

Efficient Synthesis of Gradient Solid Textures

Guo-Xin Zhang¹, Yu-Kun Lai² and Shi-Min Hu¹

Technical Report **TR-110818**,
Tsinghua University, Beijing, China

¹ Department of Computer Science and Technology, Tsinghua University, China. Email: zgx.net@gmail.com; shimin@tsinghua.edu.cn

² School of Computer Science and Informatics, Cardiff University, UK. Email: Yukun.Lai@cs.cardiff.ac.uk

Efficient Synthesis of Gradient Solid Textures

Guo-Xin Zhang[†], Yu-Kun Lai[‡] and Shi-Min Hu[†]

Abstract—Solid textures require large storage and are computationally expensive to synthesize. In this paper, we propose a novel solid representation called *gradient solids* to compactly represent solid textures, including a tricubic interpolation scheme of colors and gradients for smooth variation and a region-based approach for representing sharp boundaries. We further propose a novel approach based on this to *directly* synthesize gradient solid textures from exemplars. Compared to existing methods, our approach avoids the expensive step of synthesizing the complete solid textures at voxel level and produces optimized solid textures using our representation. This avoids significant amount of unnecessary computation and storage involved in the voxel-level synthesis while producing solid textures with comparable quality to the state of the art. The algorithm is much faster than existing approaches for solid texture synthesis and makes it feasible to synthesize high-resolution solid textures in full. Our compact representation also supports efficient novel applications such as instant editing propagation on full solids.

Index Terms—solid textures, synthesis, vector representation, gradient



1 INTRODUCTION

Textures are essentially important for current rendering techniques as they bring in richness without involving overly complicated geometry. Most previous work on texture synthesis focuses on synthesizing 2D textures, which require texture mapping with almost unavoidable distortions when they are applied to 3D objects. Solid textures represent color (or other attributes) over 3D space, providing an alternative approach to 2D textures that avoids complicated texture mapping and allows real solid objects to be represented with consistent textures both on the surface and in the interiors alike.

Due to the extra dimension, solid textures represented as attributes sampled at regular 3D voxel grids are extremely expensive to synthesize and store. To provide sufficient resolution in practice, a typical solution is to synthesize only a small cube (e.g. 128^3), and tile the cube to cover the 3D space. However, tiling may cause visual repetition (see Fig. 8). While repetitions could be alleviated with some rotations, they cannot be eliminated completely when the volumes are sliced with certain planes. Further it is possible only when the solid textures have no interaction with the underlying objects, and thus cannot respect any model features or user design intentions. To address this, previous approaches [1], [2] synthesize solid textures on demand; however, handling high-resolution solid textures is still expensive in both computation and storage.

Inspired by image vectorization, for pixels (or voxels) with dominantly smooth color variations (within each homogeneous region), vectorized graphics provide significant advantages such as being compact, resolution independent and easy-to-edit. The possibility and effectiveness of vectorizing solid textures have been recently studied in [3]. This work is essentially a 3D generalization of image vectorization, which requires

voxel-level (raster) solid textures as input and inherits similar advantages over traditional raster solid textures. It remains computationally costly and involves large intermediate storage for raster solid textures to synthesize high resolution solid textures with a nonhomogeneous spatial distribution (e.g. [2]).

In this paper, instead of first synthesizing the full voxel solid textures before vectorizing them [3], we propose a novel approach to directly synthesize vectorized solid textures from exemplars. Inspired by gradient meshes in image vectorization [4], we propose a novel *gradient solid* representation that uses a tricubic interpolation scheme for smooth color variations within a region, and a region-based approach to represent sharp boundaries with separated colors. This representation is compact, more regular than Radial Basis Functions (RBFs) [3] and thus particularly suitable for real-time rendering and efficient solid texture synthesis. Our approach can be used to vectorize input solids, which is over 100 times faster than [3] and leads to reduced approximation errors in most practical cases.

We further treat solid texture synthesis as an *optimization* process of *control points* of gradient solids to produce synthesized solids with similar sectional images as given exemplars. Compared with traditional solid texture synthesis, we have *far less* control points than voxels, leading to a much more efficient algorithm. While we solve both bitmap solid synthesis and solid vectorization together and produce solid textures with comparable quality as the state of the art, it is over 10 times faster than existing synthesis methods.

The main contributions of this paper are:

- A new *gradient solid* representation with regular structure that is compact, resolution-independent and capable of representing smooth solids and solids with separable regions.
- A novel optimization-based algorithm for *direct* synthesis of high quality solid textures vectorizing high resolution solids which is efficient both in computation and storage.

[†] Department of Computer Science and Technology, Tsinghua University, China. Email: zgx.net@gmail.com; shimin@tsinghua.edu.cn

[‡] School of Computer Science and Informatics, Cardiff University, UK. Email: Yukun.Lai@cs.cardiff.ac.uk

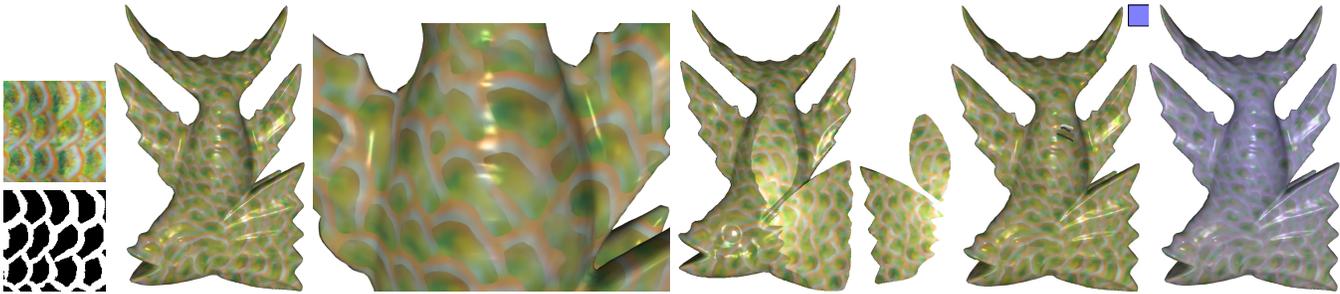


Fig. 1. High-resolution gradient solid texture synthesis and editing. From left to right: the input exemplar, the synthesized gradient solid texture following a given directional field, a closeup, internal slices and instant editing (user interaction and the output).

- Our representation also facilitates novel applications such as *instant solid editing*, as demonstrated in the paper.

To the best of our knowledge, this is the first algorithm that synthesizes vector solid textures directly from exemplars, allowing high resolution, potentially spatially nonhomogeneous solid textures to be synthesized *in full*. Various techniques have also been developed to effectively improve the quality or reduce the computational cost.

2 RELATED WORK

Our work is closely related to example based texture synthesis and vector images/textures.

Solid Texture Synthesis Texture synthesis has been an active research direction in computer graphics for many years. Please refer to [5] for a comprehensive survey of example-based 2D texture synthesis and [6] for a recent survey of solid texture synthesis from 2D exemplars.

Early work on solid texture synthesis focuses on procedural approaches [7], [8]. Since rules are used to generate solid textures, very little storage is needed. Procedural solid textures can be generated in real-time [9]. However, only restricted classes of textures can be effectively synthesized and it is inconvenient to tune the parameters. Exemplar-based approaches do not suffer from these problems, and thus received more attention. 2D exemplar images are popular due to their wide availability. Wei [10] extends non-parametric 2D texture synthesis algorithms to synthesize solid textures. An improved algorithm is proposed in [11] to generate solid textures based on texture optimization [12] and histogram matching [13]. Further extended work [14] considers k -coherent search and combined position and index histograms to improve the results. To synthesize high resolution solid textures, Dong et al. [1] propose an efficient synthesis-on-demand algorithm based on deterministic synthesis of certain windows from the whole space [15] necessary for rendering, based on the fact that only 2D slices are needed at a time for normal displays. This work is extended in [2] that introduces user-provided tensor fields as guidance for solid texture synthesis. This approach allows synthesizing solid textures with nonhomogeneous spatial distributions, thus cannot be achieved by tiling small fixed cubes.

