

Discrete Element Texture Synthesis




Figure 1: Discrete element texture synthesis. Given a small input exemplar (left within each image), our method can synthesize the corresponding output with user specified size and shape (right within each image). Our method, being data driven, can produce a variety of effects, including realistic/artistic phenomena, regular/semi-regular/irregular distribution, different number of element types, and variations in element properties including size, shape, and orientation. Case (a) is a semi-regular distribution with the input manually placed as the effect is difficult to achieve via physical simulation. The output in case (b) has very different shape and boundary conditions from the input. Case (c) contains different types of elements and different input and output shapes/boundary-conditions: the input is prepared on a planar board with the output served in a curved plate.

Abstract

A variety of natural and man-made phenomena can be characterized by repetitive discrete elements. Examples include a stack of fruits, a plate of dish, or a stone sculpture. Although certain results can be produced via existing methods based on procedural or physical simulation, these are often designed for specific applications. Some of these methods can also be hard to control.

We present discrete element texture synthesis, a data-driven method for placing repetitive discrete elements within a given large-scale structure. Our main goal is to provide a general approach that works for a variety of phenomena instead of just specific scenarios. We also want it easy to use, as the user only needs to specify an input exemplar for the detailed elements and the overall output structure, and our approach will produce the desired combination. Our method is inspired by texture synthesis, a methodology tailored for generating repetitions. However, existing texture synthesis methods cannot adequately handle *discrete* elements, often producing unnaturally broken or merged elements. Our method not only preserves the individual element properties such as color, shape, size, and orientation, but also their aggregate distributions. It also allows certain application specific controls, such as boundary constraints for physically realistic appearance. Our key idea is a new neighborhood metric for element distribution as well as an energy formulation for synthesis quality and control. As an added benefit, our method can also be applied for editing element distributions.

Keywords: discrete element, texture synthesis, editing

1 Introduction



A variety of natural or man-made phenomena can be characterized by a distinctive large scale structure with repetitive small scale elements. Some common examples include a stash of fruits or vegetables, a tiled house, or a stone sculpture. Due to the potential scale and complexity of such phenomena, it would be desirable to let the user specify only the overall structure while having automatic algorithms to produce the detailed elements.

One common method is physical simulation, for which the user specifies certain input controls (e.g. initial or boundary condition) and simply let the algorithm run course to produce the results [Baraff and Witkin 1997]. The primary advantage of physical simulation is fidelity to realism. However, such methods can be hard to control, as to produce the desired output the user might need to repeatedly tweak the input parameters. Physical simulation might not be suitable for man-made or artistic effects (e.g. see [Cho et al. 2007]). Another possibility is the procedural approach [Ebert et al. 2002]. However, procedural methods are known for their limited generality and only applicable to specific distribution (e.g. Poisson disk [Lague and Dutré 2005]) or phenomenon (e.g. rocks [Peytavie et al. 2009]). Furthermore, even though many procedural methods offer control via input parameters, tuning these to achieve the desired effects might require significant expertise.

Our main goal is to provide a general approach for placing repetitive discrete elements within a given large-scale structure. By general, we mean the approach should work for a variety of phenomena instead of tied to specific applications. We also want it easy to use, i.e. the user only needs to specify a small exemplary swath of detailed elements and the overall output shape, and the approach will produce the desired combination. In addition, we would like to produce both realistic and artistic effects simply via user specified element swath and output shape, as existing methods might not offer the kind of flexibility and control that a user desires.

We present discrete element texture synthesis, a data-driven approach for placing repetitive discrete elements within a given large-scale structure. We draw inspirations from texture synthesis [Wei et al. 2009], a methodology for producing natural repetitions from given exemplars. But unlike prior methods that might produce broken or merged elements (Figure 2), our method not only preserves the individual elements but also their aggregate distributions. Our key idea is a texture neighborhood definition for elements as well as a metric measuring neighborhood similarity. By ensuring that the input and output have similar texture neighborhoods, we are able to synthesize outputs that not only preserve the individual elements but also resemble aggregate distributions in the input exemplars. Thus, the user can easily achieve the desired result by simply supplying a proper input exemplar as well as the desired overall output size and shape. Since the user has maximum flexibility in specifying both the input exemplar and the output shape, our method is able

79 to achieve a variety of effects, including different dimensions (e.g.
80 2D or 3D), different element properties (including shapes, sizes,
81 colors) and distributions (e.g. regular/semi-regular/irregular), dif-
82 ferent number of element types (e.g. a stack of plums or a plate
83 of mixed vegetables), as well as physically realistic or artistic phe-
84 nomena (e.g. a pile of pebbles or a decorative mosaic pattern).

85 The main technical challenge of our approach is synthesizing el-
86 ment distributions. Unlike many prior texture synthesis methods
87 where the *domain* information \mathbf{p} is given (e.g. positions of pixels,
88 vertices, or voxels) and only the *range* information \mathbf{q} needs to be
89 determined (e.g. colors of pixels, vertices, or voxels), we have to
90 compute both \mathbf{p} and \mathbf{q} as part of the synthesis process. We achieve
91 this by a carefully designed metric for measuring the similarity be-
92 tween two texture neighborhoods, accounting for both \mathbf{p} and \mathbf{q} .
93 Even though there are prior methods targeting specific scenarios of
94 this general problem (e.g. 2D NPAR distribution [Ijiri et al. 2008;
95 Hurtut et al. 2009]), to our knowledge these are not for synthesizing
96 general discrete elements, e.g. 3D or physically realistic effects for
97 which our method can easily handle. We formulate our synthesis
98 algorithm as an energy optimization process [Kwatra et al. 2005]
99 to allow us not only properly determine \mathbf{p} and \mathbf{q} but also satisfy
100 certain specific application demands, such as boundary constraints
101 for physically plausible effects. We choose this optimization frame-
102 work mainly for its flexibility, as both the basic neighborhood simi-
103 larity as well as additional application-specific needs can be incor-
104 porated as individual energy terms. Adopting a familiar framework
105 of optimization also facilitates easy extension and adoption. How-
106 ever, even though optimization is a common methodology and has
107 been used in many different algorithms, the main challenges are on
108 how to properly design the individual algorithmic components for
109 discrete element synthesis.

