

# ShapeGenetics: Using Genetic Algorithms for Procedural Modeling

Karl Haubenwallner<sup>1</sup>, Hans-Peter Seidel<sup>2</sup> and Markus Steinberger<sup>2</sup>

<sup>1</sup>Graz University of Technology, Austria

<sup>2</sup>Max Planck Institute for Informatics, Saarland Informatics Campus, Germany



**Figure 1:** (a) Procedural approaches can generate models with large variety. (b) Given a high level target as a sketch (drawing at the far right), the generation can be directed towards this goal. Previous approaches are either tuned for speed (SMC, red circle and SOSMC, orange circle) and thus might not reach the desired result, or require a long time to achieve a good result (RJMCMC, green circle); our proposed solution using a genetic algorithm converges fast and achieves high quality results (blue circle), clearly matching the target the best.

## Abstract

In this paper, we show that genetic algorithms (GA) can be used to control the output of procedural modeling algorithms. We propose an efficient way to encode the choices that have to be made during a procedural generation as a hierarchical genome representation. In combination with mutation and reproduction operations specifically designed for controlled procedural modeling, our GA can evolve a population of individual models close to any high-level goal. Possible scenarios include a volume that should be filled by a procedurally grown tree or a painted silhouette that should be followed by the skyline of a procedurally generated city. These goals are easy to set up for an artist compared to the tens of thousands of variables that describe the generated model and are chosen by the GA. Previous approaches for controlled procedural modeling either use Reversible Jump Markov Chain Monte Carlo (RJMCMC) or Stochastically-Ordered Sequential Monte Carlo (SOSMC) as workhorse for the optimization. While RJMCMC converges slowly, requiring multiple hours for the optimization of larger models, it produces high quality models. SOSMC shows faster convergence under tight time constraints for many models, but can get stuck due to choices made in the early stages of optimization. Our GA shows faster convergence than SOSMC and generates better models than RJMCMC in the long run.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling —Geometric algorithms, languages, and systems

## 1. Introduction

Procedural modeling not only has a long history in computer graphics, but also sees increasing interest in recent years as the demand for detailed models and large virtual worlds is growing rapidly. Using a procedural approach, a wide variety of detailed variants of a model or family of models can be generated by altering the parameters

that control the procedural generation. Examples include vegetation models, which can generate everything from small bushes to full leaf and needle trees [PSK\*12], and building grammars, which can generate everything from dog sheds to skyscrapers [SM15]. While such generative models potentially offer technical artists the ability to create a complex object within seconds, achieving a desired result



**Figure 2:** A procedural approach that can generate random trees can be controlled to generate a model whose shadow resembles an arbitrary logo. The approach controlling the generation needs to find an appropriate tree structure and choose tens of thousands of parameters to find a suitable shape. The result was obtained with our GA in eight hours.

is often still a complicated and time-consuming process. Because the model's parameters can control recursive generation processes and small parameter adjustments can be amplified throughout the generation, the influence of a single parameter on the final model is often unpredictable and artists have to rely on excessive trial and error. Unsurprisingly, this cumbersome process turns many artists away from procedural modeling back to manual model generation.

As a remedy to this problem, the processes of procedural generation have been cast as a probabilistic inference problem [TLL\*11]. Given a procedural generator that has the expressive power to generate a suitable model and a scoring function that tells how close a model is to the desired goal, these approaches view the generation of a model as drawing a sample from a probability distribution. Naively speaking, the scoring function allows for an optimization over the entire space of possible models to find a model that fits the scoring function as closely as possible. The approach is complicated by the fact that the space of possible models is trans-dimensional, theoretically unbounded, and lies in a mixed domain space; the scoring function is in general non-linear and non-convex; and generating a model and evaluating the scoring function can take significant amounts of time. Clearly, the problem of finding a set of parameters that generates a model that matches a given target is a non-trivial problem, *e.g.*, consider the result shown in Figure 2.

Most previous approaches on probabilistic inference for procedural generation rely on Markov Chain Monte Carlo (MCMC) methods, specifically the Reversible Jump Markov Chain Monte Carlo method (RJCMC) to cope with the trans-dimensionality, to solve the problem [TLL\*11, SPK\*14, YYW\*12]. As alternative to the relatively slow converging RJCMC, a variant of the sequential Monte Carlo (SMC) method, namely stochastically-ordered sequential Monte Carlo (SOSMC) has been proposed [RMGH15]. While SOSMC, in general, improves on the convergence of RJCMC, there is a possibility that it gets stuck in bad initial conditions, especially if the procedural models become more complex. Also, SOSMC requires the scoring function to yield reasonable results for unfinished models, which cannot be guaranteed in general. Given

the mere existence of only two classes of approaches for the probabilistic inference problem in targeted procedural modeling and the increasing importance of (semi)-automated content creation, we see high potential benefits in exploring further alternatives.

In this paper, we show that genetic algorithms (GA) can be used for controlled procedural modeling. While the principles of GAs are not new, applying GA to a new domain is never without new challenges. We tackle these challenges and make the following contributions:


- We present a unifying framework that, while geared towards GA, can be used for various optimization methods like RJCMC or SOSMC, allowing for a direct performance comparison.
- We present a simple yet efficient way to encode a variety of procedural generation approaches as *genomes*, used by but not limited to GA for controlled procedural modelling.
- We show how the core operations of a GA—reproduction and mutation—can be implemented for these genomes such that only valid modeling operations are created.
- We show that our GA yields significantly better convergence rates than the previous state-of-the-art implementations, especially when more complex models are generated.

In the following, we give an overview of related work (Section 2); introduce GA for controlled procedural modeling (Section 3); present our genome representation for procedural modeling (Section 4); explain the mutation and reproduction operations (Section 5); provide details on the GA setup (Section 6); show how RJCMC and SOSMC can be modeled within the same genetic algorithm framework (Section 7); compare our GA implementation against previous state-of-the-art for different kinds of scoring functions and generation methods (Section 8); and summarize the findings (Section 9).

## 2. Related Work

**Procedural Modeling** has been a part of computer graphics for decades. Among the first procedural modeling approaches were Stiny's original shape grammars [Sti75] and Lindenmayer's L-systems [Lin68, PL91]. Starting with these early works, a variety of approaches have been proposed, including realistic trees [WP95], plants [LD98], trees interacting with their environment [MP96], split grammars for facades [WWSR03], botanic trees [OOI05], a language for complex objects [Hav05], a grammar for buildings [MWH\*06], self-organizing tree models [PHL\*09], interconnected structures [KK11], plastic trees [PSK\*12], and complex buildings [SM15]. At the same time, speeding up procedural generation on modern hardware has received increasing interest [Mag09, MBG\*12, SKK\*14a, SKK\*14b, BDK\*16]. Our approach can be used for any of the previously mentioned procedural approaches and will benefit directly from faster generations.