Alternative approaches for solid texture synthesis exist. Jagnow et al. [16], [17] propose an algorithm based on stereological analysis which provides more precise modeling of solid textures. However, their approach only works for restricted types of solid textures with well separable pieces. Lapped textures have been extended to synthesize 3D volumetric textures [18]. 3D volumetric exemplars instead of 2D image exemplars are needed as input. Solid texture synthesis has also been used for other applications. Ma et al. [19] use similar techniques for motion synthesis.

Unlike previous methods, our approach directly synthesizes *gradient solid textures* from 2D exemplars. This provides the benefits from both procedural and exemplar-based approaches: the representation is more compact and high resolution solid textures can be synthesized *in full* efficiently. The algorithm is flexible to synthesize various solid textures using 2D exemplars and follow given tensor fields if specified by the user. The whole solid textures need only to be synthesized once which reduces overall computation.

Vector Images and Vector Solid Textures Different from raster images, vector graphics use geometric primitives along with attributes such as colors and their gradients to represent the images. Due to the advantages of vector graphics, plenty of work recently focuses on generating vector representations from raster images. Recent work proposes automatic or semi-automatic approaches to high-quality image vectorization using quadrilateral gradient meshes [4], [20] or curvilinear triangle meshes for better feature alignment [21]. Diffusion curves [22] model vector images as a collection of color diffusion around curves. Some works consider combining raster images with extra geometric primitives [23], [24], [25] to obtain benefits such as improved editing and resizing.

Vector graphics have recently been generalized to solid textures [3], [26]. Compared to raster solids, vector solids have the advantages of compact storage and efficient rendering. Wang et al. [3] propose an automatic approach to vectorize given solid textures using a RBF-based representation. However, this approach relies on raster solids as input, thus an expensive raster solid texture synthesis algorithm [11] needs to be performed first if only 2D exemplars are given as input. Diffusion surfaces [26], a generalization from diffusion curves [22], was used to represent vector solids; their focus

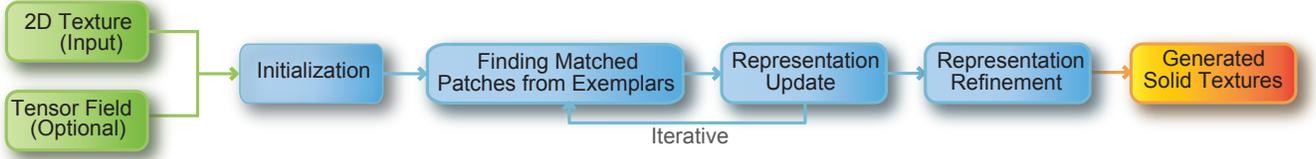


Fig. 2. Algorithm pipeline of gradient solid texture synthesis.

however is user design of solids rather than automatic generation.

Vector representation is loosely related to volume compression techniques (e.g. [27], [28]) as both consider more compact representations than raster solids. The focus of vector representation however aims at creating compact and resolution-independent representation suitable for graphics applications that produce visually similar and pleasing results even when magnified while the purpose of volume compression techniques is to reconstruct large volumes as close as possible to the original even under significant compression. Research work on volume compression tends to use blocks and block-based coding which leads to less smooth reconstruction.

We propose a novel algorithm that synthesizes gradient solids directly from 2D exemplars, bypassing intermediate bitmap solid synthesis and subsequent bitmap-to-vector conversion, leading to an efficient algorithm in both computation and storage that produces high quality solid textures. The representation although with a somewhat different aim may be useful for certain volume compression applications.

3 GRADIENT SOLID REPRESENTATION

WE give details of the gradient solid representation, allowing efficient representation of smooth regions and regions with boundaries.

3.1 Representing Smooth Regions

We first consider representing regions with smoothly varying colors. We use a $n \times n \times n$ grid of control points with axes u, v, w to represent the solid textures. At each control point (i, j, k) , we store a feature vector \mathbf{f} including r, g, b color components and additional feature channels such as the signed distance [29]. In addition, the gradients of \mathbf{f} , i.e. $\frac{d\mathbf{f}}{du}, \frac{d\mathbf{f}}{dv}, \frac{d\mathbf{f}}{dw}$ are also stored allowing flexible control of variations in 3D space. 3D tricubic interpolation with gradients [30], [31] is used to obtain the feature vector $\tilde{\mathbf{f}}$ for any voxel inside the grid. Assume that $p = 1, 2, \dots, 8$ represents the 8 control points in the cube that covers the voxel and assume second or higher order derivatives of \mathbf{f} to be zero, $\tilde{\mathbf{f}}$ at parameter (u, v, w) ($0 \leq u, v, w \leq 1$) can be evaluated as

$$\tilde{\mathbf{f}}(u, v, w) = \sum_{i,j,k=0}^3 \mathbf{a}_{ijk} u^i v^j w^k. \quad (1)$$

All the 64 coefficient vectors \mathbf{a}_{ijk} are weighted sums of 32-dimensional vectors $V = (\dots \mathbf{f}^{(p)}, \frac{d\mathbf{f}^{(p)}}{du}, \frac{d\mathbf{f}^{(p)}}{dv}, \frac{d\mathbf{f}^{(p)}}{dw} \dots)$. Using the integer weights given in [31], C^1 continuity over the whole volume is guaranteed.

The geometric positions of control points in our representation are fixed. Assuming the displacement between adjacent control points is d , the geometric position of the control point (i, j, k) is (id, jd, kd) . The displacement determines the number of voxels located within each cube of the control grid. Larger d leads to more compression while smaller d implies better capture of details. In all of our experiments we use $d = 4$ which means that the number of control points is roughly $\frac{1}{64} = 1.56\%$ of voxels.

This simple representation has several significant advantages. For any fixed point with known parameter (u, v, w) , since $u^i v^j w^k$ can be pre-computed, the expensive evaluation of Eqn. 1 can be reduced to a weighted sum of elements in V . In practice, we pre-compute these coefficients for a regular grid with 33^3 samples in each cube, with interval at $\frac{1}{8}$ voxel for accuracy. A fixed look-up table irrelevant to the input is pre-computed and stored, with $33^3 \times 32$ entries (about 4.4MB), and the interpolated feature at any space position can be computed as a linear combination of V with these prebuilt weights.

The interpolation is achieved in rendering via GPU acceleration, as detailed in Sec. 5.3. This allows efficient evaluation, particularly important as solid textures are computationally intensive. The look-up table does not need to be stored and is calculated on the fly. It is of fixed size even for very large volumes (equivalent to e.g. 512^3 or 1024^3) and in such cases becomes negligible. Compared with the RBF-based representation [3], we have regular structures suitable for texture synthesis. As demonstrated in Figs. 10 and 11, our local interpolation representation has much better color reproduction. There is no need to store the positions of control points, which further saves storage. The regularity also helps efficient direct solid texture synthesis and supports other applications such as instant editing propagation, as detailed later.

3.2 Representing Region Boundaries

If the given texture only contains gradual change of colors, the representation described in Sec. 3.1 is sufficient (e.g. Fig. 10). If the texture contains sharp boundaries that need to be preserved, a feature mask image is often used in texture synthesis as an additional component (other than color) to

better preserve structures. Similar to previous work both in 2D and 3D textures [29], [3], we assume regions can be separated using a binary mask. To represent the boundary in the solid textures, we also use a signed distance field stored at the *same* regular $n \times n \times n$ grid. We store both the signed distance \tilde{d} and its gradients $\frac{d\tilde{d}}{du}$, $\frac{d\tilde{d}}{dv}$ and $\frac{d\tilde{d}}{dw}$ and use the same tricubic interpolation as in Sec. 3.1 to calculate the interpolated signed distance \tilde{d} at each voxel. The sign of \tilde{d} indicates which side of the regions in the binary mask this voxel belongs to. Different from [3], gradients are stored in addition to the distance, and thus we process the distance field consistently with colors and represent region boundaries with flexibility. For each control point that is adjacent to at least one cube with both positive and negative distances, two feature vectors \mathbf{f}^P (positive distance) and \mathbf{f}^N (negative distance) and their gradients are stored. Any voxel with positive (or negative) distance will be evaluated using the same interpolation in Sec. 3.1 but with \mathbf{f}^P (or \mathbf{f}^N) and their gradients instead. This guarantees C^1 smoothness within each region while also allowing sharp boundaries to be produced between regions. Our gradient solid representation is easy to evaluate but also sufficient to represent various solid textures, as demonstrated in Sec. 5.