110 As an added benefit, our method can also be applied for editing
111 element distributions. Specifically, the user can change not only \mathbf{q}
112 but also \mathbf{p} of a few elements, and our method will automatically
113 propagate such changes to all other elements with similar texture
114 neighborhoods, relieving the user from the potential tedious chore
115 of manual repetitions. This editing application is possible thanks to
116 the texture neighborhood metric we developed for direct synthesis.

117 2 Previous Work

118 **Multi-scale synthesis** A variety of phenomena consists of
119 small scale repetitions within a distinctive large scale structure.
120 Such phenomena could be computed with better quality or effi-
121 ciency by applying different methods for different scales; some
122 examples include fluid turbulence [Kim et al. 2008; Narain et al.
123 2008], hair strands [Wang et al. 2009], crowds [Lerner et al. 2007;
124 Narain et al. 2009], or motion fields [Ma et al. 2009]. Our approach
125 follows this general philosophy and focuses on discrete elements.

126 **Example-based texturing** Example-based texturing is a gen-
127 eral data-driven methodology for synthesizing repetitive phenom-
128 ena (see survey in [Wei et al. 2009]). However, the basic repre-
129 sentations in most existing texture synthesis methods such as pix-
130 els [Efros and Leung 1999], vertices [Turk 2001] or voxels [Kopf
131 et al. 2007] cannot adequately represent individual or *discrete* el-
132 ements with semantic meanings, such as common objects seen in
133 our daily lives. Without a basic representation that has knowledge
134 of the discrete elements it would be very difficult to synthesize these
135 elements adequately; even though artifacts could be reduced via ad-
136 ditional constraints on top of existing methods (e.g. [Zhang et al.
137 2003; Wu and Yu 2004]), there is no guarantee that the individual
138 elements would be preserved. Thus, the synthesized textures can
139 have elements that are broken or merged with each other (Figure 2).
140 Such artifacts can be quite visible and thus better avoided.




Figure 2: Pixel versus element synthesis. The pixel synthesis result is produced by combining discrete optimization [Han et al. 2006] with texton mask [Zhang et al. 2003].

141 **Geometry synthesis** Our method is also related to geome-
142 try synthesis, especially those via texture methods such as meshes
143 [Zhou et al. 2006], models [Merrell and Manocha 2008], or ter-
144 rains [Zhou et al. 2007]. However, similar to other texture synthe-
145 sis methods these are mainly for continuous patterns and might lack
146 necessary information to preserve or control discrete elements, e.g.
147 broken elements as can be seen in Figure 5b of [Zhou et al. 2006].

148 **Element packing** There exist methods that pack a set of dis-
149 crete elements into a specific domain or shape, such as mosaic tiles
150 [Hausner 2001; Kim and Pellacini 2002] or 3D object collage [Gal
151 et al. 2007]. However, the element distributions in these methods
152 are usually determined via specific procedures or semi-manual user
153 interface, instead of from input exemplars targeted at general distri-
154 butions as in our approach.

155 **Texture element placement** Even though the majority of
156 example-based texturing methods are not suitable for discrete el-
157 ements, potential solutions have been explored by a few pioneering
158 works. However, despite the promises shown in these techniques,
159 they might fall short in certain aspects. Dischler et al. [2002] and
160 Liu et al. [2009] obtain distribution from input exemplars to place
161 2D textons, but these techniques are not designed for general dis-
162 crete elements. Barla et al. [2006] synthesized discrete elements
163 but their positions are determined by Lloyd relaxation, not from the
164 input exemplars. Ijiri et al. [2008] synthesized element positions
165 via a growth method similar to [Efros and Leung 1999] but their
166 method appears to be less general and more complex than ours,
167 e.g. dealing with only 1-ring neighborhoods and requiring trian-
168 gulation. Thus their method is sufficient for the target 2D NPAR
169 applications but probably not for more general effects such 3D or
170 physically realistic distribution. Jodoin et al. [2002] and Kim et
171 al. [2009] applied texture synthesis for generating stipple distribu-
172 tions, but not general discrete elements. Hurtut et al. [2009] took
173 into account element attributes via area and appearance analysis,
174 but only deals with static 2D non-photorealistic elements, not 3D
175 or physically-realistic phenomena. Our method is inspired by these
176 pioneering works, but aims at synthesizing discrete elements in a
177 general setting, including 2D and 3D distribution, volume and sur-
178 face synthesis, regular/semi-regular/irregular configuration, varia-
179 tions in number of element types, shapes, sizes, as well as artistic
180 and physically-realistic effects.

181 3 Core Algorithm

Given an input exemplar \mathbf{z} consisting of a set of elements with the relevant domain/position \mathbf{p} and range/attribute \mathbf{q} information, our goal is to synthesize an output \mathbf{x} that it is similar to \mathbf{z} in terms of both \mathbf{p} and \mathbf{q} . We can formulate this synthesis of discrete elements as an optimization problem [Kwatra et al. 2005] by minimizing the

following energy function:

$$E_t(\mathbf{x}; \mathbf{z}) = \sum_{s \in X^\dagger} |\mathbf{x}_s - \mathbf{z}_s|^2 \quad (1)$$

