

# Terrain Synthesis from Digital Elevation Models

Howard Zhou, Jie Sun, Greg Turk, *Member, IEEE*, and James M. Rehg, *Member, IEEE*

**Abstract**—In this paper, we present an example-based system for terrain synthesis. In our approach, patches from a sample terrain (represented by a height field) are used to generate a new terrain. The synthesis is guided by a user-sketched feature map that specifies where terrain features occur in the resulting synthetic terrain. Our system emphasizes large-scale curvilinear features (ridges and valleys) because such features are the dominant visual elements in most terrains. Both the example height field and user's sketch map are analyzed using a technique from the field of geomorphology. The system finds patches from the example data that match the features found in the user's sketch. Patches are joined together using graph cuts and Poisson editing. The order in which patches are placed in the synthesized terrain is determined by breadth-first traversal of a feature tree and this generates improved results over standard raster-scan placement orders. Our technique supports user-controlled terrain synthesis in a wide variety of styles, based upon the visual richness of real-world terrain data.

**Index Terms**—Terrain synthesis, Digital Elevation Models, terrain analysis, texture synthesis.

## 1 INTRODUCTION

THERE are numerous applications that make use of synthetic terrain and, very often, the terrain is the dominant visual element in the scene. Such applications include landscape design, flight simulators, emergency response training, battleground simulations, feature film special effects, and computer games. Over the years, graphics researchers have made considerable progress toward developing efficient methods for generating a synthetic terrain. Previous terrain synthesis work has focused on using fractal models and physical erosion models to create a realistic-appearing terrain.

With the rapid growth of computing power and development in terrain visualization techniques, the demand for more realistic terrain has increased considerably. In addition, users of terrain modeling applications want more control over the creation of new terrain. However, current terrain synthesis methods have several limitations. First, these methods provide users with little or no control over the placement of desired terrain features. Second, using the control parameters in these methods, it is difficult to generate terrain with a desired style, such as a terrain with the geological features of the Grand Canyon.

Figs. 2a and 2b show examples of fractal and erosion terrains produced by some popular commercial software, displayed as intensity-coded elevation maps. The styles of these terrains are quite unlike the natural terrain illustrated in Figs. 2c and 2d.

We present a novel example-based terrain synthesis method that addresses the need for intuitive user control over both terrain feature placement and terrain style. Our method draws upon the techniques of patch matching and

patch placement from example-based texture synthesis. In our approach, the user supplies a sketched terrain feature map (called the *sketch map*) and real terrain data (called the *example height field*), which contain the desired terrain styles. Example height fields are in the form of Digital Elevation Models (DEMs), which are available online from the US Geological Survey. Our system then automatically generates a new height field that preserves the visual style of the real terrain data and meets the feature constraints of the sketch map. Synthetic results from our approach are shown in Fig. 1.

The sketch map provides the user with an easy and intuitive way to control the synthesis process. Each map specifies the locations of important terrain features such as the bifurcation point at the center of the image in Fig. 10a. Notice that these sketches are quite coarse. In fact, the width of the brush and the pixel intensities are of little importance, as long as they follow the simple principle that darker pixels indicate lower elevations and brighter pixels indicate higher elevations. Our goal in this work is the generation of visually compelling terrain. We do not address the separate issue of whether a synthesized terrain is geologically accurate.

The starting point for our algorithm is the identification of important terrain features in both the sketch map and the example height field. Our system concentrates on large curvilinear features, such as rivers, valleys, and mountain ridges, since these are usually the most important visual elements in a large-scale terrain (as illustrated in Figs. 2c and 2d). Because the underlying terrain structures in the sketch map and the ones in the example height field are often very different, matches over large spatial scales are unlikely to be correct. However, we observe that, at small scales, common terrain features such as a bifurcation or a straight section of a valley can be found in both maps. Hence, our algorithm breaks the sketch map into small patch regions and searches through the example height field for structural feature matches.

The extraction of structural information from the sketch map is straightforward due to its simplicity. On the other

• The authors are with the GVI Center, Georgia Institute of Technology, TSRB, 85 5th Street NW, Atlanta, GA 30308.

E-mail: {howardz, sun, turk, rehg}@cc.gatech.edu.

Manuscript received 10 Oct. 2006; accepted 28 Dec. 2006; published online 17 Jan. 2007.

For information on obtaining reprints of this article, please send e-mail to: [tvcg@computer.org](mailto:tvcg@computer.org), and reference IEEECS Log Number TVCG-0184-1006. Digital Object Identifier no. 10.1109/TVCG.2007.1027.




Fig. 1. Synthesized terrains using the DEM of the Flathead National Forest Mountain Range (top image) and the DEM of the Grand Canyon (bottom image).

hand, the extraction of high-level structural information such as ridges and valleys from a real terrain height field can be difficult. Here, we draw upon work in the field of geomorphology. Specifically, we use the Profile recognition and Polygon breaking Algorithm (*PPA*) developed by Chang et al. [1] to extract structural information such as ridge and valley axes from a height field.

Once the important structural terrain features have been extracted, our algorithm proceeds in the following stages: First, curvilinear terrain features are used as constraints for the matching and alignment of patch regions in the sketch map and the example height field. Second, these features are used as constraints for matching along the overlapping boundaries between neighboring patch regions. A patch is selected based on a weighted combination of scores that assess the goodness-of-fit between a patch and its neighbors and the amount of deformation of the patch that is required. The order of patch placement is determined by a breadth-first traversal of a feature tree that is constructed from the user sketch. Each patch is placed into the output map using a combination of graph-cut seam finding and Poisson seam removal to minimize visual discontinuities.


This paper makes the following contributions:

- a method for sketch-based specification of synthesized terrain features, giving the user intuitive control over the synthesis result,
- a feature-based approach to matching and placement of large curvilinear terrain features, which makes it possible to efficiently search large terrain databases and preserve the important visual elements in the synthesis process,
- a tree-ordered patch placement algorithm as an alternative to the standard raster-scan ordering, which results in a more faithful reproduction of terrain structure in matching the user input, and
- the ability to synthesize transitions between different terrain types and incorporate multiple DEMs into a single synthesized output.



(a)

(b)



(c)

(d)

Fig. 2. Examples of terrains generated by current methods. (a) Fractal terrain (ridged multi-Perlin) generated by Terragen at  $1,025 \times 1,025$  resolution. (b) Erosion terrain generated by Bryce 5 at  $1,024 \times 1,024$  resolution. In contrast are two real elevation maps: (c) 1/3 arc second (10 m) DEM of the Flathead National Forest mountain range, Montana, and (d) 1/3 arc second (10 m) DEM of Mount Vernon, Kentucky. (DEMs courtesy of the US Geological Survey.)

We have drawn significant inspiration for our work from previous texture synthesis methods, which we review in the following section.

## 2 RELATED WORK

There are two main approaches to generating a synthetic terrain: fractal landscape modeling and physical erosion simulation. Fractal landscape modeling dates back to the pioneering work of Mandelbrot [2]. Since then, a variety of stochastic subdivision techniques have been introduced. Fournier et al. [3] introduced the random midpoint displacement technique to create fractal surfaces. Voss [4] added successive random displacement to fractional Brownian surfaces. Miller [5] proposed a square-square subdivision scheme for generating a fractal terrain and a parallel processing algorithm for rendering height fields. Lewis [6] proposed a generalized stochastic subdivision. Szeliski and Terzopoulos [7] addressed the problem of user control by combining deterministic splines and stochastic fractals into constrained fractals. Recent fractal-based approaches are reviewed in [8] and [9].