4 GRADIENT SOLID TEXTURE SYNTHESIS

Our algorithm synthesizes gradient solid textures directly from 2D exemplars, which may include optional binary masks (if sharp boundaries exist between regions). In addition, a smooth tensor field may be given to specify the local coordinate systems the exemplar images align with [2]. We use an optimization based approach to synthesize gradient solid textures, with local patches aligned to the field if given. The algorithm pipeline is summarized in Fig. 2, which involves several key steps: initialization, iterative optimization consisting of similarity patch search and gradient representation update, and the final gradient solid refinement. We will also discuss techniques to ensure efficiency in both computation and storage. If the binary mask is given, we pre-compute a signed distance field for the image with the absolute value at each pixel being the distance to the region boundary and different signs (positive or negative) for different regions. This signed distance is considered as an extra component in the feature vector \mathbf{f} [29].

4.1 Initialization

We simply start from a randomized initialization. For each control point, we randomly select a pixel from the exemplar image, and assign the feature vector at the pixel to the control point. All the gradients are initialized to zero.

4.2 Optimization-based Synthesis

Optimization is the key step in our gradient solid texture synthesis pipeline. It involves iterations of two alternating steps, namely choosing optimal patches from exemplars that

best match the current representation and updating the representation to better approximate the exemplar patches. Unlike traditional texture optimization [12], [11], we optimize the feature vectors in the control points of the gradient solids, a much more compact representation than voxels. New challenges exist due to the different nature of the representation which we will address with various technical solutions. We apply N_O iterations for each synthesis level, and use a modified coarse-to-fine strategy detailed in Sec. 4.2.3. $N_O = 3$ is sufficient and used for all the experiments in the paper.

4.2.1 Finding matched patches from exemplars

We first identify those local patches from the exemplars that best match the current gradient solid. These patches will then be used to improve the representation. Since gradient solids have much sparser control points than voxels, we randomly choose a small number N_C of check points within each cube of the grid ($N_C = 3$ provides a good balance and is used for all the examples in the paper). At each check point, we sample three orthogonal planes each with $N \times N$ samples (denoted as \mathbf{s}_x , \mathbf{s}_y and \mathbf{s}_z respectively) which are evaluated based on our representation (as illustrated in Fig. 3). A fast approximate evaluation is used in intermediate synthesis to significantly improve the performance without visually degrading the quality (see Sec. 4.2.4).

We then find three local patches from exemplars that best match these sampled patches. If all the three slices are equally important, we use three independent searches as [11]. Many practical solid textures are anisotropic and it is not possible to keep all three slices well matched with a single exemplar image. In such cases, it is known that matching two slices instead of three may lead to better results [11]. We propose a new approach that takes crossbar consistency into account, which works best when two slices are matched. Crossbars are those voxels shared by two or three slices (see Fig. 3) and inconsistent crossbars may result from independent best searches. For computational efficiency, we first search for the patch E_x from exemplars that best matches \mathbf{s}_x , as usual. We then search for the patch E_y that best matches \mathbf{s}_y from a set of N_1 candidates with the most consistent crossbar voxels as E_x . If three slices are matched, we similarly search for the best match E_z of \mathbf{s}_z from a set of N_2 candidates with the most consistent crossbars as E_x and E_y . $N_1 = 20$ and $N_2 = 50$ are used for all the experiments in the paper. This leads to improved synthesis results with better structure preservation, which shows the importance of crossbar consistency, as demonstrated in Fig. 5. While crossbar matching has been used in correction-based synthesis [1], using this in optimization based synthesis is new.

To speed up the computation, a PCA projection of the matching vectors is used [32], which effectively reduces dimensions from hundreds to 10-20 while keeping most of the energy. After this, the searches can be effectively accelerated with ANN approximate nearest neighbor library.

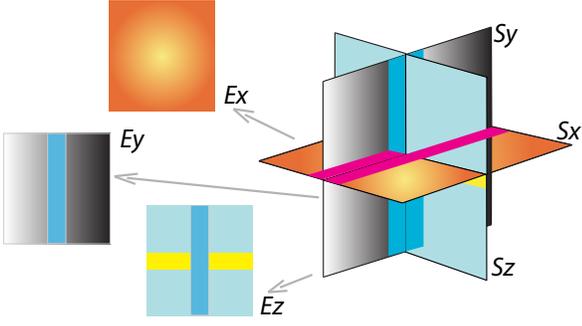


Fig. 3. Illustration of crossbars.

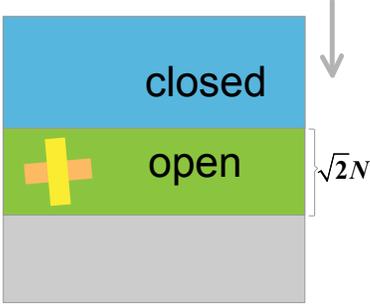


Fig. 4. Illustration of bucket reuse.

4.2.2 Representation update

Each matched patch at every check point gives $N \times N$ samples, which will be used to update the gradient solid representation. To efficiently collect samples, we conceptually build a bucket for each voxel in the grid that holds all the samples located in the voxel. After considering check points in all the cubes, each bucket may end up with none or a few samples. For buckets with more than one samples, in order to determine the feature vector, simply averaging all the samples in the bucket tends to produce blurred voxels. Previous methods [33], [11] use mean shift clustering to avoid blurring, which is expensive as all the samples in the buckets need to be preserved and clustering algorithms need to be performed many times. We propose two novel approaches to significantly improve the efficiency.

Quantization. To avoid blurring without storing all the samples in each bucket, we propose a novel approach based on vector quantization. We preprocess the given exemplar to quantize colors of all the pixels into N_T clusters. A small number of N_T (e.g. 12) is sufficient for practical textures. For the texture with a binary mask, we start from two clusters for both positive and negative regions, and iteratively allocate the new cluster to regions with most significant average quantization error, until all the N_T clusters are allocated. We use a two-pass approach in the synthesis. In the first pass, for every bucket, only the number of samples belonging to each cluster is recorded. In the second pass, we compute the average feature vector only for those samples belonging to the two *dominant* clusters (with maximum counts in the first pass). Since the dominant clusters are known before the second pass, whenever a sample is generated we test if it should be included

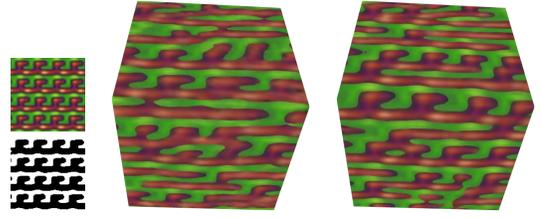


Fig. 5. Results without (left) and with (right) crossbar matching.

for averaging. Only the sum and the number of samples need to be kept which significantly saves the storage. This avoids using the computationally expensive clustering algorithm for each voxel but also significantly reduces blurring, as demonstrated in various results in Sec. 5. Using quantization in the finest level is sufficient, according to our experience.

