Depicting Terrain with Shaded Relief Maps

Mike Cammarano*
Stanford University

Abstract

Shaded relief maps provide an appealing visualization of landforms and hilly terrain. This paper considers two aspects of their construction: simplification of terrain, and effective feature depiction through exaggerated shading.

Simplification involves choosing a level-of-detail appropriate to the scale of reproduction and the desired emphasis of the map. We will model this process as a digital filtering operation applied to regularly sampled elevation data.

Hill shading techniques for indicating terrain slope have seemed natural candidates for automation using bump-mapping. However, naïve application of bump mapping to elevation data fails to reproduce the distinctive character of hill shading painted by a skilled cartographer. Examples of manual hill shading often feature idealized shading that emphasizes conceptually important terrain features like ridge lines. We treat this as a nonphotorealistic rendering problem, and propose a new shading model that attempts to model this illustrative technique.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Color, shading; I.4.3 [Image Processing and Computer Vision]: Enhancement—Filtering;

Keywords: cartography, nonphotorealistic rendering, simplification, level of detail

Figure 1: Medium-scale map with hill shading and low-detail bathymetric contours, all generated automatically.

1 Introduction

Shaded relief maps can combine an appealing aesthetic with impressively high information density, while still making their rich content readily accessible to an observer with no special training. Maps are perhaps one of the oldest forms of visualization, and pose many fundamental challenges relevant to the field as a whole. There has also been a rich scholarly tradition within cartography, allowing us to draw on techniques of visual communication that have been refined over many years. Imhof [1982] provides an excellent history of cartographic relief presentation as well as a discussion of many techniques of the discipline.

An effective map, like any effective visualization, is designed to help the user readily answer certain types of questions. For example, a map may be designed to exhibit national or regional boundaries, to illustrate a network of roads or locations of cities, or as a geographic foundation upon which to overlay weather or statistical data. In any case, when multiple layers of information are combined, each must be rendered with a level of detail and emphasis appropriate to its role in the graphic as a whole. As an example, Figure 2 shows two maps of Poland at similar scale, one intended to illustrate highways while the other delineates the national boundary in detail.

Perhaps even more critical than choice of emphasis is choice of scale. Given the limited resolution of the intended display medium, a small scale map will require at least some smoothing of contours, and perhaps even topological simplification. Figure 3 shows how the landforms of the Orkney islands might undergo dramatic topological simplification to be effectively depicted at a smaller scale.

The preceding examples should motivate the need for cartographer-directed simplification of terrain to a level of detail appropriate to the goals of the map. While much existing work in the field of Geographic Information Systems (GIS) considers simplification of vector data representing coastlines and contours, we will examine this problem in terms of digital filtering of regularly sampled elevation data.

In addition to boundaries and contours, a map may also illustrate topography using techniques like hypsometric tinting and hill shading. Hypsometric tints use hue to encode the elevation at a point, while hill shading conveys slope by (approximately) shading the land surface as if lit from a canonical direction – generally northwest. Both techniques leverage the human visual system, taking advantage of our propensity for for interpreting shape from shading, and our familiarity with colors found in nature as a diagrammatic shorthand. Terrain features depicted in this manner can generally be perceived and interpreted with little conscious effort.

*e-mail: mcammara@graphics.stanford.edu

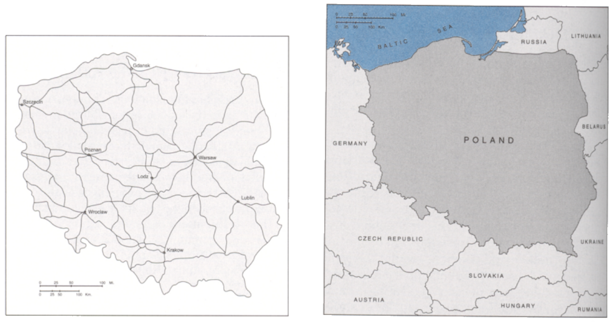

Figure 2: Two maps of Poland at similar scale. The left map is intended to portray major highways, while the right map depicts the geographic relationship to neighboring countries. Note the significant difference in detail with which the national border is represented in the two examples. Example taken from Robinson et al. [1995].