**Controlled Procedural Generation** has received increasing interest in recent years. Viewing procedural generation as an inference problem, high-level goals can be considered during generation. Taltou et al. used the reversible jump MCMC (RJCMC) algorithm, a variant of the Metropolis Hastings (MH) algorithm for controlled grammar-based procedural modeling [TLL\*11]. While their implementation achieved impressive results choosing thousands of




**Figure 3:** (a) Although quite simple, the spaceship generator can produce a variety of ships. (b) Given a volumetric target (dark shape), the generation can be optimized to fill the target. Note that the best solutions of different algorithms (GA: blue, SMC: red, SOSMC: orange, RJMCMC: green) differs not only in the size of the parts but the structure of the entire ship.

parameters for tree and building models to match target volumes and silhouettes, their MH algorithm also requires tens of thousands of iterations to converge, leading to running times of up to multiple hours. When using MCMC to only choose a few parameters, *e.g.*, to position a small number of pieces of furniture, good results can be obtained within a second [MSL\*11]. However, more complex furniture layout generation, which also require RJMCMC methods, increases the running time again to minutes or hours [YYW\*12]. In case one is not interested in a specific model, but rather a sub-family of models that are similar to a target, finding a suitable similarity measure is a non-trivial task [SPK\*14]. With such a similarity measure in place, MH can be used to choose a small number of meta-parameters for the sub-family. However, running times might still be long due to costly similarity evaluations. Tackling the rather long convergence time of the MH algorithm and its variants for controlling procedural modeling, the sequential Monte Carlo (SMC) approach can be seen as a combination of guided procedural modeling [BvMM11] with probabilistic inference. Models are not only scored when they are fully generated, but the score of early stages of the generation determines how likely it is that the generation of a model continues. Ritchie et al. proposed a stochastically ordered variant of the sequential Monte Carlo approach (SOSMC), where the execution of the possible modeling paths is permuted randomly [RMGH15]. For simple models, good results can be obtained within seconds or minutes, outperforming MH approaches when only little time is available.

Our approach can be seen as an alternative to MH and SOSMC, showing even better convergence rates than SOSMC under tight time constraints and outperforming MH in the long run.

**Evolutionary Approaches** have been used in computer graphics before. Most noteworthy is the work by Sims. He showed that evolutionary algorithms with simple mutation and random crossover operations can be used to generate 3D plant structures, images, textures, and animations [Sim91]. In his follow-up work, he described the genotypes and phenotypes of virtual creatures including their animation using evolutionary algorithms, allowing them to walk, jump, and swim [Sim94]. Sims' approach sparked a variety of approaches following the same direction, like, *e.g.*, Creature Academy [PJ08], which uses Sims' encoding schemes to evolve creatures that are capable of a variety of motions. Evolutionary



**Figure 4:** (a) The generator graph for the simple spaceship models (Figure 3) consists of three nodes: Each part can generate an instance of itself (probabilities next to edges). The body part may generate a top and a wing part. All parts are parameterized by scaling factors; the wing has an additional offset. (b,c) The derivation tree fully describes a generated model. (b) corresponds to the center model of Figure 3(a), (c) to the blue ship in Figure 3(b). Note that parameter choices are not shown here due to the size of the graphs.

algorithms have also been used to control L-systems to optimize parameters of 2D-plants [Och98], or fit basic 2D shape grammars to simple targets [OSM\*09]. They have also been applied for real world objects [FP98], enabling the design of robots that can be assembled using Lego bricks. Genetic Algorithms can also be used for shapes within shape collections, extending the variety of models in design galleries [XZCOC12]. In this 'fit and diverse' gallery, evolution is not applied on an abstract genome representation, but directly on the models by exchanging their parts using fuzzy crossover operations. While the aforementioned evolutionary approaches have been used to represent shapes and models, they either work directly on model parts or a high-level, meta-graph which allows edges between arbitrary nodes. Similarly, Martin et al. [MLCB10] used a GA to extend the variety of design gallery for a computer game, using a generative language to create models of buildings. Applying GA to inference problems for procedural generation has different requirements and a direct adoption of the previous representations is not possible. With our work we provide such a representation and underline the special requirements of controlled procedural generation with GAs.

### 3. Genetic Algorithms and Controlled Procedural Modeling

Virtually all approaches that are classified as procedural modeling, can be described in a unified way: they apply a sequence of modeling operations on objects. Each operation either alters an object or generates new objects, on which modeling operations can again be applied. These objects are in turn translated into geometry that constitutes the generated model. The operations are usually parameterized, *e.g.*, translate by vector  $\mathbf{t}$ , rotate by angles  $(\phi_x, \phi_y, \phi_z)$ , and so on. These parameters are generally chosen at random to add variations to the generated models. Also, the number and type of generated objects may be chosen randomly. Based on these observations it seems natural to capture procedural generation as directed graphs (with cycles). This fact has been explicitly noted for L-systems [Boe95], shape grammars [Pat12], and stack-based modeling languages [Hav05] before. Also languages like the one used by Ritchie et al. [RMGH15] can be translated into a directed graph by applying the principles of data flow programming [WA85]. Thus,

we do not limit our approach to any specific procedural approach, but only work with a directed graph representation, where nodes correspond to parameterized operations and the edges correspond to objects. We simply call the description of the procedural generation in their respective language a *procedural generator* and the associated graph the *generator graph*, see Figure 4(a) for a simple example.

The generator graph describes the procedural generation itself and thus all models that can be generated by that generator. The generation of a specific model can also be described by a graph or, more specifically, a tree. This tree captures all intermediate objects as they move through the generator graph, with the operations applied to them including the randomly chosen parameters. In grammar-based modeling this tree is called *derivation tree* [Sip06]. Although we do not limit our approach to grammar-based modeling, we do adopt this name. As the operations can generate a different number of objects at random, the derivation trees for different models can vary in structure and number of parameters, as shown in Figure 4(b,c). In any case, a model is fully described by its derivation tree.

For complex models, derivation trees can become very large, containing tens of thousands of nodes with as many or more parameters. This underlines that expecting a technical artist to choose these parameters by hand is infeasible. As the derivation tree describes the generation, we can view controlled procedural modeling as an optimization problem or Bayesian inference problem of derivation trees. Each possible derivation tree corresponds to a sample from the procedural generator, leading to Bayesian inference. At the same time, finding those derivation trees that maximize the scoring function creates a classical optimization problem. The problem is obviously in a mixed domain: deciding whether a node should be added to the tree is a boolean decision; the parameters describing operations in Euclidean space are in the continuous domain. Also, derivation trees of different models have different structures, which emphasizes the trans-dimensional nature of the problem.

Genetic algorithms are well suited for this kind of problem. GA is inspired by evolution and natural selection, where traits and characteristics of individuals of a species are encoded as genes. Due to the selection process individuals with successful traits get a higher chance to pass on their genes to future generations, while less successful traits tend to disappear, overall leading to a better adapted population. The basic steps of GA are outlined in Algorithm 1. For controlled procedural modeling we want each individual to represent a derivation tree. During evolution individuals can be altered (mutated) and information from different individuals can be combined (reproduction). To use GA for controlled procedural modeling we require mutation and reproduction operations that work on derivation trees while making sure the resulting derivation trees conform with the generator graph. The complexity of the optimization—working in a mixed domain and the trans-dimensionality—is hidden in these two operations. To determine the fitness of individuals, as required by a GA, we directly use the scoring function, which can also include additional constraints like the size of the derivation tree or how balanced the tree should be. However, two challenges need to be solved: (1) an encoding of the derivation tree as genome representation is required, and (2) the genetic operations for mutation and reproduction need to be described.

---

**Algorithm 1:** The basic structure of genetic algorithms is straightforward. A population of individuals are managed for a given number of generations. The best individuals are allowed to influence the next generation, by either copying them directly, mutating them, or allowing them to reproduce.