Bucket reuse. Although conceptually the number of buckets is the same as the number of voxels, i.e. $O(n^3)$, we can significantly reduce the memory requirement by bucket reuse. We update our representation in the 3D scanline order of control points. Depending on the template size N , check points more than $\frac{\sqrt{2}}{2}N$ voxels away will not produce any sample in the current bucket, where $\frac{N}{2}$ is half the template size and $\sqrt{2}$ is introduced due to rotation. As illustrated in Fig. 4 for a 2D illustration, we keep track of two references in the dominant dimension (one of the three dimensions that can be chosen arbitrarily) that mark the boundaries of the open region (from which new check points will be generated) and the closed region (where no more samples will be produced and we can safely update the gradient solid representation). In case the two-pass algorithm in quantization is used, this buffer needs to be doubled i.e. up to $2\sqrt{2}N$ span in the dominant dimension is sufficient, or the memory cost is $O(n^2N)$. This is because either pass has an affected region as we discussed and the second pass relies on the results collected from the first pass. The required buffering space does not increase with more synthesis iterations as buckets are cleared after each synthesis iteration and no further propagation as in [1] happens. Since $N \ll n$ and often constant for various examples, this effectively saves the storage by reducing the complexity from n^3 to n^2 , without any extra recomputation. This is possible, because after each iteration of optimization, only a very compact gradient solid representation is kept, while traditional solid texture synthesis requires the whole dense volume to be accessible. By using this technique, we can synthesize gradient solid textures corresponding to 1024^3 voxels within $2GB$ memory, even less than storing the voxels alone.

After obtaining the average feature vector for any bucket with at least one sample, we assign each non-empty bucket to the closest control point. The feature vector as well as gradients of the control point are updated by minimizing the fitting error in the least-squares sense. For a particular control point, assume s buckets are related with relative coordinates du_t, dv_t, dw_t and feature vector \mathbf{f}_t ($1 \leq t \leq s$), we find $\mathbf{f}_c, \frac{\mathbf{f}_c}{du}, \frac{\mathbf{f}_c}{dv}, \frac{\mathbf{f}_c}{dw}$ that

minimizes

$$E_C = \sum_{t=1}^s \left\| \mathbf{f}_c + \frac{\mathbf{f}_c}{du} du_t + \frac{\mathbf{f}_c}{dv} dv_t + \frac{\mathbf{f}_c}{dw} dw_t - \mathbf{f}_t \right\|^2. \quad (2)$$

This can be considered as a local first-order Taylor expansion of our representation which can be efficiently solved by small linear systems. This approximation is sufficient for intermediate computation and we use the accurate evaluation only in the final stage.

4.2.3 Multi-resolution synthesis

To capture features at multiple scales, a multi-resolution approach is also used in our algorithm. However, since a sparse control grid is used, reducing the resolution is not feasible as it would be too coarse in low resolutions to effectively capture details. Instead, inspired by fractional sampling [15], in each successive coarser level we keep the resolution of the control grid *unchanged* and double the spacing between sample pixels in exemplar image and voxels in the 3D space. From coarse to fine we use three levels of synthesis with $N = 9, 11, 21$ respectively. The finest level uses a significantly larger neighborhood in order to cover a few control points at minimum in our sparse representation.

4.2.4 Fast approximate evaluation

Our gradient solid representation is relatively easy to evaluate; however, in the solid texture synthesis process, many evaluations are needed. We suggest two approximations for improved performance. In the intermediate synthesis process, instead of evaluating the accurate values at each sample point, we use a first-order Taylor expansion as an approximation. For any point p whose closest control point is c with feature vector \mathbf{f}_c and its gradients $\frac{\mathbf{f}_c}{du}, \frac{\mathbf{f}_c}{dv}, \frac{\mathbf{f}_c}{dw}$, the approximate feature vector at p with relative coordinates du_p, dv_p, dw_p , can be evaluated as $\mathbf{f}_p = \mathbf{f}_c + \frac{\mathbf{f}_c}{du} du_p + \frac{\mathbf{f}_c}{dv} dv_p + \frac{\mathbf{f}_c}{dw} dw_p$. This approximation does not ensure smoothness, but only involves 3 multiplications and 3 additions for each component of the feature vector, thus only takes about 1/10 of the computation of a full evaluation. In the iterative synthesis, another approximation is to ignore the region-based calculation given in Sec. 3.2 (as if there is no separated region as in the single channel case such as Fig. 6(middle)). This may mix up voxels in different regions within *the same cube* and leads to visual degradation of the final results; the impact on the intermediate synthesis however is negligible as it is restricted to a couple of voxels due to the cube size.

4.3 Gradient solid representation refinement

As the final step, we further optimize the gradient solid representation to better represent the synthesized gradient solids.

Region separation. For solids with smooth variation of colors (e.g. Fig. 10), our algorithm does not require a binary mask

as input and can effectively reproduce the solids with a single region. For solids with sharp region boundaries that need to be preserved, we differentiate regions with positive and negative signed distances for the computation of control point parameters described in Sec. 4.2.2. For each control point, we compute positive parameters (\mathbf{f}^P and gradients) using samples with positive signed distance. Similarly, samples with negative signed distance contribute to negative parameters (\mathbf{f}^N and gradients). To improve reliability in the fitting of boundary control points, we propagate boundary samples (with neighboring samples having different signs of distance) to the near neighboring space, similar to dilation in mathematical morphology. This mainly ensures cubes near region boundaries have sufficient samples to make fitting reliable.

Control point optimization. To further improve the quality, instead of fitting with first order approximation, we minimize the fitting error of all the samples between the sample values and those interpolated using Eqn. 1. For a sample point with sampled feature vector $\hat{\mathbf{f}}_i$ located in the cube c_i with corner control points collected as V_i and parameter (u_i, v_i, w_i) , the evaluated feature vectors $\tilde{\mathbf{f}}$ are linear functions of V_i , denoted as $\mathbf{f}(V_i; u_i, v_i, w_i)$. We minimize the following quadratic energy

$$\bar{E}_C = \sum_{i=1}^{\#s\text{amp}} \|\tilde{\mathbf{f}}_i - \hat{\mathbf{f}}_i\|^2 = \sum_{i=1}^{\#s\text{amp}} \|\mathbf{f}(V_i; u_i, v_i, w_i) - \hat{\mathbf{f}}_i\|^2, \quad (3)$$

where $\#s\text{amp}$ is the number of sample points. Minimization of \bar{E}_C leads to a sparse linear system. As we have a good estimation from the previous approximation, the linear system can be effectively solved by a few iterations of gradient descent. As demonstrated in Table 1, control point optimization reduces the approximation error but also takes some extra time. Our method without this optimization is sufficiently good in many cases so it is considered as an option to tradeoff quality with speed.

4.4 Instant solid editing

3D solids are expensive to store, and also time-consuming to edit. Thanks to our compact representation, we can achieve instant solid editing, by adapting a recent development [34] in images and videos. A typical scenario for this editing is the user first draws a few strokes with different intensities indicating how strong the selected voxels will be affected by the editing. The user then selects a reference color, and voxels will be affected based on similarities in the position and the appearance (color) to those with user specifications. While the editing in [34] is generally efficient, dealing with large volumes is still relatively slow. Worse still, if the volume is in some vector representation, naive application of this method will involve converting to raster representation before editing and back to vector representation afterwards. We show as follows that our adaptation of the editing algorithm is instant with virtually *equivalent* solution; this cannot be achieved with Wang’s representation.

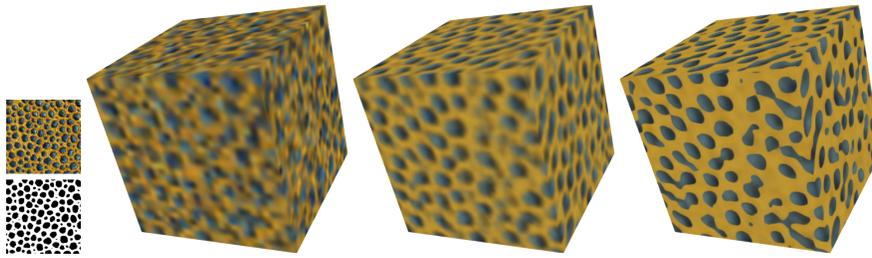


Fig. 6. Comparison of results using direct upscaling (left) and our algorithm without (middle) and with (right) region separation.