where E_t measures the similarity between the input exemplar \mathbf{z} and the output \mathbf{x} via local neighborhoods around elements. Specifically, for each output element s , we take a small set of elements near it as the texture neighborhood \mathbf{x}_s , find the most similar input neighborhood \mathbf{z}_s , and measure their distance $|\mathbf{x}_s - \mathbf{z}_s|$. We repeat this same process for each $s \in X^\dagger$, a subset of all input elements, and sum their squared differences. Our goal is to find an output \mathbf{x} with low energy value. Below we describe details about our energy formulation as well as a solver for this optimization problem. For each reference, we have summarized the algorithm in Pseudocode 1.

```

function  $\mathbf{x} \leftarrow$  DiscreteElementTextureOptimization( $\mathbf{z}$ )
  //  $\mathbf{x}$ : output distribution
  //  $\mathbf{z}$ : input exemplar
   $\mathbf{x} \leftarrow$  Initialize( $\mathbf{z}$ ) // Section 4.2
  iterate until convergence or enough # of iterations reached
    // search phase, i.e. the “M-step” in [Kwatra et al. 2005]
     $\{\mathbf{z}_s, s \in X^\dagger\} \leftarrow$  Search( $\mathbf{x}, \mathbf{z}$ )
    // assignment phase, i.e. the “E-step” in [Kwatra et al. 2005]
    Assign( $\{\mathbf{z}_s, s \in X^\dagger\}, \mathbf{x}$ )
  end
  return  $\mathbf{x}$ 

function  $\{\mathbf{z}_s, s \in X^\dagger\} \leftarrow$  Search( $\mathbf{x}, \mathbf{z}$ ) // Section 3.4
  foreach element  $s \in X^\dagger$  //  $X^\dagger$ : a subset of all output elements
     $\mathbf{x}_s \leftarrow$  output neighborhood around  $s$ 
     $\mathbf{z}_s \leftarrow$  find most similar neighborhood in  $\mathbf{z}$  to  $\mathbf{x}_s$ 
  end
  return  $\{\mathbf{z}_s\}$ 

function Assign( $\{\mathbf{z}_s\}, \mathbf{x}$ ) // Section 3.5
  foreach output element  $s \in X$ 
     $\mathbf{p}(s) \leftarrow$  weighted combination of predicted positions
      from output neighborhoods that overlap  $s$ 
     $\mathbf{q}(s) \leftarrow$  select the vote that minimizes the energy function
  end

```

Pseudocode 1: Discrete element texture synthesis.

3.1 User Inputs

In addition to the input exemplar, the user also needs to supply the following main inputs:

Neighborhood size This is the standard parameter for neighborhood-based texture synthesis [Wei et al. 2009]. The user simply specifies the spatial extent of the neighborhoods, and for each element s , we construct its neighborhood $\mathbf{n}(s)$ by taking the union of all elements within the spatial extent centered at s .

Output shape The user also needs to define the output size and shape. Our algorithm will then attempt to obey it as much as possible, i.e. filling the domain interior with elements while avoiding them spill outside the domain. The algorithm will also try to maintain similarity between input and output boundary element configurations.

Element attributes The user can also specify what kinds of element properties to consider. The element positions \mathbf{p} are mandatory, but the range attributes \mathbf{q} such as element type, geometry, and appearance could be optional depending on the target applications. See Section 3.3 for more details.

3.2 Neighborhood Metric

The neighborhood similarity metric is the core component for neighborhood-based texture synthesis algorithms [Wei et al. 2009]. For traditional texture synthesis that has fixed sample positions \mathbf{p} , this can be done easily by either simple sum-of-squared differences (SSD) of the range information \mathbf{q} (such as colors) in a regular setting (e.g. pixels or voxels) or by resampling irregular samples into a regular setting before proceeding with SSD as in the former case (e.g. mesh vertices). However, in our case, since we have to synthesize both \mathbf{p} and \mathbf{q} , we need to incorporate both of them into the neighborhood metric. Formally, let $\mathbf{n}(s)$ denote the spatial neighborhood around an element s . We measure the distance $|\mathbf{n}(s_o) - \mathbf{n}(s_i)|^2$ between the neighborhoods of two elements s_o and s_i via the following formula:

$$|\mathbf{n}(s_o) - \mathbf{n}(s_i)|^2 = \sum_{s'_o \in \mathbf{n}(s_o)} |\hat{\mathbf{p}}(s'_o) - \hat{\mathbf{p}}(s'_i)|^2 + \alpha |\mathbf{q}(s'_o) - \mathbf{q}(s'_i)|^2 \quad (2)$$

where s'_o is an element $\in \mathbf{n}(s_o)$, $s'_i \in \mathbf{n}(s_i)$ the “matching” element of s'_o (explained below), $\hat{\mathbf{p}}(s') = \mathbf{p}(s') - \mathbf{p}(s)$ (i.e. the relative position of s' with respect to the neighborhood center s), and α the relative weight between domain \mathbf{p} and range \mathbf{q} information.

Intuitively, what Equation 2 tries to achieve is (1) align the two neighborhoods $\mathbf{n}(s_o)$ and $\mathbf{n}(s_i)$, (2) match up their elements in pairs $\{(s'_i, s'_o)\}$, and (3) compute the sum of squared differences of both \mathbf{p} and \mathbf{q} among all the pairs. We determine the pairings by first identifying the pair (s'_i, s'_o) with minimum $|\hat{\mathbf{p}}(s'_o) - \hat{\mathbf{p}}(s'_i)|$, exclude them for further consideration, and repeat the process to find the next pair until $\mathbf{n}(s_o)$ runs out of elements. (We prevent $\mathbf{n}(s_i)$ from running out of elements before $\mathbf{n}(s_o)$ by not presetting its spatial extent, essentially giving $\mathbf{n}(s_i)$ an ∞ size.) We have found that the heuristic above works well in practice, and provides similar quality with a more rigorous but much slower approach that considering all possible pair matching (s'_i, s'_o) in brute force.




Figure 3: Illustration for our neighborhood metric. Each pair of matched input and output elements has not only similar relative positions (to the center element) but also similar color, shape, and orientation. Unmatched input elements are shown in black.

We use the neighborhood metric above throughout our algorithm, including both the search and the assignment steps. Specifically, in the search step, we use only the scalar distance value computed by Equation 2 to pick the most similar input neighborhood \mathbf{z}_s for each output neighborhood \mathbf{x}_s . However, in the assignment step, we need further information for the matching pairs $\{(s'_i, s'_o)\}$ in order to determine the \mathbf{p} and \mathbf{q} values for the output element s . More details can be found in Section 3.4 & 3.5.

Note that even though some prior methods have adopted similar neighborhood metric for discrete elements, they do not entirely satisfy our needs. For example, [Barla et al. 2006] uses Hausdorff distance and thus does not allow explicit control of pair-wise element matching, e.g. the need to avoid duplications, and [Ijiri et al. 2008] considers only 1-ring neighbor positions \mathbf{p} instead of our formula-

242 tion that allows not only general neighborhoods but also considers 291
 243 both \mathbf{p} and \mathbf{q} . Furthermore, our neighborhood definition does not 292
 244 require additional processing such as triangulation in [Jiri et al.
 245 2008], making our method easier to implement, especially for non-
 246 2D applications such as 3D volume or surface synthesis.

