

Fit and Diverse: Set Evolution for Inspiring 3D Shape Galleries

Kai Xu*[†]

Hao Zhang[‡]

Daniel Cohen-Or[§]

Baoquan Chen*

*Shenzhen Institutes of Advanced Technology (SIAT)

[†]National University of Defense Technology (NUDT)

[‡]Simon Fraser University

[§]Tel-Aviv University

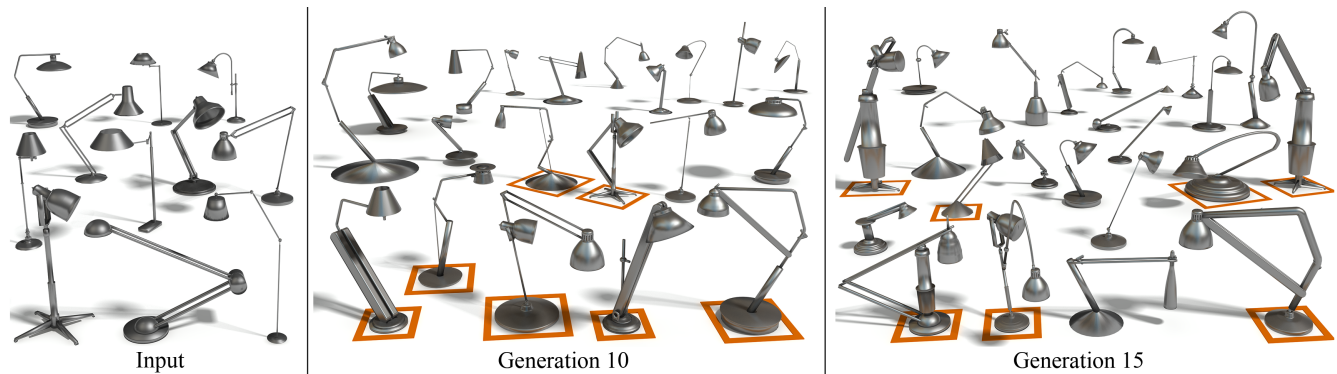


Figure 1: Set evolution starting from a small input set of lamps (left). With the set evolution “fit and diverse”, new generations of shapes are not only fit to be lamps but also exhibit significant and potentially inspiring variations.

Abstract

We introduce *set evolution* as a means for creative 3D shape modeling, where an initial population of 3D models is evolved to produce generations of novel shapes. Part of the evolving set is presented to a user as a shape gallery to offer modeling suggestions. User preferences define the fitness for the evolution so that over time, the shape population will mainly consist of individuals with good fitness. However, to inspire the user’s creativity, we must also keep the evolving set diverse. Hence the evolution is “*fit and diverse*”, drawing motivation from evolution theory. We introduce a novel part crossover operator which works at the finer-level part structures of the shapes, leading to significant variations and thus increased diversity in the evolved shape structures. Diversity is also achieved by explicitly compromising the fitness scores on a portion of the evolving population. We demonstrate the effectiveness of set evolution on man-made shapes. We show that selecting only models with high fitness leads to an elite population with low diversity. By keeping the population fit and diverse, the evolution can generate inspiring, and sometimes unexpected, shapes.

Links: [DL](#) [PDF](#) [WEB](#) [VIDEO](#) [DATA](#)

1 Introduction

In nature, no two creatures are exactly alike. During the evolution of a species, nature implicitly maintains a genetic diversity as a means for the species to better adapt to changing environments. In

product design, parallels can be drawn. Products intended to serve the same function often come in a variety of shapes and structures to adapt to the ever changing customer needs and tastes, and to inspire new designs. In computer graphics, 3D content creation remains a central and difficult problem. The ability to create a diverse set of 3D models allows one to populate virtual worlds with enriched 3D content and improve user experience.

We view 3D shape modeling as a creative task, whether for product design or scene modeling. Creativity has always been a trait bestowed to humans but not machines. An interesting question is whether a machine can assist humans in being creative and *inspire* a user during 3D modeling. One possible means to achieve this is through a design gallery which presents a variety of suggestive designs from which the user can pick the ones he or she likes the best [Marks et al. 1997]. The ensuing challenge is how to come up with intriguing suggestions which inspire creativity, rather than banal suggestions which stall the creative process.

In this paper, we introduce *set evolution* as a means for creative 3D shape modeling (Figure 1). Our goal is to create generations of novel 3D shapes starting from an initial population, where the new models can not only be adopted to populate virtual scenes but also make potentially inspiring suggestions for future creations.

Fit and diverse. During evolution, part of the evolving set is presented to a user as a *shape gallery*. User preferences define the fitness function [Sims 1994] for the evolution as he or she selects shapes from the gallery that are deemed to be plausible (a chair needs to be “chair-like”) and “liked” to breed the next generation. Through time, the shape population will mainly consist of fit individuals. However, we would also like our creations to potentially inspire the user. A key point we advocate, drawing motivation from the role diversity plays in evolution, is that, to inspire creativity, the evolving set needs to be kept *diverse*. We compromise the fitness scores of a portion of the evolving population as a means to maintain diversity. Specifically, we define a diversity measure for the evolving set and use it to control a trade-off between the counterbalancing objectives of “being fit” and “being diverse” for the set. Explicitly maintaining the diversity of the evolving set increases the potential to create surprising and inspiring suggestions.



Figure 2: The evolving population (left) consists of a diverse background set (in gray) and a fit foreground set (in gold). The gallery of shapes that is presented to the user is illustrated on the right, which consists of shapes taken from the foreground set.

Evolution of a set. We evolve an entire population of 3D models which belong to a certain class (e.g., chairs, teapots), instead of individual shapes one at a time. Rather than suggesting parts during shape composition [Chaudhuri and Koltun 2010], we offer a set of complete shapes as suggestions in each generation of the evolution (Figure 2). Throughout the process, we focus on the quality of the set rather than that of the individuals.

Evolution operators. To execute the set evolution, we develop a genetic algorithm to automatically create generations of shapes. We assume that the initial set of shapes have been pre-analyzed to possess part correspondence and built-in structural information such as inter-part symmetries. Our algorithm is built on a mutation operator and a crossover operator which correspond to part warping and part replacement, respectively. Since the fitness function only selects but does not create, the diversity of the evolving set depends on the above operators to come up with creative suggestions.

To this end, we introduce a novel part crossover operator which acts on *fine-scale* shape parts without relying on accurate correspondences between them. Throughout the evolution, the coarse- or meta-level part structure of the shapes across the whole population is maintained. The finer-level structures within the meta parts can vary significantly between different shapes, which contribute to the diversity of the set. Our part crossover builds upon the notion of *fuzzy part correspondence* to carry out *many-to-many* exchange between the fine-scale parts. Coupled with part mutation, our shape reproduction goes beyond part shuffling [Funkhouser et al. 2004; Kreavoy et al. 2007; Chaudhuri and Koltun 2010] for shape composition. The resulting creations can exhibit significant variations in the overall shape structures, even in topology (Figure 3).

