

Part-level 3D shape generation driven by user intention inference with preferential Bayesian optimization

Received: 13 October 2025

Accepted: 31 January 2026

Published online: 07 February 2026

Cite this article as: Lee S.W., Choi J. & Hyun K.H. Part-level 3D shape generation driven by user intention inference with preferential Bayesian optimization. *Sci Rep* (2026). <https://doi.org/10.1038/s41598-026-38916-7>

Seung Won Lee, Jiin Choi & Kyung Hoon Hyun

We are providing an unedited version of this manuscript to give early access to its findings. Before final publication, the manuscript will undergo further editing. Please note there may be errors present which affect the content, and all legal disclaimers apply.

If this paper is publishing under a Transparent Peer Review model then Peer Review reports will publish with the final article.

Part-Level 3D Shape Generation Driven by User Intention Inference with Preferential Bayesian Optimization

Seung Won Lee^{1,2}, Jiin Choi^{1,2}, and Kyung Hoon Hyun^{1,2,*}

¹Department of Interior Architecture Design, Hanyang University, Seoul 04763, Republic of Korea

²Human-Centered AI Design Institute, Hanyang University, Seoul 04763, Republic of Korea

*Corresponding author. E-mail address: hoonhello@hanyang.ac.kr

ABSTRACT

Advancements in generative artificial intelligence have introduced state-of-the-art models capable of producing impressive visual shape outputs. However, when it comes to supporting decisions during the three-dimensional shape creation process, prioritizing outputs that align with designers' needs over mere visual craftsmanship becomes crucial. Furthermore, designers often intricately combine three-dimensional parts of various shapes to create novel designs. The ability to generate designs that align with the designers' intentions at the part-level is pivotal for assisting designers. Hence, we introduced BOgen, a novel system that empowers designers to proactively generate and synthesize part-level three-dimensional shapes and enhances their overall user experience by reflecting designer intentions through Bayesian optimization. We assessed BOgen's performance using a study involving 30 designers. The results revealed that, compared to the baseline, BOgen fulfilled the designer requirements for three-dimensional shape part recommendations and shape exploration space guidance. BOgen assists designers in navigation and development, offering design suggestions and fostering proactive design exploration and creation during early-stage design ideation.

Introduction



Figure 1. BOgen Overview: (a) Bogen recommends user-desired chairs by inferring user's preferences; (b) BOgen allows users to generate and synthesize assemble-part-level 3D designs; (c) BOgen guides users to navigate through vast design space for design idea refinement.

Recently, there has been a surge in interest in Generative AI due to its capability to produce high-quality content with remarkable fidelity¹⁻⁴. However, 3D generative AI differs from 2D generative AI in having fewer paired text-shape data and limitations in terms of style and diversity^{5,6}, and require significant computational cost⁷. This presents many challenges and limitations in creating 3D shapes that satisfy user requirements through 3D generative AI. Nonetheless, there have been proposals for 3D generative AI models that show impressive performance in 3D generation, including capabilities such as part-level inversion, generation, interpolation, and augmentation⁸⁻¹¹. For instance, Hertz et al.⁸ introduced two types of latent concepts that represent the overall shape and fine details of 3D objects. Their model allows shape adjustments at the part-element level. Based on this, Koo et al.¹⁰ developed a diffusion-based method for learning these latent variables and introduced a model for text-guided 3D generation. With these technological advancements, labor-intensive and time-consuming 3D modeling tasks

are now more accessible to general users and can be performed using simple text-based prompting. Furthermore, these 3D generative AI enable designers to explore various rich and highly informative design variations. However, when it comes to AI models in the design process, prioritizing outputs that align with designers' needs over mere visual craftsmanship becomes even more crucial. Furthermore, designers often intricately combine parts of various designs to create novel designs. The ability to generate designs that align with the designers' intentions at the part level is pivotal for assisting designers, especially considering that 3D design processes require significant efforts¹². Recent work has also emphasized the importance of supporting part-level design decisions in assembly contexts¹³, highlighting the need for systems that can effectively capture and respond to designer intentions at this granular level.

Therefore, there is a need for methods that support designers iteratively^{14,15}, enabling them to efficiently navigate a vast design space with limited time and computational costs. This has prompted numerous studies on generative design and AI, supporting designers in efficient design exploration¹⁶⁻¹⁸. These systems utilize interactive genetic algorithms and design optimization to efficiently navigate extensive design spaces and generate optimized designs based on user inputs. Additionally, several studies have developed systems using generative adversarial network (GAN) models to create and explore desired images through various user interfaces¹⁹⁻²¹. However, despite numerous studies proposing user-friendly methods for navigating extensive design spaces, research is still limited in identifying specific designs desired by designers within these spaces, understanding part-level design intentions, and providing recommendations and guidance accordingly. Consequently, they are unable to address specific design problems and satisfy human psychological intentions because of failure to capture the designer's intention²². Guiding designers to navigate a high-dimensional design space in real time based on their preferences is a formidable challenge. To address this issue, methods that facilitate the exploration of a high-dimensional design space and estimate user intentions within that design space are required. To estimate design intentions, systems have been proposed that probabilistically interpret or optimize user behavior based on methodologies such as the Bayesian information gain framework or Bayesian optimization, providing appropriate feedback²³⁻²⁷. Despite the design intention estimation from Bayesian frameworks, the high computational cost of GenAI's 3D model generation process makes it expensive to recommend the desired 3D shapes aligned with the designer's actions. Therefore, reducing the exploration and generation costs of Bayesian models in this high-dimensional design space is critical. To address this issue, methods such as Non-negative Matrix Factorization (NMF) and variational autoencoders (VAE), which represent high-dimensional data in lower dimensions, can be utilized^{28,29}. A VAE is a deep generative model that learns the distribution of complex data in an unsupervised manner, maps it to a lower-dimensional latent space using an encoder, and generates new data resembling the original dataset from the latent variables sampled from this space. Owing to its effective learning of data distributions and ability to generate new data, the VAE has been utilized in various fields, such as structural and car design explorations^{28,30}. However, there remains a scarcity of comprehensive research on the integrated methods and interfaces that employ Bayesian models to predict user intentions within a vast design space while overcoming the high computational costs associated with part-level 3D generative AI. Furthermore, to the best of our knowledge, no study has applied dimension reduction to the latent space of 3D generative AI in part-level design synthesis and generation within the design exploration processes.

Building on prior research, this study aims to support design generation by inferring user preferences from the selection of design parts in a reduced design space termed the "exploration map." We introduce BOgen, a collaborative human-AI system designed to generate 3D models by leveraging user behavior and an exploration map, especially in scenarios that demand part-level generation and synthesis through 3D generative AI (Figure 1a,b, and c; Video Supplementary Material). BOgen identifies and delivers design outcomes containing desired parts through part-level selection. It also presents a method for creating a design exploration map that facilitates real-time interaction within the time and computational constraints. Our approach involved six steps: 1) collecting 58,750 SALAD latent data to create the exploration map; 2) developing a VAE-based 2D exploration map; 3) developing preferential Bayesian optimization (PBO) for BOgen; 4) Integrating PBO, VAE, and SALAD; 5) developing a BOgen and UIonly interface; and 6) conducting a comparative experiment with 30 designers.

Related Work

3D Generative AI in Design Exploration

Recent advances in 3D generative AI models have demonstrated remarkable capabilities for producing high-quality results^{4,8-10,31}. These models enable the generation, manipulation, and interpolation of 3D objects. For instance, Achlioptas et al.³¹ developed a model that generates point clouds from 3D GANs. Zhang et al.³² proposed Brep2Seq, which transforms B-rep models into sequences of editable parametrized feature-based modeling operations. Hui et al.⁹ introduced a wavelet-based diffusion network capable of generating, manipulating, and interpolating high-quality 3D objects at a part-level. Further advancements by Hao et al.³³ and Hertz et al.⁸ involved models that could control specific parts by decomposing the implicit shapes of objects, thereby facilitating part-level manipulations. Hertz et al.'s SPAGHETTI system notably enables part-level shape control through the use of extrinsic latent elements representing the overall shape and intrinsic latent elements representing fine details. Building on this, Koo et al.¹⁰ refined output quality and segregated the training of low-dimensional extrinsic latents

from high-dimensional intrinsic latents using a diffusion-based network. Hertz et al. and Koo et al. calculated the approximate shape of a 3D Gaussian using these extrinsic vectors to generate 3D shapes. However, the primary focus of these studies has been on perfecting the precision and quality of 3D generative outputs.

On the other hand, systems supporting design exploration through generative design based on interactive genetic algorithms and design optimization^{34–36}, as well as generative AI have been proposed^{19,37}. Hyun and Lee³⁴ proposed a GA-based framework to support designers in making decisions based on styling strategies. Ban and Hyun³⁵ suggested a framework that interpolates input sketches with those in a database to create design variations and, based on these variations, generates 3D models to support design exploration. Evirgen and Chen¹⁹ introduced GANzilla, a generative adversarial network (GAN)-based model that supports non-expert users in creating desired images using scatter-and-gather techniques. Furthermore, Zhou et al.³⁷ proposed a model enabling users to intervene in the Bayesian optimization loop, adjusting the model's exploration and exploitation to create the desired melody compositions. However, there is a clear gap in research that aligns the power of generative AI with a designer's unique intent across a myriad of design components. For instance, Shi et al.³⁸ found that generative AI models require domain-specific fine-tuning to understand both textual design intent and non-verbal intent to produce appropriate design alternatives. From a user-experience perspective, insights into how best to leverage these state-of-the-art 3D generative AIs are notably absent, especially when creating user-friendly interfaces to support the 3D generation process. While current leading studies emphasize the accuracy of the generation performance, their application in real-world design contexts remains largely unexplored. As pivotal as advancing AI performance is, devising ways for designers to utilize AI efficiently is crucial. Hence, this study seeks to bridge this gap by proposing methods and interfaces that integrate state-of-the-art 3D generative AIs into the design exploration process.

User Intention Inference and Feedback Methods

To generate objects that are meaningful and useful to designers, it is essential to enable design creation, combination, and exploration at a detailed element level that aligns with the designer's intentions. However, whether such designs are useful can change depending on the design situation, and ultimately, designers should make that decision. To support this design process, design systems employing probabilistic models like Bayesian information gain (BIG) and Bayesian optimization (BO) have been proposed^{23–25,27}. Son et al.²⁷ and Lee et al.²⁵ introduced systems that interpret user behavior probabilistically to provide design feedback based on information gain. Kadner et al.²³ proposed 'adaptifont,' a BO-based font generation system that uses the user's reading speed as an objective function. In the broader line of preference-based optimization, prior work modeled user preferences from pairwise comparisons using Gaussian processes^{39,40} and extended this idea to preferential Bayesian optimization for black-box search⁴¹. Brochu further established interactive Bayesian optimization and user preference modeling as foundational approaches for incorporating human feedback in optimization loops^{42–44}. Constrained Bayesian optimization incorporates constraint satisfaction into the optimization loop^{45,46}, providing an alternative BO formulation when feedback includes constraints in addition to preferences. Koyama and Goto²⁴ introduced preferential Bayesian optimization (PBO), which optimizes a function predicting the user's desired goal based on their slider bar manipulation information. These systems track the design process by probabilistically interpreting user actions and offering design feedback that reflects the designer's evolving intentions. However, as the information space increases, the computational cost of the BIG framework increases exponentially and BIG-based systems provide suboptimal design feedback to address this issue^{25,27,47}. Similarly, BO operates best in lower-dimensional spaces, often requiring dimension reduction due to computational limitations²³. In this context, Danhaive and Mueller²⁸ proposed a system that learns design parameters and performance using a VAE and maps them onto a 2D space to generate designs. They introduced a graph that visualizes and navigates a complex high-dimensional design space in lower dimensions. This allows users to visually grasp, compare and evaluate various design options. However, research on inferring the desired design elements from a reduced exploration map to support design creation and synthesis remains unexplored. Integrating the mentioned methods to infer user intentions, along with a streamlined exploration map approach, remains a significant challenge in proposing a 3D generative design system. Therefore, this study introduces a system that reduces the high-dimensional 3D generative latent to a more manageable exploration map, thereby guiding the design exploration process based on user preference estimations.