Physical erosion simulation is an alternative approach to synthesizing terrain details based on models of landscape formulation and stream erosion from the geomorphology community. It is often used as a refinement step after a rough height field is generated. Kelley et al. [10] first introduced a method to approximate a natural terrain by simulating the erosion of stream networks. Later, Musgrave et al. [11] then combined the fractal modeling and erosion

simulation approaches into a single framework. Recent physical erosion techniques, exemplified by [12], [13], [14], [15], and [16], have focused on improving both the physical modeling aspect and computational efficiency. With appropriately tuned parameters, these techniques can generate a realistic-appearing terrain.

Both fractal and physical erosion techniques add terrain details through procedural refinement, which often involves nonintuitive parameter tuning. Recently, Brosz et al. [17] attempted to extract high-resolution terrain details from existing DEM data and applied it to a lower resolution terrain through multiresolution analysis. In practice, their method requires both the source and target terrain to be fairly detailed and does not grant the user freedom to create an arbitrary terrain.

To provide the user with more intuitive control over the synthesized terrain, image-based alternatives were proposed by Lewis [18] and Perlin and Velho [19]. In these works, terrain was viewed as a type of texture and user control was provided through direct manipulation of the texture. The result was then interpreted as a height field to create a variety of terrain types. However, because it is difficult for a user to draw a realistic natural height field by hand, these methods typically suffer from a lack of realistic detail.

## 2.1 Commercial Software

We are not aware of any existing commercial software for terrain synthesis that employs an approach that is similar to ours. We briefly review four major commercial systems, Mojoworld, Terragen, World Machine, and Bryce, which are representative of the current state of the art. Mojoworld and Terragen both have fractal synthesis engines that generate terrain procedurally. The resulting models are very compact because they are generated on the fly. World Machine, on the other hand, offers geological erosion on top of its procedural shape and noise generator to synthesize a realistic-appealing terrain. In all of these systems, control over the synthesis result is obtained by changing the global parameters in the generation process. It is, however, difficult to set these parameters so as to generate distinctive realistic terrain types. Moreover, these systems do not support the user-specified placement of major terrain features. Bryce, in addition to having both fractal and erosion synthesis engines, also accepts as input a direct specification of the height field by the user using a painting approach (essentially the method in [18] with the addition of fractal noise). Figs. 2c and 2d are example terrains generated by such commercial systems. None of these products, however, make it possible to synthesize terrain in a specific style such as the style of the Grand Canyon.

## 2.2 Texture Synthesis

Image-based texture synthesis is the process of creating an arbitrarily large patch of texture by drawing pixels from a given example image. During the past few decades, a steady improvement in the quality of synthesized textures, both 2D and volumetric, has been achieved through an evolution from pixel-based methods [20] to nonparametric neighborhood-based methods [21], [22], [23]. The most recent patch-based techniques, exemplified by [24], [25], [26], [27], [28], [29], [30], and [31], have two common stages:

1) search in a sample texture for neighborhoods most similar to a context region and 2) merge a patch or a pixel with the (partially) synthesized output texture. Dynamic programming [24] and graph cuts [26] have been used to optimize the patch merging stage. We employ a related search-and-merge strategy that addresses the unique characteristics of terrain data.

Recently, Zhang et al. [32] introduced feature-based warping and blending techniques to synthesize progressively variant texture on arbitrary surfaces. In their work, the feature texture masks were manually extracted. Wu and Yu [27] use edges extracted from the input texture as high-level features to guide patch-based texture synthesis. We build on this earlier work in two ways. First, we employ curvilinear features to support the user sketching of the desired terrain features and efficient search for matching patches in large terrain data sets. Second, we introduce a feature analysis technique that can reliably extract global terrain characteristics such as ridges and valleys from large DEMs in a wide range of styles. We demonstrate that standard edge-finding methods are inappropriate for this task.

The image analogies framework introduced by Hertzmann et al. [25] can be used to synthesize terrain images through a texture-by-numbers approach. This work does not directly synthesize novel terrain height fields. Its application to height-field synthesis is hampered by the difficulty of guaranteeing that the local-neighborhood matching approach would preserve the extended structures such as ridges and valleys that characterize terrain style. The GPU-based texture synthesis method presented in [30] includes drag-and-drop features and synthesis magnification. These techniques were used to relocate mountains on a terrain height map. In contrast to this work, our focus is to automatically generate a complete terrain image in a particular style by matching extended terrain features with a desired user sketch. We therefore support fully automatic feature extraction and matching.

Although our example-based approach to terrain synthesis is inspired by the recent success of patch-based methods, terrain synthesis is not simply texture synthesis on height fields. The terrain generation problem can be distinguished from conventional texture synthesis in three main ways. First, a wide variety of terrain types can be characterized by a combination of global features (such as ridges and valleys), which can be reliably extracted from input terrain maps. In contrast, no such easily identifiable global features exist for general image textures. Second, terrain synthesis must be globally controllable in order to be useful for a wide range of applications and large terrain data sets (often a gigapixel or more) must be searched in order to meet the users' constraints. This stands in contrast to the canonical texture synthesis problem of "growing" a small input patch of texture into a larger output image with the same local intensity structure. We demonstrate that feature-based terrain matching can address both of these concerns. Third, many recent cut-and-merge techniques for texture synthesis exploit the fact that image textures contain many natural boundaries that provide good seams along which to merge texture elements. In contrast, there are no natural seams in terrain data. Moreover, any mismatch

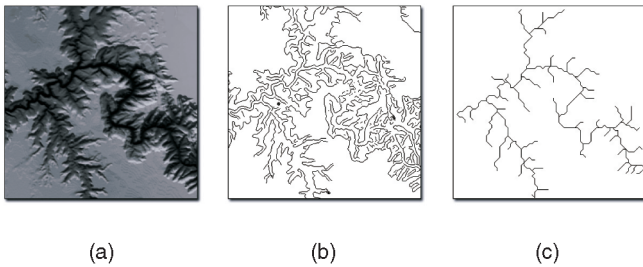


Fig. 3. Comparison of feature extraction methods. (a) Grand Canyon height field displayed as a shaded relief map. (b) Features extracted by the Canny edge detector. (c) Features extracted by PPA.

between the height fields in two adjacent patches is immediately visible. The use of blurring to remove height differences results in highly visible artifacts. We show how to combine recent graph-cut and Poisson seam removal techniques to address the problem of merging terrain patches.

### 3 FEATURE EXTRACTION

In this section, we describe our method for extracting terrain features from the example height field and from the user's sketch map. We use the term *terrain features* to mean large-scale long curvilinear features such as a river, a valley, or a mountain ridge. Such terrain features characterize the overall layout of the terrain. Fig. 3a shows an example height field, consisting of a portion of the Grand Canyon, which is depicted as a shaded relief map. This example illustrates both the existence of large scale curvilinear structures (the main river bed), as well as a rich array of secondary structures (the side canyons and other features). Our goal is to match the primary structures to a desired sketch shape while preserving the rich detail that is present in this data.