Contributions. Our contributions to creative 3D modeling can be summarized as follows:

- A set evolution which creates generations of novel 3D shapes. We focus on the quality of a set, rather than that of the individuals, and offer complete shapes as creative suggestions.
- A mechanism to keep an evolving set fit to the user’s design preferences while maintaining diversity of the set in a controlled manner. The key idea of maintaining diversity provides a source for creative suggestions and it contrasts the intuitive tendency to produce an elite population.
- A novel crossover operator based on fuzzy part correspondence which creates diverse shape (even topological) variations, while maintaining the built-in meta part structure.

We demonstrate the effectiveness of our set evolution for creative modeling of 3D man-made shapes. We show that if the set is evolved by only selecting individuals with high fitness scores, it would eventually turn into an elite population, one with a low diversity and lack of creative potential. On the contrary, by keeping the population diverse, the evolution succeeds in creative modeling. With our novel part crossover operator, we show that even when starting with a small but diverse set, the future generations can grow in size as well as diversity; see Figure 1.

2 Related works

Evolutionary modeling and design. The seminal works of Karl Sims [1991; 1994] first introduced genetic algorithms to the computer graphics community to synthesize novel creatures with desired physical behavior. In his work, both the shape and function of a creature are evolved simultaneously. Many follow-up works have appeared since, e.g., the creature academy of Pilat and Jacob [2008]. As well, Pollack et al. [1998] adopted an evolutionary framework in the design and generation of assembled objects such as robots and Jakiela et al. [1997] studied structural topology design in the context of mechanical design. The application domain of evolutionary design has spanned visual arts [Draves 2006], even music [Romero and Machado 2007]. A distinguishing feature of our problem setup is that we evolve a set simultaneously and the focus is on the quality, particularly diversity, of the set.

Peter Bentley [1999] introduced evolutionary design and collected several state-of-the-art works in the field, including evolutionary design for urban planning [Soddu and Colabella 1995] and architecture [Frazer 1995]. He [Bentley 2000] suggests that the creativity



Figure 3: Our part crossover and mutation produce significant shape variations, even topology changes. The offsprings are generated automatically and we trace the evolution paths.

of the evolutionary approach depends on the exploration of a low-level, knowledge-lean representation of solutions as it offers more flexibility for the evolution. Our work indeed exploits low-level genetic representations derived from the finer-level shape parts.

Data-driven object modeling The meta work by Funkhouser et al. [2004] on “modeling by examples” has pioneered the direction of data-driven 3D object modeling. The design task is fully controlled by the user, while a content-driven search allows geometry variations to be added via part substitution and combination. This search-and-assemble modeling paradigm has been widely received over the years, e.g., [Shin and Igarashi 2007; Kreavoy et al. 2007; Lee and Funkhouser 2008; Fisher et al. 2011]. The modeling inspirations typically come in the form of relevant part libraries as a user incrementally composes a new model, where part suggestions are driven by geometric and contextual similarity among the shape parts [Funkhouser et al. 2004; Kreavoy et al. 2007; Chaudhuri and Koltun 2010; Chaudhuri et al. 2011].

Specifically, Chaudhuri et al. [2010] developed an approach for generating data-driven part suggestions as creativity support, where they discuss the need for offering unexpected suggestions. The method is then extended to provide more semantically and stylistically compatible suggestions with probabilistic reasoning [Chaudhuri et al. 2011]. Our set evolution evolves an entire population of 3D models, instead of editing one shape at a time. Rather than suggesting parts during shape composition, we offer a set of complete shapes as suggestions in each generation of the evolution.

Design space exploration. The influential work by Marks et al. [1997] on design galleries certainly provided an inspiration for our interactive evolutionary modeling approach, whereby the user actively participates in the evolution process. Design galleries provide an interface to assist the user in selecting parameters through a visual display of random solutions. The critical component of design galleries is “dispersion”, which finds a set of parameters mapped to the most diverse representative solutions to maximize coverage of the explored space. Our creative modeling utilizes a shape gallery to steer the evolution process, where diversity also plays a key role, but with a different goal in mind.

Also falling into the category of data-driven modeling, the work by Shapira et al. [2009] allows a user to actively navigate a space of design galleries. The user does not necessarily have a mental target of what is sought. While exploring the space of solutions, the user is inspired by what he/she sees. Yang et al. [2011] present a method for exploring a space of polygonal meshes possessing the same combinatorics, where the space is characterized by non-linear constraints associated with mesh elements. Talton et al. [2009] present a collaborative design space, where a community of users define and explore a wide variety of models. This work offers a unique tool for casual users to easily create high-quality 3D models in their entirety just by navigating the design space. Over time, the space of models can grow and be re-organized for efficient reuse.

3 Set evolution

We present a set evolution technique which adopts the principles of “survival of the fittest” and “population diversity” from evolution theory. Nevertheless, our technique is not aimed to explain or mimic nature, but merely to serve as a driving mechanism that continuously develops galleries to offer the user inspiring 3D shapes. Our technique evolves an entire population of 3D models of some semantic class. In each generation, the system generates sets of new shapes from the current generation through crossovers and mutations (Section 4). A selected subset is presented via a gallery to the user who provides feedback to the system by rating them according

to his/her preference, which defines the fitness function for the evolution. Through the evolution, the set is personalized and populated with shapes that better fit to the user. At the same time, the system explicitly maintains the diversity of the population so as to prevent it from converging into an elite set; see Algorithm 1.

3.1 Interactive set evolution

We classify the evolving population of shapes into two disjoint subsets. The foreground or breeding set \mathcal{G} and background set \mathcal{B} (Figure 2). The breeding set is the set whose elements breed to produce new generations. The rest of the population is kept for the sake of diversity. In each generation a controlled portion of the background set is upgraded and included in the breeding set.

Starting from an initial population \mathcal{G}_0 , our system repeatedly generates new models for generation \mathcal{G}_{i+1} from generation \mathcal{G}_i . During the shape reproduction for \mathcal{G}_{i+1} , the system generates M_d descendants or offsprings, denoted by the set \mathcal{D} , using evolutionary operations. Among the offsprings, our system automatically selects M_f fit ones to be presented in a gallery and lets the user rate them according to his/her preference. For each model in \mathcal{D} , we either increase its fitness score if the user likes it, or remove it from the population if the user dislikes it, or otherwise simply add it to the rest of the population \mathcal{B} . The reproduction is repeated until the number of accepted new models reaches M_g , at which point we move to produce the next generation. Algorithm 1 describes the flow of our interactive set evolution. In all our experiments, we set $M_d = 16$, $M_f = 9$, and $M_g = 40$.

The initial population is a set of pre-analyzed 3D models. We describe it in detail when we present part mutation and crossover in Section 4. The set should be reasonably rich and diverse and the shapes therein sufficiently developed. Our set evolution does not explain the evolution of low level creatures into a high-level species. That said, the kind of sets we deal with are common everyday objects such as lamps, teapots, chairs, and candelabra.