BOgen

Conceptual Overview

The BOgen system comprises three main components: 1) Generative AI for 3D designs generation (Figure 2a), 2) A VAE that reduces the generative AI's latent space to two dimensions and decodes 2D latent points on the exploration map (Figure 2b and 2e), 3) a PBO model that infers user preferences in the reduced latent space (Figure 2c), and 4) an interface that enables designers to effectively generate and explore assemble-part-level designs (Figure 2d). Briefly, in the BOgen, users can synthesize and generate 3D chair designs and express their preferences for these designs in the user interface (Figure 2d). This information is mapped onto a 2D exploration map via a VAE (Figure 2b), and the PBO infers user preferences from

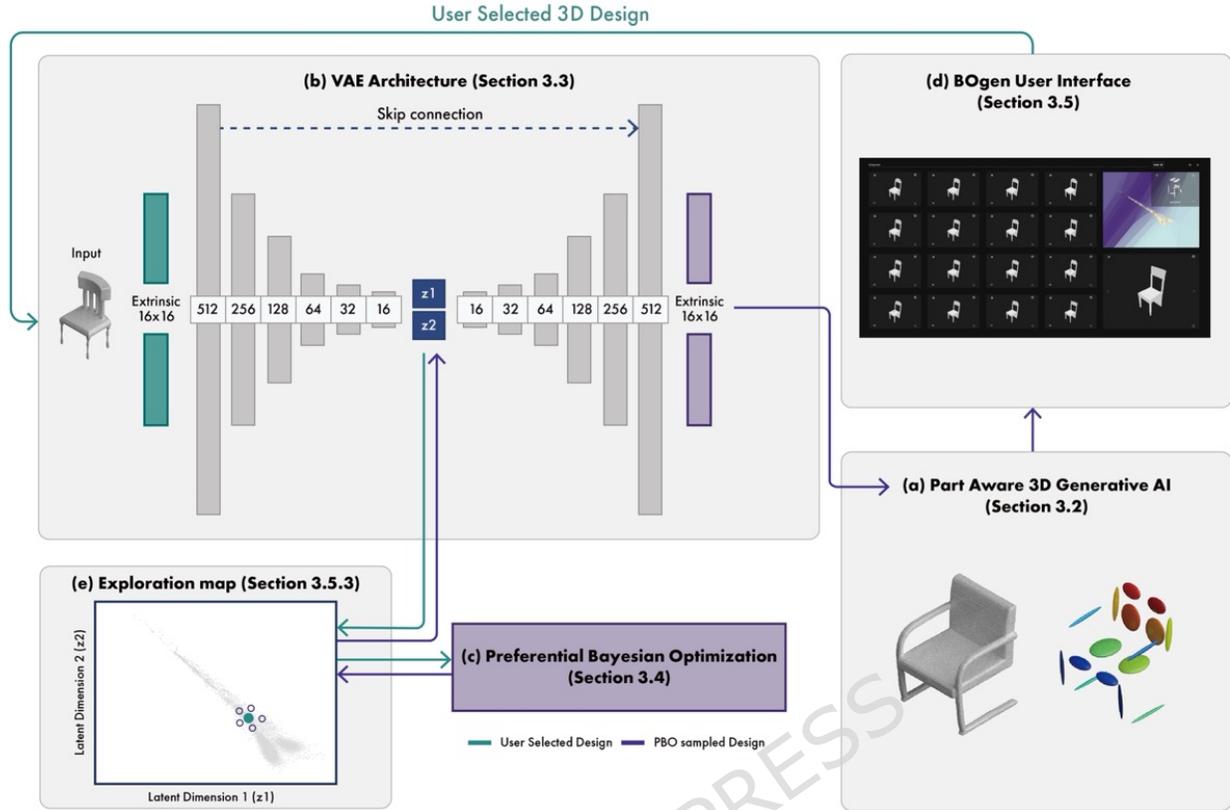


Figure 2. BOgen conceptual overview.

this exploration map (Figure 2c, e). Subsequently, the points sampled by the PBO are decoded and passed through the 3D Generative AI model (in this study, SPAGHETTI and SALAD; Figure 2a), resulting in the generation of chair designs.

Part Aware 3D Generation Models

This section introduces SALAD¹⁰, the 3D generative AI used in our research, and its basis, SPAGHETTI⁸. SPAGHETTI, proposed by Hertz et al.⁸, is a part-aware 3D generative AI model comprising three main networks: Decomposition, Mixing, and Occupancy (Figure 3a). Initially, a 3D shape is mapped to a shape-embedding Z_a . Z_a then passes through a multilayer perceptron-based (MLP)-based Decomposition Network and is encoded into part embeddings $\{p_i\}_{i=1}^N$ ($N = 16$), where each p_i represents a part of the 3D shape. Each p_i of Z_b consists of an extrinsic latent $\{e_i\}_{i=1}^N$, which represents the location, orientation, and size information of each 3D shape part, and an intrinsic latent $\{s_i\}_{i=1}^N$, which details the geometric features (Figure 3b). Specifically, extrinsic latent $e_i = \{c_i, \lambda_{i1}, \lambda_{i2}, \lambda_{i3}, u_{i1}, u_{i2}, u_{i3}, \pi_i\}$ and $e_i \in \mathbb{R}^{16}$, where $c_i \in \mathbb{R}^3$ represents the mean of the Gaussian mixture, $\lambda_i \in \mathbb{R}$ are the eigenvalues of covariance, $u_i \in \mathbb{R}^3$ are the eigenvectors, and $\pi_i \in \mathbb{R}$ represents the blending weights indicating the relative importance of each part. Based on these extrinsic latent variables, each part is computed as a spherical 3D Gaussian. Z_b passes through a Transformer-based Mixing Network, outputting contextual embedding Z_c . Finally, a 3D shape is generated through a Transformer decoder-based Occupancy Network, which uses Z_c information to predict the occupancy value $o \in [0, 1]$ for a query coordinate x . A value of $o = 1$ indicates that point x is inside the 3D shape, while a value of 0 indicates that it is outside. This occupancy value is used to reconstruct the 3D shape.

Koo et al.¹⁰ proposed SALAD, which is a two-phase diffusion model that generates both extrinsic $\{e_i\}_{i=1}^N$ and intrinsic $\{s_i\}_{i=1}^N$ latents. The first phase of the SALAD performs Transformer-based noise prediction for extrinsic noise, and the second phase predicts the conditional distribution of the intrinsic latent noise $P(\{s_i\}_{i=1}^N | \{e_i\}_{i=1}^N)$. Thus, Koo et al. generated SPAGHETTI's extrinsic and intrinsic latents through a diffusion model and input them into the Occupancy Network to enable higher-level 3D shape generation.

In summary, we generated 3D models based on SALAD and SPAGHETTI and processed extrinsic latents by reducing their dimensions to 2D to facilitate the exploration of a vast design space.

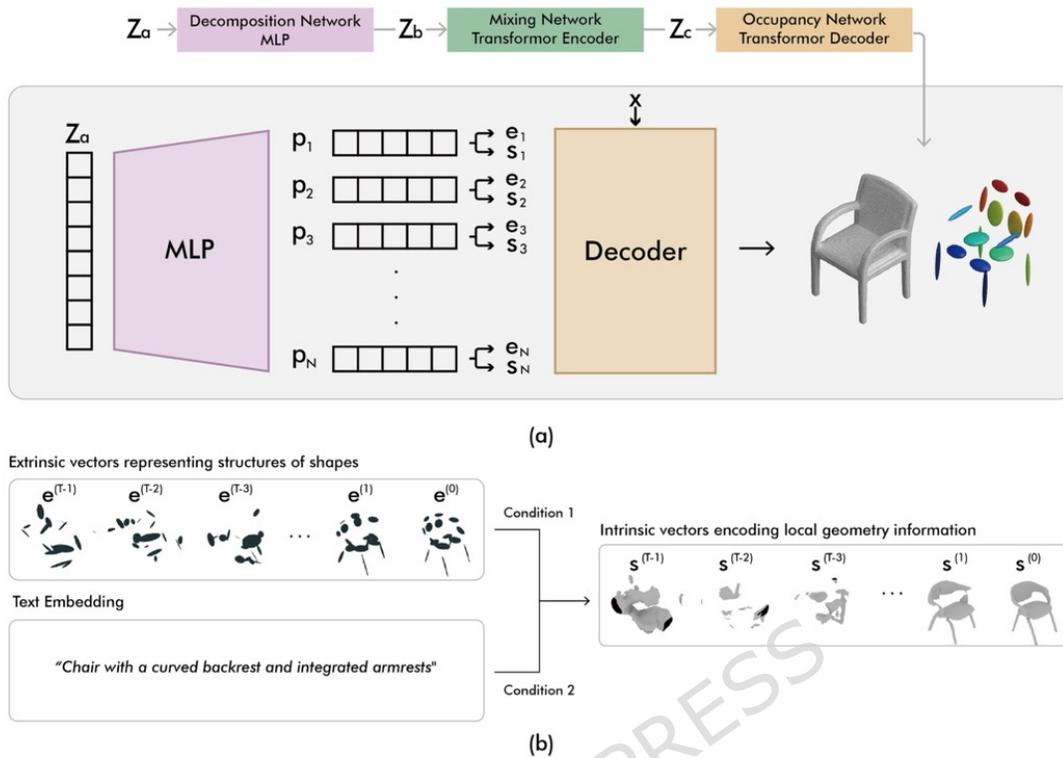


Figure 3. Part aware 3d generative model architecture: (a) three types of networks of SPAGHETTI⁸; (b) Diffusion model of SALAD¹⁰.

Dimensionality Reduction of Extrinsic Latent using VAE

The extrinsic latent consists of 16 parts, each with 16 dimensions. Although these latents effectively represent the broad shape of intricate 3D objects within a condensed dimension, adjusting each of the 16 parts to match a designer’s specifications is both laborious and time-intensive. Furthermore, visualizing this 16x16-dimensional design space for exploration is challenging. Therefore, we utilized a VAE to map the 16 × 16-dimensional extrinsic latent to a 2D space without losing detailed information and to generate an extrinsic latent based on 2D points. To utilize the VAE, we trained its encoder to represent the distribution of extrinsic latents in a lower dimension and the decoder to reconstruct points in the reduced 2D space to resemble extrinsic latents. Unlike traditional VAE models, we added skip information to the VAE architecture, where the information from the encoder’s 512-dimensional layer is directly transferred to the decoder’s 512-dimensional layer (i.e., skip connection). This direct linkage ensures that the VAE network retains information during training by transmitting the input data features from the encoder to the decoder. This allows the decoder to use the input data information during decoding, thereby enabling the generation of data with different details while maintaining the context of the input data.

We trained the VAE model with a 2-dimensional latent space for 50 epochs using the Adam optimizer (learning rate = 0.001) with a training batch size of 128. Training was conducted on an NVIDIA A6000 GPU and completed in approximately 2.7 minutes, averaging 3.3 seconds per epoch. For VAE training and evaluation, we generated 58,750 extrinsic latents using shape-referential texts employed in SALAD’s text-conditioned extrinsic diffusion model (Figure 2). Of these, 52,992 were used for training, and 5,758 were used for evaluation. The VAE reconstruction loss on the evaluation dataset was approximately 0.015 at the 50th epoch (Figure 4), indicating that the generated data closely resembled the ground truth (i.e., extrinsic latent data). The extrinsic latents mapped to 2D by the VAE encoder and the points sampled in the 2D space are termed ‘2D latent points,’ where each point comprises latent dimension 1 (z_1) and latent dimension 2 (z_2). We define a 2D space composed of these 2D latent points as an ‘exploration map.’ The data distribution mapped to the reduced 2D design space showed dense areas for chair types, such as armchairs and dining chairs (Figure 5a, b). Moving to the top left, many decorative chairs,

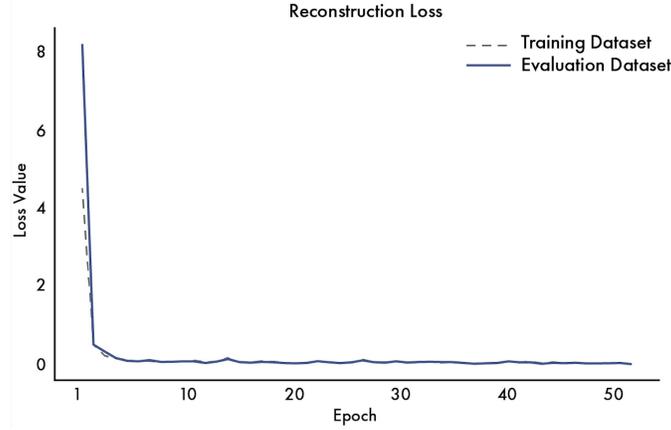


Figure 4. VAE training and reconstruction loss

folding chairs, and stools were distributed (Figure 5c, d), followed by unconventional chairs (Figure 5e). We obtained 1,981 characteristic data points from a dataset of 58,750 training and evaluation datasets using the k-means++ algorithm, which was then utilized as the final dataset.