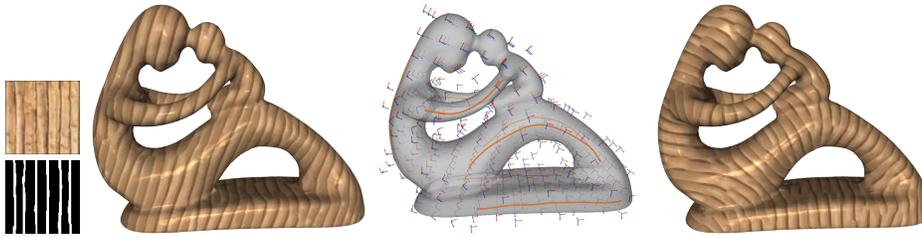


Fig. 7. Synthesized object without (a) and with the field (c); the field is given in (b).

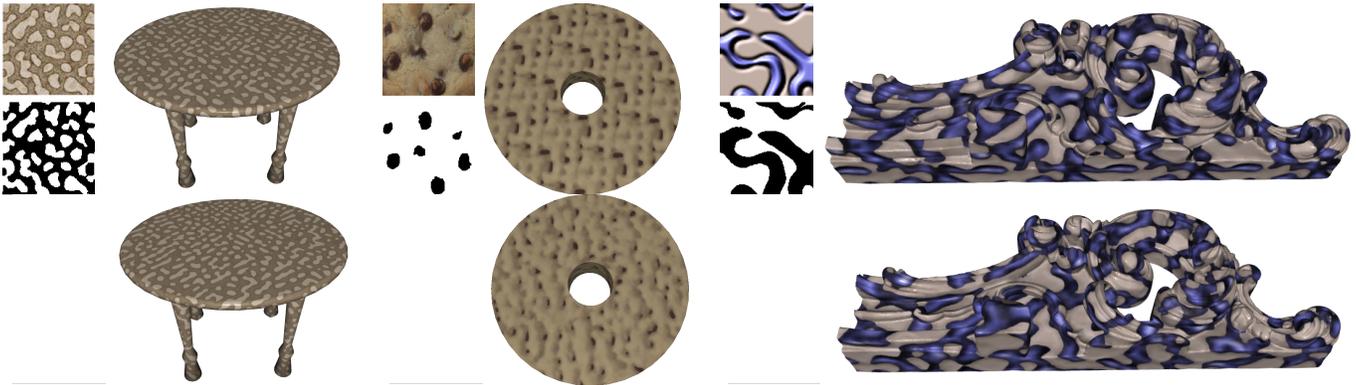


Fig. 8. Synthesized solids without fields. First row: tiled low resolution (128^3) solids. Second row: high resolution (512^3) solids.

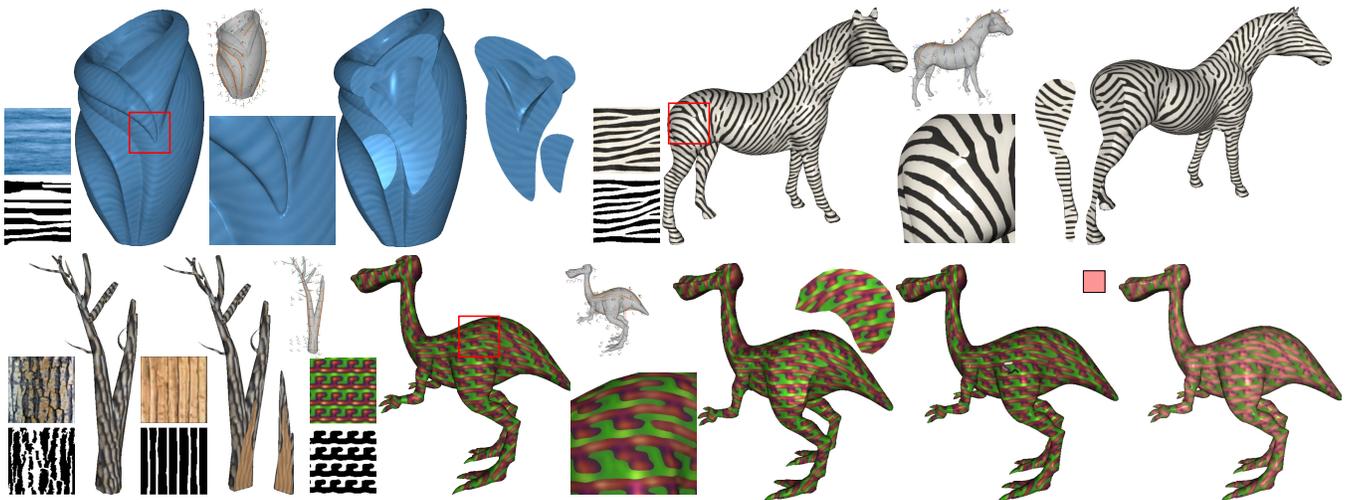


Fig. 9. Synthesized high-resolution (512 samples in the longest dimension) solids following given directional fields with our algorithm: 'vase', 'horse', 'tree' and 'dinopet' with synthesized solids, close-ups and internal slices. 'Dinopet' is turned pink with instant editing.

For each control point i with color $\mathbf{c}_i = (r, g, b)^T$ and position $\mathbf{p}_i = (x, y, z)^T$, we need to know the influence h_i . This is effectively modeled as m RBFs, the centers of which are randomly selected from the stroke voxels

$$h_i = \sum_{k=1}^m \omega_k h_{i,k} = \sum_{k=1}^m \omega_k \exp \left\{ -\alpha(\beta|\mathbf{p}_i - \bar{\mathbf{p}}_k|^2 + |\mathbf{c}_i - \bar{\mathbf{c}}_k|^2) \right\} \quad (4)$$

where $\bar{\mathbf{p}}_k$ and $\bar{\mathbf{c}}_k$ are the position and color of k -th stroke voxel selected as a RBF center. ω_k , restricted to be non-negative, can be obtained by solving a linear programming problem that minimizes the strength deviation for user specified voxels [34]. Parameters α and β control the propagation and $\alpha = 10^{-4}$, $\beta = 0.1$ work well in many cases. Assuming the reference color is \mathbf{c}_{ref} , to compute the edited gradient solids, \mathbf{c}_i and $\frac{d\mathbf{c}_i}{d\mathbf{p}_i}$ need to be updated for each control point, which can be effectively calculated as follows. We define $\mathbf{c}'_i = (1 - h_i)\mathbf{c}_i + h_i \cdot \mathbf{c}_{ref}$, and thus we have

$$\frac{d\mathbf{c}'_i}{d\mathbf{p}_i} = (1 - h_i) \frac{d\mathbf{c}_i}{d\mathbf{p}_i} - (\mathbf{c}_i - \mathbf{c}_{ref}) \left(\frac{dh_i}{d\mathbf{p}_i} \right)^T, \quad (5)$$

where

$$\frac{dh_i}{d\mathbf{p}_i} = - \sum_{k=1}^m \omega_k 2\alpha h_{i,k} \left\{ \beta(\mathbf{p}_i - \bar{\mathbf{p}}_k) + \left(\frac{d\mathbf{c}_i}{d\mathbf{p}_i} \right)^T (\mathbf{c}_i - \bar{\mathbf{c}}_k) \right\}. \quad (6)$$

The editing is demonstrated in Fig. 1 to turn a fish purple instantly. A few strokes on the fish object are drawn to indicate the effect of change and a purplish color sample is chosen (as in the box). Another example is in Fig. 9 where ‘dinopet’ is turned pink instantly with a few strokes and a pink reference color (as in the box). The editing process takes only 0.1 seconds in total with 0.01 seconds for linear programming and the remaining time for gradient solid update, providing instant feedback on such large volumes. Comparatively, direct application on raster solid of equivalent resolution takes about 1 second and naive implementation on vector solids takes a few minutes.