247 3.3 Element attributes \mathbf{p} and \mathbf{q}

248 Here, we describe more details about the element domain \mathbf{p} and
 249 range \mathbf{q} information, and how to measure their differences in Equa-
 250 tion 2. The \mathbf{p} part is relatively straightforward; it is just the element
 251 position, and we measure the difference $\mathbf{p}(s) - \mathbf{p}(s')$ between two
 252 elements s and s' via the usual Euclidean metric. The \mathbf{q} part can
 253 contain a variety of information depending on the particular applica-
 254 tion scenario. For the simplest case of point distribution, \mathbf{q} can
 255 be empty. Below is a list for more typical applications involving
 256 concrete objects as elements:

257 **Orientation** The orientation of an element is represented as a nor-
 258 malized quaternion for both 2D and 3D cases. We compute
 259 the difference between two quaternions via the standard ap-
 260 proach of taking the inverse cosine of their dot product.

261 **Geometry** Each element can have geometry with different size and
 262 shape from one another. In general, we can measure the dif-
 263 ference between two element geometries via Hausdorff dis-
 264 tance (after aligning element centers and orientations to avoid
 265 double counting their contributions).

266 **Appearance** Each element can also have different appearance at-
 267 tributes, including colors and textures. We can measure their
 268 appearance differences via color histograms.

269 **Type** In general, both the geometry and appearance are parts of the
 270 intrinsic element attributes (that remain largely invariant with
 271 respect to position and orientation). Beyond geometry and
 272 appearance, we can also consider other kinds of intrinsic ele-
 273 ment attributes depending on the specific application contexts,
 274 such as high level semantic meanings. For maximum flexibil-
 275 ity, we allow the user to specify the distance metric between
 276 intrinsic element properties. In addition, when the number of
 277 input elements is sufficiently small or can be grouped into a
 278 small number of types, we can pre-compute their intrinsic dis-
 279 tances for run time efficiency. For most of our examples, we
 280 have found it sufficient to use an integer number to identify
 281 the element type, and set the intrinsic distance to be 0 if they
 282 are the same, and or 1 if not.

283 3.4 Search Step

284 During the search step, we find, for each output element s_o , the
 285 best match input element s_i with the most similar neighborhood,
 286 i.e. minimizing the energy value in Equation 2. This search can be
 287 conducted by exhaustively examining every input element, but this
 288 can be computationally expensive. Instead, we adopt k-coherence
 289 search for constant time computation, as detailed in Section 4.3.

290 3.5 Assignment Step

\mathbf{p} assignment At the beginning of the assignment step, we have 310
 multiple input neighborhoods $\{z_{s'_o}\}$ overlapping every output ele- 311
 ment s_o , where $z_{s'_o}$ is the matching input neighborhood for output 312
 element s'_o as determined in the search step (Section 3.4) and s'_o is 313
 sufficiently close to s_o so that the spatial extent of $z_{s'_o}$ covers s_o .
 Each such $z_{s'_o}$ provides a predicted position $\tilde{\mathbf{p}}(s'_o, s_o)$ for element
 s_o :

$$314 \tilde{\mathbf{p}}(s'_o, s_o) = \mathbf{p}(s'_o) + \mathbf{p}(s_i) - \mathbf{p}(s'_i) \quad (3) \quad 315$$

where s_i/s'_i indicates the matching input element for s_o/s'_o as de-
 292 scribed in the neighborhood metric (Section 3.2). See Figure 4.




Figure 4: Illustration for the assignment step.

To minimize the energy function E_t in Equation 1, the sample posi-
 293 tion $\mathbf{p}(s_o)$ is updated as a weighted combination of all $\{\tilde{\mathbf{p}}(s'_o, s_o)\}$
 294 where $z_{s'_o}$ covers s_o :

$$295 \mathbf{p}(s_o) = \frac{\sum_{s'_o} \omega(s'_o, s_o) \cdot \tilde{\mathbf{p}}(s'_o, s_o)}{\sum_{s'_o} \omega(s'_o, s_o)} \quad (4)$$

The relative weight ω is determined as

$$300 \omega(s'_o, s_o) = \frac{1}{\alpha |s'_o - s_o| + 1} \quad (5)$$

where α is a user-specified constant. We have found it sufficient to
 301 set $\alpha = 0$ which yields Equation 4 to a simple (equal weighted)
 302 average.

\mathbf{q} assignment We assign \mathbf{q} by a simple voting scheme. For
 303 each output element s_o , we gather a set of votes $\{\mathbf{q}(s_i)\}$, where
 304 each s_i is matched to s_o for a certain overlapping neighborhood de-
 305 termined in the search step (see Figure 4). Then we choose the one
 306 that has the minimum sum of distance across the vote set $\{\mathbf{q}(s_i)\}$:

$$307 \mathbf{q}(s_o) = \arg \min_{\mathbf{q}(s_i)} \sum_{s_i \in \{s_i\}} |\mathbf{q}(s_i) - \mathbf{q}(s'_i)|^2 \quad (6)$$

308 where s_i' runs through the set of elements $\{s_i\}$ matched to s_o dur-
 309 ing the search step. Essentially, what we are trying to do is to find
 a $\mathbf{q}(s_o)$ that is closest to the arithmetic average of $\{\mathbf{q}(s_i)\}$; this
 is very similar to the use of a discrete solver [Han et al. 2006] for
 solving a least squares problem [Kwatra et al. 2005].

Discussion In the assignment steps we use blend for \mathbf{p} (Equa-
 310 tion 4) but selection for \mathbf{q} (Equation 6). (In some sense, the for-
 311 mer is analogous to the least squares solver [Kwatra et al. 2005]
 312 while the latter the discrete k-coherence solver [Han et al. 2006].)
 313 The main reason is that blend works better than selection for \mathbf{p} , but
 314 might not be suitable for all \mathbf{q} attributes. For example, the orienta-
 315 tion, shape or type information might not be meaningfully blended.
 316 Furthermore, to apply k-coherence acceleration (Section 4.3), we
 317 will have to copy instead of blending the \mathbf{q} information.

4 Advanced Features

Here, we describe several advanced features of our method beyond
 the core algorithm presented in Section 3.