---

**Function** *GeneticAlgorithm*

```

population ← { }
for i ∈ {1 ... PopulationSize} do
  append (population, newIndividual ())
for j ∈ {1 ... MaxGenerations} do
  fitness ← evaluateFitness (population)
  sort (population, fitness)
  newPopulation ← { }
  for k ∈ {1 ... Elitism} do
    append (newPopulation, population [k])
  for k ∈ {1 ... (PopulationSize- Elitism )/2} do
    par1 ← select (population,fitness)
    par2 ← select (population,fitness)
    if rand () ≤ MutationProbability then
      chld1 ← mutate (par1)
      chld2 ← mutate (par2)
    else
      chld1,chld2 ← reproduce (par1, par2)
    append (newPopulation, chld1, chld2)
  population ← newPopulation

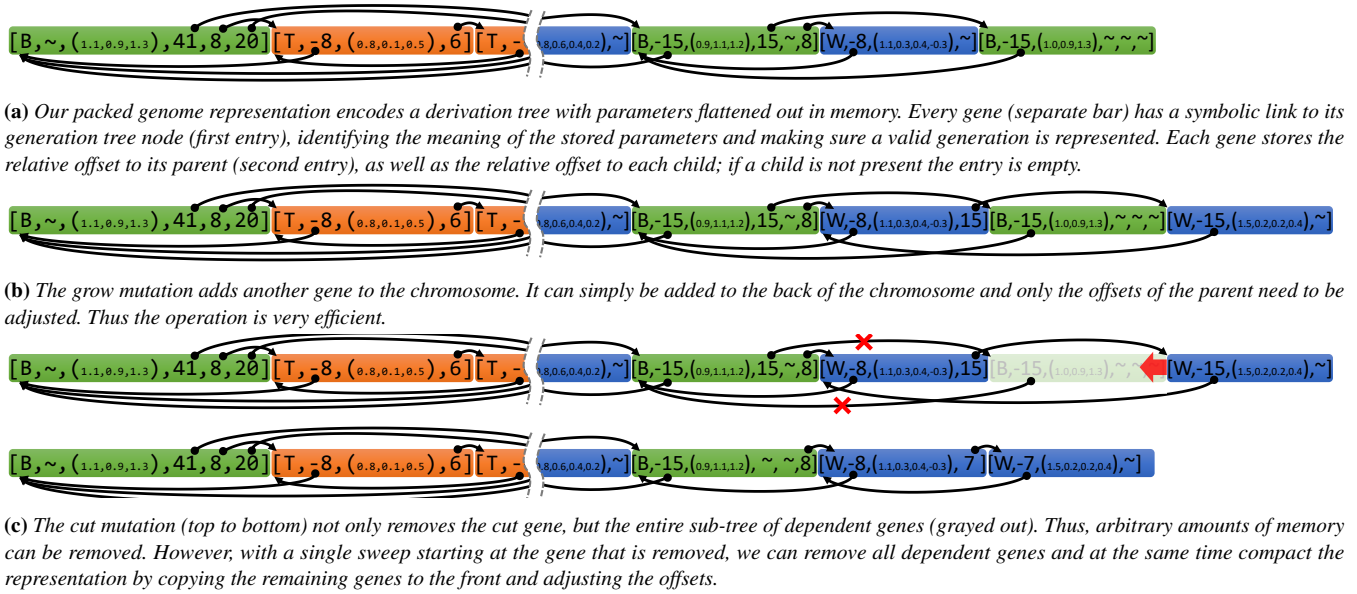
```

---

#### 4. Genome Representation for Procedural Models

The genome representation must be able to encode all possible solutions to the problem, *i.e.*, all possible derivation trees. To describe the genome representation we use the following terms: a *gene* is the smallest element of the genome; a *chromosome* is the set of genes that make up an individual. Ideally, the representation of a chromosome and hence a gene should be simple and allow for efficient genetic operations. Furthermore, the fitness evaluation requires a translation of the chromosomes into their expression, *i.e.*, into a geometric model. Thus, given the chromosome, we want to run the procedural generator replacing the random decisions with the parameters decoded from the chromosome in an efficient manner.

Because each chromosome needs to encode the derivation tree, the traditional GA approach of using bit-strings of fixed size as genome representation is not feasible. We propose to store the derivation tree directly as nodes and edges, similar to Sims' graph representation [Sim94]. Every node is described by one gene. Each gene keeps a reference to its corresponding node in the generator graph, the parameters chosen for its associated operation, and a pointer for each possible output object, *i.e.*, child node, which we set to empty if no output is generated. If an output is present, it points to the gene within the same chromosome that encodes the generated object. Furthermore, every gene stores a pointer to the gene that describes its parent node. We call the pair of pointers between a parent and child gene a *connection*. Each connection is associated with an edge in the generator graph. Note that genes are of different size, depending on which node and operation they describe. We pack all genes of



**Figure 5:** Example of our genome representation and mutations applied to the example spaceship from Figure 3(a) and 4(b)

one chromosome compactly in memory. The pointers are stored as local relative memory offsets, which allows for efficient copying of subtrees; see Figure 5(a) for an example.

Our genome representation allows for variable length chromosomes and easy insertion of new genes into the chromosome without changing the already existing entries. Furthermore, copying genes from one individual to another can be carried out efficiently due to the use of local offsets. The reference to the generator graph is required to ensure that operations on the chromosomes do not yield individuals that cannot be created by the procedural generator, *i.e.*, lie outside of the sampling space. Within the generator graph we additionally store the distribution of each random variable and probability for each output object. Thus, the GA has all required information available in the generator graph and can construct any chromosome using the genome representation.

## 5. Genetic Operations

At the core of each GA there are two genetic operations that evolve the population: mutation and reproduction. Mutation allows the GA to explore the problem space outside the already existing population by introducing random modifications to the genes of an individual, *i.e.*, generate a sample that is outside the space spanned by the individuals of the current population. Reproduction combines features from two individuals, creating individuals that contain parts of both parents and possibly yield better results than either of them. Combining traits from different high-scoring derivation trees, the GA may be able to converge faster.

**Mutation.** As the chromosomes of an individual not only encode parameters, but also the structure of the generation, we propose to use three different mutation operations: *grow*, *cut*, and *alter*.

*Grow*, outlined in Figure 5(b), appends a gene to a gene that does not generate the maximum number of children in its current form. We choose a gene  $i$  and its non-expressed child  $j$  with the probability


$$p_{grow,i,j} = \frac{p_{i,j}}{\sum_k \sum_l p_{k,l}},$$

with  $p_{i,j}$  being the probability of object  $j$  being generated by a node of type  $i$ ;  $k$  runs over all genes that do not express their maximum number of children;  $l$  runs over those non-expressed children. Thus, the probability of adding a gene is proportional to the likelihood of the object represented by the gene being generated in an uncontrolled generation. After a new gene has been added we choose its parameters randomly based on the probability distributions stored in the generator graph. The gene can simply be added to the back of the chromosome and *grow* is thus very efficient.

