Lecture notes: Range image segmentation

Segmentation for a range image follows many of the same principles as segmentation for an intensity image. However, due to the nature of the data, there are some differences.

First, the problem tends to be more objective. The goal is usually to segment the image into regions that correspond to different surfaces. For example:

range image	reflectance image	segmentation

This is a simple shape, but the segmentation is fairly objective. Each region represents a different surface of the object or background. More complicated examples are still fairly objective:

range image	reflectance image	segmentation

Second, there are two types of edges in a range image: crease and jump. Crease edges occur where surfaces join at a common geometrical boundary. Jump edges occur where one surface occludes another in the background, causing a discontinuity in depth. Several examples of both types of edges can be seen in the above images. For example, in the top image of the anvil-shape, the front edges between the bottom of the anvil and the floor are crease, while the edges between the back of the top of the anvil and the back wall are jump.

Jump edges are easier to detect due to the discontinuity, and are therefore easier to find during segmentation. Crease edges are more difficult to detect and can be subtle depending on the orientation of the surfaces on either side of the edge.

Third, the region predicate can be based upon geometry rather than intensity features. Thus, instead of grouping pixels based upon intensity or texture values, pixels can be grouped based upon surface normal orientations:

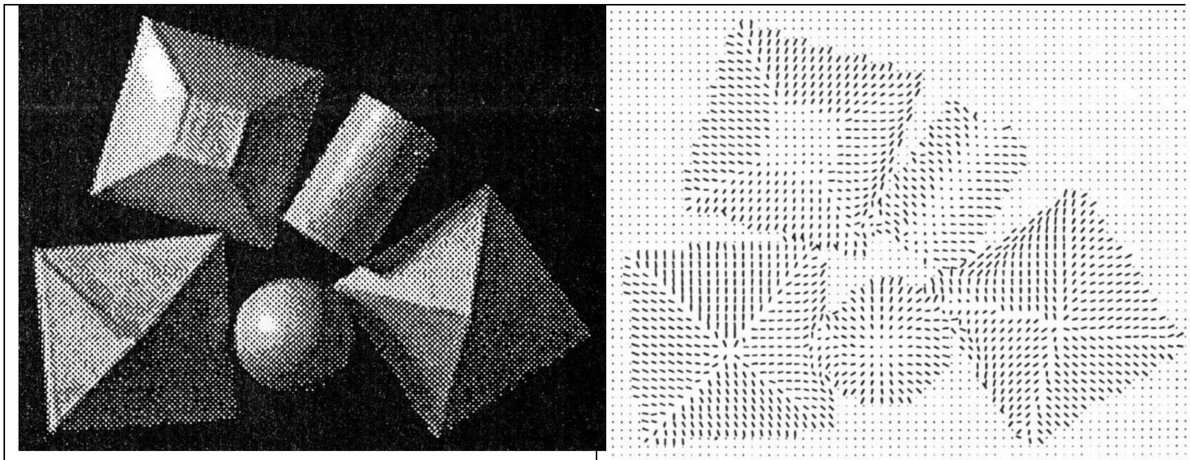

reflectance image	surface normal "needle map"

It is easy to see in this example that the surface normals share a common orientation for the same surface. Thus, the region predicate can be to require pixels in a region to have the same surface normal orientation.

This works for planar surfaces, but if the surfaces are curved, then the normals will vary across the region. This can be seen for the ball in the above example, and is further illustrated in this example:

range image	reflectance image	segmentation

In this case the segmentation is still fairly objective, but the surface normals must be grouped differently. The most common approach is to use differential geometry. See figure 4 in the Besl & Jain paper posted at the course website. The idea is to calculate changes in the surface normals across small areas in the image. These changes can be used to describe the curvature. Changes in curvature can then be used to identify region boundaries.

Curvature "type" can also be used. Figure 6 in the paper shows how different properties can be calculated for curvature (H and K, called the mean and Gaussian curvature) and used to describe different types of surfaces. Segmentation can be driven by an HK map that looks for different surface types (as opposed to thresholding changes in curvature value).

Fourth, the segmentation can be used to calculate a surface equation for each region so that the content can be described in 3D. Thus, the output is not an intensity blob; it is a 3D primitive.