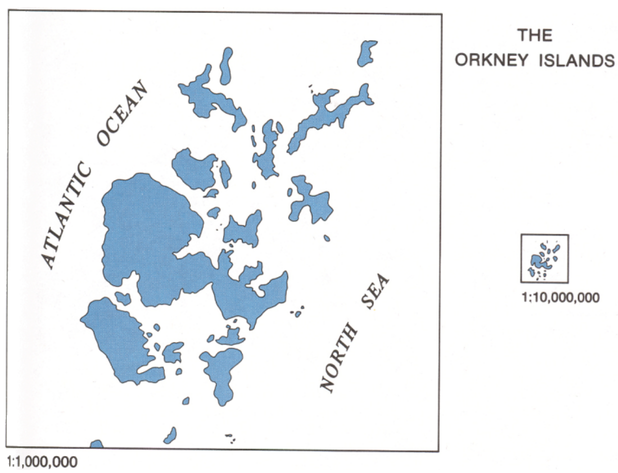

Figure 3: Orkney islands represented at two different scales. Note the substantial topological simplification of the smaller scale example. From Robinson et al. [1995].

Effective use of these perceptual cues is far from trivial, however. Again, level of detail needs to be considered, and variations in the terrain must be idealized to conform to the overall representative goals of the map. It is not desirable to depict every pebble with high-constrast shading variations, or the resulting map will be unreadable. On the other hand, features like ridge lines are very important to our understanding of terrain[1], and hill shading is often drawn so as to emphasize them strongly with a sharp transition in shading across the crest of a hill. Imhof [1982] provides many examples of hand-painted hill shading that effectively illustrate the idealization of terrain shape by a skilled practitioner. He also provides a history of hypsometric tint schemes that have been proposed for depicting elevation.

Imhof's map of Europe illustrates yet another level of detail issue in map design. Figure 4 shows a closeup of a portion of this map. Notice that the coastline is depicted in great detail, while the bathymetric contours indicating sea depths are substantially smoothed and simplified. A more detailed rendition of contours in the ocean depths would not contribute to the purposes for which the map is intended. Indeed, jagged and overly detailed ocean contours would be distracting and disruptive to the map as a whole. Hence we see a compelling example where even adjacent features in a single map may need to be presented at significantly different levels of detail.

[1] Possibly because we are accustomed to a vantage point at ground level, where we see hills in profile.

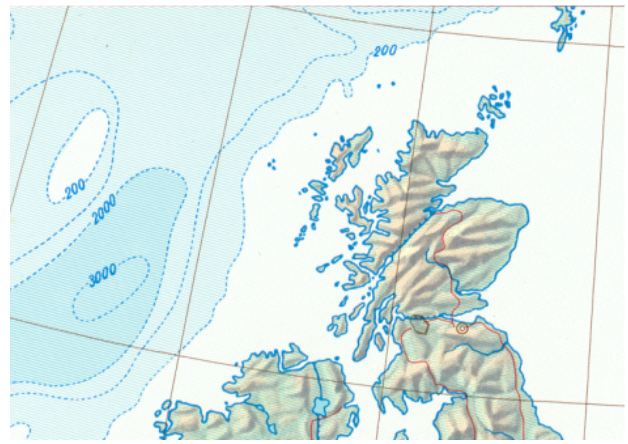

Figure 4: Detail (Scotland and Ireland) from Imhof's [1982] map of Europe. Note the difference in detail between shoreline and nearby bathymetric contours.

In this paper, we propose to treat all simplification or idealization as (possibly adaptive) digital filtering operations applied to regularly sampled elevation data. Next, inspired by research on nonphotorealistic rendering techniques to improve shape communication [Gooch and Gooch 1999], we introduce a shading technique that attempts to mimic the emphasis on ridge lines seen in hand painted hill shading. A schematic overview of our approach is illustrated in Figure 5.

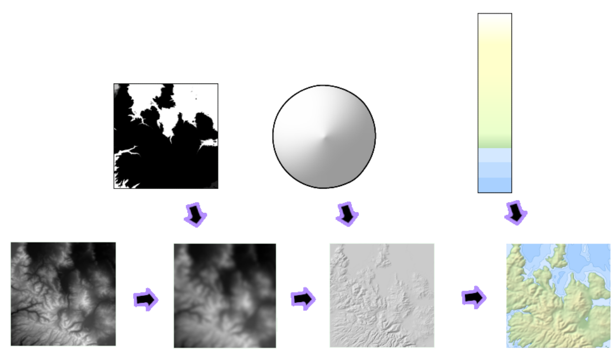

Figure 5: Schematic overview of the production of shaded relief from elevation data. From left to right: Raw elevation data is adaptively filtered based on a detail mask. Filtered elevation data is then bump-mapped, but a nonphotorealistic shading operator is applied rather than diffuse. Finally, hypsometric tints are applied and contours are drawn.