3.2 Fit and diverse

Fitness function. For each individual, the fitness score is determined by two factors: (i) an objective fit score and a subjective or

Algorithm 1: Interactive set evolution

```

Input : Initial population  $\mathcal{G}_0 = \{m_i^0\}_{i=1}^N$ ; Background set  $\mathcal{B}$ 
 $\mathcal{B} \leftarrow \mathcal{G}_0$ ;
 $i = 0$ ;
while the user is not satisfied do
   $\mathcal{G}_{i+1} = \emptyset$ ;
  while size of  $\mathcal{G}_{i+1}$  is less than  $M_g$  do
     $\mathcal{D} \leftarrow \text{Reproduce}(\mathcal{G}_i, M_d)$ ;
     $\mathcal{D} \leftarrow \text{SelectFit}(\mathcal{D}, M_f)$ ;
    foreach model  $m_s \in \mathcal{D}$  do
      if  $m_s$  is liked by the user then
        Increase the fitness of  $m_s$ ;
      else if  $m_s$  is disliked by the user then
         $\mathcal{D} = \mathcal{D} - \{m_s\}$ ;
     $\mathcal{G}_{i+1} = \mathcal{G}_{i+1} \cup \mathcal{D}$ ;
   $\mathcal{G}_{i+1} \leftarrow \text{SelectDiverse}(\mathcal{G}_{i+1} \cup \mathcal{G}_i \cup \mathcal{B}, M_n)$ ;
   $\mathcal{B} \leftarrow \text{SelectDiverse}(\mathcal{B} \cup \mathcal{G}_{i+1}, M_b)$ ;
   $\mathcal{G}_i \leftarrow \mathcal{G}_{i+1}$ ;
   $i = i + 1$ ;

```




Figure 4: Part crossover showing many-to-many exchanges at the sub-part level. Parts involved in crossover are marked in red.

personalized score. The objective score is meant to quickly reject unsuccessful offsprings, avoiding presenting them to the user in the first place. Identifying unsuccessful offsprings is a non-trivial task. In our implementation, we simply estimate how likely the object is able to stand upright. Specifically, we compute the projection of the model’s center of mass and test whether it is within the convex hull of its supporting points, based on upright orientation [Fu et al. 2008]. That measure, denoted by f_s , is a binary value where 1 implies that the object can stand well and 0 otherwise. The rejection mechanism is rather conservative and the user can always come in and reject any shape when it is presented in the gallery.

The subjective term is a continuous value recording the model’s history of being liked by the user. Whenever a user likes the model, its subjective term, denoted by f_l , is doubled. During reproduction, the subjective term is propagated to the offsprings by:

$$f_l(m_d) = \max\{f_l(m_d), f_l(m_1)p_1 + f_l(m_2)p_2\}, \quad (1)$$

where $p_i, i = 1, 2$, is the percentage of parts selected from parent m_i during the reproduction of m_d . The initial value for f_l is set to be 1. The final fitness function is defined as: $f = f_s f_l$.

Diversity control. To control the diversity of generation \mathcal{G}_{i+1} , we refine its content as follows. The refined set consists of M_n models selected from \mathcal{G}_{i+1} , \mathcal{G}_i , and \mathcal{B} , with relative portions 80%, 15% and 5%, respectively, in the descending order of fitness scores for each of the three sets. To control the storage space while keeping diversity, we also perform diversity control for the background set by selecting the top M_b most diverse models from $\mathcal{B} \cup \mathcal{G}_{i+1}$ and remove the rest. In all experiments, $M_b = 120$ and $M_n = 30$.

As a means to measure diversity between shapes, we resort to the Light Field Descriptor (LFD) [Chen et al. 2003] as a similarity measure. First, we embed the LFD descriptors of all models in the evolved set into 3D Euclidean space using Multidimensional Scaling (MDS). To select the most diverse models, we rely on farthest point sampling in MDS space. Specifically, we first select the point which is farthest from the center and then repeatedly select points which have the farthest average distance from the selected set, until the desired number of diverse models have been selected.

4 Part mutation and crossover

A basic consideration for set evolution is to ensure meaningful offsprings while avoiding as much as possible invalid ones, so as to alleviate the user’s effort when making selections in the shape gallery. This is achieved in our approach by storing pre-analyzed

shape structures in the initial set and preserving the stored structures throughout the evolution. More importantly, the “fit and diverse” characteristic of our evolution requires the reproduction operators, mutation and crossover, to also generate significant variations. This is achieved by allowing random mutation and crossover of shape parts and enabling crossover of a finer granularity of parts.

4.1 Shape representation

Each shape in the initial set is pre-segmented with each part enclosed by a controller, either a cuboid or a generalized cylinder (GC) [Zheng et al. 2011]. Symmetry and proximity relations between the parts per shape are pre-analyzed and stored as part of the controller representation. The set of shapes have pre-established correspondence at their coarse-level components, which we refer to as the *meta parts*, e.g., a leg or back of a chair. Each meta part may have a finer level of parts, the *sub-parts*, e.g., the back of a chair may be formed by several smaller components (Figure 4). From now on, parts refer to sub-parts unless otherwise noted.

The evolutionary operations, mutation and crossover, are performed directly on the controllers. When a new shape is created, it inherits the controllers from its parents, as well as the meta part correspondence and structural constraints, when appropriate. Our structure-preserving reproduction scheme builds on the component-wise controller framework of Zheng et al. [2011]. However, to allow more degrees of freedom in part mutation and crossover, we ignore controller constraints related to orthogonality and coplanarity.

Two additional types of controller constraints are added to accommodate the more drastic shape variations our reproduction operations allow. In some cases, part replacements and significant deformations of the controllers may compromise the geometric coherence of a new model, e.g., connected parts are detached or ground touching parts are off-ground. We detect physical connectivity between parts in the input set and propagate the connectivity to subsequent generations. Enforcing the connectivity constraint leads to a “snapping” between two detached parts. Important to man-made shapes is ground support. The second constraint ensures that ground touching parts remain so during evolution. Specifically, the contact regions with the ground are identified in the input set. If in an offspring model, such regions are off-ground, they are snapped to the ground. Certainly, additional consideration of appropriate physical or geometric constraints is also possible.

4.2 Part crossover

Crossover happens between two (parent) shapes and they produce two new shapes (the offsprings). With a given probability, currently set at 5%, an offspring goes through a mutation right after crossover. A crossover involves exchange of parts between the parents and it shares the same goal as shape modeling via part re-composition [Funkhouser et al. 2004; Kreavoy et al. 2007; Chaudhuri and Koltun 2010]. However, there are two significant differences. First, our crossover does not occur only between corresponding parts; in fact, the finer-level sub-parts may not possess accurate correspondence — only the meta parts do. To this end, we introduce *fuzzy part replacement* based on *fuzzy part correspondence*. Second, our part exchange is *many-to-many* instead of one-to-one, which explains the use of the term “fuzzy”; see Figure 4.