*Cut*, outlined in Figure 5(c), removes a random gene and the sub-tree that follows it. The probability of choosing a gene  $i$  is

$$p_{cut,i} = \frac{1 - p_{j,i}}{\sum_k \sum_l (1 - p_{k,l})},$$

where  $k$  runs over all genes that have children and  $l$  over those expressed children. Thus, genes are removed with a probability that is inverse proportional to the likelihood that the object represented by the gene is generated in the uncontrolled production. Note that objects which are generated with a probability of one are never cut and thus no invalid derivation trees are generated. If there is no gene with a probability  $p_{cut,i} \neq 0$ , no cut is applied. Although cut may involve the removal of multiple genes which are spread across the chromosome, the operation can be completed efficiently, sweeping over the chromosome, moving non-removed genes further to the front essentially replacing removed ones and adjusting the local offsets. The operation can therefore be completed in  $\mathcal{O}(G)$ , where  $G$  is the chromosome length. This guarantees that the operation remains efficient even for more complex derivation trees.



**Figure 6:** (a) Our crossover operation chooses one random gene connection in the left parent (red highlight), identifies all matching connections from the right parent by consulting the generator graph (red edges) and chooses one at random (red highlight). (b) Exchanging the subtrees (yellow and purple) yields two offspring.

*Alter* changes the parameter values of a random gene. We choose a gene with parameters from the chromosome with equal probability and replace all its parameters by drawing a new random sample from the parameter's distribution as described in the generator graph. *Alter* is obviously the simplest and most efficient mutation.

**Reproduction.** Our reproduction operation is an adapted version of the single-point-crossover, which creates a pair of children from a pair of parents. In its basic form the single-point-crossover selects a random crossover point in both parent chromosomes and creates the offspring by swapping the genes at the crossover point. This strategy cannot be directly applied to chromosomes describing a derivation tree since the results may be incompatible with the generator graph.

Our proposed crossover operation selects a random gene connection from the first parent, chooses another random connection from all *compatible* connections in the second parent, and exchanges the genes including all successors at these connections. The set of compatible connections contains those that describe the same edge in the generator graph. In this way, we can guarantee that the offspring are compatible with the generator. The crossover operation is outlined in Algorithm 2 and Figure 6. With our genome representation the crossover operation is efficient, since compatible connections can easily be selected based on the reference to the generator graph and copying genes from a parent can be viewed as combination of cut and grow operations, which are both efficient as well.

## 6. Genetic Algorithm Setup

**The selection method** is one of the central parts of a GA. It determines which individuals are allowed to reproduce and influence the next generation. The selection is usually done by selecting individuals according to their fitness, with some margin for stochasticity. A purely deterministic selection method would select the best individuals only, limiting the exploration of the problem space to the proximity of the fittest individual, while a purely random method would lead to an entirely undirected exploration decreasing the probability of finding a good solution considerably. There exist several selection methods that take the fitness into account while allowing for random deviation. The most widely used ones are *roulette wheel* and *k-tournament* selection [BT96].

---

**Algorithm 2:** Our proposed crossover operation is computationally efficient, as every offspring is a valid chromosome due to the explicit links to the generator graph. Thus, crossover involves only drawing two random numbers and copying genes.

---

```

Function reproduce(parent1, parent2)
  for  $i \in \{1 \dots \text{MaxRetries}\}$  do
    C1  $\leftarrow$  connections (parent1)
    e1  $\leftarrow$  selectUniformly (C1)
    type  $\leftarrow$  generatorGraphEdge (C1)
    C2  $\leftarrow$  findMatching (parent2, type)
    if C2 is empty then
      continue
    e2  $\leftarrow$  selectUniformly (C2)
    child1  $\leftarrow$  cloneUntil (parent1, e1)
    child1  $\leftarrow$  cloneFrom (parent2, e2)
    child2  $\leftarrow$  cloneUntil (parent2, e2)
    child2  $\leftarrow$  cloneFrom (parent1, e1)
  return child1, child2
return parent1, parent2

```

---

Roulette wheel selection chooses an individual  $I$  with a probability  $p_I$  proportional to the individual's fitness  $f_I$ :

$$p_I = \frac{f_I}{\sum_j f_j},$$

where  $j$  runs over all individuals of the current generation. The name stems from the informal description as a roulette wheel, where the size of each spot relates to the fitness of the individual occupying it.

$k$ -tournament selection starts by selecting  $k$  individuals at random with equal probability. From that group, the one with the highest fitness is selected, *i.e.*, letting the individuals fight in a tournament.  $k$ -tournament selection can be implemented efficiently, while  $k$  controls the selection pressure. In our approach roulette wheel selection is used to choose one parent and  $k$ -tournament the other.

**A good initial population** is important to introduce enough variety into the evolution. We create the initial population starting with individuals that contain an empty chromosome. For each individual we repeatedly apply the grow mutation until their chromosomes reach a randomly chosen size. The GA converges fastest if the average chromosome length of the initial population is comparable to the final solution. However, as crossover and mutation change the chromosome length, even far off initial populations converge. As the grow mutation is very efficient, generating the chromosomes of the initial population has a low computational cost.

**Elitism** is a technique that ensures that future generations contain the best individuals of previous generations. Using elitism, the individuals with the highest fitness are copied unchanged to the new generation. However, these individuals can still also be selected as parents for reproduction. This ensures that the quality of the solution never decreases and increases convergence rates as good solutions are available for reproduction more often. However, if the population size is too small, elitism can lead to stagnation as the GA cannot explore different solutions. We use elitism in all our experiments.

## 7. RJMCMC and SOSMC as GA

As SMC, SOSMC [RMGH15], and RJMCMC [TLL\*11] for controlled procedural modeling work on individual particles or chains that can be viewed as individual (partially finished) procedural generations, we argue that they can be implemented in a GA framework. The following description is an outline of how to implement these approaches, details can be found in the original papers.

*SMC* is arguably the simplest of the three approaches. The particles used in SMC can be seen as individuals, each starting with an empty chromosome. SMC does not allow for interaction between particles and thus can be described with mutations only. Ritchie et al. sample particles based on their fitness, which we can provide in terms of a selection method. As mutation operations, we only require an adjusted grow mutation: It is not allowed to choose any non-expressed child, but has to follow the order they are found in the chromosome. Additionally, the next child is only expressed with a probability equal to the original generation (as found in the generator graph). If it is not expressed, it is marked as such in the chromosome and will never be expressed. The mutation continues until a new geometric output object is generated, showing the same behavior as the original SMC [RMGH15].

*SOSMC* can be implemented the same way as SMC with the difference that the mutation chooses a random non-expressed child and does not have to follow the order in the chromosome. This step again continues until a new geometric output object is generated. Note that the state of the chromosome can be seen as the structure storing the stochastic futures [RMGH15].

*RJMCMC* is arguably more complex. Each production chain corresponds to an individual. In the default setting, no reproduction is supported. The selection method picks each individual once, moving mutated versions of each individual to the next generation. The diffusion operations correspond to a sequence of alter mutations followed by a scoring function evaluation, which determines whether the mutation is accepted. The jump moves can be viewed as mutations that combine a cut with grow mutations on top of the cut gene. Furthermore, the initialization of the newly found subtree does not draw new random variables, but copies the ones from the cut subtree (new values are drawn if there are not enough for copying). The mutation is accepted based on the fitness ratio between the mutated and original individual, and the probability of the jump move.