2 Related Work

Naturally, this project draws a great deal of inspiration from techniques of manual relief illustration [Imhof 1982]. However, this section will mainly consider existing applications of computers to simplification and shaded relief construction. Imhof presents a brief survey of early research into automating hill shading with computers. He observes that map design typically requires design decisions regarding which features to emphasize which cannot readily be automated. Notable among these are choice of level of detail for specific features, and local variations in shading. Hopefully the examples from Section 1 have already illustrated that such decisions must be based on the communicative intent of the map. However, we can endeavor to create interactive tools that will efficiently support the cartographer's work.

In recent years, the field of Geographic Information Systems has grown rapidly. Through the efforts of the United States Geological Survey (USGS), vast amounts of elevation survey data and other cartographic information has been made available in standardized digital formats. Digital cartographic data has generally been divided into two categories:

- Vector data, such as the Digital Line Graphs from USGS. Used to describe cartographic features like regional boundaries, roads and transportation lines, rivers, and elevation contours.
- Raster data, notably the Digital Elevation Maps from USGS. Generally, any type of data sampled on a regular grid. In addition to DEMs, which encode elevation, other raster data might include aerial or satellite images.

There has been extensive research into computer-assisted simplification of cartographic data in vector form. For example, a call for papers inspired by the International Cartographic Association Working Group on Map Generalization led to special issues of both *GeoInformatica* [Geo 1999] and *Cartography and Geographic Information Science* [CaG 1999] addressing the topic. While much of this generalization work usually applies to vector data, it seems beneficial to notice that contours are implicitly represented in sampled elevation data. Applying simplification operations in this raster domain will naturally simplify the implicit contours, and trivially permits contours to undergo radical topological change, which is cumbersome at best when operating with vector data.

Bump mapping [Blinn 1978] is widely used for automating production of shaded relief from elevation data. For example, general purpose raster editing tools like Photoshop can be used to manipulate elevation data and then render it as a bump map. Tom Patterson has written a number of articles and tutorials describing the use of this process as used at the U.S. National Park Service [Patterson]. Many GIS software packages also provide bump mapped hill shading from elevation data. A recent survey of desktop software for relief shading [Dorn 2002] seems to suggest that the common packages are all performing bump mapping substantially as Blinn described it, although there are a few customizations specific to the needs of shaded relief.

One problem with straightforward application of bump mapping to elevation data is terracing due to quantization in the data, and gaps where data is missing. MacDEM attempts to address both problems by filtering with a 3×3 box filter. MacDEM also includes a "low elevation haze filter", which attenuates shading variation at low elevations [Dorn 2002].

In addition to bump-mapped shaded relief, MicroDEM provides an analytical display of slope by coloring gradient direction according to a quantized color wheel, as shown in Figure 6 [Dorn 2002].

We believe there is still considerable room for refinement of these techniques. In particular, improving filtering of elevation data and enhancing bump mapping with nonphotorealistic shading models for feature emphasis.

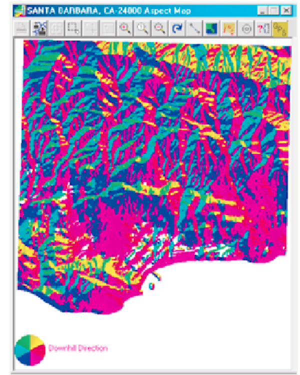

Figure 6: MicroDEM can produce a nonphotorealistic diagram of downhill direction quantized to a six-hued color wheel.

3 Filtering Elevation Data

Some care is required when filtering elevation data to address existing sampling artifacts and avoid introducing new ones. The problem of terracing due to quantization of elevation levels has already been mentioned. It manifests itself as stairstep shading artifacts in regions with mild slope, as shown in Figure 7.