5 RESULTS AND DISCUSSIONS

Our algorithm is useful for either direct synthesis of solid textures, or vectorizing input solids. We carried out our experiments on a computer with $2 \times 2.26\text{GHz}$ CPU and NVIDIA GTS 450 GPU.

5.1 Solid texture synthesis

Our algorithm directly synthesizes more compact and resolution-independent gradient solid textures from 2D exemplars. Solids with comparable quality to the state of the art can be synthesized, as shown in Figs. 1, 4-8, 12. As for other CPU-based algorithms that focus on synthesizing full solids of a 128^3 cube, the typical reported times have been tens of minutes, e.g. [11] uses 10-90 minutes (without tensor fields) and [19] (a CPU-based implementation similar to [1] with direction fields considered) reported about 30

minutes with a single core. Our results are vector solid textures which are resolution-independent. For simplicity, we consider solid textures with equivalent detail resolution to raster solid textures when certain resolution is referred to in the following discussion. Our current implementation, after about 10 seconds preprocessing of the input exemplar (which is the same for arbitrarily sized output volumes), takes only 13 seconds. Even counting the different performance of CPUs, our algorithm is over 10 times faster. Due to the compactness in representation and the technique for memory reuse, we can synthesize high-resolution solid textures in full. 512^3 solids can be synthesized within 15 minutes (Fig. 8). Other examples throughout the paper with 512 samples in the longest dimension take 3-7 minutes, while the example in Fig. 12 at resolution of 1024 takes 35 minutes and within 900MB memory. Region separation is not needed if the input texture does not contain sharp boundaries, as the ‘vase’ and ‘tree’ examples in Fig. 9. In these examples, the binary mask is used only as part of the feature vector, not for region separation. The ‘tree’ example shows that our synthesis algorithm can be generalized to synthesize solids with different exemplars covering different spaces, mimicking the real structure of a tree.

We demonstrate the effectiveness of our algorithm with various examples. Although our method uses a rather sparse set of control points, they are much more expressive than voxels at the same resolution. An example is given in Fig. 6. The left result is synthesized with [2] (using a proportionally downsized exemplar image as input) and looks sensible at original 32^3 resolution. We use tricubic interpolation to upscale the volume to 128^3 and clear artifacts appear indicating that 32^3 volume is not sufficient to capture the structure of the solid. Our results with also 32^3 control points are significantly better and sharp region boundaries can be recovered with region separation. Tiling small cubes such as of 128^3 size to cover the whole space is commonly used, due to the prohibitively expensive computation with most previous algorithms. Synthesizing high resolution solids is essential to avoid visual repetition (as demonstrated by ‘table’, ‘cake’ and ‘statue’ in Fig. 8) or produce solids following certain direction fields (see Fig. 9). A comparison of results without or with the field is given in Fig. 7. High resolution solid textures with 512 and 1024 samples in the longest dimension are shown in Figs. 1 and 12, respectively. Note that in all the results we synthesize the full solids rather than only the visible voxels [1], [2]. This is preferred since in many applications objects are synthesized once but rendered many times on lower-end systems. Our representation makes rendering algorithm both efficient and simple to implement.

5.2 Vectorization of solid textures

Our approach can also be used for solid texture vectorization. In this application, we take each voxel as a sample and produce gradient solids with the method in Sec. 4.3 if optimization is used or otherwise a first order approximation as in Sec. 4.2.4. We perform comparative experiments on the same computer, using the code directly from [3]. 5,000 RBFs are used to

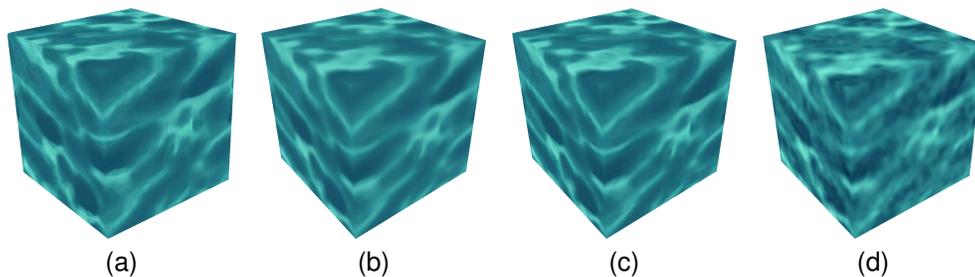


Fig. 10. Solid vectorization of the input volume ‘caustic’ without binary mask. (a) input volume; (b)(c) our results without and with further optimization; (d) result using [3].

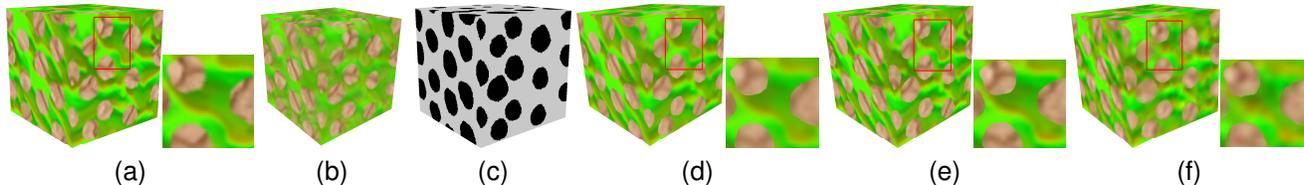


Fig. 11. Solid vectorization results with a binary mask. (a) input volume ‘balls’; (b) input volume rendered in transparency; (c) input mask; (d) vectorized solid with our algorithm without optimization; (e) our result with optimization; (f) result using [3].

provide sufficient flexibility, more than most examples in [3], for fair comparison. Although our algorithm is highly parallel, we only use a single core for fair comparison. Detailed running times and fitting errors are given in Table 1. Whilst pixel-wise error measurement before and after vectorization may not be the best criterion perceptually, it is widely used in image vectorization. For most solids suitable for vectorization, our method produces results with lower per-pixel error and avoids the spotty artifacts caused by the use of RBFs. Although RBFs seem to be more flexible, unless using a (potentially impractically) large number of RBFs for relatively complicated input, spotty artifacts are reasonable reflection of radial bases and large approximation errors result. We also experimented with varying RBFs from 3,000 to 5,000 but the approximation errors in our experiments only drop marginally. Wang’s algorithm may also get stuck at suboptimal solutions due to the highly non-linear nature. Our vectorization does not suffer from these problems and is much simpler to optimize as only sparse linear systems need to be solved.

Our method without control point optimization is on average 500 times faster and has much reduced reconstruction error and better color reproduction than [3], as shown in Figs. 10 and 11 as well as Table 1. If the optional control point optimization is used, the error can be further reduced at a small cost. This shows that we currently achieve interactive performance for vectorization of moderate sized volumes. Direct synthesis of gradient solid textures requires many times of intermediate vectorization and evaluation and it would become impractically slow without the speedup. Since the algorithm is highly parallel, a parallel GPU-based implementation may further improve the speed.

We use regions to represent sharp boundaries (Fig. 11) but our method can deal with input solids that cannot be naturally

separated into multiple regions (see Fig. 10). In this case, no binary mask image needs to be provided. A more thorough evaluation on the whole dataset provided by [11] with 21 solid textures shows that for more than 75% of the examples, especially those examples more suitable for vectorization (with lower approximation for both methods), our method outperforms [3] in fitting error (see the accompanied supplementary material for detailed statistics).

We quantize each value with 8 bits, and our representation, without further careful coding, takes only 6.5% (without region separation) or 15% (with region separation) of the voxel solids, while 17%-26% is reported in [3]. The size of the look-up table is not considered because it does not depend on input and thus does not need to be stored in external files and its size is fixed and small enough to be kept in current graphics card without any problem.

