perspective.

535 3.6 Latent space navigation

536 As the last section shows, it is possible to build visualizations that summarize the latent space and show
 537 both the geometric diversity and the performance of the designs it contains. Such visualizations offer great
 538 immediate snapshots of the high-performance regions compressed by the PVAE. However, the PVAE needs
 539 not be constrained to the production of static visualizations but can also be explored interactively. This is
 540 particularly easy given the 2-dimensionality of the latent space, and designers can navigate the latent space
 541 using a simple interface with a three-dimensional view showing the current design geometry, an interactive
 542 performance map that designers can move through using with the mouse cursor, and a knob, slider, or even
 543 text field to adjust the performance condition. Figure 18 shows a prototype of such an interface with the
 544 latent space represented by a 2.5D surface and an additional parallel coordinate plot indicating the decoded
 545 original design vector.




Figure 18: Prototype interface for latent space exploration.

546 Figure 19 shows a potential exploration path that a designer may take through the latent space in such an
 547 interface and the designs that would be explored along such a path. The visualization highlights how
 548 navigating through the latent space yields diverse high-performing designs, and the parallel coordinate plot
 549 in particular further shows the smoothness and high nonlinearity of the latent space. In addition, it shows
 550 the influence of the tunable performance condition. Compared to a manual exploration of the original design
 551 space, which would require sequentially adjusting 36 sliders or knobs, this mode of exploration allows
 552 designers to generate design variations by controlling only 3 variables, the first of which, the performance
 553 condition, allows them to tighten or relax their performance requirements.




Figure 19: Exploration path in the synthetic latent design space. This visualization highlights a specific exploration path on the performance landscape of the latent design space. Renderings on the right display designs (accompanied by their performance scores in kg/m²) on the exploration path for $p = 0$ and $p = 1$. The parallel coordinate plot on the left shows how the original design variables change as one navigates along the S-shaped exploration path for $p = 0$, with colors indicating the corresponding location of each of the designs on the path.

555 **4 Conclusion**

556 This research contributes a method to build and train deep, performance-conditioned latent variable models
557 that pack complex design spaces into continuous, smooth, low-dimensional design subspaces. Design
558 subspace learning offers a new paradigm for performance-informed design exploration that is neither
559 optimization nor random or undirected and that provides a navigable cartography of otherwise unwieldy
560 design spaces.

561 *4.1 Future Work*

562 There are many interesting and potentially impactful directions for future work related to this research.
563 First, it would be compelling to adapt design subspace learning to rule-based design spaces which have a
564 variable number of parameters and are notoriously hard to control. Second, it would be powerful to illustrate
565 the use of the contribution in multi-objective design contexts, where structural considerations might conflict
566 with other performance objectives typical to architecture, such as daylight autonomy or energy use
567 intensity. In some ways, this can be straightforward: the research presented here can be directly on custom
568 composite objective functions that combine and weigh multiple and potentially divergent performance
569 metrics in a single score, though there are probably more nuanced ways to marry the proposed method with
570 existing multi-objective design approaches. Finally, this research discusses ways low-dimensional design
571 subspaces may be integrated simply as explorable and interactive maps in design tools. Previous work has
572 shown that the nature of the tools we use influences the way we design and that better tools, with more
573 integrated and interactive interfaces, yield better design outcomes [34], [35]. Future work will be devoted
574 to further studying how the interfaces proposed here impact design outcomes in user studies compared to
575 undirected design exploration.

576 *4.2 Concluding remarks*

577 Design subspace learning offers a new paradigm for performance-informed design exploration that is
578 neither optimization nor undirected search and that provides a navigable cartography of otherwise unwieldy
579 design spaces. Design subspace learning provides ways to explore more design solutions that perform well
580 and explicitly gives designers control over how much performance matters and allows them to negotiate
581 the tension between functional requirements and intangible human factors. It can be used to power
582 intelligent interfaces that foster the exploration and discovery of high-performing designs. Because these
583 interfaces can act as a more natural and intuitive layer of understanding between human designers and
584 complex design spaces, design subspace learning overcomes several important and fundamental limitations
585 to many existing computational design methods and has the potential to broaden the adoption of
586 performance-informed design processes.

587 **Acknowledgments**

588 This material is based upon work supported by the National Science Foundation under Grant No. 1854833.

589 **References**

- 590 [1] T. Kohonen, "Self-organized formation of topologically correct feature maps," *Biological Cybernetics*, vol.
591 43, no. 1, pp. 59–69, Jan. 1982, doi: 10.1007/BF00337288.
- 592 [2] T. Kohonen, *Self-Organizing Maps*. Berlin, Germany: Springer-Verlag Berlin Heidelberg, 2001, ISBN: 978-
593 3-642-56927-2.
- 594 [3] J. Harding, "Dimensionality Reduction for Parametric Design Exploration," in *Advances in Architectural*
595 *Geometry*, 2016, pp. 274–286, doi: 10.3218/3778-4.
- 596 [4] E. Fuchkina, S. Schneider, S. Bertel, and I. Osintseva, "Design Space Exploration Framework: A modular