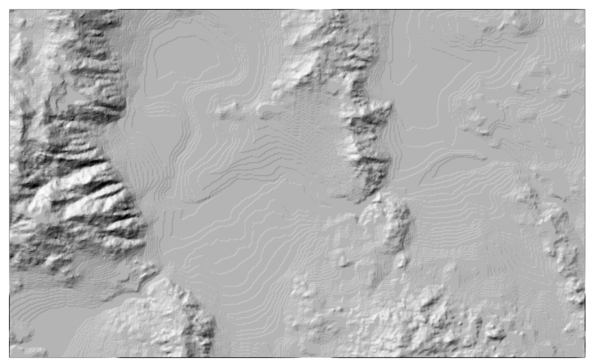

Figure 7: Quantization of elevation data leads to terracing artifacts in the bump mapped image.

Also note that the use of bilinear interpolation for magnification of elevation data will lead to another shading artifact. The elevation detail it produces will be piecewise linear, hence having constant slope over each region of interpolation. Since bump-mapping depends only upon local elevation gradients, the resulting hill shading would not benefit from the interpolation at all. Higher order filtering is required to provide reconstruction that is continuous in first derivative. This suggests that smoothness is a desirable property for filtering operations on elevation data. Since there is no evident reason why two terrains that are equivalent up to rotation should be simplified differently, radial symmetry seems like a desirable trait as well. Accordingly, we will use a Gaussian kernel as our primary filter, since it provides these properties. It is also separable, which is convenient for implementation.

To provide varying levels of detail across the map (and potentially even use a different level of terrain detail for contouring than for hill shading), we want to be able to vary the width of the the filter kernel applied. In the interest of allowing rapid interactive simplification of large terrain, we don't attempt to perform this filtering with large, variable-size kernels on the fly. Instead, we precompute a series of progressively smoother versions of the terrain by filtering with linearly increasing kernel widths. This is similar to a Gaussian pyramid, except that the filter radius increases linearly rather than exponentially and we do not decimate the pyramid levels. By linearly interpolating between successive levels of smoothing, we can rapidly provide a desired degree of simplification at any point of the terrain, and have the potential to adjust it interactively. Figure 16 shows a series of shaded images of synthetic terrain at increasing levels of detail. The same level of detail is used uniformly throughout each image, for both shading and contours.

Note that in the highest level of detail, the bathymetric contours have increased in detail along with the rest of the image. As suggested by Figure 4, this level of detail is probably excessive and inappropriate for bathymetric contours. However, by locally adjusting the smoothness level based on a mask of sea regions like that shown in Figure 8, we can reduce the detail in the bathymetric contours to a more reasonable level, as depicted in Figure 1. This process is fully automatic.

Figure 8: Detail mask for sea regions.

4 Shading

We determine the base color for a map pixel based on the hypsometric tint associated with its elevation in a lookup table. The tint scheme applied is roughly based on an example from Imhof [1982], and can easily be modified by changing the lookup table. This base color is multiplied by a shading coefficient determined using bump-mapped hill shading.

As in Schlag [1994], we estimate the partial derivatives of elevation in x and y by convolving the elevation map with the masks:

$$i_x = \begin{pmatrix} -1 & 0 & 1 \\ -1 & 0 & 1 \\ -1 & 0 & 1 \end{pmatrix}$$

and

$$i_y = \begin{pmatrix} 1 & 1 & 1 \\ 0 & 0 & 0 \\ -1 & -1 & -1 \end{pmatrix}$$

And accordingly estimate the normal as $N = (-i_x, -i_y, 1)$ (although we will first allow an arbitrary scaling of the z component so the user can vary the apparent depth of the scene). A typical bump mapping scheme would now apply lambertian diffuse shading with a large ambient term. Following the convention of Gooch and Gooch [1999], we can illustrate a shading model by showing how it appears on a lit sphere. Assuming oblique lighting incident from the northwest (in keeping with cartographic tradition), basic diffuse+ambient shading would yield an effect like that shown in the upper left of Figure 9.

If we have moderately smoothed the terrain as in our medium-scale example, this shading model may cause the softened terrain to appear overly blurry or indistinct, with a very gradual transition from light to shadow across the top of a ridge. To accentuate the shading variations, we can sharpen the diffuse term by applying a sinusoid to smoothly exaggerate the shading variations. This shading model is illustrated in the upper right of Figure 9. The "sharpened-diffuse" model is used for Figures 1 and 16.