5.3 Rendering

While our current *synthesis* implementation is CPU-based, gradient solids are *rendered* efficiently with commodity GPUs. For each visible pixel, we obtain the interpolated texture coordinate using the vertex shader and evaluate the color with Eqn. 1 using the fragment shader; the colors and gradients at control points are stored as textures for efficient GPU access. The color at any continuous position is calculated by linear interpolation of entries in the look-up table described in Sec. 3 through hardware-supported texture fetches, a commonly used technique for realtime rendering. For solid textures with binary masks, the relevant set of feature vectors is used based on the evaluated signed distance. This is both efficient and ensures accuracy as accurate values are obtained at 8^3 times higher resolution than raster solids. This linear interpolation is

example	error			time		
	Wang's	ours w/o opt.	ours w/ opt.	Wang's	ours w/o opt.	ours w/ opt.
Fig. 10 ('caustics')	6.14	2.62	1.50	8 min 25 sec	1.08 sec	5.10 sec
Fig. 11 ('balls')	8.59	5.83	4.21	24 min 21 sec	2.67 sec	11.46 sec

TABLE 1

Statistics of solid vectorization results.

accurate at $\frac{1}{8}$ voxel resolution and is sufficiently close to the real function such that no visible artifact is produced, even when extremely magnified. In most practical applications, a precomputed look-up table with $\frac{1}{4}$ voxel resolution is sufficient, which leads to a look-up table with $17^3 \times 32$ entries and taking less than 0.6MB storage. To avoid jagged boundaries when gradient solids with two regions are rasterized, we use similar antialiasing technique as in [3]. The idea is for pixels close to boundaries, colors evaluated with both positive and negative regions are linearly blended.

Our representation has similar real-time rendering performance as [3]. To make a fair comparison, in the performance measurement, we have disabled mipmapping for [3] and enabled antialiasing for both methods. For a 128^3 solid with a mesh containing 70K vertices rendered at 1024×768 resolution, our average frame rate is 80 fps and the rendering algorithm from [3] achieves on average 75 fps. High resolution solid textures in this paper are rendered with 30-60 fps. The slightly lower frame rates are due to the relatively complicated geometry and large textures with lower cache performance.

5.4 Discussions and Limitations

Although we can represent sharp boundaries with regions, similar to Wang et al. [3] using a single distance field we cannot in general recover sharp boundaries if more than two regions touch. An example is given in Fig. 13. The input solid (a) can be vectorized with our algorithm producing the reconstructed solid (b) with close-up (c). Sharp boundaries between triangles cannot be preserved with the single binary mask. Compared with [3], our blurring effects are much more local. If such blur is not acceptable, our algorithm needs to be modified to be augmented with another distance field to separate adjacent triangle pairs, as shown in (d) (a close-up view). Another limitation is although our fitting error is usually lower than Wang et al. [3] for typical input, fine details within a region may not be fully reproduced; this however is a limitation for virtually all the vectorization methods. To simulate fine details of textures without excessive storage, the approximation error at any position is modeled as a Gaussian distribution. Assume for each position \mathbf{x} , and an arbitrary channel c (including r , g , b) with a sample pixel value $\hat{p}_c(\mathbf{x})$ and corresponding reconstructed value from the vector representation $\tilde{p}_c(\mathbf{x})$, the residual $r_c(\mathbf{x}) = \hat{p}_c(\mathbf{x}) - \tilde{p}_c(\mathbf{x})$ is a Gaussian distribution with probability p satisfying

$$p(r_c(\mathbf{x}) = y) = G(0, \sigma_c(\mathbf{x})) = \frac{1}{\sqrt{2\pi\sigma_c(\mathbf{x})^2}} \exp\left\{-\frac{y^2}{2\sigma_c(\mathbf{x})^2}\right\}, \quad (7)$$

where $\sigma_c(\mathbf{x})$ is the standard deviation and y an arbitrary value. We optimize $\sigma_c(\mathbf{x})$ such that $p(r_c(\mathbf{x}) = \hat{p}_c(\mathbf{x}) - \tilde{p}_c(\mathbf{x}))$ is

maximized, which is worked out as $\sigma_c(\mathbf{x}) = |\hat{p}_c(\mathbf{x}) - \tilde{p}_c(\mathbf{x})|$. For efficiency, σ_c is also compactly represented using our vector representation, treating as an additional channel. When rendering at any position, the residual r is randomly sampled from the distribution. To ensure consistent result, a position determined hash function as Perlin noise [8] is used. With similar lookup table based GPU acceleration, extra computation can be efficiently done, keeping the rendering algorithm realtime (current implementation with 30% – 50% of the original fps). An example is shown in Fig. 14 where richer details are recovered without losing the benefits of vector representation such as resolution independence.

6 CONCLUSIONS AND FUTURE WORK

IN this paper, we propose a novel gradient solid representation for compactly representing solids. We also propose an efficient algorithm for direct synthesis of gradient solid textures from 2D exemplars. Our algorithm is very efficient in both computation and storage, compared with previous voxel-level solid texture synthesis methods and thus allows high-resolution solid textures to be synthesized in full. The algorithm can be generalized to take 3D solids as exemplars which will also benefit from the compactness of our representation. The representation is also potentially useful for accelerating volume processing. We have demonstrated instant editing of large volumes, and we would like to explore other applications such as efficient volumetric rendering and manipulation of (solid) textures (e.g. [35], [36]) in the future. Our current implementation of the synthesis algorithm is purely CPU based. The algorithm is highly parallel and we expect to implement this on the GPU to further improve the performance. Our rendering implementation can be further augmented with mipmapping for adaptive scaling especially minification and texture composition to produce richer fractal-like boundaries, using techniques similar to those in [3].

REFERENCES

- [1] Y. Dong, S. Lefebvre, X. Tong, and G. Drettakis, "Lazy solid texture synthesis," *Computer Graphics Forum*, vol. 27, no. 4, pp. 1165–1174, 2008.
- [2] G.-X. Zhang, S.-P. Du, Y.-K. Lai, T. Ni, and S.-M. Hu, "Sketch guided solid texturing," *Graph. Mod.*, vol. 73, no. 3, pp. 59–73, 2011.
- [3] L. Wang, K. Zhou, Y. Yu, and B. Guo, "Vector solid textures," in *Proc. ACM SIGGRAPH*, 2010, p. Article 86.
- [4] J. Sun, L. Liang, F. Wen, and H.-Y. Shum, "Image vectorization using optimized gradient meshes," *ACM Trans. Graph.*, vol. 26, no. 3, p. Article 11, 2007.

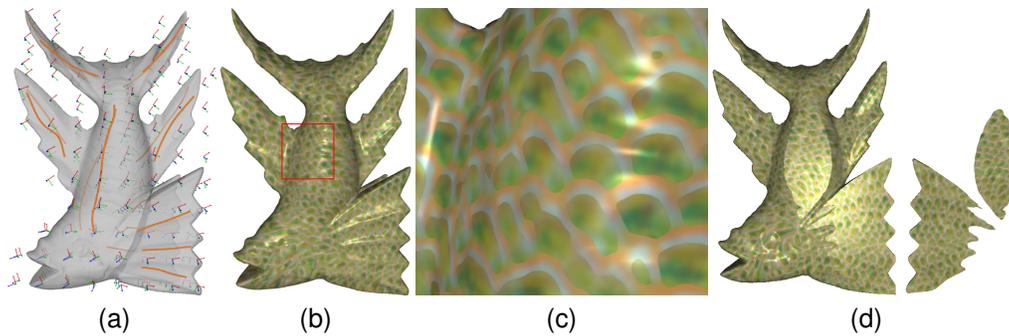


Fig. 12. Synthesized high-resolution solid texture (1024 samples in the longest dimension) with a field. (a) input user specified tensor field; (b) synthesized solid texture; (c) close up (d) internal slices.