- 597 approach to flexibly explore large sets of design variants of parametric models within a single environment,”
598 in *Proceedings of the 36th eCAADe Conference*, A. Kepczynska-Walczak and S. Bialkowski, Eds. Lodz,
599 Poland, 2018, pp. 367–376.
- 600 [5] L. Fuhrmann, V. Moosavi, P. O. Ohlbrock, and P. Dacunto, “Data-Driven Design: Exploring new Structural
601 Forms using Machine Learning and Graphic Statics,” Sep. 2018, Accessed: Jun. 15, 2020. [Online]. Available:
602 <http://arxiv.org/abs/1809.08660>.
- 603 [6] N. C. Brown and C. T. Mueller, “Design variable analysis and generation for performance-based parametric
604 modeling in architecture,” *International Journal of Architectural Computing*, vol. 17, no. 1, pp. 36–52, Mar.
605 2019, doi: 10.1177/1478077118799491.
- 606 [7] Z. X. Conti and S. Kaijima, “Enabling Inference in Performance-Driven Design Exploration,” in *Humanizing*
607 *Digital Reality*, 2017, pp. 177–188, doi: 10.1007/978-981-10-6611-5.
- 608 [8] T. Karras, T. Aila, S. Laine, and J. Lehtinen, “Progressive Growing of GANs for Improved Quality, Stability,
609 and Variation,” *arXiv preprint arXiv:1710.10196*, Oct. 2017.
- 610 [9] D. P. Kingma and M. Welling, “Auto-encoding variational bayes,” in *2nd International Conference on*
611 *Learning Representations (ICLR), Conference Track Proceedings*, 2014, Accessed: Feb. 04, 2020. [Online].
612 Available: <http://arxiv.org/abs/1312.6114>.
- 613 [10] D. J. Rezende, S. Mohamed, and D. Wierstra, “Stochastic backpropagation and approximate inference in deep
614 generative models,” in *Proceeding of the 31st International Conference on Machine Learning*, 2014, pp.
615 1278–1286, Accessed: Jan. 24, 2020. [Online]. Available: <http://proceedings.mlr.press/v32/rezende14.pdf>.
- 616 [11] K. Sohn, H. Lee, and X. Yan, “Learning Structured Output Representation using Deep Conditional Generative
617 Models,” in *NIPS’15: Proceedings of the 28th International Conference on Neural Information Processing*
618 *Systems*, C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, Eds. Montreal, Canada: MIT
619 Press, 2015, pp. 3483–3491.
- 620 [12] N. Umetani and Nobuyuki, “Exploring generative 3D shapes using autoencoder networks,” in *SA ’17:*
621 *SIGGRAPH Asia 2017 Technical Briefs*, 2017, pp. 1–4, doi: 10.1145/3145749.3145758.
- 622 [13] A. Burnap, Y. Liu, Y. Pan, H. Lee, R. Gonzalez, and P. Y. Papalambros, “Estimating and Exploring the
623 Product Form Design Space Using Deep Generative Models,” in *Proceedings of the ASME 2016 International*
624 *Design Engineering Technical Conferences and Computers and Information in Engineering Conference*,
625 2016, pp. 1–13, doi: 10.1115/DETC2016-60091.
- 626 [14] S. Carter and M. Nielsen, “Using Artificial Intelligence to Augment Human Intelligence,” *Distill*, Dec. 2017,
627 doi: 10.23915/distill.00009.
- 628 [15] J. Wu, C. Zhang, T. Xue, W. T. Freeman, and J. B. Tenenbaum, “Learning a Probabilistic Latent Space of
629 Object Shapes via 3D Generative-Adversarial Modeling,” in *30th Conference on Neural Information*
630 *Processing Systems (NIPS 2016)*, 2016, pp. 82–90.
- 631 [16] P. Achlioptas, O. Diamanti, I. Mitliagkas, and L. Guibas, “Learning Representations and Generative Models
632 for 3D Point Clouds,” in *Proceedings of the 35th International Conference on Machine Learning*, 2018, pp.
633 40–49.
- 634 [17] J. Wu, Y. Wang, T. Xue, X. Sun, W. T. Freeman, and J. B. Tenenbaum, “MarrNet: 3D shape reconstruction
635 via 2.5D sketches,” in *Proceedings of the 31st International Conference on Neural Information Processing*
636 *Systems*, 2017, pp. 540–550.
- 637 [18] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, “Image-to-Image Translation with Conditional Adversarial
638 Networks,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, Nov. 2017,
639 pp. 1125–1134, doi: 10.1109/CVPR.2017.632.