Preferential Bayesian Optimization

PBO Algorithm

BOgen employs a PBO to predict and provide the user with the desired data based on their preferences. However, estimating the designer's desired part from existing extrinsic (16×16 dimensions) or intrinsic latent (16×512 dimensions) is computationally inefficient. Therefore, to solve this problem, we propose a method that applies a PBO based on a reduced exploration map via a VAE to estimate and suggest a user's desired design in real-time. The PBO used in our study followed that of Koyama and Goto²⁴. First, the objective function of the PBO is the function g which predicts the user's desired design. For instance, if the user is presented with n chairs and prefers chair number 4, we represent this observation as $d = [x^4 > x^1, x^2, x^3, x^5, \dots, x^n]$. This means x^4 represents the 'preferred design', while the rest are 'other designs'. With multiple observations of d , it's denoted as $D = [d_1, d_2, \dots, d_n]$. For each data point x^i , a preference or 'goodness' value, denoted by g^i , is assigned, where i ranges from 1 to N . These values, g^i , are represented as $\mathbf{g} = [g^1, \dots, g^N]^T$. For the preference-based objective function g , preferential modeling uses the data created or preferred by the user to model d , the relative preference among the remaining provided data, using the Bradley-Terry-Luce model⁴⁸.

$$P(d|\mathbf{g}) = \frac{\exp(g^i)}{\sum_{j=1}^N \exp(g^j)}. \quad (1)$$

The Bradley-Terry-Luce (BTL) model was used to model the relative preferences among the given data based on user preferences. The likelihood of multiple observation data, given the preference values g , is calculated as the product of individual likelihoods, represented by $P(D|\mathbf{g}) = \prod_i P(d_i|\mathbf{g})$. However, the $g(x)$ values for unobserved data points x remain unknown. To predict these values, a Gaussian process (GP) was used, which assumes that all data points share a common probability space⁴⁹. Based on this assumption, the relative preference d calculated through the BTL is used to estimate $g(x)$, reflecting the user's desired design. The maximum a posteriori (MAP) methodology is then applied to obtain an optimal estimate for these $g(x)$ values.

$$\mathbf{g}^{map} = \operatorname{argmax} P(\mathbf{g}|D) = \operatorname{argmax} P(D|\mathbf{g})P(\mathbf{g}). \quad (2)$$

Essentially, GP estimates various possible probability distributions for $g(x)$ based on the observed data, and MAP identifies the distribution with the highest posterior probability that most closely aligns with the observed data. Specifically, in Eq 2, $P(D|\mathbf{g})$ represents the posterior probability of $g(x)$ values, given the observed data D , expressed as the product of $P(D|\mathbf{g})$ and $P(\mathbf{g})$. $P(D|\mathbf{g})$ denotes the conditional probability of D given the $g(x)$ values (i.e., the probability of occurrence of D when $g(x)$ values are given), and $P(\mathbf{g})$ represents the prior distribution of $g(x)$ values, which, under GP assumptions, is a Gaussian distribution. Using the estimated $g(x)$ values, the predicted distribution of an unseen data point x can be calculated as follows:

$$g(x) \sim N(\mu(x), \sigma^2(x)). \quad (3)$$

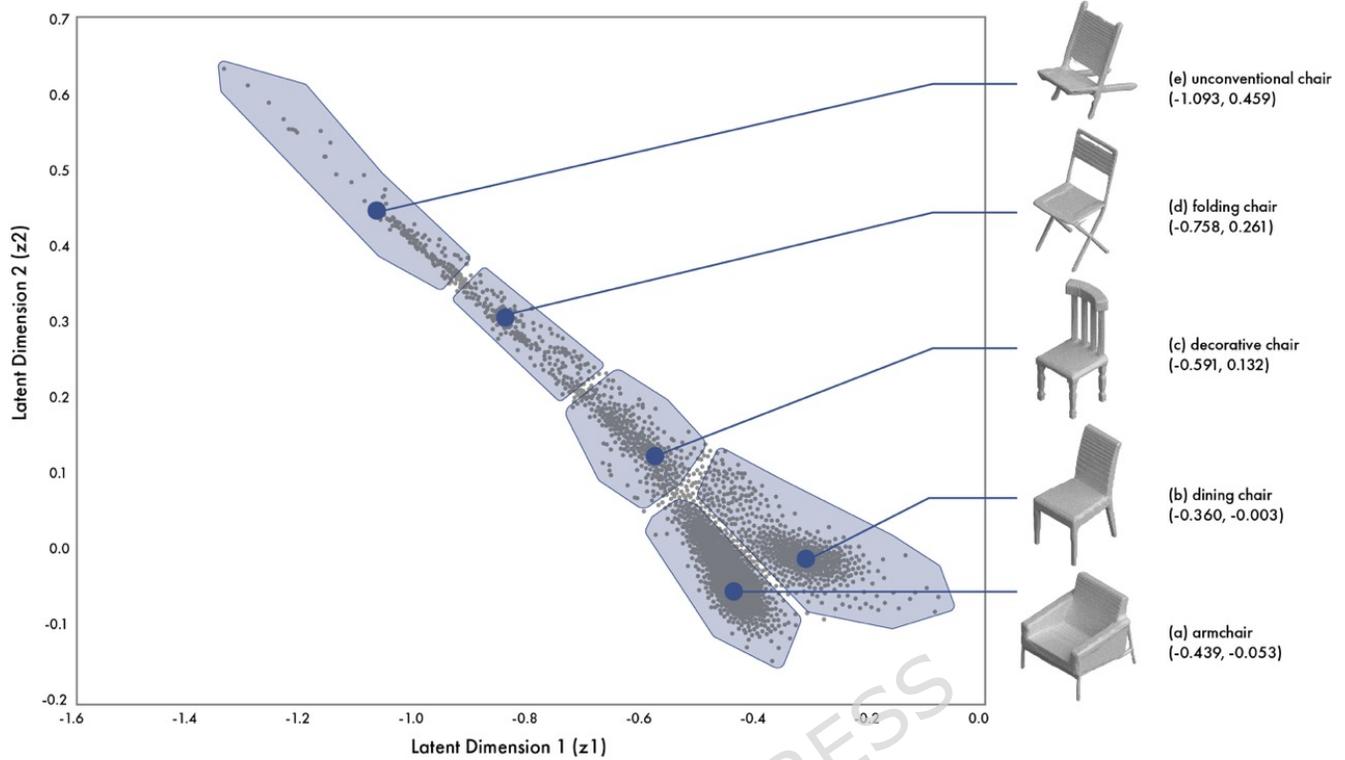


Figure 5. Data distribution of exploration map which is divided by five clusters and examples of the chairs corresponding to the 2D latent points.

This stems from the Gaussian Process (GP) model, where $g(x)$ is defined by a normal distribution. The mean, ($\mu(x)$), represents the predicted user preference for design option x , and the variance, $\sigma^2(x)$, quantifies the uncertainty associated with this prediction²⁴. Specifically, a high $\mu(x)$ in the normal distribution of a particular data point x implies a high $g(x)$ value, indicating a high probability of user preference. Conversely, a large $\sigma^2(x)$ value suggests greater uncertainty in GP's prediction, indicating less confidence about that particular data point x . The Acquisition function used to provide suggestions to the user based on these estimated $\mu(x)$ and $\sigma(x)$ employs the Gaussian process upper confidence bound (GP-UCB⁵⁰):

$$a^{GP-UCB} = \mu(x) + \beta \sigma(x). \quad (4)$$

The GP-UCB is an acquisition function based on a Gaussian process. The β parameter in GP-UCB (Equation 4) controls the trade-off between exploration and exploitation: higher β values prioritize exploring uncertain regions of the design space, while lower values focus on exploiting regions with known high preferences. Following Koyama and Goto²⁴, we adopted $\beta = 0.5$. PBO recommends optimal suggestions to the user by selecting the points with the highest acquisition values based on the mean and variance of the preference-based objective function. To generate multiple samples simultaneously, we employed a sequential greedy batch selection technique proposed by Schonlau et al.⁵¹. After selecting each point by maximizing the acquisition function (Equation 4), only the variance function $\sigma^2(x)$ is updated, while the mean function $\mu(x)$ remains unchanged because the actual goodness value is unknown at this stage. Subsequent points are then sampled by maximizing the acquisition function computed with the original mean and updated variance. This approach ensures spatial diversity among batch samples by reducing the acquisition value near already-selected points. Smaller batch sizes provide higher-quality recommendations per sample but limit diversity, while larger batch sizes increase diversity at the cost of later samples potentially having lower acquisition values. The batch size was set to $k = 16$.

Preferential Bayesian Optimization for BOgen

Generating samples with VAE. We need to predict user preferences for the desired chair designs, which requires defining an information space for chair designs. Conventionally, design information spaces are discrete design objects^{25,27} or feature vector spaces²³. Defining design objects as information spaces is useful for predicting which design references match a user's intended design; however, expanding or reducing the information space is challenging and requires a pre-built design reference

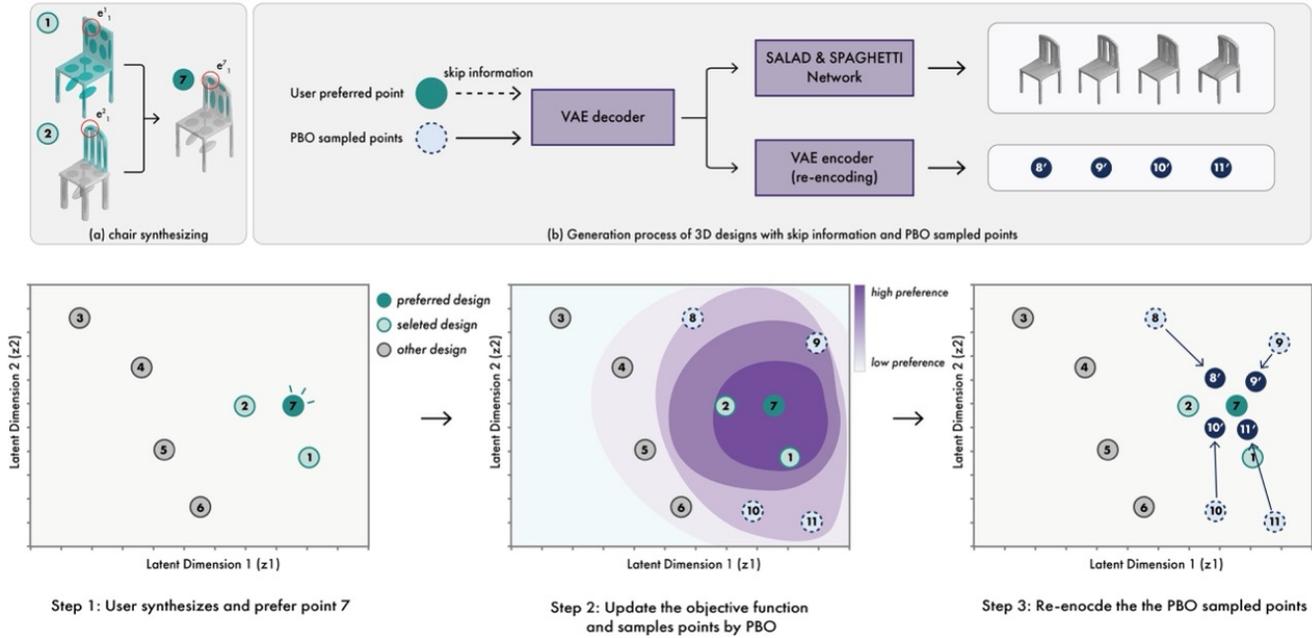


Figure 6. The process of updating the preference objective function with user preference.

database, making it less suitable for design generation scenarios. By contrast, defining a vector space as a design space allows the prediction of the intended vector based on user actions, and if a feature vector can be reconstructed into a design, it enables design generation. Therefore, a viable information space that can be selected is the extrinsic latent space. If we can interpret user intent in this space and find the intended extrinsic latent, we can generate a complete 3D chair object through SALAD's 'intrinsic diffusion network' and SPAGHETTI's 'Occupancy Network.' However, the extrinsic latent is 16×16 -dimensional, and computational feasibility is nearly impossible at the level required for real-time interaction and design exploration through BO. Hence, we incorporated the VAE and PBO into the existing SALAD and SPAGHETTI frameworks. The VAE encoder learns the extrinsic latent distribution and represents it in 2D space, whereas the decoder reconstructs points sampled in 2D space into the form of extrinsic latents. PBO performs preference optimization in a reduced 2D space by sampling the optimal intended 2D latent points. We set the sampling boundary of the PBO to encompass the 2D latent-space range of the VAE. The range for the first latent dimension was set between 0 and 1.5, while the second dimension's range was between -0.1 and 0.7.