If we magnify a small region of the elevation data with a smooth interpolation, we may find that the undersampling of elevation gives insufficient shading detail. See Figure 11. Ideally, we would like small scale maps to exhibit the same crisp shading seen in Figure 10 [Imhof 1982].

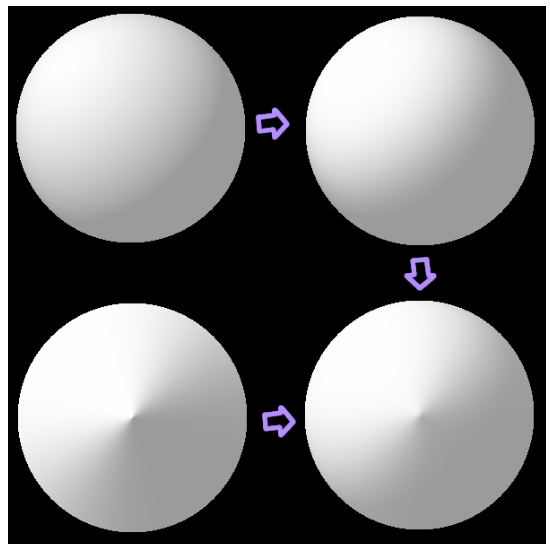

Figure 9: Four shading models, illustrated on shaded spheres. Top row: Diffuse + ambient, Sharpened diffuse + ambient. Bottom Row: Radial gradient, Hybrid (Ridge shading).

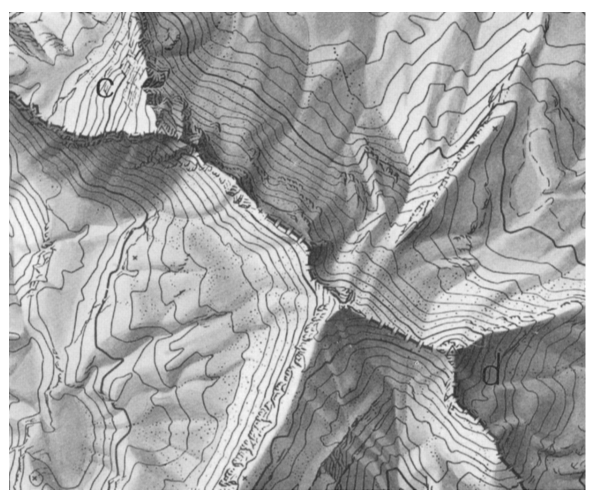

Figure 10: Example of relief shading at small scale. From Imhof. Notice the sharp transition from light to shadow across ridge lines.

Figure 11: Magnified region of map, with sharpened diffuse shading.

Figure 12: Magnified region of map, with ridge shading.

To maximally emphasize ridge lines, we could adopt the radial gradient shading function (shown in the lower left of Figure 9). This shading function is discontinuous across the upward-facing normal, where it has a singularity. This ensures an abrupt transition in shading across ridge lines. However, having a shading continuity about the upward facing normal has unappealing side effects – it introduces jarring, high-contrast shading transitions in gently rolling flat terrain. Overall, the radial gradient deviates far enough from familiar lighting models that the hill shading it produces is no longer readily perceived as a shaded relief surface.

By smoothing the discontinuity about the upward normal and forming a weighted blend with the sharpened diffuse shading function, however, we produce a hybrid shading function that greatly enhances ridge lines, while still resembling a shaded surface. The effect of this shading model on the magnified terrain is shown in Figure 12. Note that this shading model sacrifices detail and contrast on the slopes in order to emphasize the ridge lines. This makes it well suited to small-scale relief maps where the terrain varies gradually. On medium-scale relief, the contrast enhancement provided by this method would be excessive, and the sharpened diffuse model remains an appropriate choice.

5 Contours

All the images we have shown from our system include contours. These are readily extracted from the elevation data, essentially by "marching squares". In the current implementation there is a one-to-one mapping from elevation samples to pixels, and contours are drawn by coloring pixels adjacent to an elevation crossing. Coloring all pixels adjoining a contour equivalently would produce an objectionable aliasing artifact, as 45 degree lines would no longer appear to have the same thickness as axis-aligned lines – see Figure 13B. To address this, we use an efficient "approximate antialiasing" technique driven by a lookup table. There are 16 possible configurations of edge-crossings between a pixel and its four nearest neighbors. Ignoring reflections and rotations, there are five distinct configurations, as shown in Figure 14. The contribution to a pixel from nearby contour lines is weighted according to this local edge-crossing configuration. The relative weights of the edge and corner cases are chosen so that lines of varying slope are rendered with the same apparent width (Figure 13C).