While the described setup models the basic reversible jump RJMCMC algorithm, further optimizations are required to achieve good results [TLL\*11]. These optimizations can also be modeled in a GA. *Non-terminal selection* is added by altering the probability with which the mutations choose genes. *Parallel tempering* adds a temperature to each individual. The temperature changes the acceptance probability and—implementing a reproduction operation—exchanges the temperature of two individuals depending on their fitness. *Delayed rejection* extends the jump mutation with additional alter mutations, which allows a mutated individual to adjust before being rejected. Finally, *annealing* can be added on top of all rejection tests, slowly reducing the chance of accepting a worse individual.

## 8. Evaluation

To evaluate the performance of GA for controlled procedural modeling, we compare it to the previous state-of-the-art for a variety of scenes, complexity levels, and types of scoring functions. As evaluation platform we used an Intel Core i5-4570 @ 3.2GHz with 8GB of memory and an NVIDIA Geforce GTX 970. For running the procedural generation, we used a custom, multi-threaded, template-meta programming C++ generator that works on abstract shapes. This generator can be used for L-systems, shape grammars, and custom modeling languages. To obtain fast generation speeds, the generator graph descriptions is input to the C++ compile step, specializing the generator for the approach at hand.

**Volumetric targets** are a common way to specify high-level goals for a procedural generator. To this aim, we take any number of volume-describing target models  $\mathbf{T}_i$  and weights  $w_i$  and  $\tilde{w}_i$  as input, voxelize them and compute the scoring function  $s_{vol}$  for a generated model  $\mathbf{M}$ :

$$s_{vol}(\mathbf{V}_M) = \sum_i \left( \left( \sum_{v \in \mathbf{V}_i} w_i \cdot \mathbf{V}_M(v) \right) + \left( \sum_{v \in \mathbf{V}_M} \tilde{w}_i \cdot (1 - \mathbf{V}_i(v)) \right) \right),$$

where  $\mathbf{V}_M$  is the voxelized representation of  $\mathbf{M}$  and  $\mathbf{V}_i$  is the voxelized representation of  $\mathbf{T}_i$ . In this way, volumes that should be filled or avoided can be specified setting positive or negative  $w_i$  and  $\tilde{w}_i$ .

For efficient implementation, we compute the maximum bounds among all  $\mathbf{T}_i$  and combine the voxels of all targets into a regular voxel grid  $\mathbf{V}_{comb}$ , where each voxel corresponds to the sum over the weights of all targets. During evaluation, we voxelize the generated model, compute the overlap with  $\mathbf{V}_{comb}$  and sum up all weights. For all voxels that fall outside the bounds of the voxel grid, we add  $\sum \tilde{w}_i$  to the score. To perform the voxelization of each  $\mathbf{M}$  efficiently, we use a GPU voxelizer implemented in CUDA.


During voxelization it is easy to detect self intersections of the model components. Thus, we take another optional weight  $w_{self}$  which is subtracted from the overall model score every time a voxel is hit multiple times. In this way, self intersections of models can be avoided completely by setting  $w_{self}$  to a high weight, or punished slightly in case a few self intersections are acceptable.

**Images as targets** are also a common way to specify high-level goals. We take any number of floating point images  $I_i$  and camera parameters  $C_i$  as input and compute the scoring function  $s_{img}$  for a generated model  $\mathbf{M}$ :

$$s_{img}(\mathbf{M}) = \sum_i \left( \sum_{p \in I_i} project_{C_i}(\mathbf{M})(p) \cdot I_i(p) \right),$$

where  $project_{C_i}(\mathbf{M})(p)$  corresponds to projecting  $\mathbf{M}$  using the camera parameters  $C_i$  and sampling the image at position  $p$ . Regions that are very important to be hit by the generated model should have high positive pixel values, areas that should be avoided high negative pixel values, and areas that do not matter zero values.


For efficient evaluation of  $s_{img}$ , we render the generated model into a black and white texture using OpenGL. Then, we run a compute shader to multiply the rendering result with  $I_i$  and perform a parallel reduction of the obtained values on the GPU.



**Figure 7:** Development of the mean fitness values and standard deviations over time for the scenes shown in Figure 1 and 3. For a simple scene and short optimizations times as in the spaceship scenes, SOSMC and even SMC show good performance, while RJMCMC takes considerably longer to achieve good results. In complex scenarios, like the tree sketch, SMC and SOSMC only show good performance in the beginning, while RJMCMC slowly converges. GA in both cases shows very good early convergence rates and continues to improve over time.

**The test scenes** used recreate scenes from previous work [TLL\* 11, RMGH15]: a spaceship generator with models shown in Figure 3 and the generator graph in Figure 4(a); a tree generator with models shown in Figure 1, 2, and 9(a) (using a different foliage geometry), and generator graph in Figure 8(a), and a city generator with models shown in Figure 9(b-d) and generator graph in Figure 8(b). The spaceship generator is straight forward, as each new object occupies its own space and every node generates geometry. The tree generator is still relatively simple, however, there are many possible ways for branches reaching the same spot. Thus a variety of solutions and scenarios with self intersection are possible. Every node generates geometric output. The city generator has two recursive nodes which generate empty building lots and essentially iterate over the scene. Only the building node is responsible for creating output geometry. Hence, decisions made during city generation might not directly lead to visible differences and a possibly large number of empty lots are required to extend the city towards areas where buildings should be set up. The targets and scoring functions for the scenes are: a sketch for the silhouette of the generated model from a fixed perspective in Figures 1, 2 and 9(b,c), a volumetric target to fill for Figures 3 and 9(d), and a volumetric target to avoid for Figure 9(a).

**Results** obtained during optimization for the models presented in Figure 1 and Figure 3 are shown in Figure 7. The results obtained for the remaining test cases are shown in Figure 9. The RJMCMC method for Metropolis procedural modeling includes the parallel tempering, nonterminal selection and annealing optimization described by Talton et al. [TLL\* 11]. We use their proposed parallel tempering factor that assigns a 1% acceptance probability in the coldest chain to mutations that have a 70% chance in the hottest. As annealing factor we use 1.1. SMC and SOSMC correspond to the approaches described by Ritchie et al. [RMGH15]. GA corresponds to the genetic algorithm proposed by us. In all cases, we run them in our framework, as described in Section 7. GA uses a mutation probability of 30%, divided into 13.5%, 13.5%, and 3.0% for grow, cut, and alter, respectively for tree and spaceship scenes and a probability of 40%, divided into 12.5%, 10% and 18.0% respectively