BOgen User Interface

Updating Preference-based Objective Function. We regard the 2D latent points of the chairs for which users have explicitly shown preferences or requested PBO suggestions (i.e., they want to see similar designs) as the *preferred design* (Figure 6-step 1). The 2D latent points of all other chairs were considered to be the *other designs* (Figure 6- step 1). If the chair preferred by the user or for which a PBO suggestion was requested was a combination of more than one chair, those chairs were excluded from the other designs. Specifically, as shown in Figure 6, if a user prefers a synthesized '7' chair in the step 1 and requests a PBO based on this, the preferred design is '7' chair and the other designs are chairs '3', '4', '5', '6'. Based on this, the objective function was updated. The PBO then samples new points based on the updated objective function, where each point is decoded into extrinsic latent points using skip information and re-encoded (Figure 6b) into 2D latent points. This VAE encode and decode method, linked with PBO and based on skip information, enables the prediction of which vector space in 2D the user is interested in, even if they only select the 3D Gaussian parts of a chair design. In addition, the original 2D latent points sampled by the PBO (i.e., points before encoding) were updated as other designs and excluded from the next update. This was performed to prevent the PBO from continuously sampling at the same points.

User interface main screen

The BOgen system interface is shown in Figure 7. Initially, users can input their desired chair into the prompt (Figure 7a) employed by SALAD's conditional text-guided extrinsic diffusion model, enabling them to simultaneously observe up to 16 different variations of chairs. The sequence of each card is displayed at the bottom left. Additionally, a 2D exploration map located at the top right facilitates the navigation of designs, allowing users to explore beyond those shown on the cards (Figure 7c). Users can mark the designs of interest using a button on the top left of each card (Figure 7b). Such marked designs are

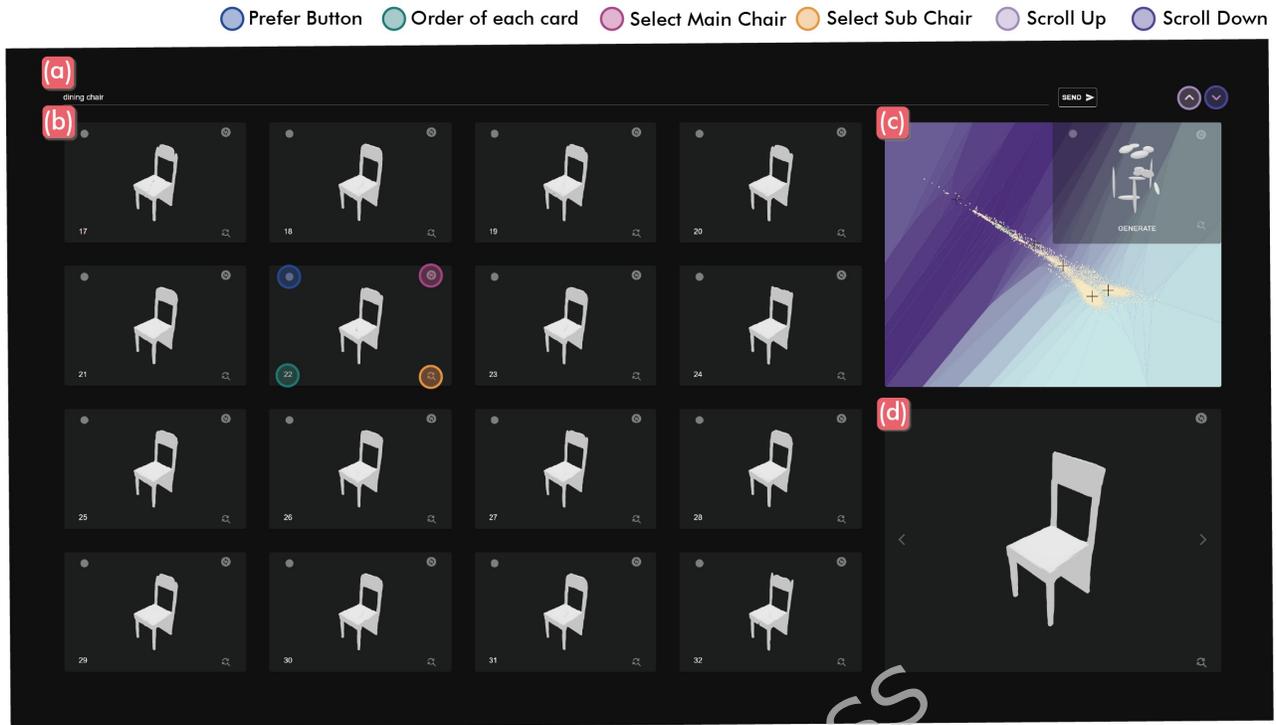


Figure 7. BOgen System Overview: (a) prompt input field; (b) generated chairs by prompt or PBO; (c) exploration map; and (d) added synthesized chair.

considered the *preferred designs* and are excluded from subsequent updates and *other designs*. Upon selecting the main and sub designs through the generated chairs and exploration map, a screen pops up, allowing for the synthesis of different parts. The resulting design can then be saved (Figure 7d). When users scroll down, the next 16 cards display the chairs sampled by the PBO based on the user’s recent synthesis. The VAE latent of this design is deemed the *preferred design*, whereas the VAE latent of *other designs* that are not part of the synthesis is considered as *other designs*.

Synthesis screen

When the main chair and subchair are selected on the BOgen main screen, a synthesis screen appears (Figure 8). Within the synthesized chair, parts of the main chair (Figure 8a) and subchair (Figure 8b) can be selected individually. Upon clicking the “Interpolate” button, changes in the 3D Gaussian function based on the main chair were observed. For example, if a user interpolates the 1st part of the 3D Gaussian of the main chair A (e_1^A) with the 1st part of the 3D Gaussian of sub-chair B (e_1^B), the newly synthesized 1st part of the 3D Gaussian of chair C (e_1^C) is $(e_1^A + e_1^B)/2$, as explained in Figure 6a. The 3D Gaussians of the remaining parts remain unchanged and are the original ones from the main chair A. By pressing the generate button, the shape of the synthesized chair can be previewed (Figure 8c), and by clicking on “Add,” it can be saved to the card located at the bottom right of the main screen (Figure 8d).

Exploration Map

Designs can be explored through the ‘*exploration map*,’ a 2D latent space reduced by VAE, with areas divided using Voronoi diagrams to distinguish the different areas on the map, located at the top right of the main screen. The color of the area was updated (from purple to sky blue) based on the probability of the chair being selected (Figure 9a). Each area of the map is marked with a “+” sign to indicate the characteristics of the designs distributed (Figure 9b). Hovering over a point representing a design on the exploration map displays it in a 3D Gaussian form at the top right (Figure 9c). Designs explored on the map can also be selected as either the main or sub-design and synthesized in the same manner as before. The 16 chairs displayed on the cards were numbered (marked) on the exploration map (Figure 9d). In addition, clicking on the area of the map highlights it as fluorescent lime, indicating the exploration area (Figure 9e).

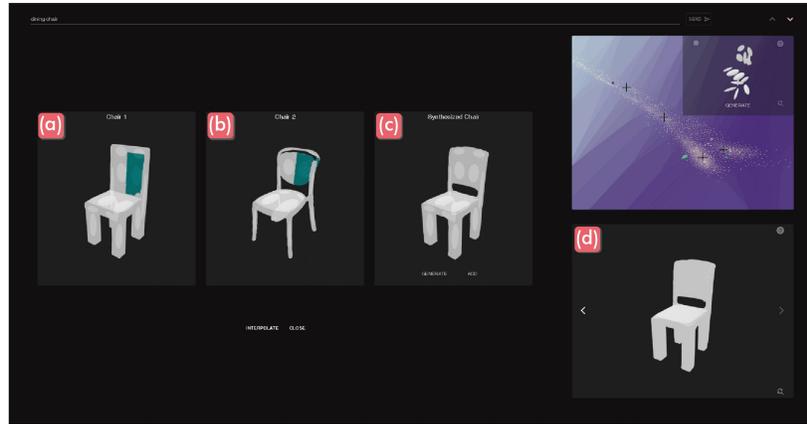


Figure 8. BOgen chair synthesizing interface: (a) main chair; (b) sub-chair; (c) preview of synthesized chair; and (d) saved synthesized chair.

Implementation and Results

Experimental Design

To evaluate BOgen’s performance, we recruited 30 experimental participants (including 20 females), all of whom had majors related to design or had professional design experience in the field. (Age: Mean = 23.03; Max = 27; and Min = 19). This study was approved by the Institutional Review Board of Hanyang University (IRB approval number: HYUIRB-202304-024-4). All methods were performed in accordance with the relevant guidelines and regulations. Informed consent was obtained from all participants prior to the experiment. The experiment followed a within-subject design and proceeded in the following order: introduction, consent, and tutorial (about 10 minutes) → Brief 1 (about 30 minutes) → Brief 2 (about 30 minutes) → In-depth interview (about 20 minutes). The participants performed two tasks (BOgen and UI only). Here, the UI only served as a baseline interface for comparison, excluding the exploration map and PBO-based recommendation features from the BOgen system (Figure 10). The remaining synthesis function and prompt-based creation used the same model. In the exploration task scenario, the participants briefly designed a chair to fit into a designated empty space within a given interior scene (e.g., a living room or a hotel lobby). The aim was to observe how participants, as designers, explored and selected chair designs during the early stages of the design process. To achieve this, participants were required to use only two systems, BOgen and UI only, to investigate the effect of the exploration map and PBO-based recommendations. To control for potential order effects, we employed a fully counterbalanced design: 15 participants used BOgen first followed by UI only, while the remaining 15 participants used UI only first followed by BOgen. The design stimuli were also counterbalanced across conditions. Each task had a maximum duration of 30 min, although participants could conclude earlier if they found a design that satisfied their needs. The system used in the experiment was developed using Python 3.9. All participants used a high-performance client-server system (client: React on Windows OS with Intel Core i9 10980XE, 64 GB RAM; server: Python Flask on Linux OS with AMD Ryzen Threadripper Pro 3995WX, 256 GB RAM). During each experiment, the extrinsic and 2D latent points of all the generated and explored chairs were logged. After each session, a 7-point scale survey was conducted to assess the influence of the system on the design process. After all the sessions were completed, we collected opinions on the overall experience of both systems through interviews.

Experimental Analysis Metrics

In our experimental analysis of BOgen, we focused on four key metrics along with surveys and in-depth interviews for a comprehensive evaluation. Four quantitative analysis metrics were utilized: 1) the mean uncertainty evaluates the model’s ability to accurately capture and adapt to user preferences, indicating the model’s confidence level in its predictions; 2) the mean probability of the preferred chair assesses how well the system identifies designs that align with the user’s preferences, reflecting the system’s predictive accuracy; 3) the explored area captures the extent of the user’s exploration within the design space, highlighting the breadth of their design exploration; and 4) the number of clusters indicates the density and distribution of user design exploration, providing insights into the patterns of their exploration behavior. We performed paired-sample t-tests on these four quantitative metrics, and the Wilcoxon signed-rank test for ordinal survey responses.

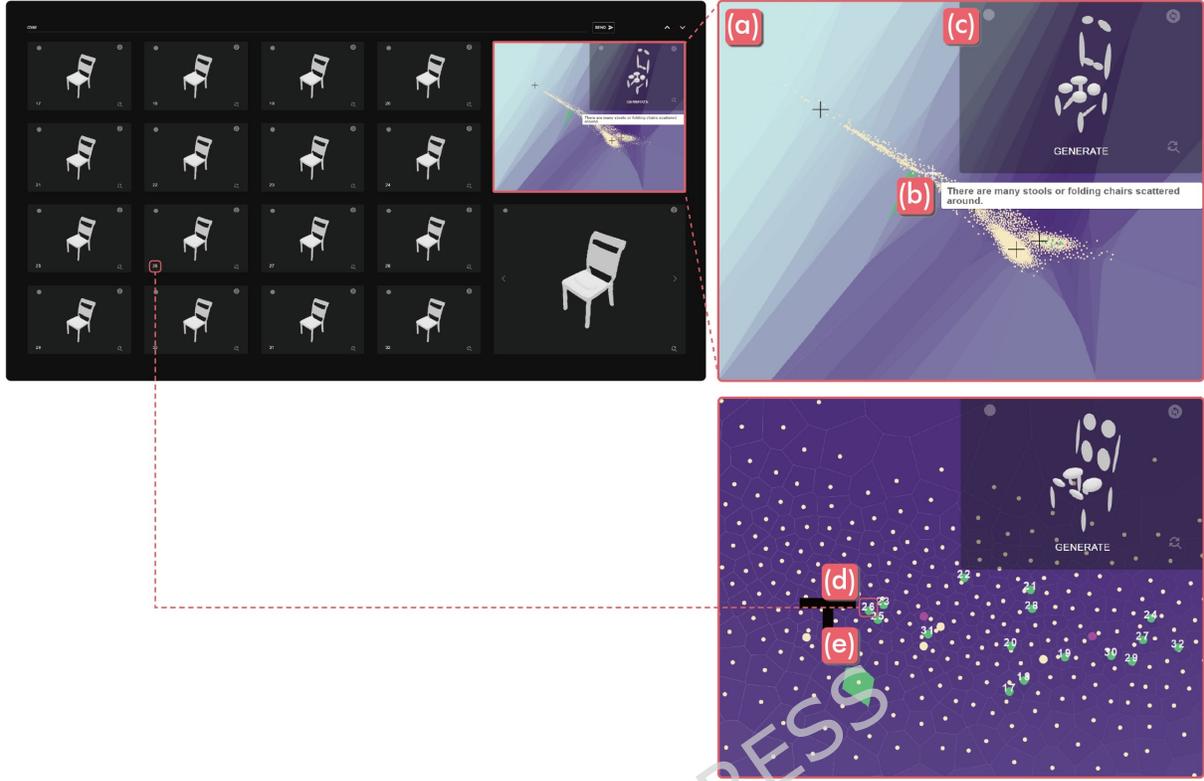


Figure 9. Exploration Map: (a) exploration map; (b) brief explanation of chairs; (c) 3D gaussian of a hovered design point ;(d) order of each card; and (e) exploration area mark.