It may seem that terrain features could be identified through edge detection. Closer scrutiny, however, reveals several problems in applying edge detection to complex terrain data. First, terrain features are characterized by the local extrema of the height field, whereas edge detection methods are based on the locally maximal derivatives of the image. As a consequence, the application of edge detection to terrain data results in spurious features due to local height variations, as shown in the upper right corner of the image in Fig. 3b.<sup>1</sup> In addition, terrain data has characteristic branching structures that are not always handled correctly by standard edge detection algorithms (for example, in the lower right corner of the image in Fig. 3b).

We draw upon work in the geomorphology literature to identify large-scale terrain features. In particular, we have adapted the PPA, developed by Chang et al. [1], to our task. We briefly summarize the method in this section and refer the reader to their paper for additional details.

PPA can extract either ridges or valleys. Here, we will only describe ridge finding, with the necessary modification to finding valleys being understood. First, each grid point (pixel in the height field) is visited to determine if it is a candidate for

1. Searching for edges at multiple scales cannot easily solve this problem as perceptually important terrain structures are not guaranteed to persist over scale.




Fig. 4. Chang's PPA. (a) Profile recognition. (b) Target connection. (c) Polygon breaking. (d) Branch reduction.

being on a ridge. The grid points in eight outward paths are examined (the yellow points in Fig. 4a), and the current point is marked as a candidate if the height dips below the central point by more than a threshold<sup>2</sup> along any of these paths. All such candidates are then connected by a segment to all other adjacent candidates (the red circles in Fig. 4b). When one segment crosses another, the lower elevation segment is cancelled (shown as dotted segments in Fig. 4b). The polygons are then broken into dendritic line patterns by repeatedly eliminating the least important segment (the remaining segment with the lowest height, shown as dotted segments in Fig. 4c) for each closed polygon. This process terminates when there are no more closed polygons. After the polygon breaking process, shorter branches are eliminated (dotted segments in the Fig. 4d).

The output of PPA for a sample terrain is shown in Fig. 3c. Notice that the valley axes coincide with human perception of the important curvilinear features in the height field. Our system uses PPA to identify features both in the example height field and in the user's sketch map. The gray-scale values in the sketch map are treated as elevation values. The output of PPA is a collection of line segments that are connected and that form long chains along ridges or valleys. These chains of segments are the basis for feature matching between the user's sketch and regions in the given height field. Once the chains of segments have been identified, they are analyzed to form two classes of features: isolated features and curvilinear (path) features. *Isolated features* are branch points and endpoints. *Curvilinear features* are long chains of segments that connect isolated features. The feature matching process described next makes use of these two feature categories.

A key property of the PPA is that the extracted features form a tree (a forest in general) since all closed polygons are broken in the analysis process. This property allows us to use tree traversal to order the placement of patches during synthesis, as described next.

### 4 FEATURE-BASED PATCH MATCHING AND PLACEMENT

Our synthesis process creates a new height field by extracting patches from the example height field and placing them in an output height field in a manner that is dictated by the user's sketch. Typical patch sizes are  $80 \times 80$  pixels (this is determined by the spatial scale of the example terrain data and the detail of the result desired by the user). Patch selection and placement is performed in two stages: *feature patch matching and placement* and *nonfeature patch placement*.

2. We typically set the threshold to one-half of the elevation range of the input height field.




Fig. 5. Examples of terrain patches for each feature type. (a), (b), and (c) Branch point, endpoint, and path patch, respectively. (d), (e), and (f) Corresponding patches after feature extraction. Border colors identify feature types in Fig. 6b.

The first stage, feature matching and placement, locates both isolated features, branch points and endpoints, and nonisolated features, curvilinear features (*paths*) in the sketch map. See Fig. 5 for examples of each of these feature types. Patches are found that match these features and these patches are placed in the output height field.

If we treat the isolated features extracted by the PPA as nodes, and curvilinear features (paths) as edges of a graph, then this graph is guaranteed to be acyclic. Our algorithm follows a breadth-first search order to match and align the patch regions. It first picks an isolated feature (usually a branch if it is present) as the root, it then traverses down through the graph one edge at a time until every edge reachable from the root is covered. This process is illustrated in Figs. 6c, 6d, and 6e. Note that this is a departure from traditional texture synthesis methods that follow a rigid patch placement order (for example, left to right and top down).

Finally, all the as-yet unfilled regions in the output height field are filled using regions from the input that do not contain any strong features. The final result is depicted in Fig. 6f. In both stages, several candidate patches are considered for placement and the patch that is selected for use is the one that gives the least cost in terms of its match with the user's sketch and its overlap with already placed patches. (Patch placement invokes a graph-cut algorithm to select which pixels should be deposited in the output height field and this will be described in Section 5.1). In Sections 4.1 and 4.2, we describe these matching and placement operations in more detail.

#### 4.1 Feature Patch Matching and Placement

As described above, isolated features in the user's sketch are either branch points or endpoints in the segment graph. Fig. 6b contains three branch points and five endpoints. The system analyzes the user's sketch map and identifies all such isolated features. One by one, each such isolated feature is examined and a list of candidate matches is formed from the isolated features in the example height field. We will consider branch points and endpoints in turn.

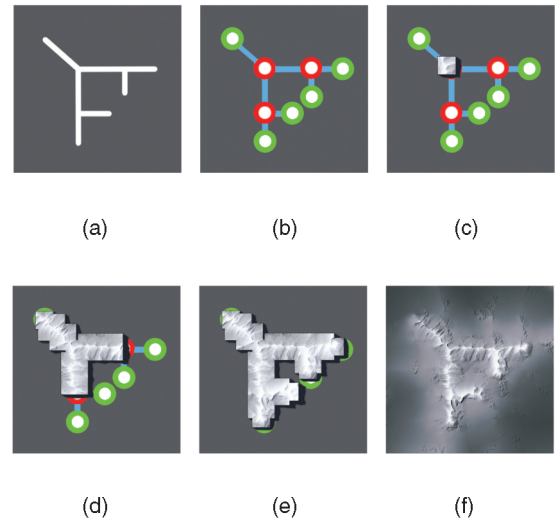


Fig. 6. Illustration of patch placement order. (a) Sample sketch map. (b) Tree structure returned by PPA analysis. Branch-point features and endpoint features are connected by curvilinear path features. (c) The root patch is placed first. (d) Breadth-first traversal guides the placement of additional patches. (e) Once tree traversal is complete, begin placing nonfeature patches. (f) Final result.

##### 4.1.1 Branch Points

For branch-point matching, the degree (or valence) of the branch points must match and the angles of the outgoing segment chains must be similar. For determining the quality of angle matching, only  $d$  possible alignments have to be considered for a degree  $d$  branch point. For instance, as shown in Fig. 5d, the degree 3 branch patch would have three alignments  $(1, 2, 3) \rightarrow (1, 2, 3), (2, 3, 1), (3, 1, 2)$  (the mirror image of this patch is treated as a new patch and its alignment is a separate process). Because there are typically not many such branch points in a given example height field, testing against all candidate matches is extremely fast.