4.1 Synthesis control

Even though texture synthesis can automatically produce a station-
 ary output, for realistic effects, it is usually desirable to control cer-
 tain global aspects of the synthesis process. This synthesis control
 has appeared in prior methods, e.g. controllable [Lefebvre and

318 Hoppe 2005] or globally-varying [Zhang et al. 2003; Wei et al. 369
319 2008] synthesis. Here, we describe several synthesis controls that 370
320 we have found useful in producing our results. 371

321 **Overall shape** Given a user specified output domain shape 372
322 (Section 3.1), we would like the synthesis process to comply with 373
323 this as much as possible, i.e. put elements inside instead of out- 374
324 side the overall shape, and transfer boundary/interior output ele- 375
325 ments from similarly configured boundary/interior input elements. 376
326 We can achieve this with a density map \mathbf{c} that shaped as the out- 377
327 put domain with values within the range $[0, 1]$, where higher value 378
328 indicates larger probability of element appearance. 379

In the search step, we find the input neighborhood \mathbf{z}_s that minimizes 380
not only the usually texture (dis)similarity $|\mathbf{x}_s - \mathbf{z}_s|^2$ but also an 381
additional term $\lambda|\mathbf{c}_s - \mathbf{z}_s|^2$, where λ is a relative weight and \mathbf{c}_s the 382
sampled density value of \mathbf{x}_s . Specifically, 383

$$|\mathbf{c}_s - \mathbf{z}_s|^2 = \sum_{s'_i \in \mathbf{z}_s} |c(s'_i) - 1|^2 \quad (7)$$

329 where $\{c(s'_i)\}$ are sampled density \mathbf{c} values at positions 384
330 $\{\mathbf{p}(s_o) + \mathbf{p}(s'_i) - \mathbf{p}(s_i), s'_i \in \mathbf{z}_s\}$. Essentially, we shift the entire 385
331 input neighborhood \mathbf{z}_s to the center location $\mathbf{p}(s_o)$ and sample \mathbf{c} at 386
332 the shifted element positions. 387

333 **Local orientation** The user can also optionally specify a local 388
334 orientation of the output texture so that the output patterns are 389
335 aligned with the user choice instead of the default global coordi- 390
336 nate frame. This allows the production of more interesting results, 391
337 e.g. oriented flow patterns as in [Ijiri et al. 2008]. Algorithmically, 392
338 this can be easily achieved by using the local instead of the global 393
339 frame at each element during each step of our algorithm, including 394
340 neighborhood metric, search, assignment, and initialization. Note 395
341 that the incorporation of local frames into a texture optimization 396
342 framework has been done in prior methods, e.g. [Ma et al. 2009].

343 **Constraints** For certain application scenarios it might be desir- 397
344 able to maintain specific constraints, e.g. minimize penetrations for 398
345 physical elements or avoid elements floating in the mid air. Even 399
346 though texture synthesis cannot completely guarantee all these con- 400
347 straints, it can usually be tuned to produce visually plausible results. 401
348 For inter-penetration, we have found that minimizing neighborhood 402
349 dissimilarity in Equation 2 would also lead to less penetrations. For 403
350 other kinds of constraints, we have found it effective to restrain the 404
351 kinds of input elements that can be transferred to the constrained 405
352 output regions. (This is a commonly used method in texture syn- 406
353 thesis, e.g. for volumetric layers [Owada et al. 2004].) For example, 407
354 to reduce the chance of elements floating in the mid air, during the 408
355 search step we only select input floor elements for output floor ele- 409
356 ments. During the assignment step, we maintain the vertical eleva- 410
357 tion of these floor elements to be invariant while minimizing other 411
358 energy terms as described in Section 3.5. Similar mechanisms can 412
359 be applied to other kinds of constraints, as we will show in Sec- 413
360 tion 5. 414

361 4.2 Initialization 415

362 **White noise** This is perhaps the simplest and most flexible ini- 416
363 tialization method, by randomly copying elements from the input to 417
364 the output domain. One downside of such a white noise initializa- 418
365 tion, though, is that it may require an excessive number of iterations 419
366 to converge via our optimization procedure. It could also get stuck 420
367 in a local minimum, causing unsatisfactory element distribution in 421
368 certain regions of the output. 422

Patch copy To address the deficiencies of white noise initial- 423
ization, we have found another strategy, patch copy, which works 424
quite well. Patch-base synthesis has demonstrated to be effective 425
for image textures (see e.g. [Liang et al. 2001; Efros and Freeman 426
2001] and the survey in [Wei et al. 2009]). Here, we apply a simi- 427
lar method for initialization. We first divide the input exemplar 428
and output region into uniform grids, with each grid cell correspond- 429
ing to a patch of elements, and then randomly copy input patches 430
into output grids, just like patch-based image synthesis. In addition, 431
when copying patches we take into account the user controls (Sec- 432
tion 4.1), such as aligning patches with local orientations as well 433
as preferring input patches with similar boundary conditions to the 434
output region. 435

436 4.3 Acceleration by k-coherence 437

438 Since our method copies the \mathbf{q} information from input to output ele- 439
ments, we can apply k-coherence [Tong et al. 2002] throughout our 440
entire algorithm. The main difference between our method and the 441
original k-coherence method is that we have to deal with irregularly 442
placed samples. However, this problem has been addressed in the 443
context of irregular mesh vertices [Han et al. 2006], and we could 444
adopt a similar strategy here. Specifically, during the pre-process, 445
we can build a similarity-set for each input sample via our search- 446
step as described in Section 3.4. At run-time, we build the candi- 447
date set by collecting the similarity sets from all the neighboring 448
elements, with the offset part properly computed by the recorded 449
element pairs (Section 3.2). 450

451 5 Results 452

453 5.1 Element distribution 454

455 Our method can produce a variety of element distributions with dif- 456
ferent attributes, such as dimensionality (2D/3D), volume/surface 457
synthesis, regular/semi-regular/irregular distribution, number of el- 458
ement types, variations in element size/shape/color/texture, out- 459
put domain size/shape/orientation, and artistic/realistic phenomena. 460
Since our method is data driven, we can handle all these by simply 461
using different input exemplars and output domains. We wish to 462
emphasize that the input and output specifications are more or less 463
de-coupled, i.e. the same input exemplar can be used for different 464
output domains, and vice versa (see Figure 6). This is a key factor 465
facilitating easy and flexible usage of our method. 466