Mean Uncertainty

We evaluated all the 2D latent points generated by the user, encompassing the initial 1,981 points, as the information space. We established the average value of uncertainty in that space, denoted as $\sigma^2(x)$, as the mean uncertainty. The mean uncertainty serves as a crucial metric, reflecting the model's confidence in capturing and adapting to the user's evolving preferences over time. Calculating the mean uncertainty involved deriving the average uncertainty across all information spaces after executing a PBO update. Therefore, a lower mean uncertainty indicates a high level of confidence in the model's predictions within the information space. Specifically, for the i -th PBO update, the uncertainty for the information space can be expressed as the average of $\sigma^2(x_i)$ over all x in the space. This is computed as $meanuncertainty = \frac{1}{N} \sum_{i=1}^N \sigma^2(x_i)$, where N represents the total number of points in the information space.

Mean Probability of preferred chair

The mean probability of the preferred chair represents the average of the prior probabilities (i.e., $\mu(x)$) of the preferred designs, representing how well the system aligns with the user's preferences. A high mean probability value reflects a strong alignment with user preferences, indicating the system's effectiveness in identifying preferred designs. Suppose only one design is preferred while the user uses the system, and it results in a prior probability value of 0. In that case, it suggests that the system may not fully understand the user's preferences yet. Mean probability is calculated as the average of the $\mu(x)$ values for each preferred design prior to its PBO update. For example, if the designer preferred x^1 design in 4th PBO iteration state, the $\mu(x_4^1)$ is the prior probabilities. This can be mathematically expressed as $meanprobability = \frac{1}{N} \sum_{i=1}^N \mu(x_i)$, where N represents the total number of preferred designs.

Explored Area

We defined the 2D latent points of the chairs that users searched for, generated, and hovered over as the 2D latent points explored by the users. The explored area is a significant indicator of a user's exploratory behavior, showing the breadth and diversity of the design elements investigated. The explored area is determined by the convex hull of the 2D latent points, representing the breadth of the design explored by the user on the exploration map. Mathematically, if we denote the set of explored 2D latent points as $X = x^1, x^2, \dots, x^n$, where each x^i is a point in the latent space, the area of the explored region, A , can be calculated as the area of the convex hull formed by the points in X . This is expressed as $A = Area(convexhull(X))$.

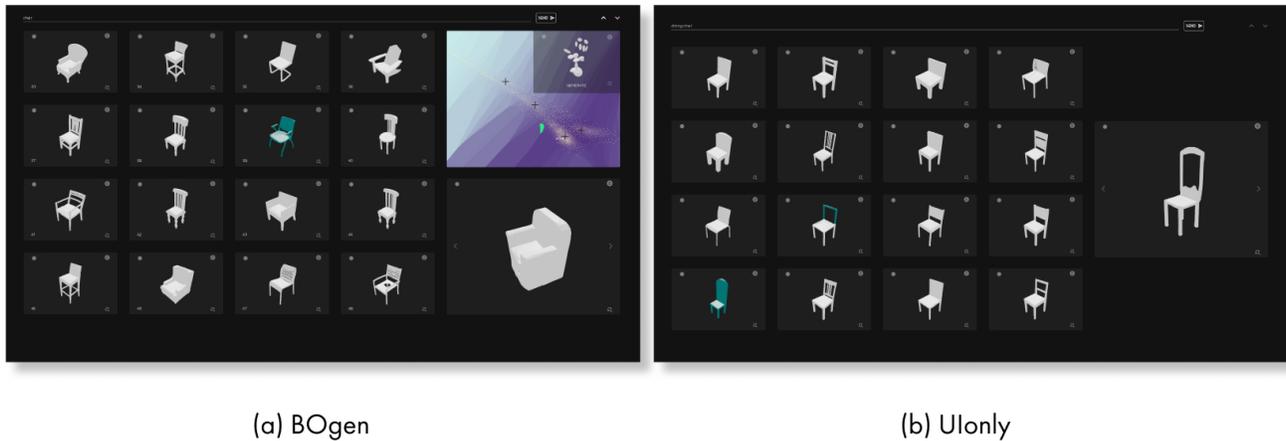


Figure 10. System Interface for (a) BOgen; and (b) UIonly.

Number of Cluster

To measure the density of the 2D latent points explored by users, we utilized the DBSCAN algorithm. DBSCAN is renowned for clustering based on point density. The number of clusters identified by DBSCAN is critical because it provides insight into the patterns and structures within the user's exploration, indicating areas of concentrated interest. Before using DBSCAN, one must determine the ϵ and MinPts value. The ϵ value serves as a parameter to find neighboring points within the ϵ distance from a given data point. Too small a value categorizes many data points as noise, while too large a value can merge different clusters. A common approach to finding an appropriate ϵ is to identify the "elbow point" in a k-neighbor graph. This point corresponds to a sudden change in the graph's slope, where the curvature is maximized. By setting the ϵ based on this "elbow point," one can optimize clustering performance. Based on the 2D latent points explored in each system by users, we set an elbow point with a maximum curvature value of 0.05 or less. Ultimately, the average of these 60 elbow points (number of experiment participants * number of systems) was set to 0.039 for DBSCAN's ϵ . MinPts was set to 3.

Survey and In-depth interview

To assess if BOgen delivered designs with the user's preferred parts, assisted in design development and navigation, and provided useful designs, we surveyed the following four questions on a 7-point Likert scale: Q1) The system effectively suggested designs containing my desired parts; Q2) The system effectively clarified my design goals; Q3) The system assisted in navigating towards my preferred design direction; and Q4) The system provided useful design suggestions. Additionally, an in-depth interview inquired about the impact of each system on the design process, the pros and cons experienced during system usage, and their opinions.

Results and Discussions

Based on the experimental results, when compared with the UIonly system, BOgen exhibited three main findings: 1) guiding design space navigation for design idea refinement, 2) recommendations for designers to find prominent design parts, and 3) facilitating designers to explore dense and diverse design spaces.

Guiding Design Space Navigation for Design Idea Refinement

During the design process, designers continuously refine their design ideas and navigate the design space to search for and integrate the desired design parts. Therefore, effectively guiding the design space to infer the designer's preference, especially those with desired parts, is a crucial task. To accomplish this task effectively, a system needs to better capture user preferences and discern the preferred designs. We employed the mean uncertainty and mean probability metrics to measure the system's capability to reflect user preferences and identify preferred designs. Consequently, BOgen displayed a significantly lower mean uncertainty (BOgen: 0.070; UIonly: 0.237; $p < 0.001$; Figure 11a) and a significantly higher mean probability (BOgen: 0.092; UIonly: 0.043; $p < 0.001$; Figure 11b) compared to the UIonly system. This indicates that BOgen captured user preferences more accurately and had greater confidence in determining the chairs preferred by the users. Furthermore, for survey questions Q2 and Q3, BOgen scored significantly higher (Figure 12; Q2 : $Mean_{BOgen} = 5.8, Mean_{UIonly} = 4.833, p < 0.001$; Q3 : $Mean_{BOgen} = 5.766, Mean_{UIonly} = 4.8, p < 0.001$). The majority of users (90%) expressed that they could better navigate towards their desired design direction and refine their ideas with the aid of BOgen's exploration map and PBO-based recommendations. Some notable examples of user feedback include the following:

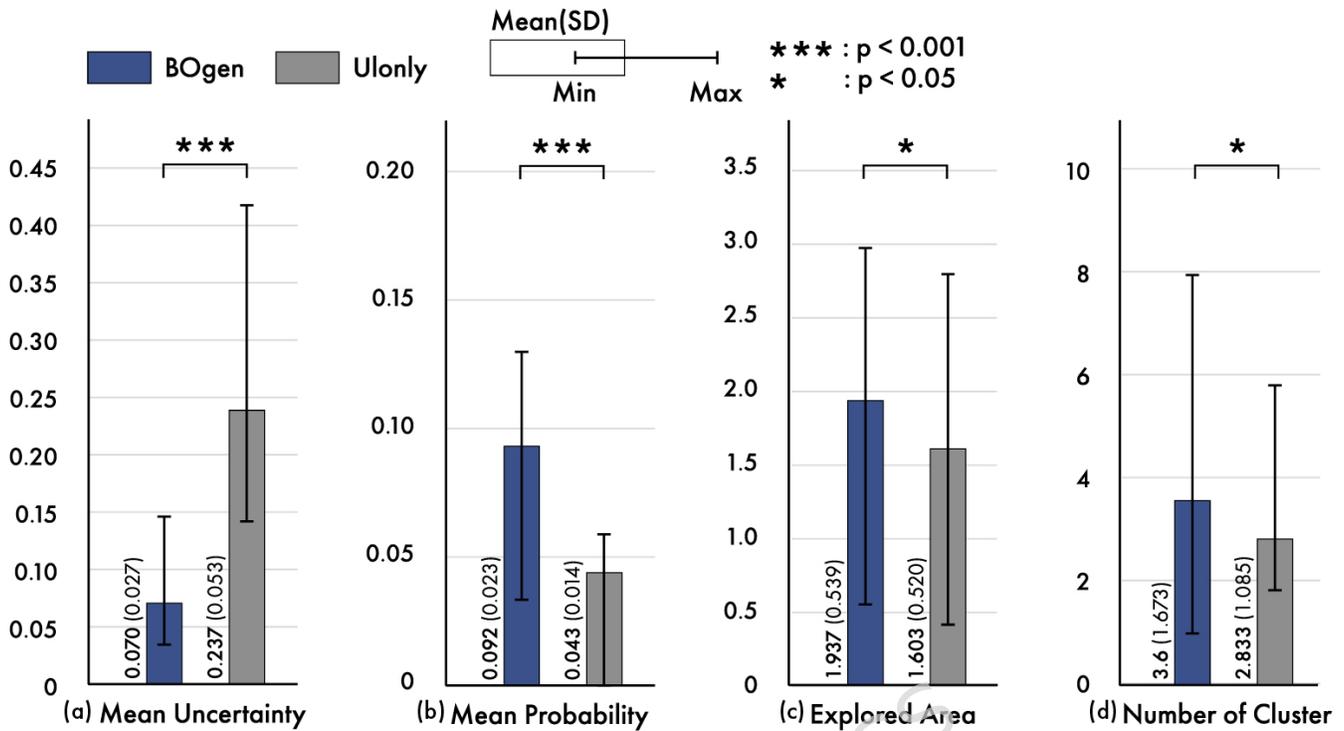


Figure 11. Results of (a) mean uncertainty; (b) mean probability; (c) explored area; and (d) number of clusters.

- P1: "With the system (BOgen), the map offered unexpected design inspirations and shapes. I could view something unique on the map and then scroll to see various variations with (PBO) suggestions, which was helpful."
- P7: "There were usable designs in the (PBO) recommendations. The map, with its color information, helped me look at both areas of interest and those that I had not previously considered, aiding in synthesizing elements and invoking images."
- P24: "I had a design intent, and the map distribution was helpful when specifying it. I was looking for a unique shape, and it was great not having to search but to have a guide. A similar recommendation function (PBO) is also helpful for seeking details. The ideation process benefited from observing several similar details."

In summary, these results demonstrate that BOgen captured user preferences more effectively, made accurate predictions about user-preferred designs, and guided design-space navigation for refining design ideas.