Overview. Given two parent shapes S and T , we only describe how to produce a crossover from S into T ; crossover in the other direction is similar. First, we select subsets of parts R_S in S and R_T in T to perform part exchange or replacement. The choices are randomized but aim to ensure a high likelihood that the offspring




Figure 5: Color-coded visualization of fuzzy part correspondence (red color: larger FPC value; yellow: small).

contains at least some sub-parts in each meta part to have a sufficient coverage of the major, i.e., coarse-level, components of each shape. Next, R_T is removed from shape T and parts in R_S are placed into T . Part placement is dictated by fuzzy part correspondence, essentially a *confidence* measure for replacement between any two parts. We progressively place parts from R_S into T in decreasing order of confidence. The confidence of placing a part is not defined based on the part alone, but also based on the confidence in the part’s neighbors, which form the *context*.

Fuzzy part correspondence. We wish to judge sub-part replaceability by functionality and not shape alone. With a focus on man-made shapes, we rely on the spatial position of a part in the context of the whole shape as a rough way of modeling the part’s functionality. For example, the legs, seats, arm rests, and backs of chairs are all spatially arranged in a predictable manner. To define fuzzy part correspondence, or FPC for short, we first globally align all the shapes in the set. To account for varying part scales, we rely on the method of Xu et al. [2010] to factor out part proportion variations in the set. After aligning the shapes, we characterize each part by its tightest oriented bounding box (OBB). Finally, the FPC measure between two parts p and q is defined as $\theta(p, q) = 1.0 - d(p, q)/\ell$, where $d(p, q)$ is the Hausdorff distance between the OBB’s of p and q and ℓ is the average diagonal length of the OBB’s of all the (whole) input shapes. Figure 5 visualizes FPC values between some parts taken from a chair set and a lamp set. In a concurrent work, Kim et al. [2012] propose a method for computing fuzzy correspondences between feature points over a set of shapes via diffusion maps.

Context-based part substitution. Given the subset of parts R_S from source shape S and the target shape T , we place one part from R_S into T at a time based on FPC of the parts. Let $p \in R_S$ be a part to be placed in. Due to structural discrepancies, p may not have a sufficiently confident target in T based on FPC. Thus, instead of




Figure 6: Illustration of context-based part placement. The (vertical) red part can be placed in the other shape since the contexts, matching parts in (a), are sufficiently replaceable based on fuzzy part correspondence. A topology change is introduced (b).

accounting for FPC between a source and target part alone, we also resort to a *context-based* approach. That is, p can be placed into T as long as the neighbors of p in R_S have sufficiently confident counterparts in T . This is illustrated in Figure 6, where we see that a topological variation is possible.

To execute context-based part placement, first, the set of *boundary parts* for R_S is substituted into shape T . A part p is substituted in this step if it is adjacent to some part in R_S and there is a part $q \in T$ that is sufficiently replaceable, i.e., $\theta(p, q) > \theta_b$. The placement of the boundary parts forms the initial context. Then we place a part $p \in R_S$ if p has a counterpart $q \in T$ where the FPC $\theta(p, q)$ between p and q is sufficiently high, $\theta(p, q) > \theta_r$. Throughout our experiments, the threshold values are set as $\theta_b = 0.8$ and $\theta_r = 0.96$. Finally, we consider contexts. Each unprocessed part $p \in R_S$ is assigned a confidence value, which is the sum of FPC values of p ’s neighbors in the updated shape T . The neighbors include boundary parts as well as parts already substituted in. We form a priority queue in decreasing confidence for the remaining parts in R_S and update it after a part goes into T . Figure 7 shows a short sequence demonstrating this.

Finding crossover subsets. We consider three criteria when selecting the subsets R_S and R_T for crossover: 1) to avoid spatially large overlap among the parts in the offspring, a part in R_S should not have large overlap with any part in $T \setminus R_T$; 2) to preserve symmetry constraints, parts belonging to the same symmetry group should be selected (or not selected) simultaneously; 3) the key consideration is that the parts in the offspring as a whole should sample all the semantic parts of the shapes. The first two criteria are fairly




Figure 7: Step-by-step illustration of context-based part crossover. The crossover (part) subsets of both shapes S and T are marked in red (left). Our method first places the boundary parts (marked yellow) of S into T (Step 1). Then parts in the crossover subset of S which have sufficiently replaceable counterparts in T are placed (Step 2). Under the constraint of all the placed parts, each unprocessed part is placed in T in decreasing confidence (Step 3).



Figure 8: Importance values for two shape sets. Higher importance is assigned to parts performing more major functionality.

straightforward to fulfill. For the last, we adopt a stochastic sampling approach inspired by [Merrell et al. 2011] which randomly selects parts for R_S and R_T so that their combined importance exceeds a threshold I_{imp} .

The importance value of a part represents the extent of spatial overlap the part shares with other shape parts in the set. The overlaps are estimated after all the shapes are globally aligned as in the case for computing FPC. The intuition is that parts that are estimated as important according to this measure are more likely to be important semantically. For example, for the set of chairs, we would expect a leg part to possess higher importance than an auxiliary bar between two adjacent legs, since all chairs must have legs but not necessarily the auxiliary bars. Figure 8 shows a few models with the importance of their parts visualized. In our experiments, we set the threshold I_{imp} to be 1.2 times the averaged total part importance of all models in the set. It is worth noting that the threshold can be set conservatively, resulting in more parts than necessary to be selected. However, the FPC-driven replacement step would still produce an adequate filter to obtain a proper part crossover.

Constraint inheritance. After placing R_S into the target shape T , we rebuild the controller constraints in the offspring model. Ground support constraints are carried over by individual parts. Symmetry groups are inherited from parent shapes only when all controllers sharing the same symmetry group are inherited during crossover. Both proximity and connectivity constraints are inherited in the same way. If two parts in the offspring come from the same parent, the constraint is simply maintained. For two parts originated from different parents, the constraint is inherited based on FPC. Specifically, let p be a new part placed in T , a proximity or connectivity relation between p and $q \in T$ is established if q has a neighboring part $r \in T$ that is sufficiently replaceable by p , that is $\theta(p, r) > \theta_n$. In our experiments, we set $\theta_n = 0.8$.

4.3 Part mutation

Part mutation happens to an individual shape. It is achieved by randomly selecting and deforming a small number of, typically one to three, controllers and then performing structure optimization on the whole shape with the mutated controllers as constraints. In order to obtain random yet meaningful mutations, we exploit the availability of a set of shapes. Specifically, we rely on fuzzy part correspondence to collect a set of similar controllers and construct a deformation space from them. Mutation of any controller in the set is carried out by random sampling in that space.