467 **Input exemplar properties** Using input exemplars with differ- 468
ent properties, our method can produce a variety of different results 469
as shown in Figure 1 & Figure 6. We begin with the simplest but 470
also very common case of one type of elements, e.g. Figure 1a, 471
6a, and 6e. But even such one-element-type distributions may have 472
certain properties that cannot be easily captured by procedural or 473
physical simulation methods. For example, the user might prefer to 474
arrange a stack of plums in a near-regular configuration (Figure 1a), 475
or a collection of carrots in specific orientations (Figure 6f, 6g, and 476
6h). Notice that these examples cannot be produced by physical 477
simulation (e.g. dropping objects until they come to rest) as the out- 478
puts are unlikely to reach the desired user intention. One possibility 479
is to manually place the elements, but this could quickly become 480
very tedious for sufficiently large outputs. Using our method, the 481
user only needs to manually place a small input exemplar and our 482
method will automatically produce the desired output. The bananas 483
(Figure 6a) present another interesting case due to their unique long 484
and curvy shapes. For this case, we generated the input via physical 485
simulation to show that our method can produce visually realistic 486
outputs via physically validated input. More interesting distribution 487
can be produced by multiple types of elements with different sizes 488

and shapes, e.g. a dish containing corns, diced carrots, and green beans (Figure 1c).

Output domain properties In addition to the input exemplar properties like element type and distribution, the user can also specify the output domain properties, including size, shape, and orientation field, to achieve different effects. Beyond physically plausible shapes like a stack, a box, a pile, or a bowl as shown in Figure 1 and 6, the user can also specify a more complex or interesting shape as a sculpture (Figure 1b), a tai-chi pattern (Figure 6f), a knot (Figure 6h), or a building (Figure 6k). Our method can also be applied to both volume (e.g. Figure 1) and surface/shell (e.g. Figure 6k, 6l, and 6m) synthesis. Note that these results span both physically realistic as well as artistic effects. As noted in [Cho et al. 2007], physical simulation might produce output distributions that look flat or boring. To produce visually more appealing effects, it is often desirable to have the output in a physically unstable or implausible configuration. Cho et al. [2007] achieved this via certain ad-hoc approaches, e.g. stopping physical simulation in the middle prior to completion (Figure 10 in [Cho et al. 2007]) or using repeated skimming and an up-side-down collision mesh (Figure 15 in [Cho et al. 2007]). Our method can easily produce the desired effect in a more principled and more controllable manner by simply using the proper output domains.




Figure 5: Boundary condition comparisons. Shown here are the profile views for the texture in Figure 1c.

Boundary handling Properly boundary handling is important to produce satisfactory results for certain discrete element textures that exhibit different distributions near and away from the overall boundaries, e.g. floor or box sides. Our experimental results indicate that these boundary conditions can be adequately handled by our control mechanisms described in Section 4.1. Without such mechanisms, the synthesis results might exhibit poor boundary conditions, as shown in Figure 5. We wish to emphasize that our method does not require all possible output boundary configurations to be present in the input exemplar; as shown in Figure 1 and 6, even though the output can contain different boundary shapes and orientations not present in the simpler input exemplars, the combination of local orientation and boundary handling can still produce satisfactory results.

5.2 Distribution editing

As an added benefit, our method can also be applied for editing discrete element textures, for not only individual element properties \mathbf{q} but also their distributions \mathbf{p} . All these can be achieved by the very same algorithms that we have built for synthesizing discrete element textures, especially the neighborhood metric. Texture editing

has been shown to be useful for a variety of application scenarios (see e.g. [Brooks and Dodgson 2002; Matusik et al. 2005; Zhou et al. 2006; Liu et al. 2009; Cheng et al. 2010]). Our method follows this line of thinking, but can achieve certain effects that may benefit from the explicit knowledge of the discrete elements.

Figure 7 demonstrates a potential example. Given an input pattern consisting of discrete elements, we aim to use our method to edit the element properties \mathbf{q} and distributions \mathbf{p} to produce more versatile effects. The user may simply select a typical element, performs some edits, and our method will automatically propagate relevant edits to all other elements with similar neighborhoods to the user interacted element. Note that without our automatic propagation, it would be quite tedious for the user to manually repeat the same edits to all relevant elements.

5.3 Usage and parameters

Input preparation Unlike other texture synthesis applications where the input exemplars can be obtained directly (e.g. downloading an image), for discrete element textures the user would have to do some work to produce the input exemplars, including both the individual elements and their distribution. For the results shown in this paper, we prepare the elements via standard modeling tools (e.g. Maya) and distribute them either manually or by simple simulation. For the modeling part, we have found it sufficient to make just one element for each type and the quality seems to work quite well for human perceptions [Ramanarayanan et al. 2008]. If additional element prototypes are desired, we have found it sufficient to slightly perturb the prototype element properties (e.g. geometry or color) via procedural noise. For the distribution part, since the input exemplar is usually quite small, manual placement seems quite feasible (e.g. the inputs for Figure 1a, 6e & 6i). It is also possible to use simple physical simulation for the input distribution for more random or physically realistic effects, even for outputs that might not be easy to produce via simulation (e.g. Figure 1b).

Parameters Similar to prior texture synthesis methods, one of the most important parameters is the neighborhood size. In our results we have found it sufficient to use a neighborhood size containing roughly 1- to 3-ring neighbors ($\sim 3^n$ to 7^n neighborhood in n -D pixel synthesis) depending on whether the pattern is more stochastic or structured. Other important parameters include α (for Equation 2) and λ (for Equation 7), for which we set to be of the same order of magnitude as the average distance between elements. (For example, if the average element distance is 0.01 we just set α and $\lambda \in [0.005 \ 0.05]$.)

Regarding speed, our current implementation takes about seconds to minutes to generate each result, containing number of elements in the range 500 \sim 2000. We have found this fast enough to produce results shown in the paper, even though we have not attempted any further speed optimization beyond the basic k-coherence introduced in Section 4.3.

6 Limitations and Future Work

Even though our method can produce visually plausible results, it cannot guarantee certain domain specific properties, e.g. complete obedience to physical laws like gravity or shape penetration. If such properties need to be more strongly enforced, one possibility is to add them as extra energy terms into our current framework.

