Exercises for

Chapter 10: Volume Visualization

1 EXERCISE 1

Volume visualization is a technique that generates 2D images of 3D scalar vector fields by considering the entire information present in such a field, in contrast to *e.g.* slicing and iso-surfacing techniques, which first select a *subset* of the field data points, and then generate the final image from such a subset. Along these lines, discuss the possibility to generate volume visualizations of *vector fields*. Which are the classical vector-visualization techniques (or elements thereof) which you would keep unchanged, and why? Which are the differences in your solution due to the fact you have a vector, rather than scalar, field input?

Hints: Consider first which of the existing 2D vector visualization techniques you would like to extend to 3D, *e.g.*, oriented glyphs, stream objects, direction color-coding, or image-based techniques.

2 EXERCISE 2

Most volume rendering techniques have two components: a so-called *ray function* and a separate *classification* function. Describe these two components, and also the way they interact with each other to yield the final volume-rendered image.

3 EXERCISE 3

Describe the difference, in terms of observed results, between using volume rendering with the maximum intensity projection (MIP) and the average intensity projection. Which of these two functions is analogous to an X-ray image of the input dataset, and under which conditions?

4 EXERCISE 4

Volumetric lighting, or shading, is a technique that brings additional realism to a volume-rendered image, by adding effects such as highlights on the apparent surfaces in the image. Can volumetric lighting be used with all types of ray functions? If so, argue why. If not, describe which ray functions admit volumetric shading.

5 EXERCISE 5

Volumetric shading requires the computation of an apparent surface normal at each voxel in the dataset which is visited by a ray. However, there are no explicit surfaces present in our volume. Describe how, given a scalar volume, this apparent surface normal is evaluated.

6 EXERCISE 6

An efficient acceleration strategy for image-order volume rendering is the *early ray termination* (ERT). Describe this strategy. Can ERT be applied for both back-to-front and front-to-back image-order ray casting? Can this technique be applied in conjunction with any ray function (MIP, average intensity, distance to value, isosurface, compositing)? If so, argue why. If not, explain for which ray functions ERT can be applied and for which not, and why.

7 EXERCISE 7

Two main interaction types exist between the interpolation of scalar values along a ray in volume rendering, and the application of transfer functions, or classification (scalar to color mapping): *preclassification* first classifies and next interpolates; *postclassification* first interpolates and then classifies. Which technique is, under no specific assumptions about the dataset and transfer functions, better?

8 EXERCISE 8

Object-order techniques are a variant of volume rendering. In contrast to image-order techniques which cast a ray through each pixel and accumulate values as the ray traverses the volume, object-order techniques pre-slice the volume using a set of parallel planes to create 2D textures, and create the final result by blending these textures. Consider now an object-ordering volume rendering of a dataset, where the view vector is parallel to one of the coordinate axes of the dataset. Would this rendering exhibit any decreased quality as compared to an image-order rendering of the same dataset and from the same viewpoint?

9 EXERCISE 9

Imagine an experimental object-order volume rendering technique where, instead of using planar textured slices, we would render the scalar volume data by extracting a high number of densely-spaced isosurfaces, and render them, one isosurface at a time, with suitable transfer functions mapping their values to color and/or transparency. Would this technique render qualitatively better results than standard object-order or image order techniques? If so, argue why. If not, explain why not.

10 EXERCISE 10

Consider a 3D scalar volume sampled on an uniform grid of N^3 cubic cells. To generate a visual representation of a contour surface from this volume, we can use either isosurfaces, such as implemented by the marching cubes algorithm, or volume rendering, using ray casting with an isosurface ray function. Let us assume that we want to render a single contour surface (for a single scalar value), and this surface is to be rendered fully opaque. Which technique has a lower worst-case complexity as a function of the number N of grid cells? Separately, which technique needs to visit, on average, fewer voxels?

11 EXERCISE 11

A user runs a ray-casting volume renderer to visualize an isosurface of some scalar field, using the isosurface ray function. By turning the ray casted volume around its center c, the user next records a very large number of N snapshots of the images I_i, $1 \leq i \leq N$, shown by the volume renderer, each snapshot being taken from a different viewpoint. For all N viewpoints v_i, the camera is placed at uniformly-spread positions over a sphere centered at c, looking towards c, and uses an orthographic projection and a fixed rendering-window size, large enough to see the entire isosurface. The color of the rendered isosurface in all views is different than the image background. The user next stores the set $V = \{I_i, v_i\}$.

Given V, describe an algorithm to reconstruct a binary volumetric representation of the visualized isosurface, *i.e.* a binary volume B whose foreground voxels are contained inside the rendered isosurface. Next, detail the types of isosurface shapes for which the reconstruction would fail from recovering the correct 3D shape of the isosurface.

End of Exercises for
Chapter 10: Volume Visualization