To perform warping of a candidate patch to fit the user's sketch, first a set of control points  $\{P_i\}$  must be identified for the patch. These control points consist of the location of the branch point itself, plus those places where each outgoing path intersects a circle that is inscribed in the patch (see Fig. 5d). For instance, a patch with a three-way branch will have a total of four control points. Corresponding control points  $\{P'_i\}$  are also defined for the patch from the sketch map. Now, we desire a continuous coordinate transformation that maps the control points  $\{P_i\}$  to  $\{P'_i\}$  and that gives the minimal amount of distortion. We use a thin-plate spline (TPS) interpolant [33] for this purpose because it works well with few constraints and it introduces minimal distortion as measured by the integral bending norm. We use two separate TPS functions to form a coordinate transformation that maps any position in the original height field to its interpolated location in the warped height field. The best  $k$  branch-point patches with the lowest deformation energy are the candidates for further matching.




Fig. 7. (a) Height profiles perpendicular to the path are stored at uniformly sampled points along the path. The height profile at point  $P$  that joins two paths is linearly interpolated from the profiles stored at  $P2$  and  $P3$ . (b) An example of the 5-point height profile that is stored at  $P3$ .

### 4.1.2 Endpoints

Endpoint matching is straightforward because all endpoints have similar curvilinear features: a relatively short segment chain going out of the endpoint, which can be aligned easily. Thus, we select all endpoint patches (Fig. 5e) as candidates for further matching without applying warping in this case.

### 4.1.3 Path Features

All of our curvilinear features are chains of line segments (from PPA) that stretch between two isolated features. Our system finds matching patches along such curvilinear features while traversing these chains of line segments in the user’s sketch. The system travels along these chains in steps that are one-half a patch in size, laying down a patch with each step. As with branching patches, the candidate patches along a curvilinear feature are deformed to better fit the user’s sketch. There are always three control points,  $P_0$ ,  $P_1$ , and  $P_2$ , in the candidate patch (Fig. 5f), and three corresponding points,  $P'_0$ ,  $P'_1$ , and  $P'_2$ , in the user’s sketch that determine the warp. The outer control points  $P_1$  and  $P_2$  are located where the path crosses the inscribed circle of the patch. The central point  $P_0$  is the midpoint in the chain within a patch, which is analogous to the central control point in a branch patch.

All of the candidate patches are ranked according to a combination of matching criteria. The best candidate patch is merged in the output height field. We use the following cost terms:

- $c_d$ . This refers to the deformation energy from the TPS warping (we refer the readers to [33]). The deformation energy measures how well the candidate patch terrain structure matches the sketch map constraints. Although TPS warping can warp the patch into the desired configuration in most cases (except for degenerate cases such as when three control points are collinear), a large deformation results in a noticeable distortion and is penalized here.
- $c_g$ . This refers to the graph-cut score (see Appendix A). The graph-cut seam cost is an indication of how well the candidate patch matches the already merged patches in the overlap regions (graph cut automatically handles the case where one patch overlaps with multiple patches and the overlapping can be disjoint). A higher graph-cut score means the seam is more noticeable and harder to remove by Poisson editing.




Fig. 8. (a) Synthesis result using raster-scan patch placement order. (b) Synthesis result using tree traversal.

- $c_f$ . This refers to the feature dissimilarity cost. The dissimilarity cost measures how closely features along the joining paths are matched.<sup>3</sup> We uniformly sample the curvilinear path and store a coarse height profile perpendicular to the path (our current implementation uses a 7-point profile) at each sampled point (Fig. 7). When the path joins the other path at point  $P$ , the height profile at  $P$  is linearly interpolated from the profiles stored at its two nearest neighbors,  $P2$  and  $P3$ . Height profiles from joining paths are compared using the sum of squared differences (SSD).
- $c_i$ s. This refers to other user specified constraint(s) such as the height constraint or the path smoothness constraint.

The total cost is then a linear combination of all of the matching costs  $c = \alpha_d c_d + \alpha_g c_g + \alpha_f c_f + \sum_{i=1}^n \alpha_i c_i$ . Note that the  $\alpha$ s serve not only as weights, but also as normalization coefficients for the costs. In our current system, we set  $\alpha_d = 1,000$ ,  $\alpha_g = 1$ , and  $\alpha_f = 3$  for most test cases. Changing these coefficients changes the emphasis on the matching criteria and results in different synthesis results. After curvilinear feature matching and breadth-first search order placement, the output height field looks like the one in Fig. 6e.

Fig. 8 shows a comparison between raster-scan and tree-ordered patch placement using the sketch map in Fig. 10a and the example height map of Fig. 10b. The inner structure of the  $\lambda$  symbol is garbled in the raster scan result (Fig. 8a). This is because the row-based scan locks onto the boundary of the circle immediately, creating constraints on future matching that prevent good matches in the interior. As Fig. 8b shows, both the  $\lambda$  and the circle can be reproduced if the most difficult element (the center branch point) is matched first.

## 4.2 Nonfeature Patch Placement

Once the isolated features and curvilinear features have been placed, the empty areas in the output height field are those places without strong features. A feasible way to fill these areas is to copy patches that match the pixels that have already been placed. To do this, our system “grows out” from the already filled-in areas. Specifically, square patch positions in coarsely spaced increments (for example, every 100 pixels horizontally and vertically) are filled in descending order

3. The graph cut cost gives an overall measure of the overlap quality, but it weights every pixel in the overlap region equally. It is difficult to integrate the feature dissimilarity measure into the graph cut algorithm; therefore, a separate feature dissimilarity cost term is introduced into our system to place emphasis on features.




Fig. 9. Illustration of patch placement and seam removal. (a) Matching patches are identified via graph cut. (b) Patch placement results in a seam. (c) Poisson seam remover yields the final output.

according to the area of the overlapping region with the already synthesized height field.

The system looks for a high-quality match between such a patch and the already synthesized height field based on the SSD of the overlapping regions and selects  $k$  candidates. It then finds the best match according to the combination of the SSD score and the graph-cut cost in merging. The best patch is placed and the system continues traversing the output height field looking for overlaps. The process terminates when all pixels in the output map have been filled.

The SSD-based search to find the patch that best matches the already synthesized output height field can be accelerated using Fast Fourier Transforms (FFTs) [34], [35]. Using the FFT-based approach, the matching cost can be computed in  $O(n \log(n))$  time, where  $n$  is the number of pixels in the source height field. This is in comparison to  $O(n^2)$  time for a naive SSD implementation.

## 5 PATCH MERGING

Our system combines two techniques to assure smooth transitions between patches that are placed in the output height field. The first of these is the graph-cut technique [26], [36], which finds good seams between already placed pixels and the pixels from a region that is in the process of being placed. The second procedure solves a discrete Poisson equation [37] to create more gentle transitions between the existing pixel elevations and the pixels from a newly placed patch.

### 5.1 Graph-Cut Optimal Seam Finder

The graph-cut algorithm finds a seam in the overlapping region between patches that determines which pixels will be kept in the final image. This is accomplished by solving a max-flow/min-cut graph problem that minimizes the cost of mismatched elevations across the cut. Edges in the graph represent connections between adjacent pixels and they are given weights based on elevation differences. For some patches, we want to insist that particular groups of pixels from a given patch should be included, such as the central pixels in a branch feature patch. This is accomplished by setting some of the edge weights infinitely high.

As mentioned earlier, the cost of the final seam from the graph-cut algorithm is used by the system to select among candidate patches for representing a given feature. Details on the use of the graph-cut algorithm for seam finding can be found in [26].