Our approach synthesizes element distribution only but not the individual elements, for which we rely on user inputs. It will be interesting to devise methods that can more automatically obtain the individual elements, e.g. 2.1d textons [Ahuja and Todorovic 2007], vector primitives [Hurtut et al. 2009], or even 3D geometry.




Figure 6: Element distribution. The input exemplars are shown as smaller images, with the corresponding synthesis results shown on larger ones. Each exemplar in (a), (e), and (i) is used to produce multiple outputs with different sizes, shapes, or orientation fields. The same output model is used to produce different results in (l) and (m) via different exemplars in (i) and (j).

532 We also rely on the user input for the overall output shape. On one
 533 hand this provides the flexibility for the users to choose whatever
 534 shapes they like, but on the other hand it may be a nuisance if the
 535 users do not feel like doing so. For the latter case it would be in-
 536 teresting to apply more automatic methods to determine the output
 537 shape, e.g. what would the output shape be for a pile of potatoes?

538 We have only tried to apply our method to static but not dynamic el-
 539 ement distributions. Based on texture optimization, we believe that
 540 our basic framework can be applied for frame coherent animation
 541 effects as in [Kwatra et al. 2005; Kyriakou and Chrysanthou 2008].
 542 The really interesting issue here is on what kinds of input exam-
 543 ples to specify; dynamic inputs would be easier for our method to
 544 work with, but static inputs might be more convenient and practical
 545 to obtain.

546 References

547 AHUJA, N., AND TODOROVIC, S. 2007. Extracting texels in 2.1d natural
 548 textures. *ICCV 0*, 1–8.

549 BARAFF, D., AND WITKIN, A. 1997. Physically based modeling: Princi-
 550 ples and practice. In *SIGGRAPH '97 Course*.

551 BARLA, P., BRESLAV, S., THOLLOT, J., SILLION, F., AND MARKOSIAN,
 552 L. 2006. Stroke pattern analysis and synthesis. In *EUROGRAPH '06:
 553 EUROGRAPH 2006 papers*, vol. 25, Citeseer, 663–671.

554 BROOKS, S., AND DODGSON, N. 2002. Self-similarity based texture edit-
 555 ing. In *SIGGRAPH '02: SIGGRAPH 2002 Papers*, 653–656.

556 CHENG, M.-M., ZHANG, F.-L., MITRA, N. J., HUANG, X., AND HU,
 557 S.-M. 2010. Finding approximately repeated scene elements for image
 558 editing. In *SIGGRAPH '10*.

559 CHO, J. H., XENAKIS, A., GRONSKY, S., AND SHAH, A. 2007. Course
 560 6: Anyone can cook: inside ratatouille’s kitchen. In *SIGGRAPH 2007
 561 Courses*.

562 DISCHLER, J., MARITAUD, K., LÉVY, B., AND GHAZANFARPOUR, D.
 563 2002. Texture particles. *Computer Graphics Forum 21*, 401–410.

564 EBERT, D. S., MUSGRAVE, K. F., PEACHEY, D., PERLIN, K., AND WOR-
 565 LEY, S. 2002. *Texturing & Modeling: A Procedural Approach*. Morgan
 566 Kaufmann.

567 EFROS, A. A., AND FREEMAN, W. T. 2001. Image quilting for texture
 568 synthesis and transfer. In *SIGGRAPH '01: SIGGRAPH 2001 Papers*,
 569 341–346.




Figure 7: Discrete element texture editing. Inspired by a real-world example (a), we aim to enhance the pattern quality from the input (b). Since the original pattern is a bit boring with only little dots, the user first changes one element shape. Our method then automatically propagates that change to all other elements with similar neighborhoods. The user then goes on to edit other element properties, including both positions p and attributes q such as color, size, and shape.