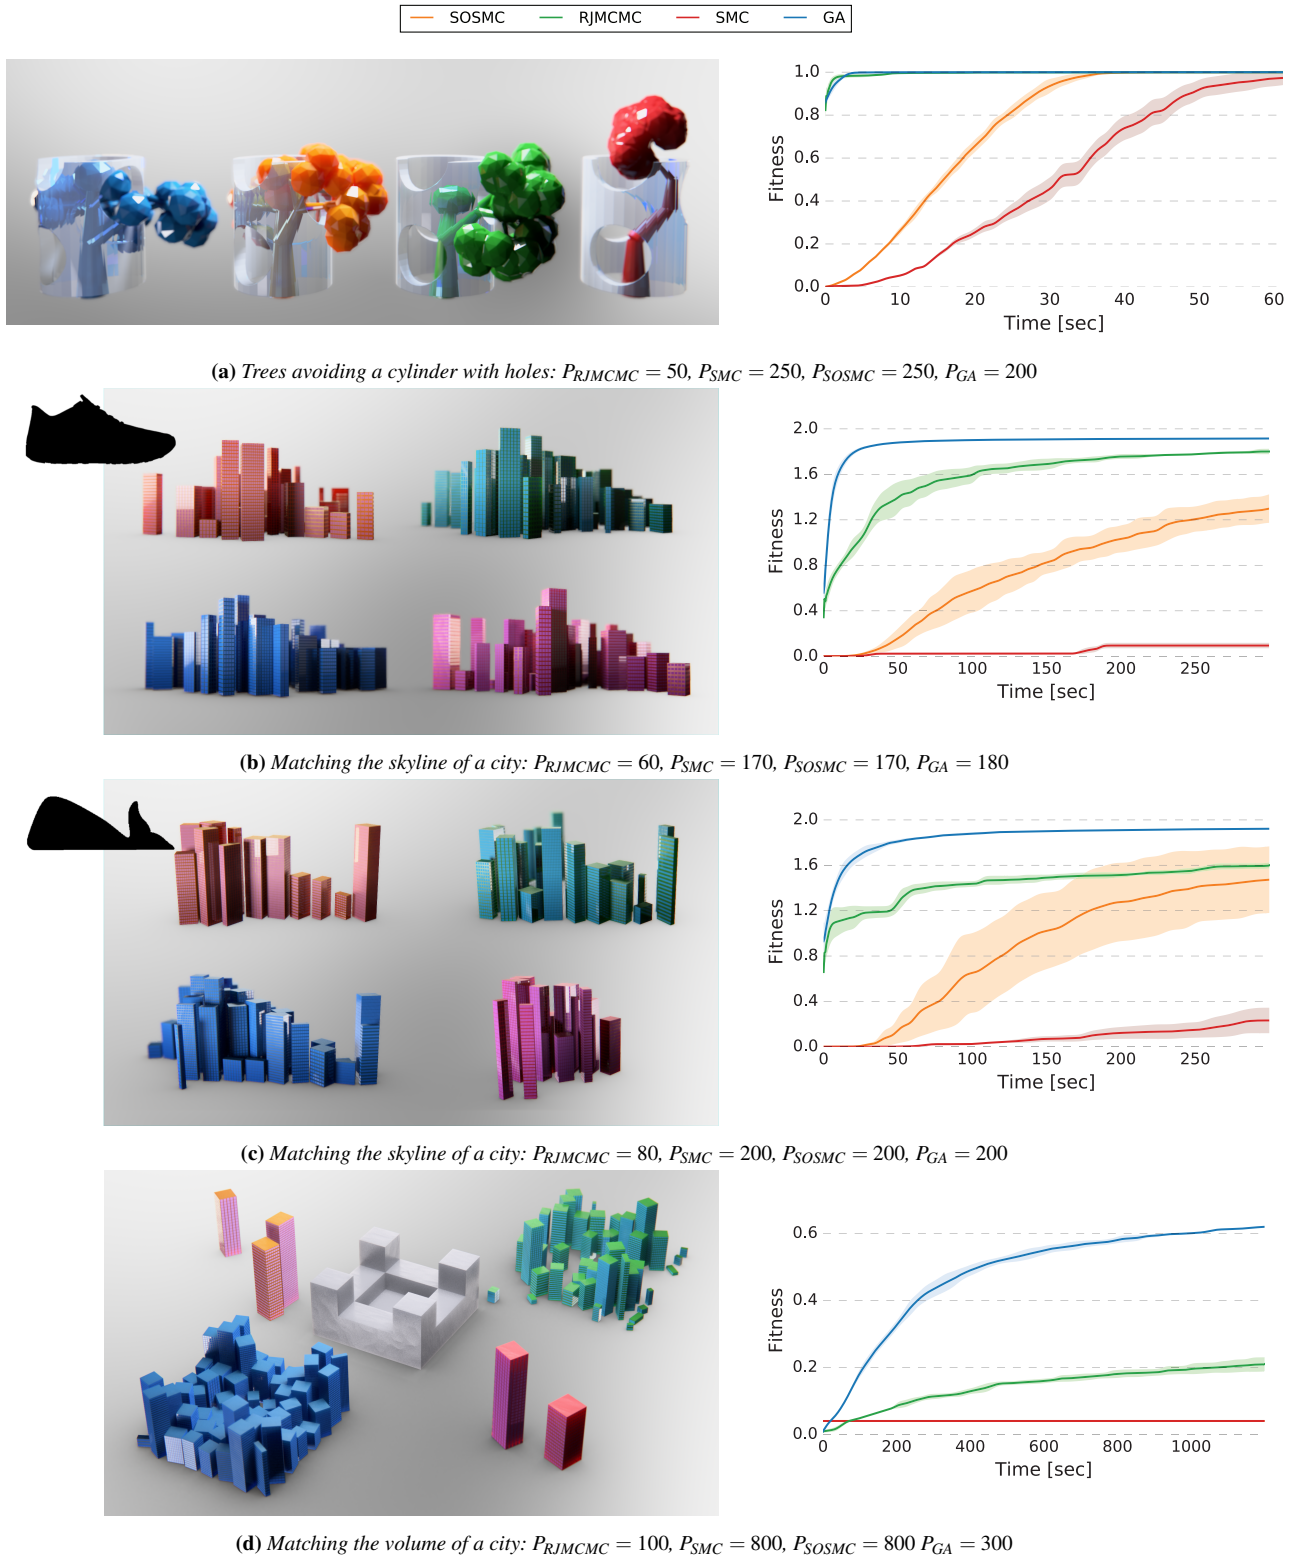


**Figure 8:** (a) The generator graph for the tree scene can output a trunk with a radius and length; the trunk and branches can have up to three sub-branches, which are each controlled by normal and inclination angles, radial scaling, length and offset; the foliage is always added as extension of a final branch. (b) In the city generator only the building node directly generates visible objects; the blue nodes recursively iterate over the space of possible building lots; the building can be rotated, and its size and height is controllable.

for the city scenes. For tournament selection we use  $k = 10$  and an elitism of one fifth of the population size. The optimal fitness value is determined by the weights used in the fitness function. Depending on the target and the expressiveness of the generator, it might not be reachable at all. However, the theoretical optimal fitness is 1.0 for the tree sketch and the volumetric city and tree scenes, and 2.0 for the spaceship and city shape scenes. We tried to tune the population size for all approaches to achieve the best possible results and report them as  $P_{RJMCMC}$ ,  $P_{SMC}$ ,  $P_{SOSMC}$ , and  $P_{GA}$  alongside the test results. Note that SMC and SOSMC work better with larger population size while RJMCMC works more efficiently with smaller populations; GA usually works best with a value in between.