### 5.2 Poisson Seam Remover

After the system has placed a new patch into the output map using graph cut, a discontinuity may still be visible along the seam (Fig. 9b). We further improve elevation matching across a seam by adjusting the heights according

to artificially set gradient values at the seams. We do this by solving a Poisson equation similar to the manner in which Pérez et al. perform pixel matching across seams in a color image [37].

Our elevation adjustment stage first translates the elevation values in the overlap region (typically one-third of the patch size) into gradient values. Then, the gradient values across the seam are artificially set to zero. Finally, a Poisson equation is solved to find the best fit set of elevations to these adjusted gradient values. The result of this process is a new set of elevations that is much smoother at the seam (Fig. 9c). Our system performs such elevation adjustments locally, invoking the method one time per placed patch. Performing elevation adjustment locally avoids the challenge of solving huge matrix equations. For more details, we refer the reader to Appendix B.

## 6 TERRAIN SYNTHESIS RESULTS

We used DEMs from the US Geological Survey terrain data for the results shown in Figs. 1, 10, 11, 12, 13, and 15. The origin, resolution, and size of the data are listed in Table 1.

Our example height fields range from  $1,200 \times 1,200$  to  $4,097 \times 4,097$  samples at a height resolution of 16 bits. These large terrain maps pose significant search problems for a traditional texture synthesis approach. At our patch size of  $80 \times 80$  pixels, there are approximately 16 million possible matches in a given  $4,000 \times 4,000$  map. However, as a result of our feature-based analysis, we can filter this total set down to approximately 600 match evaluations for a given candidate position to be filled, which dramatically reduces the computational cost. Our method can synthesize a representative terrain in approximately 5 to 6 minutes on an Intel Pentium 4 2.0 GHz processor with a 2 Gbyte memory. All of the results are rendered using Planetside's Terragen terrain rendering system with procedural textures (determined by height and slope) overlaid on top of the terrain geometry.

Figs. 10 and 11 illustrate the variation in output that can be obtained from the same sketch using our terrain synthesis approach. Fig. 10d shows the synthesis result ( $1,000 \times 1,000$ ) for the Mount Jackson terrain, and the sketch map for this example is shown in Fig. 10a. Although this result was created from many patches, both the interior  $\lambda$  and the outer circle are formed by unbroken mountain ridges. Also note the characteristic curtainlike folds in the sides of the mountains. Fig. 10h shows the synthesis results ( $2,000 \times 2,000$ ) using the Grand Canyon terrain. The sketch map in Fig. 10e used for this example is the inverse of that in Fig. 10a so that valleys are selected instead of mountains. Here, the  $\lambda$  and the circle are formed from joined pieces of the Colorado River, yet the seam locations cannot be detected. Although guided by the user's sketch, this terrain




Fig. 10. Multiple terrain synthesis results with the sketched half-life symbol. (a) User sketch. (b) Mount Jackson. (c) Synthesis result. (d) Rendered terrain synthesized from an elevation map of Mount Jackson, Colorado. (e) User sketch. (f) Grand Canyon. (g) Synthesis result. (h) Rendered terrain synthesized from an elevation map of the Grand Canyon.

retains the rich water-carved features of the original data. Fig. 11d shows the synthesis result ( $1,000 \times 1,000$ ) for the Mount Vernon terrain, and Fig. 11h shows the synthesis results ( $2,000 \times 2,000$ ) using the DEM from Flathead National Forest mountain region. Fig. 12d shows the synthesis result ( $1,000 \times 800$ ) for the Mount Vernon terrain with a user sketched Chinese character for “water.” These

results illustrate the ability of our system to work with a wide variety of terrain types with very different characteristics while maintaining the salient features of the input sketch map. Fig. 15 provides close-up views of the Grand Canyon and Flathead range synthesis results. Fig. 13a shows a rendering of the Grand Canyon height map. We extract the canyon feature from this map using PPA to