- 640 [19] S. Chaillou, “ArchiGAN: a Generative Stack for Apartment Building Design,” Harvard Graduate School of
641 Design, 2019.
- 642 [20] D. Ritchie, K. Wang, and Y. Lin, “Fast and Flexible Indoor Scene Synthesis via Deep Convolutional
643 Generative Models,” in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*,
644 Nov. 2019, pp. 6175–6183, doi: 10.1109/CVPR.2019.00634.
- 645 [21] K. Wang, M. Savva, A. X. Chang, and D. Ritchie, “Deep convolutional priors for indoor scene synthesis,”
646 *ACM Transactions on Graphics*, vol. 37, no. 4, pp. 1–14, Jul. 2018, doi: 10.1145/3197517.3201362.
- 647 [22] M. E. Yumer, P. Asente, R. Mech, and L. B. Kara, “Procedural Modeling Using Autoencoder Networks,”
648 *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology - UIST '15*, pp.
649 109–118, 2015, doi: 10.1145/2807442.2807448.
- 650 [23] T. Wortmann and T. Schroepfer, “From Optimization to Performance-Informed Design,” in *SIMAUD '19:*
651 *Proceedings of the Symposium on Simulation for Architecture and Urban Design*, 2019, pp. 1–8.
- 652 [24] G. E. P. Box and K. B. Wilson, “On the Experimental Attainment of Optimum Conditions,” *Journal of the*
653 *Royal Statistical Society. Series B (Methodological)*, vol. 13, Wiley, pp. 1–45, 1951, doi: 10.2307/2983966.
- 654 [25] S. A. Renganathan, R. M. and, and J. Ahuja, “Enhanced data efficiency using deep neural networks and
655 Gaussian processes for aerodynamic design optimization,” *arXiv*, Aug. 2020, Accessed: Jan. 31, 2021.
656 [Online]. Available: <http://arxiv.org/abs/2008.06731>.
- 657 [26] H. Ma, X. Hu, Y. Zhang, N. Thuerey, and O. J. Haidn, “A Combined Data-driven and Physics-driven Method
658 for Steady Heat Conduction Prediction using Deep Convolutional Neural Networks,” *arXiv*, May 2020,
659 Accessed: Jan. 31, 2021. [Online]. Available: <http://arxiv.org/abs/2005.08119>.
- 660 [27] M. A. Mohamad and T. P. Sapsis, “Sequential sampling strategy for extreme event statistics in nonlinear
661 dynamical systems,” *Proceedings of the National Academy of Sciences of the United States of America*, vol.
662 115, no. 44, pp. 11138–11143, Oct. 2018, doi: 10.1073/pnas.1813263115.
- 663 [28] T. P. Sapsis, “Output-weighted optimal sampling for Bayesian regression and rare event statistics using few
664 samples,” *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 476,
665 no. 2234, p. 20190834, Feb. 2020, doi: 10.1098/rspa.2019.0834.
- 666 [29] P. Pandita, I. Billionis, and J. Panchal, “Bayesian Optimal Design of Experiments For Inferring The Statistical
667 Expectation Of A Black-Box Function,” *Journal of Mechanical Design, Transactions of the ASME*, vol. 141,
668 no. 10, Jul. 2018, doi: 10.1115/1.404393.
- 669 [30] R. Agrawal, C. Squires, K. Yang, K. Shanmugam, and C. Uhler, “ABCD-Strategy: Budgeted Experimental
670 Design for Targeted Causal Structure Discovery,” Feb. 2019, Accessed: Mar. 22, 2020. [Online]. Available:
671 <http://arxiv.org/abs/1902.10347>.
- 672 [31] C. Preisinger and M. Heimrath, “Karamba—A Toolkit for Parametric Structural Design,” *Structural*
673 *Engineering International*, vol. 24, pp. 217–221, 2014, doi: 10.2749/101686614X13830790993483.
- 674 [32] G. Hinton, “Neural Networks for Machine Learning: Lecture 6,” 2012.
675 <http://www.cs.toronto.edu/~hinton/coursera/lecture6/lec6.pdf> (accessed Aug. 07, 2020).
- 676 [33] N. Zhao, M. Aldrich, C. Reinhart, and J. Paradiso, “A Multidimensional Continuous Contextual Lighting
677 Control System Using Google Glass,” in *BuildSys '15: Proceedings of the 2nd ACM International Conference*
678 *on Embedded Systems for Energy-Efficient Built Environments*, 2015, pp. 235–244, doi:
679 10.1145/2821650.2821673.
- 680 [34] G. Tsai and M. Yang, “How It Is Made Matters: Distinguishing Traits of Designs Created by Sketches,
681 Prototypes, and CAD,” in *Proceedings of the ASME 2017 International Design Engineering Technical*
682 *Conferences and Computers and Information in Engineering Conference*, 2017, vol. 7, doi:

683 10.1115/DETC2017-68403.

684 [35] E. Burnell, M. Stern, A. Flocks, and M. C. Yang, "Integrating design and optimization tools: A designer
685 centered study," in *Proceedings of the ASME Design Engineering Technical Conference, 2017*, vol. 7, doi:
686 10.1115/DETC2017-68307.

687

688

689