Given a controller c , we collect all controllers \hat{c} satisfying the FPC threshold $\theta(c, \hat{c}) > \theta_m$ into a set Ψ_c ; we set $\theta_m = 0.8$ in our experiments. From Ψ_c , we construct a deformation space for c in a similar way to [Ovsjanikov et al. 2011]. First, we define a shape descriptor for each controller in Ψ_c . We then perform PCA on the set of shape descriptors. In the 2D PCA space, we compute a deforma-



Figure 9: Six mutations on the chair model on the left. Controllers undergone mutations are highlighted in red. Structure optimization of the controllers ensures coherence of the resulting models.

tion vector for a controller $\hat{c} \in \Psi_c$ as the displacement vector of \hat{c} 's shape descriptor from the center of the space. Then we project all the deformation vectors using 2D PCA again and compute the minimum and maximum projected values along each dimension, forming a bounding box of allowed deformations. A random mutation of controller c then corresponds to a deformation vector randomly sampled within the bounding box. See Figure 9 for an example.

5 Results

We present results of our set evolution. Reports from preliminary user studies are also provided to evaluate our method both qualitatively and quantitatively. All results were obtained from real interactive sessions with graduate student participants.

Set evolution. The evolution of four sets of man-made shapes, lamps, chairs, candelabra, and TV-shaped aliens, are shown. For each set, we invited a participant to run our program and govern the evolution in a “fit and diverse” manner. Figures 1, 10, and 15 show randomly selected subsets of the foreground/breeding sets at several generations. Shapes that are marked out were those picked by the user as unexpected or interesting. We observe that even with a fairly small input set (ranging in size from 11 to 24), our tool is able to generate populations of diverse yet meaningful shapes, thanks to the part crossover and mutation operators we designed, as well as to diversity control throughout the set evolution.

Diversity control. Figure 11 shows that a lack of diversity control would lead to the generation of an elite set. For the perfume bottle set, our participant preferred heart- or diamond-shaped bottles, which directed the evolution into an elite set of bottles with heart and diamond shapes. The same happened to the aliens where the user preferred cat-like creatures. Contrast Figure 11(b), an elite set, to a “fit and diverse” one shown in Figure 15.

To quantitatively evaluate the effect of diversity control, we measure the diversity of a set of models by the standard deviation of LFD. We record the diversity values throughout the interactive set evolution carried out by a participant. In Figure 12, we plot the diversity of the breeding set over all the generations for both with and without diversity control. Evidently, “fit and diverse” leads to higher degrees of set diversity than just “fit”.

Parameters and statistics. All the experiments on set evolution were conducted with the same parameter setting as described in the preceding sections. Table 1 shows some statistics and timing for set evolution. Preprocessing of the initial input set includes meta part correspondence, upright orientation, global alignment which