The spaceship and tree sketch scenes (Figure 1, 3 and 7) outline the performance of the approaches well. For simple models like the spaceships with volume evaluation, where every node of the generator graph leads to a geometric object, SMC and SOSMC perform





**Figure 9:** Mean fitness values and standard deviation obtained for different test scenes plotted over time. SMC and SOSMC struggle during the beginning of the volumetric tree scene (a), as they run the relatively costly voxelization on many non-finished models. The image city scene (b,c) is also difficult for them as the generator graph is more complex. They fail to improve after the first few generations for the volumetric city scene (d). RJMCMC works reasonably well in all cases but shows slow convergence rates. GA in all cases shows the best convergence behavior among all tested approaches. Note that the fitness values are not normalized and depend on the chosen weights.

quite well. As they require the evaluation of the scoring function for every step it takes some time until a first model of sufficient size is generated (about 4s). While SMC slightly outperforms SOSMC during this first seconds, SOSMC outperforms SMC in the longer run. However, both of them plateau pretty soon, as they cannot ‘undo’ their initial choices. RJMCMC behavior is the opposite, since the initial state contains already full models, which score reasonably well. However, improving them takes a long time, falling below SMC and SOSMC. Then again, the improvement is steady and in the long run RJMCMC eventually outperforms SMC and SOSMC. GA also starts with a full initial population, from which it quickly increases the score, matching SOSMC during its best phase. However, in the long run GA keeps improving, similarly to RJMCMC. Note that the generation of the initial population hardly costs any time, as repeatedly calling the grow mutation has virtually no cost. For more complex targets, like the tree sketch image evaluation, SMC and SOSMC only show a good performance over the first seconds. Afterwards, RJMCMC and GA significantly improve the result, while SOSMC and SMC are stuck with the early parameter choices. GA in this case significantly outperforms RJMCMC. The GA solution after 3min achieves the same score as RJMCMC after 20min; SMC and SOSMC never reach this score. In this case GA profits strongly from the fact that it can copy partial solution from one of the three branches to the other.

The SMC and SOSMC methods have problems with complex generators, since choices made in the early iterations reflect in the performance during later stages. While this can be mitigated with a larger sample size at the cost of longer calculation time for each iteration, these complex generators usually also require a large number of geometric objects for good solutions. Which, due to the nature of SMC and SOSMC, require many iterations, therefore increasing the required calculation time even further. This makes choosing the right number of samples highly problem dependent, and especially cumbersome when aiming to achieve the best result within a limited computational budget.

We observe comparable behaviors in the other test cases shown in Figure 7. The large number of volumetric scoring function evaluations for uncompleted models reduce the early performance of SMC and SOSMC (a). More complex generators, like the city (b,c), which contain nodes that do not generate any visible objects are troublesome for the convergence of SMC and SOSMC. However, after these issues have been overcome, SOSMC at least can catch up with RJMCMC, before it plateaus. RJMCMC on the other hand always slowly and steadily converges. For the volumetric city scene (d), the weights were chosen such that wrongly placed buildings are punished severely. This removes the advantage of the initial population for RJMCMC and GA. But even so they manage to improve the score, with GA significantly outperforming RJMCMC, while SMC and SOSMC do not succeed to improve upon the solution after the first couple of generations. GA in almost all cases seems to combine the best of all worlds. Starting with a random population there is enough information to recombine good initial solutions which quickly increase the score. Furthermore, in contrast to SOSMC, GA continues to improve similarly to RJMCMC. Overall GA achieved the best results in all test cases during initial convergence and in the long run.

## 9. Discussion and Future Work

We have shown that genetic algorithms can be used for controlled procedural modeling in a variety of procedural approaches. Using our compact tree representation to encode genomes and linking this tree to the generator graph allows for efficient mutation and reproduction operations while making sure that all resulting chromosomes are valid. Comparing our GA to the state-of-the-art, we found that using a GA combines the best characteristics of previous approaches. The initial convergence is better than SOSMC and in the long run it significantly outperforms RJMCMC, yielding the best convergence in all stages of optimization. We attribute that to GA being able to combine the best features from an entire population.

We also found that GA shows good performance independently of the complexity of the used generators and scoring function. Also, in our testcases it was rarely necessary to tune the parameters to achieve good results, which hints at good stability of the approach. A limitation we experienced is that by copying parts of the genome, similar features can be duplicated in the model, *e.g.*, the remains of a heart shape are visible in the left branch of the GA result in Figure 1. Furthermore, our implementation of elitism is somewhat susceptible to ‘sample impoverishment’, leading to almost identical copies of the same individual in the elite population, but we did not notice a diminished performance caused by this.

In the future we will increase the efficiency of our approach further, by using a GPU-based generator. This will also increase the efficiency of the proposed scoring function evaluations which already use the GPU. Furthermore, we believe that combining SOSMC, RJMCMC, and GA and automatically choosing the best approach might even increase convergence further, *i.e.*, using SOSMC to create an initial population, switching to GA after a few iterations and then switching to RJMCMC in the long run. Considering our description of SOSMC and RJMCMC within our GA framework this seems to be possible. Finally, the major challenge we see is tackling full inverse procedural modeling, where the generator consists of the entire set of modeling operations and the goal is given by a fully detailed object or a scan of a real world object. We believe genetic algorithms have the potential to help make full inverse procedural modeling possible.

## Acknowledgments

We would like to thank Martin Kenzel for generating the final renderings for the paper. This research was supported by the Max Planck Center for Visual Computing and Communication.

Our approach can be downloaded from <https://github.com/crest01/ShapeGenetics>.

## References

- [BDK\*16] BOECHAT P., DOKTER M., KENZEL M., SEIDEL H.-P., SCHMALSTIEG D., STEINBERGER M.: Representing and scheduling procedural generation using operator graphs. *ACM Trans. Graph.* 35, 6 (Nov. 2016), 183:1–183:12. doi:10.1145/2980179.2980227. 2
- [Boe95] BOERS E. J. W.: Using L-Systems as graph grammar: G2L-Systems, 1995. 3