Note that these implicitly represented contours behave reasonably when the elevation data undergoes smoothing, simplification, and interpolation, without requiring any attention to topological change as in vector-based schemes. Of course, the implicit contours could still be easily extracted into a vector form for applications that require it. Figure 15 shows the magnified terrain region with hypsometric contours included.

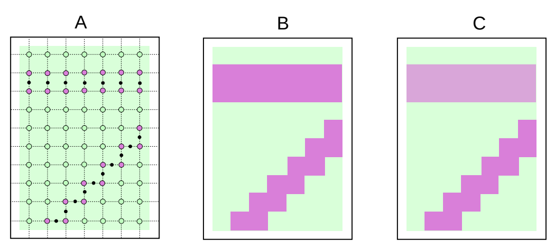

Figure 13: Lines drawn on a pixel grid. A) Grid of elevation samples. Black dots mark the edges between neighboring samples through which a contour passes. B) Weighting all pixels adjoining the line equally makes produces lines that seem to vary in thickness depending on their slope. C) Approximate antialiasing based on local edge-crossing configurations avoids this.

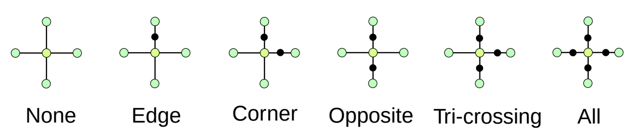

Figure 14: Possible edge-crossing configurations about a pixel.

Figure 15: Magnified region of map, with ridge shading and contours.

6 Summary and Future Work

We have suggested that simplification operations can be expressed on raster elevation data more conveniently than on vector data. Additionally, we suggest a simple method for precomputing terrain at progressive levels of simplification, which could permit local levels of detail to be efficiently manipulated interactively. To better emphasize ridge lines in a shaded relief, we introduce a nonphotorealistic shading model that depends on downhill direction. Finally, we use a simple image-based technique for drawing contour lines that preserves apparent line width without the expense of full antialiasing.

There are a number of interesting directions for further work. It is likely that other filters besides the Gaussian might be desirable for simplification. A filter that selectively preserved some sharp features, like sharp-angled corners, would be very useful. Also, the use of nonphotorealistic shading techniques for emphasizing shape features should be further explored.

References

BLINN, J. F. 1978. Simulation of wrinkled surfaces. In *Computer Graphics (Proceedings of SIGGRAPH 78)*, ACM Press / ACM SIGGRAPH, Computer Graphics Proceedings, Annual Conference Series, ACM, 286–292.

1999. Cartography and Geographic Information Science, vol. (26)1, American Congress on Surveying and Mapping. Special Issue on Computational Processes for Map Generalization.

DORN, D. V., 2002. Evaluation of desktop shaded relief software. http://www.reliefshading.com/software/index.html, May. Presented at Workshop of the ICA Commission on Mountain Cartography.

1999. GeoInformatica, vol. (2)4, Kluwer Academic Publishers. Special Issue on Computational Processes for Map Generalization.

GOOCH, A. A., AND GOOCH, B., 1999. Using non-photorealistic rendering to communicate shape. ACM SIGGRAPH 1999 Course Notes, Course on Non-Photorealistic Rendering., July.

IMHOF, E. 1982. *Cartographic Relief Presentation*. Walter de Gruyter & Co.

PATTERSON, T. Nacis shaded relief homepage. http://www.nacis.org/cp/cp28/resources.html.

ROBINSON, A. H., MORRISON, J. L., MUEHRCKE, P. C., KIMERLING, A. J., AND GUPTILL, S. C. 1995. *Elements of Cartography*. John Wiley & Sons.

SCHLAG, J. 1994. Fast embossing effects on raster image data. In *Graphics Gems IV*. AP Professional, 433–437.

Figure 16: Progressively higher level of details for a synthetic terrain. Each image uses a consistent level of detail throughout, for both contours and shading.