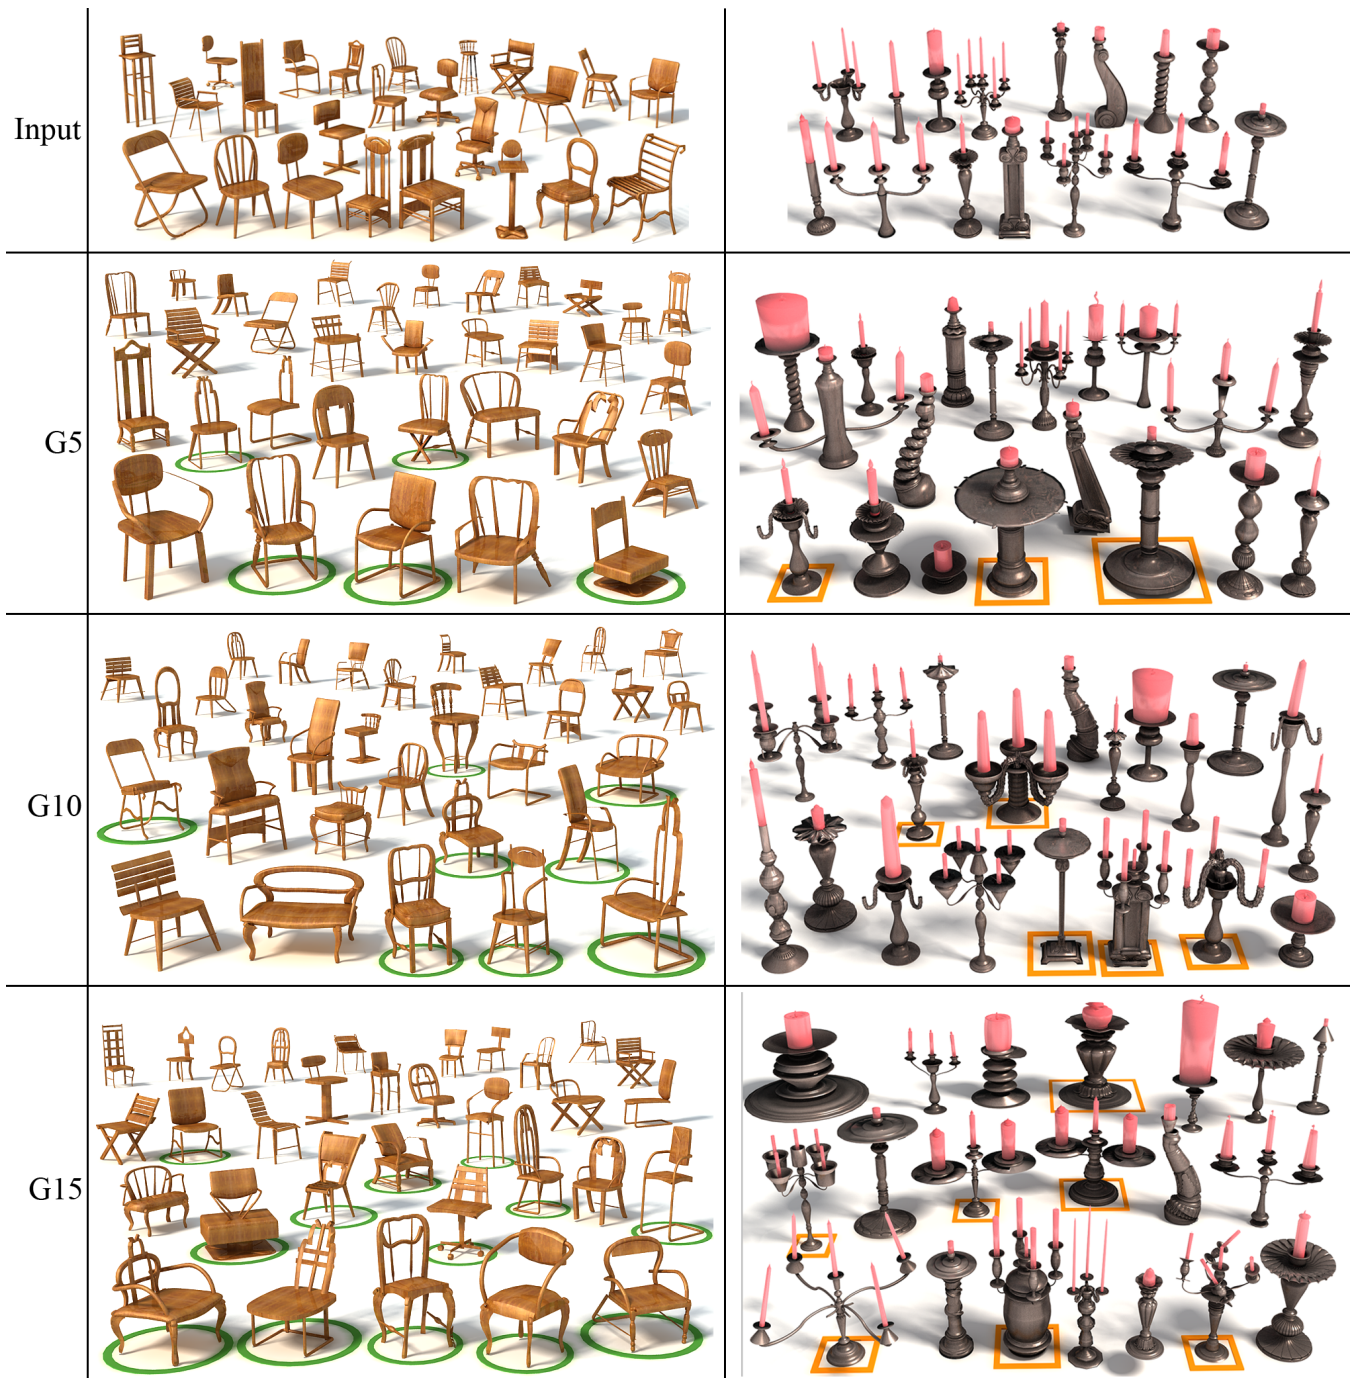


Figure 10: Evolutions of a chair set (left) and a candelabrum set (right). The entire input sets are shown. We show randomly selected shapes from the gallery in three generations. Shapes marked are those identified as unexpected/interesting by the participants.

factors out part proportions, and controllers fitting and preanalysis. The preprocessing time of the aliens set is not reported since the meta part correspondence was difficult to produce automatically and hence was manually specified. The percentages of valid shapes demonstrate the ability of our mutation and crossover operators in creating mostly valid shapes, as judged by human users. The low percentage for the alien set can possibly be attributed to the users' unfamiliarity of what makes a valid TV-like alien.

Preliminary user study. We conducted an informal user study to evaluate the effectiveness of our set evolution in generating unexpected 3D shapes. Sixteen graduate students from the graphics labs at SIAT and NUDT were invited to run the evolution of two sets: chairs and lamps. For each set and in each generation, 18 shapes were generated and presented to the user (9 at a time) and he or she was asked to vote on each shape as either unexpected or not. We assumed that the users have sufficient familiarity with the input set. We then computed the average percentage of unexpected shapes in each generation over all the users. In Figure 13, we plot the aver-

Set	#input	#part	prep_t	repr_t	%valid
Perfumes	16	4	25m	0.1s	92%
Teapots	15	4	20m	0.24s	86%
Chairs	24	17	40m	0.4s	75%
Lamps	11	10	15m	0.17s	81%
Candelabra	15	5	20m	0.16s	90%
Aliens	12	10	N/A	0.12s	47%

Table 1: Various statistics from our set evolution experiments. #input denotes size of the input set. #part denotes the average number of (sub-)parts per shape in the set. prep_t denotes the time (in minutes) spent on preprocessing. repr_t denotes the time (in seconds) needed, on average, to reproduce an offspring. %valid denotes the percentage of valid models for the shapes presented in the gallery over 30 generations; validity was evaluated through voting by a number of participants.

age percentage of unexpected shapes collected over the generations for both “fit and diverse” and “fit” only. The trends show that with only “fit”, the set may become an elite group which then contains less number of unexpected shapes (blue curves dipping down after 15 generations or so). On the other hand, the degree of unexpectedness of the evolving set is well-maintained under diversity control.

In Figure 14, we plot the average percentage of “dislike” votes collected from five participants, with both “fit and diverse” and “fit” only. The plots show that shape classes with more complex structures (e.g. chairs and aliens) have higher rates of dislikes, compared to simpler shapes such as perfume bottles. Moreover, the dislike rate for “fit and diverse” usually increases with more generations, due to the fact that the interactive user selection cannot reject all implausible or disliked shapes in the background set. Obviously, evolutions with only “fit” have slower growth of the dislike rates. To prevent the dislike rate from increasing too fast, we clean up the background set in every 30 generations by randomly removing half of its members which are disliked or unmarked.

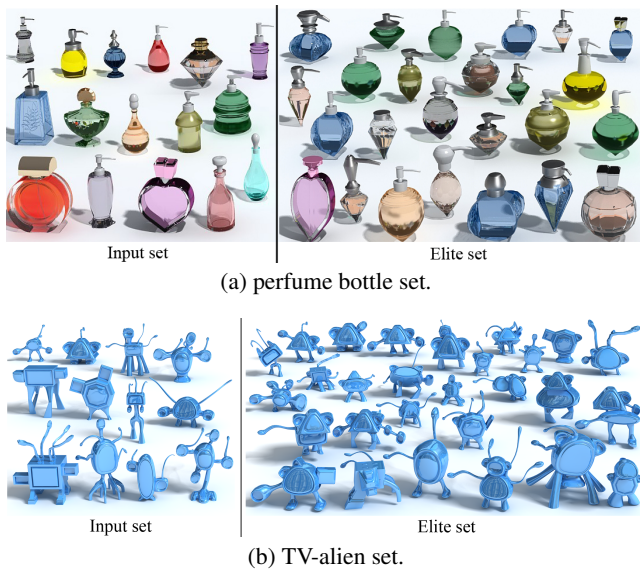


Figure 11: Elite sets generated due to a lack of diversity control. (a) Preferences given to heart- or diamond-shaped bottles. (b) Preferences given to cat-like creatures.




Figure 12: Plots of set diversity, measured as the standard deviation of LFD, over number of generations. “Fit and diverse” (red) leads to more diversity than just “fit” (blue).




Figure 13: Plot of percentage of unexpected shapes, as judged by humans, over the generations. Unexpectedness decreases without diversity control (blue) but is maintained by “fit and diverse”.

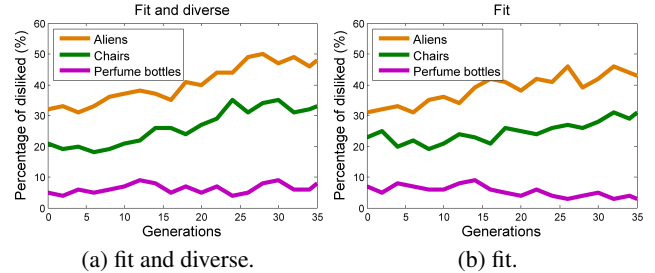


Figure 14: Plot of percentage of disliked shapes over the generations. The number of dislike models in each generation is obtained by averaging votes from five participants.

6 Discussion, limitations, and future work

We present a set evolution method designed to generate sets of 3D novel shapes to inspire the user and to assist in the creative modeling process. The focus of this work is on the set, rather than the individuals. The evolution keeps the population of shapes diverse, with a distinguishing foreground subset fit to the user tastes or preferences, and portions of the background models, not as fit, but ensuring that some unexpected models would emerge in future generations. The high-level notions of fitness and diversity are both biologically motivated, as are the low-level reproduction mechanisms we use, namely, part mutation and crossover.