- 570 EFROS, A. A., AND LEUNG, T. K. 1999. Texture synthesis by non- 618
571 parametric sampling. In *ICCV '99*, 1033. 619
- 572 GAL, R., SORKINE, O., POPA, T., SHEFFER, A., AND COHEN-OR, D. 620
573 2007. 3d collage: expressive non-realistic modeling. In *NPAA '07: 621*
574 *Proceedings of the 5th international symposium on Non-photorealistic 622*
575 *animation and rendering*, 7–14. 623
- 576 HAN, J., ZHOU, K., WEI, L.-Y., GONG, M., BAO, H., ZHANG, X., AND 624
577 GUO, B. 2006. Fast example-based surface texture synthesis via discrete 625
578 optimization. *Vis. Comput.* 22, 9, 918–925. 626
- 579 HAUSNER, A. 2001. Simulating decorative mosaics. In *SIGGRAPH '01: 627*
580 *SIGGRAPH 2001 Papers*, 573–580. 628
- 581 HURTUT, T., LANDES, P.-E., THOLLOT, J., GOUSSEAU, Y., DROUILLET- 629
582 HET, R., AND COEURJOLLY, J.-F. 2009. Appearance-guided synthesis 630
583 of element arrangements by example. In *NPAA '09*, 51–60. 631
- 584 IJIRI, T., MECH, R., IGARASHI, T., AND MILLER, G. 2008. An example- 632
585 based procedural system for element arrangement. In *EUROGRAPH 633*
586 '08: *EUROGRAPH 2008 papers*, vol. 27, 429–436. 634
- 587 JODOIN, P.-M., EPSTEIN, E., GRANGER-PICHÉ, M., AND OSTRO- 635
588 MOUKHOV, V. 2002. Hatching by example: a statistical approach. In 636
589 *NPAA '02: Proceedings of the 2nd international symposium on Non- 637*
590 *photorealistic animation and rendering*, 29–36. 638
- 591 KIM, J., AND PELLACINI, F. 2002. Jigsaw image mosaics. In *SIGGRAPH 639*
592 '02: *SIGGRAPH 2002 Papers*, 657–664. 640
- 593 KIM, T., THÜREY, N., JAMES, D., AND GROSS, M. 2008. Wavelet turbu- 641
594 lence for fluid simulation. In *SIGGRAPH '08: SIGGRAPH 2008 Papers*, 642
595 50:1–6. 643
- 596 KIM, S., MACIEJEWSKI, R., ISENBERG, T., ANDREWS, W. M., CHEN, 644
597 W., SOUSA, M. C., AND EBERT, D. S. 2009. Stippling by example. In 645
598 *NPAA '09*. 646
- 599 KOPF, J., FU, C.-W., COHEN-OR, D., DEUSSEN, O., LISCHINSKI, D., 647
600 AND WONG, T.-T. 2007. Solid texture synthesis from 2d exemplars. In 648
601 *SIGGRAPH '07: ACM SIGGRAPH 2007 papers*, 2. 649
- 602 KWATRA, V., ESSA, I., BOBICK, A., AND KWATRA, N. 2005. Texture opti- 650
603 mization for example-based synthesis. In *SIGGRAPH '05: SIGGRAPH 651*
604 *2005 Papers*, 795–802. 652
- 605 KYRIAKOU, M., AND CHRYSANTHOU, Y. 2008. Texture synthesis based 653
606 simulation of secondary agents. In *Motion in Games*, 1–10. 654
- 607 LAGAE, A., AND DUTRÉ, P. 2005. A procedural object distribution func- 655
608 tion. *ACM Trans. Graph.* 24, 4, 1442–1461. 656
- 609 LEFEBVRE, S., AND HOPPE, H. 2005. Parallel controllable texture syn- 657
610 thesis. In *SIGGRAPH '05: ACM SIGGRAPH 2005 Papers*, 777–786. 658
- 611 LERNER, A., CHRYSANTHOU, Y., AND LISCHINSKI, D. 2007. Crowds 659
612 by example. *Computer Graphics Forum* 26, 655–664. 660
- 613 LIANG, L., LIU, C., XU, Y.-Q., GUO, B., AND SHUM, H.-Y. 2001. Real- 661
614 time texture synthesis by patch-based sampling. *ACM Trans. Graph.* 20, 662
615 3, 127–150. 663
- 616 LIU, Y., WANG, J., XUE, S., TONG, X., KANG, S. B., AND GUO, B. 664
617 2009. Texture splicing. In *Pacific Graphics '09*. 665
- 618 MA, C., WEI, L.-Y., GUO, B., AND ZHOU, K. 2009. Motion field texture 666
619 synthesis. In *SIGGRAPH Asia 2009*, 110:1–8. 667
- 620 MATUSIK, W., ZWICKER, M., AND DURAND, F. 2005. Texture design 668
621 using a simplicial complex of morphable textures. In *SIGGRAPH '05: 669*
622 *ACM SIGGRAPH 2005 Papers*, ACM, New York, NY, USA, 787–794. 670
- 623 MERRELL, P., AND MANOCHA, D. 2008. Continuous model synthesis. In 671
624 *SIGGRAPH Asia '08: ACM SIGGRAPH Asia 2008 papers*, 1–7. 672
- 625 NARAIN, R., SEWALL, J., CARLSON, M., AND LIN, M. 2008. Coupling 673
626 physically-based and procedural methods for animating turbulent fluids. 674
627 In *SIGGRAPH Asia '08: ACM SIGGRAPH Asia 2008 Papers*. 675
- 628 NARAIN, R., GOLAS, A., CURTIS, S., AND LIN, M. 2009. Aggregate 676
629 dynamics for dense crowd simulation. In *SIGGRAPH Asia '09*. 677
- 630 OWADA, S., NIELSEN, F., OKABE, M., AND IGARASHI, T. 2004. Volumetric 678
631 illustration: designing 3d models with internal textures. In *SIG- 679*
632 *GRAPH '04: ACM SIGGRAPH 2004 Papers*, 322–328. 680
- 633 PEYTAVIE, A., GALIN, E., MERILLOU, S., AND GROSJEAN, J. 2009. 681
634 Procedural generation of rock piles using aperiodic tiling. In *Pacific 682*
635 *Graphics '09*. 683
- 636 RAMANARAYANAN, G., BALA, K., AND FERWERDA, J. A. 2008. Per- 684
637 ception of complex aggregates. In *SIGGRAPH '08: ACM SIGGRAPH 685*
638 *2008 papers*, 60:1–10. 686
- 639 TONG, X., ZHANG, J., LIU, L., WANG, X., GUO, B., AND SHUM, H.-Y. 687
640 2002. Synthesis of bidirectional texture functions on arbitrary surfaces. 688
641 In *SIGGRAPH '02: SIGGRAPH 2002 Papers*, 665–672. 689
- 642 TURK, G. 2001. Texture synthesis on surfaces. In *SIGGRAPH '01: SIG- 690*
643 *GRAPH 2001 Papers*, 347–354. 691
- 644 WANG, L., YU, Y., ZHOU, K., AND GUO, B. 2009. Example-based 692
645 hair geometry synthesis. In *SIGGRAPH '09: SIGGRAPH 2009 Papers*, 693
646 56:1–9. 694
- 647 WEI, L.-Y., HAN, J., ZHOU, K., BAO, H., GUO, B., AND SHUM, H.-Y. 695
648 2008. Inverse texture synthesis. In *SIGGRAPH '08: ACM SIGGRAPH 696*
649 *2008 papers*, 1–9. 697
- 650 WEI, L.-Y., LEFEBVRE, S., KWATRA, V., AND TURK, G. 2009. State of 698
651 the art in example-based texture synthesis. In *Eurographics '09 State of 699*
652 *the Art Report*, 93–117. 700
- 653 WU, Q., AND YU, Y. 2004. Feature matching and deformation for texture 701
654 synthesis. In *SIGGRAPH '04: ACM SIGGRAPH 2004 Papers*, 364–367. 702
- 655 ZHANG, J., ZHOU, K., VELHO, L., GUO, B., AND SHUM, H.-Y. 2003. 703
656 Synthesis of progressively-variant textures on arbitrary surfaces. In *SIG- 704*
657 *GRAPH '03: ACM SIGGRAPH 2003 Papers*, 295–302. 705
- 658 ZHOU, K., HUANG, X., WANG, X., TONG, Y., DESBRUN, M., GUO, B., 706
659 AND SHUM, H.-Y. 2006. Mesh quilting for geometric texture synthesis. 707
660 In *SIGGRAPH '06: ACM SIGGRAPH 2006 Papers*, 690–697. 708
- 661 ZHOU, H., SUN, J., TURK, G., AND REHG, J. M. 2007. Terrain synthesis 709
662 from digital elevation models. *IEEE Transactions on Visualization and 710*
663 *Computer Graphics* 13, 4, 834–848. 711