Recommending to Designers to Find Prominent Design Parts

Designers' intentions evolve as they progress through the design process. Hence, it is essential to develop designs that respond to these changing needs. In this regard, BOgen scored statistically significantly higher in surveys Q1 and Q4 compared to the Ulonly system (Figure 12; Q1 : $Mean_{BOgen} = 5.6, Mean_{Ulonly} = 4.233, p < 0.001$; Q4 : $Mean_{BOgen} = 5.8, Mean_{Ulonly} = 4.6, p < 0.001$). This indicates that BOgen more effectively recommended designs that included the desired design parts and provided valuable suggestions to users. To further investigate these results, we analyzed users' sequential exploration paths. For instance, P3 was searching for a chair with a rounded backrest. After searching "round table chair", P3 explored the 2D latent points of the resulting chairs. P3 found a chair corresponding to sequence 1 in figure 13a and combined it with the original search result, creating a chair in sequence 2. After that, they continued exploring and merging chairs. Subsequently, P3 searched "backrest chair" and combined the resulting design with their previously developed design. During the interview, P3 mentioned, "(In BOgen) I was looking for a chair with a rounded backrest. I explored around the dots on the map ... I went on to find the desired part by looking at the dots on the map." On the other hand, with the Ulonly system, P3 said (Figure 13b), "I had a hard time searching because I could only think of a few keywords. When I searched for a sofa, it didn't match my desired design, so I had to synthesize more." Observing Figure 13c, P5 started by searching for a stool, and while exploring within those results, in sequence 3, P5 merged their design with a chunky armchair and then chose a chair based on the PBO recommendation. In sequence 4, P5 found a folding chair on the exploration map and combined it, receiving another PBO recommendation and selecting a chair in sequence 5. P5 shared in an interview, "Because I only knew a limited type of chairs, I found using the map

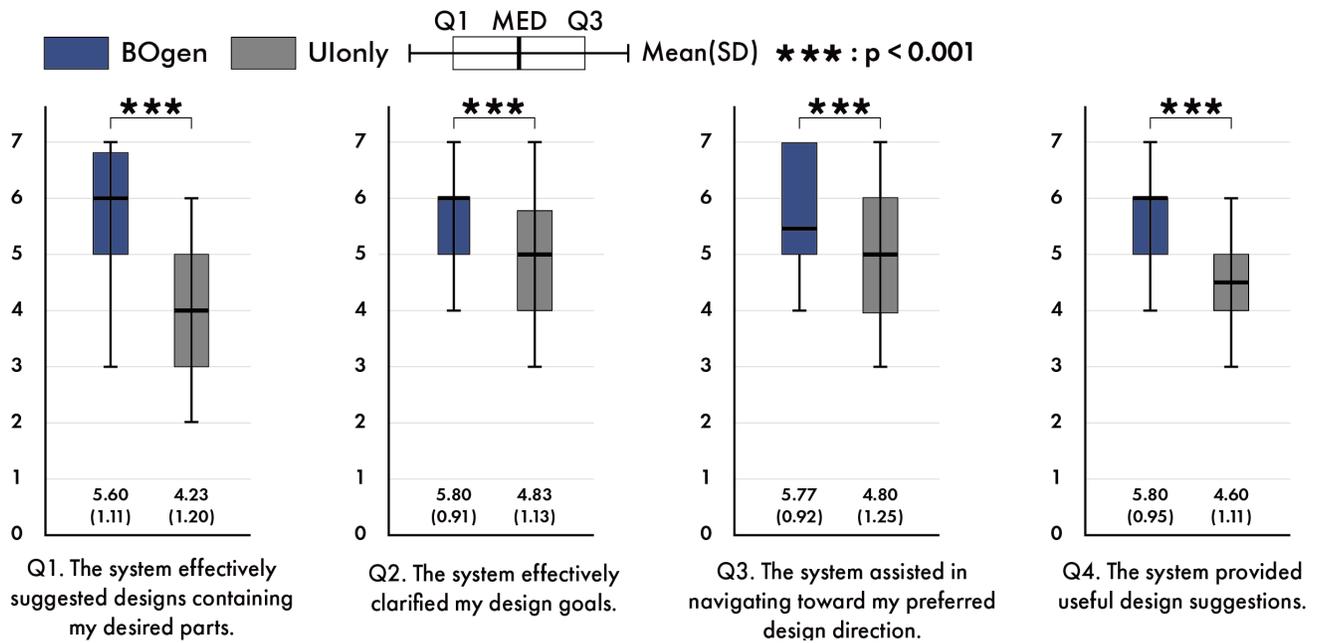


Figure 12. Results of survey.

in the system (BOGen) more helpful than just keywords (UOnly in Figure 13d)." Lastly, as seen in Figure 13e-f, P11 began by exploring the 2D latent points of their initial search result on the map. Particularly in sequence 3, P11 found desired chairs on the exploration map and kept developing their design. P11 said, "There were many unexpected discoveries from the map or (PBO) recommendations. The process was faster, and I could pull out many designs." To sum up, BOGen was superior to UOnly in effectively suggesting the design elements users were looking for and presenting them with prominent designs.

Facilitating Designers to Explore Dense and Diverse Design Space

In the design context, as one iteratively explores diverse designs and converges on a specific design, the design process becomes novel⁵². In this regard, it is imperative to understand how densely a user has explored the design space and how diverse the exploration has been. We measured the area of the design space explored by users in each system and the number of clusters they explored. First, the exploration area in BOGen was significantly larger, approximately 120% more extensive than that in UOnly (BOGen: 1.937, UOnly: 1.603, $p < 0.05$; Figure 11c). Next, BOGen produced more clusters (approximately 127%) than the UOnly system with our DBSCAN parameter setting (BOGen: 3.6, UOnly: 2.833, $p < 0.05$; Figure 11d). In BOGen, users synthesize designs by crossing multiple clusters, notably. For example, in figure 14a, looking at P6's use case with the BOGen system, P6 explored a total of four clusters. P6 developed designs by crossing over clusters 1 and 3 and clusters 1 and 2 (①+②+③, ④+⑤+⑥). Importantly, designs developed from cluster 1, 2, and 3 were synthesized with cluster 4 before ending the experiment (④+⑤+⑥+⑦+⑧+⑨). In contrast, in figure 14b for the UOnly system, designs were synthesized only within cluster 1, with designs centered around the synthesis of ① and ②, limited to a narrow area. P6's interview reflected: "(In BOGen) I searched in broad categories first, then worked in detail based on the map. Some of the design options were synthesized to refine the design idea. With the map, I felt that I explored a more diverse design space than with the other systems (UOnly). Without the map (UOnly), it felt limiting and repetitive." The results indicated that through BOGen, users explored a broader, denser, and more varied design space. These advantages of BOGen were also evident in in-depth interviews with other participants:

- P7: "In the system (UOnly), I could only see a limited range of chairs that appeared in the search, so I wasn't sure about the designs available. However, in the other system (BOGen), through the map and its colors, I could check areas of interest and even areas that I did not care about, allowing me to see the overall designs."
- P12: "In the system (BOGen), I could quickly and easily check what I wanted within a set of similar designs. The designs I had seen previously were marked on the map, which helped me think about exploring similar designs."
- P27: "The addition of a map allowed me to check distributed points to see chair features. When ideas were limited to searches alone, new ideas were discovered in less interesting areas (lighter parts) on the map. In addition, since similar

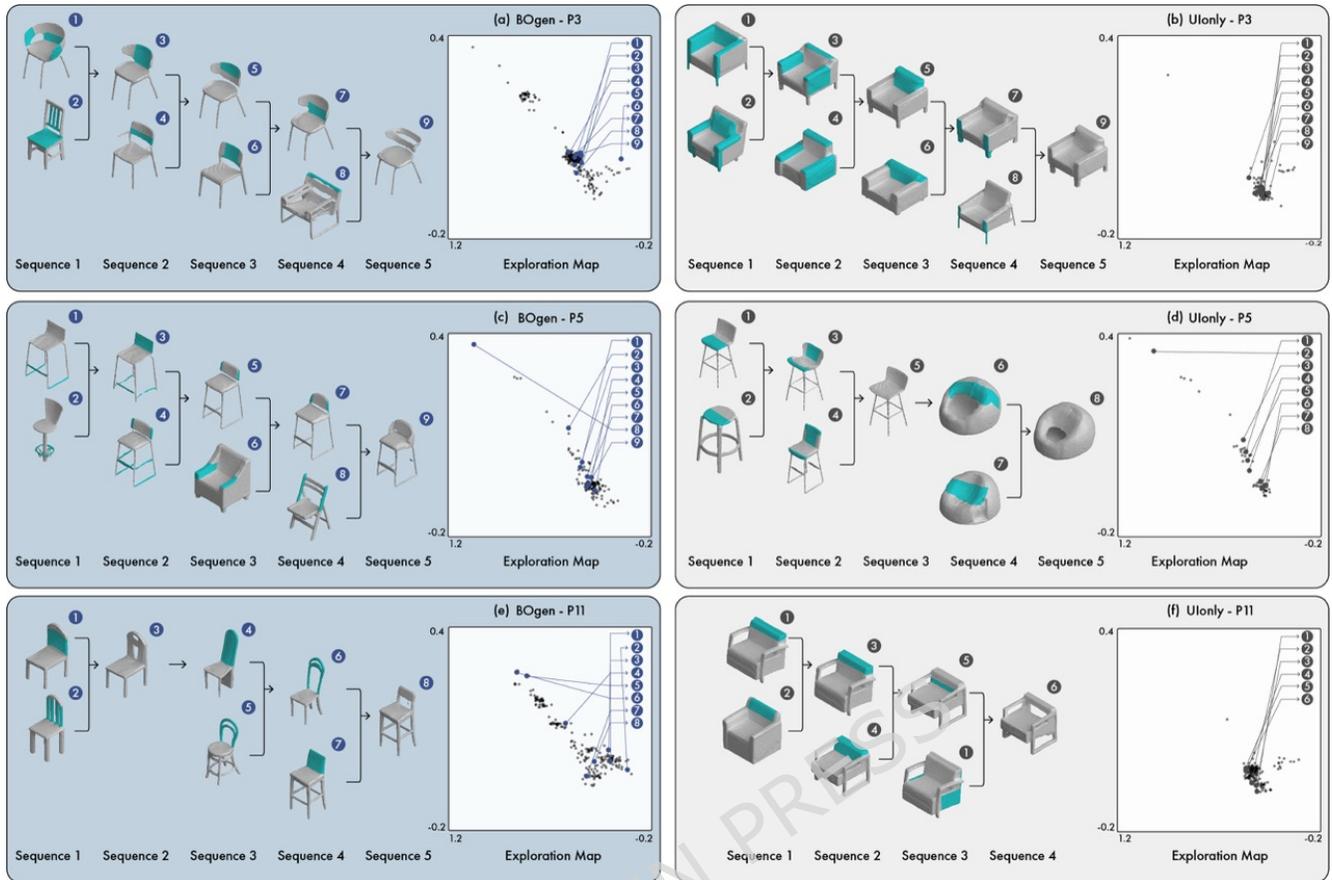


Figure 13. Sequential exploration path of BOgen and UIonly: (a) BOgen of P3; (b) UIonly of P3; (c) BOgen of P5; (d) UIonly of P5; (e) BOgen of P11; and (f) UIonly of P11.

designs were closely distributed, I could find the design I wanted without searching multiple times and verify new design ideas."

To summarize, BOgen not only guides the navigation of the design space for refining design ideas but also assists designers in finding prominent design parts. Additionally, it enables designers to delve into a dense and varied design space, thereby ensuring that users are more proactive in their design exploration.

Characterization of UIonly Baseline Limitations

Based on our experimental findings and participant interviews, we consolidate the key limitations observed in the UIonly baseline condition, which ablated the exploration map and PBO-based recommendation features.

First, users were constrained to exploring designs only within their existing vocabulary of text prompts. Without the exploration map, participants could only generate designs that they could verbally articulate, limiting discovery to their prior knowledge of chair types and design terminology. Several participants reported difficulty in searching due to their limited vocabulary of chair-related keywords.

Second, users had difficulty discovering unexpected or novel design possibilities. The UIonly condition restricted exploration to prompt-driven generation, preventing serendipitous encounters with designs outside the user's initial conceptualization. Participants noted that unexpected discoveries, which frequently occurred through the exploration map or PBO recommendations in BOgen, were absent in the UIonly condition.

Third, exploration patterns were narrower and less diverse. As reported in the quantitative analysis, the UIonly condition resulted in significantly smaller explored areas and fewer clusters compared to BOgen. Participants' exploration behaviors exemplified this pattern: designs in the UIonly condition tended to be synthesized within narrowly confined regions, whereas BOgen enabled cross-cluster synthesis spanning multiple distinct design regions.