Part crossover and mutation. We have devised a part exchange mechanism, the crossover, that goes beyond simple shuffling of the major components of a set of shapes. The major components are typically associated with meanings or functionalities that are to be preserved throughout the evolution. The smaller-scale parts within the major components are more stylistic in nature, but they are the main reason for the perceived shape complexity and variability in a set. Our technique operates at a finer granularity of part struc-

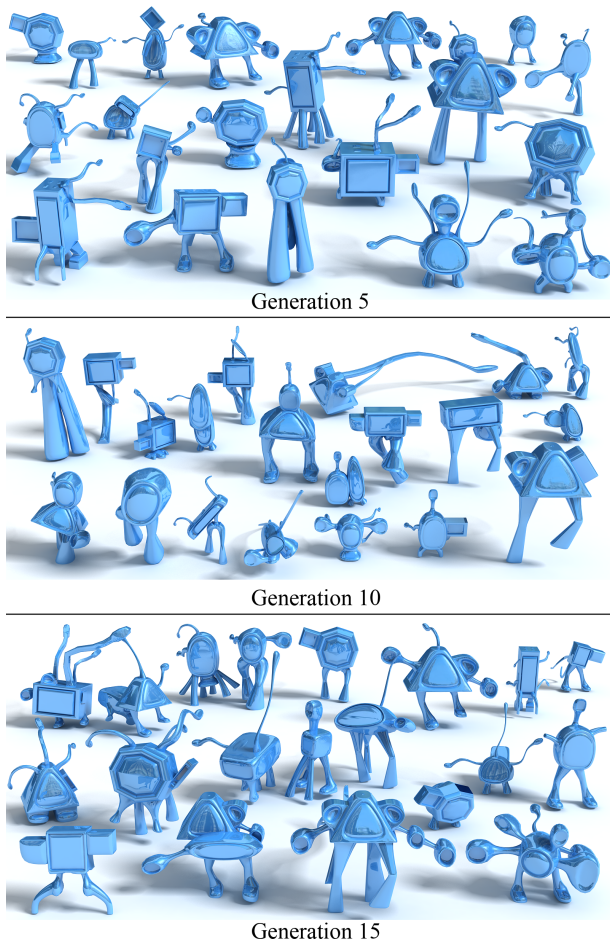


Figure 15: Evolution of a set of TV-like alien creatures. The input set can be located in Figure 11(b), where we find an elite set produced with a lack of diversity control to contrast with the more diverse sets shown in this figure. Since the whole set of creatures are rather unfamiliar themselves, voting from the participants did not reveal particularly unexpected individuals. Nevertheless, significant shape variations enabled by our part crossover and mutation operators are not difficult to spot.

tures via fuzzy correspondence applied at the sub-part level, where meaningful correspondence is not always clear. Part mutation also adds a great deal to the variability of shape forms. The two reproduction operators together contribute to significant shape variations which leads to the diversity of the population.

Initialization. Our current shape reproduction mechanism is still rather limited compared to what happens during biological evolution; it cannot start from a set consisting of highly primitive models and progressively evolve it into a richly diverse set consisting of complex and advanced shapes. We require that the evolution be initialized with a sufficiently “developed” set, i.e., a set of shapes possessing high-level structural information and correspondence at the meta part level. In our setting, any evolving shape is regarded as a combination and interpolation/extrapolation of the initial set of shape parts. Although we consider this as a limitation, it is intriguing to ask whether an algorithm can ever create a novel shape, not from the geometry latent in the initial set. The genetic similarity and apparent dissimilarity in appearance, say between a peacock and an elephant, suggest that this is possible through a long and complex

process. It would likely require at first a shape representation with much finer granularity, like a true “shape DNA” that can evolve. We would like to explore along this direction in future work.

Generative scheme. Our method evolves a set to fit, and we coined the term “fit and diverse”. It is interesting to note that if one ignores the fitness and is content with just remaining “diverse”, the evolution is reduced to a generative scheme, which creates variations from a given set of examples, e.g., [Lin et al. 2011; Jain et al. 2012]. Moreover, when the user does not have a clear fitness function in mind, he or she can explore a shape space [Shapira et al. 2009; Talton et al. 2009; Yang et al. 2011; Kim et al. 2012]. In our setting, the space does not only provide a set of control points for interpolated or extrapolated exploration, but it is dynamic and reproduces new generations along the evolutionary path designed by the user. We plan to develop our current evolution framework further into the shape exploration setting as future work.

Collaborative set evolution. In our current implementation, the foreground subset is evolved to fit the preferences of a particular user. This can be extended to accommodate a community of users by evolving a number of foreground sets. If the users express rather similar taste, it can be considered as a trend. Our system has the potential to evolve a trendy collection, without losing the capability to diverge again over time into a new trend. This multi-user fitness is related to the collaborative design space framework in [Talton et al. 2009]. We leave that for future work as well.

To inspire or not. The ultimate and difficult question of whether our set evolution produces suggestions that truly inspire user creativity is yet to be answered. Our preliminary user study only asks for the identification of unexpected shapes and it only points to the potential to inspire. Indeed, uncommon segments of the diverse set contribute to evolution along unexpected paths to generate surprises, which form an essential source for creativity [Chaudhuri and Koltun 2010]. However, a formal evaluation requires a carefully designed used study and we leave that for future work.

Limitations. In addition to the limitations mentioned above, there is still plenty of room to improve our current approach. First, the capability of our reproduction operators is limited by the underlying structure representation. We adopted the component-wise controllers [Zheng et al. 2011] which only handle cuboid and generalized cylinder types of shape parts. Second, due to the stochastic nature of the crossover operator, highly implausible shapes can still be produced but they are typically filtered out by the user as unfit. In terms of low-level geometry, parts still may not be stitched well to form a watertight shape and structure modification at the sub-part level may result in unnatural looking shapes. Third, our measure of diversity is limited to the shape signature and similarity metric employed. More advanced choices would likely yield improved results. Finally, user feedback in the evolution is only limited to liking or disliking of individual shapes. While there is the gain at simplicity, more fine-grained user feedback, e.g., at the sub-part level, will lead to better control of the set evolution.

Future work. In addition to the possible future works mentioned above, we would like to pay more attention to the quality of the evolved individuals — higher-quality shapes lead to higher-quality descendants. Also of interest is to allow evolution over sets that belong to different but relevant semantic classes to generate interesting hybrids. Finally, a more ambitious attempt would be to go beyond functionality preservation via geometry evolution and enable the evolution to discover new functionalities. This may simply be a natural consequence of more aggressive part mutations and crossovers when a quantitative leap becomes a qualitative one.