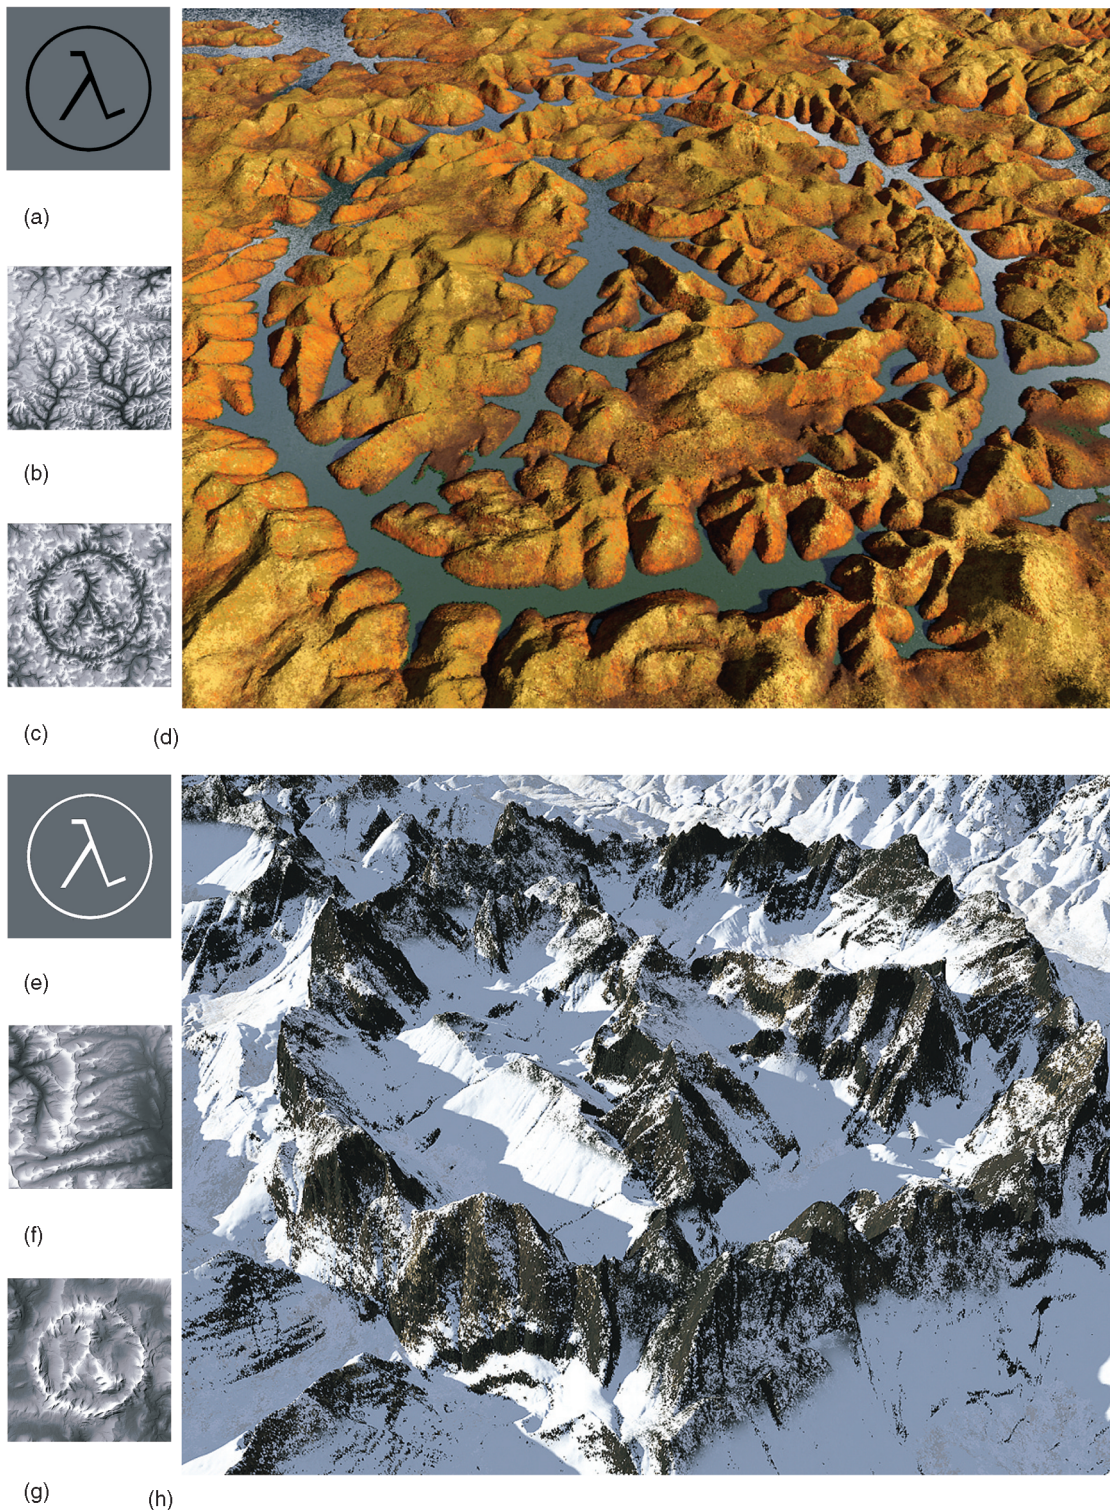


Fig. 11. Multiple terrain synthesis results with the sketched half-life symbol. (a) User sketch. (b) Mount Vernon. (c) Synthesis result. (d) Rendered terrain synthesized from an elevation map of Mount Vernon, Kentucky. (e) User sketch. (f) Flathead range. (g) Synthesis result. (h) Rendered terrain synthesized from an elevation map of Flathead National Forest mountain range, Montana.

obtain a sketch map. By synthesizing the sketch in the style of the Puget Sound DEM, we obtain a “mountain” version ( $4,000 \times 2,000$ ) of the Grand Canyon, illustrated in Fig. 13b.

### 6.1 Synthesis with Multiple Terrain Styles

A larger example, illustrated in Fig. 14, shows the generation of a 3D map of Middle Earth using several different terrain

styles. A simple sketch map (Fig. 14b) was created by the user from an artist rendering of the map of Middle Earth (Fig. 14a). The user assigned different terrain styles to each part of the sketch map (Table 2). In this example, the user first placed Mount St. Helen at the location of Mount Doom in Mordor (a localized feature). Our system then synthesized the rest of the terrain ( $3,470 \times 2,996$ ) automatically and merged it seamlessly




Fig. 12. Terrain synthesis result with the sketched Chinese character for “water.” (a) User sketch. (b) Mount Vernon. (c) Synthesis result. (d) Rendered terrain synthesized from an elevation map of Mount Vernon, Kentucky.

with Mount Doom. Note that the user sketch includes both a variety of mountain ranges and a river valley (the Dead Marshes, near the northwest corner of Mordor). The coast, lakes, and rivers were created using a simple mask, which was extracted from the original map. The synthesized height map (Fig. 14c) was rendered using Planetside’s Terragen terrain rendering system with both procedural and masked textures to give it an artistic feel. This example illustrates the ability of our system to combine multiple terrain styles into one synthesized terrain and to blend seamlessly among them.

Please view the movie (DivX encoded) that accompanies this paper at <http://www.computer.org/tvcg/archives.htm>. Additional results may be found at <http://www.cc.gatech.edu/howardz/terrain>.

## 7 CONCLUSION

We have demonstrated that example-based texture synthesis methods can be successfully adapted to the domain of terrain synthesis. The result is a new level of visual realism in a user-controllable synthesized terrain. Our approach leverages the fact that useful terrain features can be extracted from height fields using analysis techniques from the geomorphology community.

We have introduced a tree-ordered patch placement algorithm that is based on a breadth-first traversal of a feature tree. Our results demonstrate that this placement method is superior to standard raster-scan placement orders.

We have demonstrated the ability to synthesize terrain in widely differing styles while retaining control over the positioning of major terrain features. Our system is based on an intuitive sketch-based interface for specifying the desired terrain characteristics. We believe ours is the first

system to make full use of an example-based approach in the domain of terrain synthesis.


Our method suffers from the same limitations as all example-based methods. In particular, the quality of the final synthesis depends upon the richness of the available terrain data. If the terrain features desired by the user cannot be found in the example height field, then it will not be possible to produce the desired result. For instance, our method will perform poorly if the terrain data is from a desert region where few significant curvilinear features are present. The method can also perform poorly when the curvilinear feature pattern is extremely complicated. An additional issue is the need to specify the patch size, which depends in turn upon the resolution and scale of the example terrain.

We plan to extend our current method in several ways. First, it would be interesting to give the user greater control over the synthesized terrain by incorporating additional constraints into the sketch map. For example, we plan to provide the ability to specify a desired elevation at a specific position. This could be accomplished by constraining the Poisson solver in conjunction with the matcher. We are also investigating the enforcement of  $C^1$  continuity in addition to our current enforcement of  $C^0$  continuity. We intend to explore the joint synthesis of elevation and texture maps for rendering and the interactive control of terrain synthesis.

## APPENDIX A

### GRAPH-CUT SEAM FINDING

The graph-cut algorithm finds the minimum cost seam (according to some matching quality measure) in the overlapping region between patches that determines which



(a)



(b)

Fig. 13. The Grand Canyon turned into a mountain range. Features extracted from the Grand Canyon DEM are used as the user sketch to synthesize a mountain range following the structure of the Grand Canyon from an elevation map of the Puget-Sound-style mountain range.

pixels will be kept in the final image. We choose the matching quality measure defined in [26], which is a measure of intensity (elevation) difference between the pairs of pixels. For example, in Fig. 16, let  $s$  and  $t$  be two adjacent pixel positions in the overlap region  $\Omega$  between patch  $A$  and  $B$ . Let  $a(s)$  and  $b(s)$  be the elevation at the position of the patches, respectively. Then, the matching quality measure  $M$  between the two adjacent pixels at position  $s$  and  $t$  from patches  $A$  and  $B$  is defined to be

$$M(s, t, A, B) = |a(s) - b(s)| + |a(t) - b(t)|. \quad (1)$$

The graph shown in Fig. 16 has one node per pixel in the overlap region between patches. The weight of the edge connecting the adjacent pixel nodes  $s$  and  $t$  is set to equal the matching quality cost  $M(s, t, A, B)$ . We use two additional

nodes  $A$  and  $B$  to represent the old and new patches. The edges that connect pixels in  $\Omega$  and nodes  $A$  and  $B$  are set to have infinitely high weights indicating that these pixels are constrained to come from one particular patch. In Fig. 16, pixels 1, 2, and 3 have to come from patch  $A$  and pixels 7, 8, and 9 from  $B$ . To determine which patch each of the pixels 4, 5, and 6 will come from, we solve a max-flow/min-cut graph problem that minimizes the cost of mismatched elevations across the cut. The red line shows the minimum cut (the elevation difference between the two patches along the cut is minimum). In the overlap region, pixels 5 and 6 will be copied from the old patch  $B$  since they are still connected to node  $B$ . Likewise, pixel 4 will be copied from  $A$ . The cost  $c_g$  of the minimum cut  $\mathcal{C}$  is defined in terms of the matching quality measure as

$$c_g = \sum_{\substack{\langle s,t \rangle \in \mathcal{C} \\ s,t \in \Omega}} M(s, t, A, B). \quad (2)$$