These limitations suggest that text-prompt-only interfaces, while functional for known design targets, may inadvertently constrain the breadth of design exploration during early-stage ideation where designers benefit from exposure to unexpected

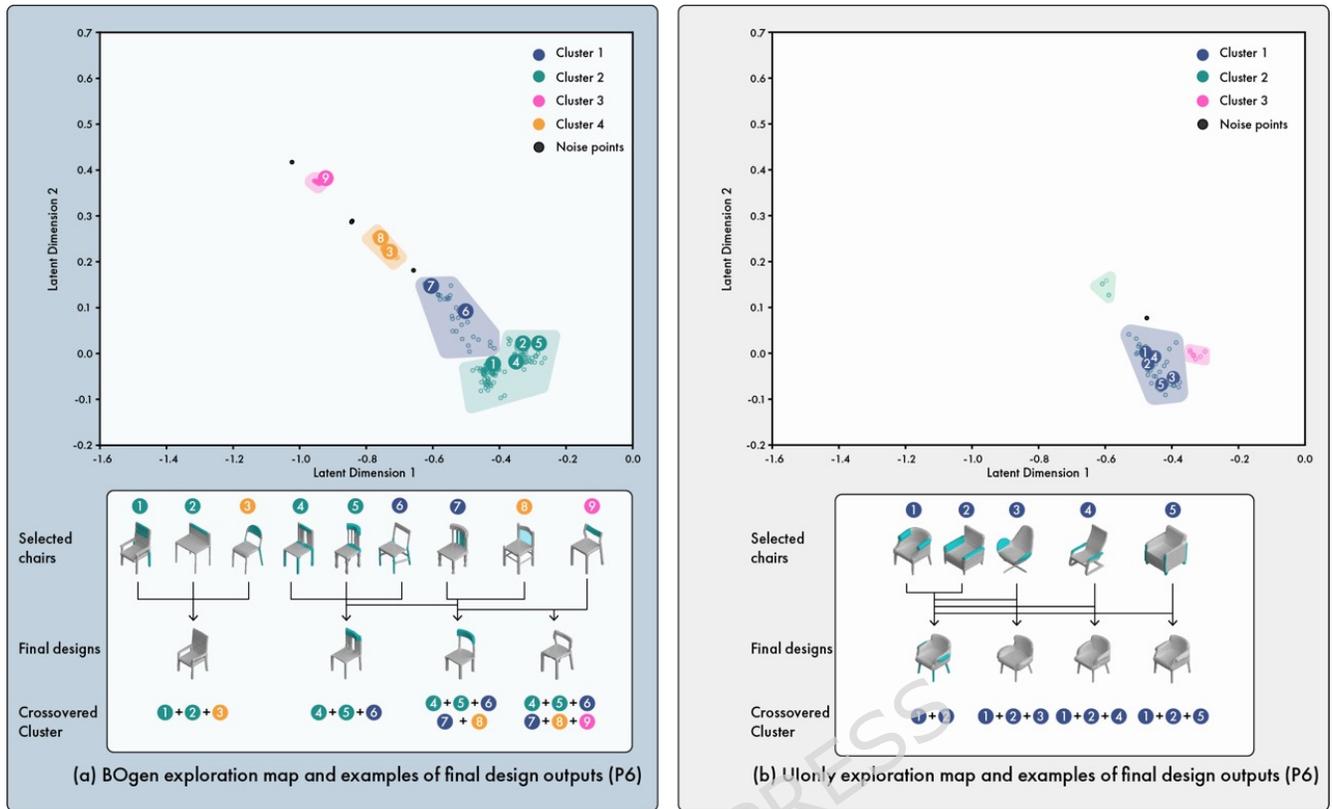


Figure 14. Exploration map and examples of final design outputs of P6: (a) BOGen; and (b) UOnly.

alternatives. The exploration map and PBO-based recommendations in BOGen address these limitations by providing visual navigation of the design space and proactive suggestions based on inferred preferences, respectively.

Implication

Recent advancements in 3D generative artificial intelligence (AI) have empowered designers and general users to easily explore high-quality design variations. However, compared to 2D generative AI, 3D generative AI has complex and challenging characteristics^{5,6}. Therefore, the generation and synthesis of assemble-part-level designs that satisfy user requirements have significant implications. Fundamentally, design is a combination of elements through which designers assess their appropriateness and creativity. In this context, our experimental results offer crucial insights into both the design process and design industry.

Initially, interfaces supporting iterative design processes for 2D and 3D generative AI were not sufficiently developed. To tackle this, we created a UOnly interface to support the evolution of designers' thought processes. Building on this, we enhanced the system with BOGen, implementing an information model that offers recommendations based on user actions. However, under the UOnly condition, users were limited to generating and synthesizing within their knowledge range, leading to a superficial exploration of the design. Conversely, BOGen enables designers to explore denser and more diverse design spaces, creating designs with prominent design elements based on assemble-part-level recommendations. Particularly in the UOnly situation, designers perceived the design space as only what the generative AI could create based on prompts, potentially missing unexplored design possibilities and limiting the discovery of desired part elements. However, BOGen's exploration map enables designers to understand the entire design space of generative AI, visualize their design preferences, and accurately recommend assemble-part-level designs, thereby fostering a more proactive design process.

Furthermore, 3D design contains significantly more information than 2D design, providing designers with deeper insights. Specifically, 3D design objects enable simulations such as structural and topology simulations, production cost, and aerodynamic evaluations, which are limited to 2D designs. Thus, BOGen has the potential to support early-stage design exploration in the design industry. It identifies preferred areas in a complex 3D chair design space based on designer preferences and generates assemble-part-level 3D chair design outputs. However, it should be noted that the current system optimizes for visual preferences and does not incorporate mechanical feasibility constraints; designs generated through BOGen would require additional validation for structural soundness and manufacturability before proceeding to fabrication stages. Although modeling

a desired 3D shape is a challenging task even for experts, BOgen simplifies this complex and difficult process using 3D generative AI. By reflecting designer preferences, BOgen creates feasible and plausible 3D chair design outcomes using a user-preferred design (i.e., skip information). While this approach has potential to support design prototyping and digital fabrication processes, enabling general users to explore user-customized 3D design concepts, additional engineering validation would be necessary to ensure structural integrity and manufacturability before physical production.

Limitations and Future Work

This study has several limitations that suggest directions for future research. First, BOgen optimizes for designers' visual and aesthetic preferences at the part level, but does not incorporate mechanical or structural feasibility constraints. A visually preferred design may not necessarily be mechanically sound or manufacturable. Integrating structural simulation or manufacturability analysis into the preference optimization process would enable end-to-end support from design exploration to fabrication.

Second, while our methodology—integrating VAE-based dimensionality reduction with PBO for part-level 3D generation—is not inherently domain-specific, our evaluation focused on chair design. Validating BOgen's effectiveness on other object categories supported by part-aware 3D generative models (e.g., tables, airplanes) remains an important direction for future work.

Finally, regarding experimental design, although we employed a counterbalanced design to mitigate order effects, carryover effects between conditions cannot be entirely eliminated in within-subject designs. Additionally, our UIonly baseline was designed to isolate the combined effect of the exploration map and PBO recommendations; however, this approach does not allow us to disentangle the individual contributions of each component.

Conclusions

In this study, we introduce the UIonly and BOgen systems designed to assist in the development of designer processes using 3D generative AI. We emphasize BOgen, which empowers designers to actively generate, synthesize, and explore part-level 3D chair designs. It utilizes preferential Bayesian optimization and a variational autoencoder, incorporating user behavior to enhance the design experience and enable efficient generation and synthesis. We effectively reduced the 3D generative AI latent space to a navigable 2D latent space, which is an exploration map for efficiently inferring the users' desired chair models.

From our experimental results, the PBO-based recommendation and exploration map features of BOgen showed three main findings in comparison to the UIonly. First, BOgen supports design-space navigation for the development of design ideas by significantly predicting the uncertainty of the information space and the probability of the selected chair designs. Second, BOgen successfully recommends the specific chair parts desired by the designer by supporting sequential design exploration, reflecting the changing intentions of the designer. Finally, BOgen supports designers by exploring and synthesizing denser and more diverse designs.

This paper presents a validated framework for managing high-dimensional design processes, with a particular focus on the initial stages of 3D generative AI applied to chair designs. Although our methodology initially concentrated on chair design, it demonstrates the potential for broader applicability in a more expansive information space, particularly for 3D digital asset generation and synthesis with dimensions similar to 16 (parts) \times 16. Specifically, dimension reduction through VAE and the PBO sampling and update process for user-preferred designs, linked with skip information, extends to various other design industries, including vehicles and human faces, showcasing its versatility in a wide array of 3D digital assets. By exploring the impact of 3D Generative AI on the design process, this study has significant implications for the design industry. BOgen, enables the generation of customized designs and part-level recommendations in a complex 3D shape space, thereby reflecting designer preferences. This approach allows a broader and more diverse exploration of the design space compared to the traditional UIonly approach, offering creative design combinations that were previously unexplored. Such methods effectively assist designers in evolving their design ideas and simplifying complex 3D shape tasks that are limited to 2D designs. In the broader context of the design industry, BOgen contributes to supporting early-stage design exploration and prototyping, potentially reducing iteration costs in the conceptual design phase, while realizing user-customized designs. This represents a significant shift in the design industry, improving the efficiency and creativity of 3D shape processes for both designers and users. Finally, BOgen allows users to explore designs more proactively. In the future, by integrating various modalities including prompting, sketching, and eye-gaze-based interactions, we can better understand user intentions and support a more designer-centric process in 3D generative AI.

Data Availability

The data underlying this article will be shared on reasonable request to the corresponding author.