Acknowledgements. We would first like to thank the anonymous reviewers for their valuable feedback. Thanks also go to Chao Lai and Shuai Lin from NUDT for their help with making the video and to the artists from SIAT, especially Jiacheng Ren and Qifeng Wei, for the modeling and rendering efforts. This work is supported in part by grants from NSFC (61161160567, 61025012), National 863 Program (2011AA010503), Shenzhen Science and Innovation Program (CXB201104220029A, JC201005270329A), NSERC (No. 611370), and the Israel Science Foundation.

References

- BENTLEY, P. J. 1999. *Evolutionary Design by Computers*. Morgan Kaufman Publishers.
- BENTLEY, P. J. 2000. Exploring component-based representations - the secret of creativity by evolution? In *Proc. of Int. Conf. on Adaptive Computing in Design and Manufacture*, 161–172.
- CHAUDHURI, S., AND KOLTUN, V. 2010. Data-driven suggestions for creativity support in 3D modeling. *ACM Trans. on Graph (Proc. of SIGGRAPH Asia)* 29, 6, 183:1–9.
- CHAUDHURI, S., KALOGERAKIS, E., GUIBAS, L., AND KOLTUN, V. 2011. Probabilistic reasoning for assembly-based 3d modeling. *ACM Trans. on Graph (Proc. of SIGGRAPH)* 30, 35:1–10.
- CHEN, D.-Y., TIAN, X.-P., SHEN, Y.-T., AND OUHYOUNG, M. 2003. On visual similarity based 3D model retrieval. *Computer Graphics Forum (Special Issue of Eurographics)* 22, 3, 223–232.
- DRAVES, S. 2006. The electric sheep and their dreams in high fidelity. In *Proc. NPAR*, 7–9.
- FISHER, M., SAVVA, M., AND HANRAHAN, P. 2011. Characterizing structural relationships in scenes using graph kernels. *ACM Trans. on Graph (Proc. of SIGGRAPH)* 30, 4, 34:1–11.
- FRAZER, J. 1995. *An Evolutionary Architecture*. Architectural Association Publications.
- FU, H., COHEN-OR, D., DROR, G., AND SHEFFER, A. 2008. Upright orientation of man-made objects. *ACM Trans. on Graph (Proc. of SIGGRAPH)* 27, 42:1–7.
- FUNKHOUSER, T., KAZHDAN, M., SHILANE, P., MIN, P., KIEFER, W., TAL, A., RUSINKIEWICZ, S., AND DOBKIN, D. 2004. Modeling by example. *ACM Trans. on Graph (Proc. of SIGGRAPH)* 23, 3, 652–663.
- JAIN, A., THORMÄHLEN, T., RITSCHER, T., AND SEIDEL, H.-P. 2012. Exploring shape variations by 3D-model decomposition and part-based recombination. *Computer Graphics Forum (Special Issue of Eurographics)* 31, 2, to appear.
- JAKIELA, M. J., AND DUDA, J. 1997. Generation and classification of structural topologies with genetic algorithm speciation. *Journal of Mechanical Design* 119, 1, 127–130.
- KIM, V. G., LI, W., MITRA, N., DIVERDI, S., AND FUNKHOUSER, T. 2012. Exploring collections of 3d models using fuzzy correspondences. *ACM Trans. on Graph (Proc. of SIGGRAPH)* 31, to appear.
- KREAVOY, V., JULIUS, D., AND SHEFFER, A. 2007. Model composition from interchangeable components. In *Proc. of Pacific Conference on Computer Graphics and Applications*, 129–138.
- LEE, J., AND FUNKHOUSER, T. 2008. Sketch-based search and composition of 3D models. In *Proc. of SBIM*, 97–104.
- LIN, J., COHEN-OR, D., ZHANG, H., CHENG, L., SHARF, A., DEUSSON, O., AND CHEN, B. 2011. Structure-preserving re-targeting of irregular 3D architecture. *ACM Trans. on Graph* 30, 6, 183:1–10.
- MARKS, J., ANDALMAN, B., BEARDSLEY, P. A., FREEMAN, W. T., GIBSON, S., HODGINS, J. K., KANG, T., MIRTICH, B., PFISTER, H., RUMEL, W., RYALL, K., SEIMS, J., AND SHIEBER, S. M. 1997. Design galleries: a general approach to setting parameters for computer graphics and animation. In *Proc. of SIGGRAPH*, 389–400.
- MERRELL, P., SCHKUFZA, E., LI, Z., AGRAWALA, M., AND KOLTUN, V. 2011. Interactive furniture layout using interior design guidelines. *ACM Trans. on Graph (Proc. of SIGGRAPH)* 30, 4, 87:1–10.
- OVSJANIKOV, M., LI, W., GUIBAS, L., AND MITRA, N. J. 2011. Exploration of continuous variability in collections of 3d shapes. *ACM Trans. on Graph (Proc. of SIGGRAPH)* 30, 4, 33:1–10.
- PILAT, M. L., AND JACOB, C. 2008. Creature academy: A system for virtual creature evolution. In *IEEE Congress on Evolutionary Computation*, 3289–3297.
- POLLACK, J., AND FUNES, P. 1998. Evolutionary body building: Adaptive physical designs for robots. *Artificial Life* 4, 337–357.
- ROMERO, J., AND MACHADO, P. 2007. *The Art of Artificial Evolution*. Springer.
- SHAPIRA, L., SHAMIR, A., AND COHEN-OR, D. 2009. Image appearance exploration by model-based navigation. *Computer Graphics Forum (Special Issue of Eurographics)* 28, 2, 629–638.
- SHIN, H., AND IGARASHI, T. 2007. Magic canvas: interactive design of a 3D scene prototype from freehand sketches. In *Proc. of Graphics Interface*, 63–70.
- SIMS, K. 1991. Artificial evolution for computer graphics. In *Proc. of SIGGRAPH*, 319–328.
- SIMS, K. 1994. Evolving virtual creatures. In *Proc. of SIGGRAPH*, 15–22.
- SODDU, C., AND COLABELLA, E. 1995. Recreating the city's identity with a morphogenetic urban design. In *Proc. of Int. Conf. on Making Cities Livable*, 5–9.
- TALTON, J. O., GIBSON, D., YANG, L., HANRAHAN, P., AND KOLTUN, V. 2009. Exploratory modeling with collaborative design spaces. *ACM Trans. on Graph (Proc. of SIGGRAPH)* 28, 5, 167:1–10.
- XU, K., LI, H., ZHANG, H., COHEN-OR, D., XIONG, Y., AND CHENG, Z. 2010. Style-content separation by anisotropic part scales. *ACM Trans. on Graph (Proc. of SIGGRAPH Asia)* 29, 5, 184:1–10.
- YANG, Y.-L., YANG, Y.-J., POTTMANN, H., AND MITRA, N. J. 2011. Shape space exploration of constrained meshes. *ACM Trans. on Graph (Proc. of SIGGRAPH Asia)* 30, 124:1–12.
- ZHENG, Y., FU, H., COHEN-OR, D., AU, O. K.-C., AND TAI, C.-L. 2011. Component-wise controllers for structure-preserving shape manipulation. *Computer Graphics Forum (Special Issue of Eurographics)* 30, 2, 563–572.