TABLE 1  
Input DEM Locations, Resolution, and Sizes

Origin	Resolution	Size
Mount Jackson, CO	10m	1620 × 1620
The Grand Canyon	30m	4080 × 2760
Mount Vernon, KY	10m	1200 × 1200
Flathead Range, MT	10m	4097 × 4097
Puget Sound, WA	10m	4097 × 4097


For some patches, we want to insist that particular groups of pixels from a given patch should be included, such as the central area surrounding a branch-point feature or along a path. This is accomplished by setting the edge weights connecting those areas to the nonoverlapping region infinitely high.



(a)

(b)

(c)



(d)

Fig. 14. A 3D map of Middle Earth synthesized from multiple elevation maps. (a) Map of Middle Earth. (b) User sketch with masked water system. (c) Synthesis result. (d) Rendered terrain synthesized from multiple elevation maps.

**APPENDIX B  
POISSON SEAM REMOVAL**

Even using the graph-cut approach to choose where to join two patches, elevation differences may still be visible. For example, in Fig. 9a, the overlap region  $\Omega$  is the new height field from our graph-cut seam finder. Though optimum, seam  $\mathcal{C}$  is still visible in the new height field (Fig. 9b). We

will describe how we remove this seam through our Poisson elevation adjustment method.

Let  $\Omega$  be a closed subset of  $\mathbb{R}^2$  with boundary  $\partial\Omega$ . Let  $f$  be the elevation value in the overlap region; hence,  $f$  is a scalar function defined over  $\Omega$ . Our elevation adjustment stage first translates the elevation values in  $\Omega$  into gradient vector fields  $\mathbf{v} = (u, v)$ , with  $u = \frac{\partial f}{\partial x}$  and  $v = \frac{\partial f}{\partial y}$ . To remove the height difference between pixels across the seam  $\mathcal{C}$ , the




Fig. 15. Close-up views of synthesized terrain from the Grand Canyon (top images) and Flathead National Forest mountain range (bottom images).

TABLE 2  
Middle Earth Regions

Middle Earth region	Elevation map
Mordor	Puget Sound, WA
Gondor, Rhun, Ered	Rocky Mountain, CO
Moria	Mount Jackson, CO
Dead Marshes	Mount Vernon, KY
Iron	Flathead Range, MT




Fig. 16. Finding the optimal seam in the overlapping region with graph cut.

gradient values along the seam are artificially set to zero. However, the resulting gradient vector field  $\mathbf{v}'$  is most likely no longer conservative ( $\text{curl}(\mathbf{v}') \neq 0$ ); in other words, it is no longer the gradient of any scalar function. Here, the Poisson methodology comes into play because it allows nonconservative vector fields to be used to reconstruct a plausible elevation field. As shown in [37], we can find the best fit set of elevations  $f'$  to the adjusted vector field  $\mathbf{v}'$  by solving the Poisson equation with Dirichlet boundary conditions.

The result of this process is a new set of elevations that change very little at the seam, as shown in Fig. 9c.

## ACKNOWLEDGMENTS

The authors would like to thank Yet-Chung Chang for his discussions on the PPA algorithm and Yuri Boykov and Vladimir Kolmogorov for making their implementation of the maxflow algorithm publicly available on the Web. Thanks are also due to Vivek Kwatra, Matthew Mullin, Huamin Wang, Spencer Charles Brubaker, Matthew Flagg, Pei Yin, Jianxin Wu, Yifan Shi, and Ping Wang for their helpful comments and suggestions. This work was supported by US National Science Foundation (NSF) Grants ITR-0205507 and CCF-0204355 and an NSF Graduate Research Fellowship (to Howard Zhou). The authors would also like to thank the reviewers whose excellent comments led to the current paper.

## REFERENCES

- [1] Y.-C. Chang, G.-S. Song, and S.-K. Hsu, "Automatic Extraction of Ridge and Valley Axes Using the Profile Recognition and Polygon-Breaking Algorithm," *Computer and Geosciences*, vol. 24, no. 1, pp. 83-93, 1998.
- [2] B.B. Mandelbrot, *The Fractal Geometry of Nature*. W.H. Freeman, 1982.
- [3] A. Fournier, D. Fussler, and L. Carpenter, "Computer Rendering of Stochastic Models," *Comm. ACM*, vol. 25, no. 6, pp. 371-384, 1982.
- [4] R.F. Voss, "Random Fractal Forgeries," *Fundamental Algorithms for Computer Graphics*, R.A. Earnshaw, ed., 1985.
- [5] G.S.P. Miller, "The Definition and Rendering of Terrain Maps," *Proc. 13th Ann. Conf. Computer Graphics and Interactive Techniques (SIGGRAPH '86)*, vol. 20, no. 4, pp. 39-48, 1986.
- [6] J.P. Lewis, "Generalized Stochastic Subdivision," *ACM Trans. Graphics*, vol. 6, no. 3, pp. 167-190, 1987.