References

1. Ramesh, A., Dhariwal, P., Nichol, A., Chu, C. & Chen, M. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125* **1**, 3, DOI: <https://doi.org/10.48550/arXiv.2204.06125> (2022).
2. Rombach, R., Blattmann, A., Lorenz, D., Esser, P. & Ommer, B. High-resolution image synthesis with latent diffusion models. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, DOI: <https://doi.org/10.1109%2Fcvpr52688.2022.01042> (IEEE, 2022).
3. von Rütte, D., Fedele, E., Thomm, J. & Wolf, L. Fabric: Personalizing diffusion models with iterative feedback. *arXiv preprint arXiv:2307.10159* DOI: <https://doi.org/10.48550/arXiv.2307.10159> (2023).
4. Zheng, X.-Y. *et al.* Locally attentional SDF diffusion for controllable 3d shape generation. *ACM Transactions on Graph.* **42**, 1–13, DOI: <https://doi.org/10.1145%2F3592103> (2023).
5. Liu, V., Vermeulen, J., Fitzmaurice, G. & Matejka, J. 3dall-e: Integrating text-to-image ai in 3d design workflows. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference, DIS '23, 1955–1977*, DOI: [10.1145/3563657.3596098](https://doi.org/10.1145/3563657.3596098) (Association for Computing Machinery, New York, NY, USA, 2023).
6. Sanghi, A. *et al.* Clip-sculptor: Zero-shot generation of high-fidelity and diverse shapes from natural language. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 18339–18348, DOI: [10.1109/CVPR52729.2023.01759](https://doi.org/10.1109/CVPR52729.2023.01759) (IEEE Computer Society, Los Alamitos, CA, USA, 2023).
7. Park, J. & Kang, N. Bmo-gnn: Bayesian mesh optimization for graph neural networks to enhance engineering performance prediction. *J. Comput. Des. Eng.* **11**, 260–271, DOI: [10.1093/jcde/qwae102](https://doi.org/10.1093/jcde/qwae102) (2024). <https://academic.oup.com/jcde/article-pdf/11/6/260/61212383/qwae102.pdf>.
8. Hertz, A., Perel, O., Giryas, R., Sorkine-Hornung, O. & Cohen-Or, D. Spaghetti: Editing implicit shapes through part aware generation. *ACM Transactions on Graph. (TOG)* **41**, 1–20, DOI: <https://doi.org/10.1145/3528223.3530084> (2022).
9. Hui, K.-H., Li, R., Hu, J. & Fu, C.-W. Neural wavelet-domain diffusion for 3d shape generation. In *SIGGRAPH Asia 2022 Conference Papers*, DOI: <https://doi.org/10.1145%2F3550469.3553394> (ACM, 2022).
10. Koo, J., Yoo, S., Nguyen, M. H. & Sung, M. Salad: Part-level latent diffusion for 3d shape generation and manipulation. *arXiv preprint arXiv:2303.12236* DOI: <https://doi.org/10.48550/arXiv.2303.12236> (2023).
11. Lee, H., Lee, J., Kim, H. & Mun, D. Dataset and method for deep learning-based reconstruction of 3d cad models containing machining features for mechanical parts. *J. Comput. Des. Eng.* **9**, 114–127, DOI: [10.1093/jcde/qwab072](https://doi.org/10.1093/jcde/qwab072) (2021). <https://academic.oup.com/jcde/article-pdf/9/1/114/41988785/qwab072.pdf>.
12. Nyamsuren, P., Lee, S.-H., Hwang, H.-T. & Kim, T.-J. A web-based collaborative framework for facilitating decision making on a 3d design developing process. *J. Comput. Des. Eng.* **2**, 148–156, DOI: [10.1016/j.jcde.2015.02.001](https://doi.org/10.1016/j.jcde.2015.02.001) (2015). <https://academic.oup.com/jcde/article-pdf/2/3/148/33133610/j.jcde.2015.02.001.pdf>.
13. Auyeskhani, U. *et al.* Virtual reality-based assembly-level design for additive manufacturing decision framework involving human aspects of design. *J. Comput. Des. Eng.* **10**, 1126–1142, DOI: [10.1093/jcde/qwad041](https://doi.org/10.1093/jcde/qwad041) (2023). <https://academic.oup.com/jcde/article-pdf/10/3/1126/52600312/qwad041.pdf>.
14. Son, K. & Hyun, K. H. Designer-centric spatial design support. *Autom. Constr.* **137**, 104195, DOI: <https://doi.org/10.1016/j.autcon.2022.104195> (2022).
15. Son, K., Lee, S. W., Yoon, W. & Hyun, K. H. Creativesearch: Proactive design exploration system with bayesian information gain and information entropy. *Autom. Constr.* **142**, 104502, DOI: <https://doi.org/10.1016/j.autcon.2022.104502> (2022).
16. Pandey, K., Chevalier, F. & Singh, K. Juxtaform: interactive visual summarization for exploratory shape design. *ACM Trans. Graph.* **42**, DOI: [10.1145/3592436](https://doi.org/10.1145/3592436) (2023).
17. Averkiou, M., Kim, V. G., Zheng, Y. & Mitra, N. J. Shapesynt: Parameterizing model collections for coupled shape exploration and synthesis. *Comput. Graph. Forum* **33**, 125–134, DOI: [10.1111/cgf.12310](https://doi.org/10.1111/cgf.12310) (2014).
18. Matejka, J. *et al.* Dream lens: Exploration and visualization of large-scale generative design datasets. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, CHI '18, 1–12*, DOI: [10.1145/3173574.3173943](https://doi.org/10.1145/3173574.3173943) (Association for Computing Machinery, New York, NY, USA, 2018).
19. Evirgen, N. & Chen, X. A. Ganzilla: User-driven direction discovery in generative adversarial networks. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology, UIST '22*, DOI: [10.1145/3526113.3545638](https://doi.org/10.1145/3526113.3545638) (Association for Computing Machinery, New York, NY, USA, 2022).

20. Evirgen, N. & Chen, X. A. Ganravel: User-driven direction disentanglement in generative adversarial networks. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI '23, DOI: [10.1145/3544548.3581226](https://doi.org/10.1145/3544548.3581226) (Association for Computing Machinery, New York, NY, USA, 2023).
21. Dang, H., Mecke, L. & Buschek, D. Ganslider: How users control generative models for images using multiple sliders with and without feedforward information. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, CHI '22, DOI: [10.1145/3491102.3502141](https://doi.org/10.1145/3491102.3502141) (Association for Computing Machinery, New York, NY, USA, 2022).
22. Wu, Y., Ma, L., Yuan, X. & Li, Q. Human-machine hybrid intelligence for the generation of car frontal forms. *Adv. Eng. Informatics* **55**, 101906, DOI: <https://doi.org/10.1016/j.aei.2023.101906> (2023).
23. Kadner, F., Keller, Y. & Rothkopf, C. AdaptiFont: Increasing individuals' reading speed with a generative font model and bayesian optimization. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, DOI: <https://doi.org/10.1145%2F3411764.3445140> (ACM, 2021).
24. Koyama, Y. & Goto, M. BO as assistant: Using bayesian optimization for asynchronously generating design suggestions. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*, DOI: <https://doi.org/10.1145%2F3526113.3545664> (ACM, 2022).
25. Lee, S. W., Kim, H., Yi, T. & Hyun, K. H. Bigaze: An eye-gaze action-guided bayesian information gain framework for information exploration. *Adv. Eng. Informatics* **58**, 102159, DOI: <https://doi.org/10.1016/j.aei.2023.102159> (2023).
26. Liu, W., d'Oliveira, R. L., Beaudouin-Lafon, M. & Rioul, O. Bignav: Bayesian information gain for guiding multiscale navigation. In *Proceedings of the 2017 CHI conference on human factors in computing systems*, 5869–5880, DOI: <https://doi.org/10.1145/3025453.3025524> (2017).
27. Son, K., Kim, K. & Hyun, K. H. BIGexplore: Bayesian information gain framework for information exploration. In *CHI Conference on Human Factors in Computing Systems*, DOI: <https://doi.org/10.1145%2F3491102.3517729> (ACM, 2022).
28. Danhaive, R. & Mueller, C. T. Design subspace learning: Structural design space exploration using performance-conditioned generative modeling. *Autom. Constr.* **127**, 103664, DOI: <https://doi.org/10.1016%2Fj.autcon.2021.103664> (2021).
29. Choi, J. & Hyun, K. H. Typeface network and the principle of font pairing. *Sci. Reports* **14**, 30820, DOI: <https://doi.org/10.1038/s41598-024-81601-w> (2024).
30. Umetani, N. Exploring generative 3d shapes using autoencoder networks. In *SIGGRAPH Asia 2017 Technical Briefs*, SA '17, DOI: [10.1145/3145749.3145758](https://doi.org/10.1145/3145749.3145758) (Association for Computing Machinery, New York, NY, USA, 2017).
31. Achlioptas, P., Diamanti, O., Mitliagkas, I. & Guibas, L. Learning representations and generative models for 3d point clouds. In *International conference on machine learning*, 40–49, DOI: <https://proceedings.mlr.press/v80/achlioptas18a.html> (PMLR, 2018).
32. Zhang, S. *et al.* Brep2seq: a dataset and hierarchical deep learning network for reconstruction and generation of computer-aided design models. *J. Comput. Des. Eng.* **11**, 110–134, DOI: [10.1093/jcde/qwae005](https://doi.org/10.1093/jcde/qwae005) (2024). <https://academic.oup.com/jcde/article-pdf/11/1/110/57275777/qwae005.pdf>.
33. Hao, Z., Averbuch-Elor, H., Snavely, N. & Belongie, S. DualSDF: Semantic shape manipulation using a two-level representation. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, DOI: <https://doi.org/10.1109%2Fcvpr42600.2020.00765> (IEEE, 2020).
34. Hyun, K. H. & Lee, J.-H. Balancing homogeneity and heterogeneity in design exploration by synthesizing novel design alternatives based on genetic algorithm and strategic styling decision. *Adv. Eng. Informatics* **38**, 113–128, DOI: <https://doi.org/10.1016/j.aei.2018.06.005> (2018).
35. Ban, S. & Hyun, K. H. 3d computational sketch synthesis framework: Assisting design exploration through generating variations of user input sketch and interactive 3d model reconstruction. *Comput. Des.* **120**, 102789, DOI: <https://doi.org/10.1016/j.cad.2019.102789> (2020).
36. Khan, S. & Awan, M. J. A generative design technique for exploring shape variations. *Adv. Eng. Informatics* **38**, 712–724, DOI: <https://doi.org/10.1016/j.aei.2018.10.005> (2018).
37. Zhou, Y., Koyama, Y., Goto, M. & Igarashi, T. Interactive exploration-exploitation balancing for generative melody composition. In *26th International Conference on Intelligent User Interfaces*, IUI '21, 43–47, DOI: [10.1145/3397481.3450663](https://doi.org/10.1145/3397481.3450663) (Association for Computing Machinery, New York, NY, USA, 2021).

38. Shi, M., Seo, J., Cha, S. H., Xiao, B. & Chi, H.-L. Generative ai-powered architectural exterior conceptual design based on the design intent. *J. Comput. Des. Eng.* **11**, 125–142, DOI: [10.1093/jcde/qwae077](https://doi.org/10.1093/jcde/qwae077) (2024). <https://academic.oup.com/jcde/article-pdf/11/5/125/59140299/qwae077.pdf>.
39. Chu, W. & Ghahramani, Z. Preference learning with gaussian processes. In *Proceedings of the 22nd international conference on Machine learning*, 137–144, DOI: <https://doi.org/10.1145/1102351.1102369> (2005).
40. Eric, B., Freitas, N. & Ghosh, A. Active preference learning with discrete choice data. *Adv. neural information processing systems* **20**, DOI: [10.5555/2981562.2981614](https://doi.org/10.5555/2981562.2981614) (2007).
41. González, J., Dai, Z., Damianou, A. & Lawrence, N. D. Preferential bayesian optimization. In *International Conference on Machine Learning*, 1282–1291, DOI: <https://doi.org/10.48550/arXiv.1704.03651> (PMLR, 2017).
42. Brochu, E. *Interactive Bayesian optimization: learning user preferences for graphics and animation*. Ph.D. thesis, University of British Columbia (2010). DOI: <https://dx.doi.org/10.14288/1.0051462>.
43. Brochu, E., Cora, V. M. & De Freitas, N. A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv preprint arXiv:1012.2599* DOI: <https://doi.org/10.48550/arXiv.1012.2599> (2010).
44. Brochu, E., Brochu, T. & De Freitas, N. A bayesian interactive optimization approach to procedural animation design. In *Proceedings of the 2010 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, 103–112, DOI: [10.5555/1921427.1921443](https://doi.org/10.5555/1921427.1921443) (2010).
45. Gelbart, M. A., Snoek, J. & Adams, R. P. Bayesian optimization with unknown constraints. *arXiv preprint arXiv:1403.5607* DOI: <https://doi.org/10.48550/arXiv.1403.5607> (2014).
46. Hernández-Lobato, J. M., Gelbart, M. A., Adams, R. P., Hoffman, M. W. & Ghahramani, Z. A general framework for constrained bayesian optimization using information-based search. *J. Mach. Learn. Res.* **17**, 1–53, DOI: <https://doi.org/10.48550/arXiv.1511.09422> (2016).
47. Liu, W., Rioul, O., Mcgreneire, J., Mackay, W. E. & Beaudouin-Lafon, M. Bigfile: Bayesian information gain for fast file retrieval. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, 1–13, DOI: <https://doi.org/10.1145/3173574.3173959> (2018).
48. Bradley, R. A. & Terry, M. E. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika* **39**, 324, DOI: <https://doi.org/10.2307/2334029> (1952).
49. Seeger, M. Gaussian processes for machine learning. *Int. journal neural systems* **14**, 69–106, DOI: [10.1142/S0129065704001899](https://doi.org/10.1142/S0129065704001899) (2004).
50. Srinivas, N., Krause, A., Kakade, S. M. & Seeger, M. W. Information-theoretic regret bounds for gaussian process optimization in the bandit setting. *IEEE Transactions on Inf. Theory* **58**, 3250–3265, DOI: <https://doi.org/10.1109/2Ftit.2011.2182033> (2012).
51. Schonlau, M., Welch, W. J. & Jones, D. R. Global versus local search in constrained optimization of computer models. In *Institute of Mathematical Statistics Lecture Notes - Monograph Series*, 11–25, DOI: <https://doi.org/10.1214/2Flnms%2F1215456182> (Institute of Mathematical Statistics, 1998).
52. Goldschmidt, G. Linkographic evidence for concurrent divergent and convergent thinking in creative design. *Creat. research journal* **28**, 115–122, DOI: <https://doi.org/10.1080/10400419.2016.1162497> (2016).

Funding

This work was supported by the Technology Innovation Program (RS-2025-02317326, Development of AI-Driven Design Generation Technology Based on Designer Intent) funded by the Ministry of Trade, Industry & Energy (MOTIE, Korea) and National Research Foundation of Korea (NRF) grant funded by the Korean government (MSIP: Ministry of Science, ICT and Future Planning) (RS-2023-00208542).

Author contributions statement

S.W.L. led the conceptualization and methodology design, developed and prepared the original draft of the manuscript. J.C. contributed to writing, reviewing, and editing the manuscript. K.H.H. supervised the project and participated in writing, reviewing, and editing. All authors reviewed and approved the final version of the manuscript.

Competing interests

The authors declare no competing interests.

ARTICLE IN PRESS