- [BT96] BLICKLE T., THIELE L.: A Comparison of Selection Schemes Used in Evolutionary Algorithms. *Evolutionary Computation* 4, 4 (1996), 361–394. doi:10.1162/evco.1996.4.4.361. 6
- [BvMM11] BENES B., ŠTAVA O., MĚCH R., MILLER G.: Guided procedural modeling. *Computer Graphics Forum* 30, 2 (2011), 325–334. doi:10.1111/j.1467-8659.2011.01886.x. 3
- [FP98] FUNES P., POLLACK J.: Evolutionary body building: Adaptive physical designs for robots. *Artif. Life* 4, 4 (Oct. 1998), 337–357. doi:10.1162/106454698568639. 3
- [Hav05] HAVEMANN S.: *Generative mesh modeling*. PhD thesis, University of Braunschweig-Institute of Technology, 2005. 2, 3
- [KK11] KRECKLAU L., KOBELT L.: Procedural modeling of interconnected structures. *Computer Graphics Forum* 30, 2 (2011), 335–344. doi:10.1111/j.1467-8659.2011.01864.x. 2
- [LD98] LINTERMANN B., DEUSSEN O.: A modelling method and user interface for creating plants. *Computer Graphics Forum* 17, 1 (1998), 73–82. 2
- [Lin68] LINDENMAYER A.: Mathematical models for cellular interactions in development ii. simple and branching filaments with two-sided inputs. *Journal of theoretical biology* 18, 3 (1968), 300–315. 2
- [Mag09] MAGDICS M.: Real-time generation of l-system scene models for rendering and interaction. In *Proceedings of the 25th Spring Conference on Computer Graphics* (New York, NY, USA, 2009), SCCG '09, ACM, pp. 67–74. doi:10.1145/1980462.1980478. 2
- [MBG\*12] MARVIE J.-E., BURON C., GAUTRON P., HIRTZLIN P., SOURIMANT G.: GPU Shape Grammars. *Computer Graphics Forum* (2012). doi:10.1111/j.1467-8659.2012.03201.x. 2
- [MLCB10] MARTIN A., LIM A., COLTON S., BROWNE C.: *Evolving 3D Buildings for the Prototype Video Game Subversion*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2010, pp. 111–120. URL: [http://dx.doi.org/10.1007/978-3-642-12239-2\\_12](http://dx.doi.org/10.1007/978-3-642-12239-2_12), doi:10.1007/978-3-642-12239-2\_12. 3
- [MP96] MĚCH R., PRUSINKIEWICZ P.: Visual models of plants interacting with their environment. In *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 1996), SIGGRAPH '96, ACM, pp. 397–410. doi:10.1145/237170.237279. 2
- [MSL\*11] MERRELL P., SCHKUFZA E., LI Z., AGRAWALA M., KOLTUN V.: Interactive furniture layout using interior design guidelines. *ACM Trans. Graph.* 30, 4 (July 2011), 87:1–87:10. doi:10.1145/2010324.1964982. 3
- [MWH\*06] MÜLLER P., WONKA P., HAEGLER S., ULMER A., VAN GOOL L.: Procedural modeling of buildings. *ACM Trans. Graph.* 25, 3 (July 2006), 614–623. doi:10.1145/1141911.1141931. 2
- [Och98] OCHOA G.: *On genetic algorithms and lindenmayer systems*. Springer Berlin Heidelberg, Berlin, Heidelberg, 1998, pp. 335–344. URL: <http://dx.doi.org/10.1007/BFb0056876>, doi:10.1007/BFb0056876. 3
- [OOI05] OKABE M., OWADA S., IGARASHI T.: Interactive design of botanical trees using freehand sketches and example-based editing. *Computer Graphics Forum* 24, 3 (2005), 487–496. doi:10.1111/j.1467-8659.2005.00874.x. 2
- [OSM\*09] O'NEILL M., SWAFFORD J. M., MCDERMOTT J., BYRNE J., BRABAZON A., SHOTTON E., McNALLY C., HEMBERG M.: Shape grammars and grammatical evolution for evolutionary design. In *Proceedings of the 11th Annual Conference on Genetic and Evolutionary Computation* (2009), GECCO '09, pp. 1035–1042. 3
- [Pat12] PATOW G.: User-friendly graph editing for procedural modeling of buildings. *IEEE Computer Graphics and Applications* 32, 2 (March 2012), 66–75. doi:10.1109/MCG.2010.104. 3
- [PHL\*09] PALUBICKI W., HOREL K., LONGAY S., RUNIONS A., LANE B., MĚCH R., PRUSINKIEWICZ P.: Self-organizing tree models for image synthesis. In *ACM SIGGRAPH 2009 Papers* (New York, NY, USA, 2009), SIGGRAPH '09, ACM, pp. 58:1–58:10. doi:10.1145/1576246.1531364. 2
- [PJ08] PILAT M. L., JACOB C.: Creature academy: A system for virtual creature evolution. In *2008 IEEE Congress on Evolutionary Computation (IEEE World Congress on Computational Intelligence)* (June 2008), pp. 3289–3297. doi:10.1109/CEC.2008.4631243. 3
- [PL91] PRUSINKIEWICZ P., LINDENMAYER A.: The algorithmic beauty of plants (the virtual laboratory). 2
- [PSK\*12] PIRK S., STAVA O., KRATT J., SAID M. A. M., NEUBERT B., MĚCH R., BENES B., DEUSSEN O.: Plastic trees: Interactive self-adapting botanical tree models. *ACM Trans. Graph.* 31, 4 (July 2012), 50:1–50:10. doi:10.1145/2185520.2185546. 1, 2
- [RMGH15] RITCHIE D., MILDENHALL B., GOODMAN N. D., HANRAHAN P.: Controlling procedural modeling programs with stochastically-ordered sequential monte carlo. *ACM Trans. Graph.* 34, 4 (July 2015), 105:1–105:11. doi:10.1145/2766895.2, 3, 7, 8
- [Sim91] SIMS K.: Artificial evolution for computer graphics. *SIGGRAPH Comput. Graph.* 25, 4 (July 1991), 319–328. doi:10.1145/127719.122752. 3
- [Sim94] SIMS K.: Evolving virtual creatures. In *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 1994), SIGGRAPH '94, ACM, pp. 15–22. doi:10.1145/192161.192167. 3, 4
- [Sip06] SIPSER M.: *Introduction to the Theory of Computation*, vol. 2. Thomson Course Technology Boston, 2006. 4
- [SKK\*14a] STEINBERGER M., KENZEL M., KAINZ B., MÜLLER J., PETER W., SCHMALSTIEG D.: Parallel generation of architecture on the GPU. *Computer Graphics Forum* 33, 2 (2014), 73–82. doi:10.1111/cgf.12312. 2
- [SKK\*14b] STEINBERGER M., KENZEL M., KAINZ B., WONKA P., SCHMALSTIEG D.: On-the-fly generation and rendering of infinite cities on the GPU. *Computer Graphics Forum* 33, 2 (2014), 105–114. doi:10.1111/cgf.12315. 2
- [SM15] SCHWARZ M., MÜLLER P.: Advanced procedural modeling of architecture. *ACM Trans. Graph.* 34, 4 (July 2015), 107:1–107:12. doi:10.1145/2766956.1, 2
- [SPK\*14] STAVA O., PIRK S., KRATT J., CHEN B., MÄŽCH R., DEUSSEN O., BENES B.: Inverse procedural modelling of trees. *Computer Graphics Forum* 33, 6 (2014), 118–131. doi:10.1111/cgf.12282. 2, 3
- [Sti75] STINY G. N.: *Pictorial and Formal Aspects of Shape and Shape Grammars and Aesthetic Systems*. PhD thesis, 1975. 2
- [TLL\*11] TALTON J. O., LOU Y., LESSER S., DUKE J., MĚCH R., KOLTUN V.: Metropolis procedural modeling. *ACM Trans. Graph.* 30, 2 (Apr. 2011), 11:1–11:14. doi:10.1145/1944846.1944851. 2, 7, 8
- [WA85] WADGE W. W., ASHCROFT E. A.: *LUCID, the Dataflow Programming Language*. Academic Press Professional, Inc., San Diego, CA, USA, 1985. 3
- [WP95] WEBER J., PENN J.: Creation and rendering of realistic trees. In *Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 1995), SIGGRAPH '95, ACM, pp. 119–128. doi:10.1145/218380.218427. 2
- [WWSR03] WONKA P., WIMMER M., SILLION F., RIBARSKY W.: Instant architecture. *ACM Trans. Graph.* 22, 3 (July 2003), 669–677. doi:10.1145/882262.882324. 2
- [XZCOC12] XU K., ZHANG H., COHEN-OR D., CHEN B.: Fit and diverse: Set evolution for inspiring 3d shape galleries. *ACM Trans. Graph.* 31, 4 (July 2012), 57:1–57:10. doi:10.1145/2185520.2185553. 3
- [YYW\*12] YEH Y.-T., YANG L., WATSON M., GOODMAN N. D., HANRAHAN P.: Synthesizing open worlds with constraints using locally annealed reversible jump mcmc. *ACM Trans. Graph.* 31, 4 (July 2012), 56:1–56:11. doi:10.1145/2185520.2185552. 2, 3