- [7] R. Szeliski and D. Terzopoulos, "From Splines to Fractals," *Proc. 16th Ann. Conf. Computer Graphics and Interactive Techniques (SIGGRAPH '89)*, pp. 51-60, 1989.
- [8] D.S. Ebert, F.K. Musgrave, D. Peachy, K. Perlin, and S. Worley, *Texturing and Modeling: A Procedural Approach*. Morgan Kaufmann, 2002.
- [9] C. Dachsbacher, *Interactive Terrain Rendering: Towards Realism with Procedural Models and Graphics Hardware*, <http://www.opus.uni-erlangen.de/opus/volltexte/2006/354/>, 2006.
- [10] A.D. Kelley, M.C. Malin, and G.M. Nielson, "Terrain Simulation Using a Model of Stream Erosion," *Proc. 15th Ann. Conf. Computer Graphics and Interactive Techniques (SIGGRAPH '88)*, pp. 263-268, 1988.
- [11] F.K. Musgrave, C.E. Kolb, and R.S. Mace, "The Synthesis and Rendering of Eroded Fractal Terrains," *Proc. 16th Ann. Conf. Computer Graphics and Interactive Techniques (SIGGRAPH '89)*, vol. 23, no. 3, pp. 41-50, 1989.
- [12] P. Roudier and B.P.M. Perrin, "Landscapes Synthesis Achieved through Erosion and Deposition Process Simulation," *Computer Graphics Forum*, vol. 12, no. 3, p. 375, Aug. 1993.
- [13] N. Chiba, K. Muraoka, and K. Fujita, "An Erosion Model Based on Velocity Fields for the Visual Simulation of Mountain Scenery," *J. Visualization and Computer Animation*, vol. 9, no. 4, pp. 185-194, 1998.
- [14] K. Nagashima, "Computer Generation of Eroded Valley and Mountain Terrains," *The Visual Computer*, vol. 13, nos. 9-10, pp. 456-464, 1997.
- [15] B. Benes and R. Forsbach, "Layered Data Representation for Visual Simulation of Terrain Erosion," *Proc. 17th Spring Conf. Computer Graphics (SCCG '01)*, p. 80, 2001.
- [16] B. Neidhold, M. Wacker, and O. Deussen, "Interactive Physically Based Fluid and Erosion Simulation," *Proc. Eurographics Workshop Natural Phenomena*, 2005.
- [17] J. Brosz, F.F. Samavati, and M.C. Sousa, "Terrain Synthesis By-Example," *Proc. First Int'l Conf. Computer Graphics Theory and Applications (GRAPP '06)*, 2006.
- [18] J.-P. Lewis, "Texture Synthesis for Digital Painting," *Proc. 11th Ann. Conf. Computer Graphics and Interactive Techniques (SIGGRAPH '84)*, pp. 245-252, 1984.
- [19] K. Perlin and L. Velho, "Live Paint: Painting with Procedural Multiscale Textures," *Proc. 22nd Ann. Conf. Computer Graphics and Interactive Techniques (SIGGRAPH '95)*, pp. 153-160, 1995.
- [20] D.J. Heeger and J.R. Bergen, "Pyramid-Based Texture Analysis/Synthesis," *Proc. 22nd Ann. Conf. Computer Graphics and Interactive Techniques (SIGGRAPH '95)*, pp. 229-238, 1995.
- [21] A.A. Efros and T.K. Leung, "Texture Synthesis by Non-Parametric Sampling," *Proc. Int'l Conf. Computer Vision*, pp. 1033-1038, 1999.
- [22] L.-Y. Wei and M. Levoy, "Fast Texture Synthesis Using Tree-Structured Vector Quantization," *Proc. 27th Ann. Conf. Computer Graphics and Interactive Techniques (SIGGRAPH '00)*, pp. 479-488, 2000.
- [23] M. Ashikhmin, "Synthesizing Natural Textures," *Proc. Symp. Interactive 3D Graphics*, pp. 217-226, 2001.
- [24] A.A. Efros and W.T. Freeman, "Image Quilting for Texture Synthesis and Transfer," *Proc. 28th Ann. Conf. Computer Graphics and Interactive Techniques (SIGGRAPH '01)*, pp. 341-346, 2001.
- [25] A. Hertzmann, C.E. Jacobs, N. Oliver, B. Curless, and D.H. Salesin, "Image Analogies," *Proc. 28th Ann. Conf. Computer Graphics and Interactive Techniques (SIGGRAPH '01)*, pp. 327-340, 2001.
- [26] V. Kwatra, A. Schödl, I. Essa, G. Turk, and A. Bobick, "Graphcut Textures: Image and Video Synthesis Using Graph Cuts," *ACM Trans. Graphics, Proc. 30th Int'l Conf. Computer Graphics and Interactive Techniques (SIGGRAPH '03)*, vol. 22, no. 3, pp. 277-286, 2003.
- [27] Q. Wu and Y. Yu, "Feature Matching and Deformation for Texture Synthesis," *ACM Trans. Graphics, Proc. 31st Int'l Conf. Computer Graphics and Interactive Techniques (SIGGRAPH '03)*, vol. 23, no. 3, pp. 364-367, 2004.
- [28] P. Bhat, S. Ingram, and G. Turk, "Geometric Texture Synthesis," *Proc. Eurographics Symp. Geometry Processing*, 2004.
- [29] A. Lagae, O. Dumont, and P. Dutré, "Geometry Synthesis by Example," *Shape Modeling Int'l*, 2005.
- [30] S. Lefebvre and H. Hoppe, "Parallel Controllable Texture Synthesis," *ACM Trans. Graphics, Proc. 32nd Int'l Conf. Computer Graphics and Interactive Techniques (SIGGRAPH '05)*, pp. 777-786, Aug. 2005.

- [31] S. Lefebvre and H. Hoppe, "Appearance-Space Texture Synthesis," *ACM Trans. Graphics, Proc. 33rd Int'l Conf. Computer Graphics and Interactive Techniques (SIGGRAPH '06)*, vol. 25, no. 3, pp. 541-548, 2006.
- [32] J. Zhang, K. Zhou, L. Velho, B. Guo, and H.-Y. Shum, "Synthesis of Progressively-Variant Textures on Arbitrary Surfaces," *ACM Trans. Graphics*, vol. 22, no. 3, pp. 295-302, 2003.
- [33] F.L. Bookstein, "Principal Warps: Thin-Plate Splines and the Decomposition of Deformations," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 11, no. 6, pp. 567-585, 1989.
- [34] C. Soler, M.-P. Cani, and A. Angelidis, "Hierarchical Pattern Mapping," *ACM Trans. Graphics*, vol. 21, no. 3, pp. 673-680, July 2002.
- [35] S.L. Kithau, M.S. Drew, and T. Möller, "Full Search Content Independent Block Matching Based on the Fast Fourier Transform," *Proc. IEEE Int'l Conf. Image Processing (ICIP '02)*, pp. 669-672, 2002.
- [36] Y. Boykov, O. Veksler, and R. Zabih, "Fast Approximate Energy Minimization via Graph Cuts," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 23, no. 11, pp. 1222-1239, Nov. 2001.
- [37] P. Pérez, M. Gangnet, and A. Blake, "Poisson Image Editing," *ACM Trans. Graphics, Proc. 30th Int'l Conf. Computer Graphics and Interactive Techniques (SIGGRAPH '03)*, vol. 22, no. 3, pp. 313-318, 2003.



**Howard Zhou** received the BS degree in engineering and applied science from the California Institute of Technology in 2002 and the MS degree in computer science from the Georgia Institute of Technology in 2005. He is currently working toward the PhD degree under the supervision of Dr. James M. Rehg and Dr. Greg Turk. His research interests include computer vision, computer graphics, and machine learning.



**Jie Sun** received the BE degree in computer science and engineering from Zhejiang University, China, in 2002 and the MS degree in computer science from the Georgia Institute of Technology in 2005, where he is currently working toward the PhD degree under the supervision of Dr. James M. Rehg and Dr. Aaron Bobick. His research interests include computer vision, machine learning, and computer graphics.



**Greg Turk** received the PhD degree in computer science from the University of North Carolina, Chapel Hill, in 1992. He is currently an associate professor at the Georgia Institute of Technology, where he is a member of the College of Computing and the Graphics, Visualization, and Usability Center. His research interests include computer graphics, computer vision, and scientific visualization. He is a member of the IEEE and the IEEE Computer Society.



**James M. Rehg** received the PhD degree in electrical and computer engineering from the Carnegie Mellon University. He is an associate professor at the College of Computing, Georgia Institute of Technology. He is a member of the Graphics, Visualization, and Usability Center and codirects the Computational Perception Laboratory. His research interests are computer vision, computer graphics, machine learning, and robotics. He is a member of the IEEE and the IEEE Computer Society.

► **For more information on this or any other computing topic, please visit our Digital Library at [www.computer.org/publications/dlib](http://www.computer.org/publications/dlib).**