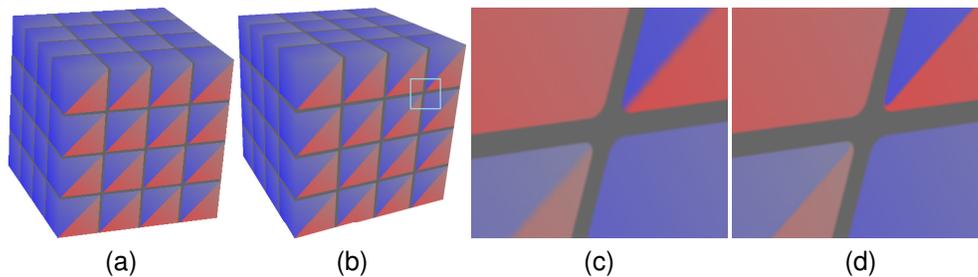


Fig. 13. An example that a single distance field is not sufficient to fully recover sharp boundaries. (a) input solid; (b) vectorized solid with a single distance field; (c) close-up of (b); (d) vectorized solid with an additional distance field to recover sharp boundaries (close-up).

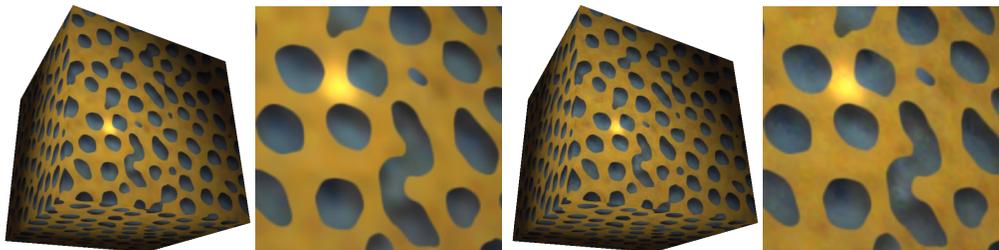


Fig. 14. Vector solid textures without (left) and with added details (right).

- [5] L.-Y. Wei, S. Lefebvre, V. Kwatra, and G. Turk, "State of the art in example-based texture synthesis," in *Eurographics State-of-Art Report*, 2009.
- [6] N. Pietroni, P. Cignoni, M. A. Otaduy, and R. Scopigno, "Solid-texture synthesis: a survey," *IEEE Computer Graphics and Applications*, vol. 30, no. 4, pp. 74–89, 2010.
- [7] D. R. Peachey, "Solid texturing of complex surfaces," in *Proc. ACM SIGGRAPH*, 1985, pp. 279–286.
- [8] K. Perlin, "An image synthesizer," in *Proc. ACM SIGGRAPH*, 1985, pp. 287–296.
- [9] N. A. Carr and J. C. Hart, "Meshed atlases for real-time procedural solid texturing," *ACM Trans. Graph.*, vol. 21, no. 2, pp. 106–131, 2002.
- [10] L.-Y. Wei, "Texture synthesis from multiple sources," in *SIGGRAPH 2003 Sketch*, 2003.
- [11] J. Kopf, C.-W. Fu, D. Cohen-Or, O. Deussen, D. Lischinski, and T.-T. Wong, "Solid texture synthesis from 2D exemplars," *ACM Trans. Graph.*, vol. 26, no. 3, p. Article 2, 2007.
- [12] V. Kwatra, I. Essa, A. Bobick, and N. Kwatra, "Texture optimization for example-based synthesis," *ACM Trans. Graph.*, vol. 24, no. 3, pp. 795–802, 2005.
- [13] D. J. Heeger and J. R. Bergen, "Pyramid-based texture analysis/synthesis," in *Proc. ACM SIGGRAPH*, 1995, pp. 229–238.
- [14] J. Chen and B. Wang, "High quality solid texture synthesis using position and index histogram matching," *The Visual Computer*, vol. 26, no. 4, pp. 253–262, 2010.
- [15] S. Lefebvre and H. Hoppe, "Parallel controllable texture synthesis," *ACM Trans. Graph.*, vol. 24, no. 3, pp. 777–786, 2005.
- [16] R. Jagnow, J. Dorsey, and H. Rushmeier, "Stereological techniques for solid textures," in *Proc. ACM SIGGRAPH*, 2004, pp. 329–335.
- [17] R. Jagnow, J. Dorsey, and H. Rushmeier, "Evaluation of methods for approximating shapes used to synthesize 3d solid textures," *ACM Trans. Applied Perception*, vol. 4, no. 4, p. Article 24, 2008.
- [18] K. Takayama, M. Okabe, T. Ijiri, and T. Igarashi, "Lapped solid textures: filling a model with anisotropic textures," *ACM Trans. Graph.*, vol. 27, no. 3, p. Article 53, 2008.
- [19] C. Ma, L.-Y. Wei, B. Guo, and K. Zhou, "Motion field texture synthesis," *ACM Trans. Graph.*, vol. 28, no. 5, p. Article 110, 2009.
- [20] Y.-K. Lai, S.-M. Hu, and R. R. Martin, "Automatic and topology-preserving gradient mesh generation for image vectorization," *ACM Trans. Graph.*, vol. 28, no. 3, p. Article 85, 2009.

- [21] T. Xia, B. Liao, and Y. Yu, "Patch-based image vectorization with automatic curvilinear feature alignment," *ACM Trans. Graph.*, vol. 28, no. 5, p. Article 115, 2009.
- [22] A. Orzan, A. Bousseau, H. Winnemöller, P. Barla, J. Thollot, and D. Salesin, "Diffusion curves: a vector representaiton for smooth-shaded images," *ACM Trans. Graph.*, vol. 27, no. 3, p. Article 92, 2008.
- [23] W. Barrett and A. S. Cheney, "Object-based image editing," *ACM Trans. Graph.*, vol. 21, no. 3, pp. 777–784, 2002.
- [24] J. Tumblin and P. Choudhury, "Bixels: Picture samples with sharp embedded boundaries," in *Proc. Eurographics Symposium on Rendering*, 2004, pp. 186–196.
- [25] D. Pavić and L. Kobbelt, "Two-colored pixels," *Computer Graphics Forum*, vol. 29, no. 2, pp. 743–752, 2010.
- [26] K. Takayama, O. Sorkine, A. Nealen, and T. Igarashi, "Volumetric modeling with diffusion surfaces," *ACM Trans. Graph.*, vol. 29, no. 6, p. Article 180, 2010.
- [27] P. Ning and L. Hesselink, "Fast volume rendering of compressed data," in *Proc. IEEE Visualization*, 1993, pp. 11–18.
- [28] B.-L. Yeo and B. Liu, "Volume rendering of dct-based compressed 3d scalar data," *IEEE Trans. Vis. Comp. Graph.*, vol. 1, no. 1, pp. 29–43, 1995.
- [29] S. Lefebvre and H. Hoppe, "Appearance-space texture synthesis," *ACM Trans. Graph.*, vol. 25, pp. 541–548, 2006.
- [30] J. Ferguson, "Multivariable curve interpolation," *J. ACM*, vol. 11, no. 2, pp. 221–228, 1964.
- [31] F. Lekien and J. Marsden, "Tricubic interpolation in three dimensions," *J. Numerical Methods Engin.*, vol. 63, pp. 455–471, 2005.
- [32] A. Hertzmann, C. E. Jacobs, N. Oliver, B. Curless, and D. H. Salesin, "Image analogies," in *Proc. ACM SIGGRAPH*, 2001, pp. 327–340.
- [33] Y. Wexler, E. Shechtman, and M. Irani, "Space-time completion of video," *IEEE Trans. PAMI*, vol. 29, no. 3, pp. 463–476, 2007.
- [34] Y. Li, T. Ju, and S.-M. Hu, "Instant propagation of sparse edits on images and videos," *Computer Graphics Forum*, vol. 29, no. 7, pp. 2049–2054, 2010.
- [35] H. Fang and J. C. Hart, "Textureshop: Texture synthesis as a photograph editing tool," in *ACM SIGGRAPH*, 2004, pp. 354–359.
- [36] J. Lu, A. S. Georghiadis, A. Glaser, H. Wu, L.-Y. Wei, B. Guo, J. Dorsey, and H. Rushmeier, "Context-aware textures," *ACM Trans. Graph.*, vol. 26, no. 1, p. Article 3, 2007.