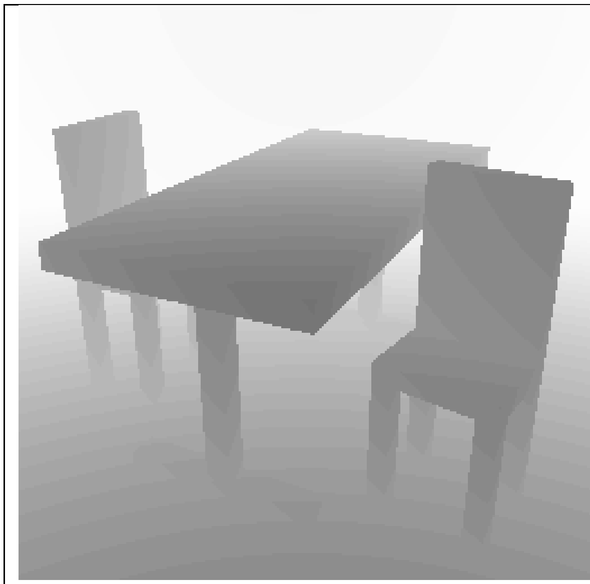

range image

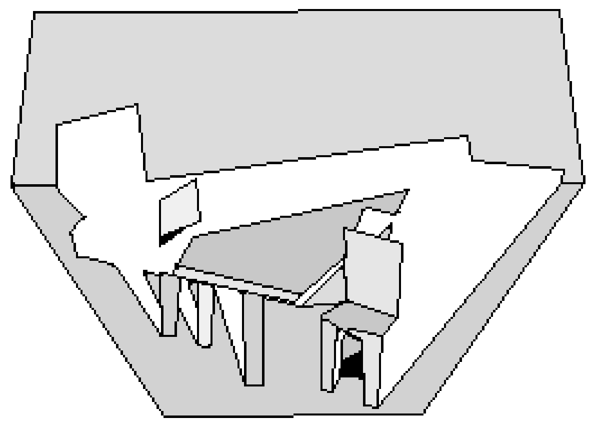

surface primitives

From a single image the object/scene geometry can only be partially recovered. The goal is to take multiple images from multiple viewpoints to fully recover the geometry and build a model:

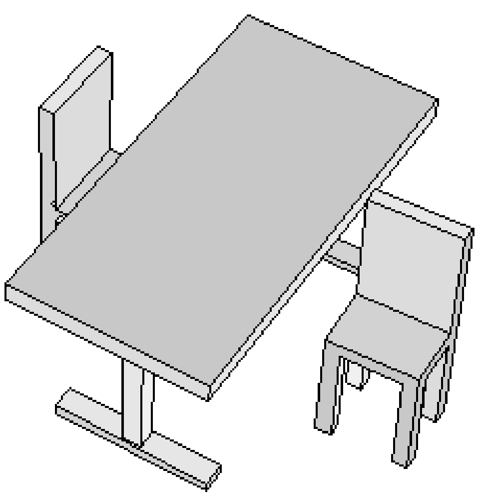

3D model (one viewpoint)

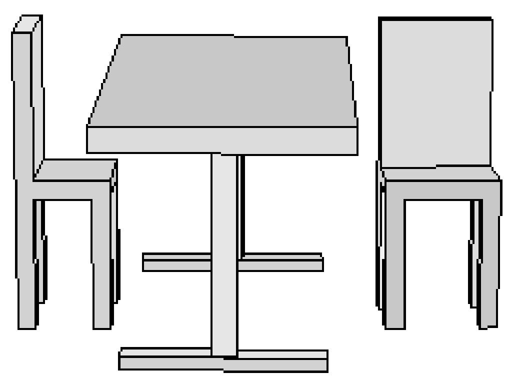

3D model (another viewpoint)

Surface normals can be calculated using a window of pixels and doing some fitting, or using the differential geometry techniques for curved surfaces. They can also be approximated simply by taking cross products:

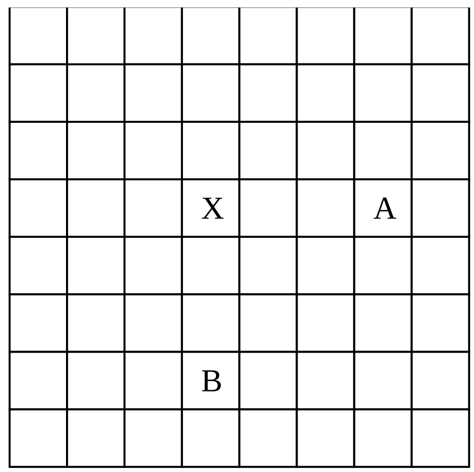

The surface normal at pixel X can be calculated by taking the cross product of (B-X) x (A-X), where A, B and X are the 3D coordinates of those pixels. The pixel distances B-X and A-X must be large enough to be robust to noise and quantization, but small enough to capture the local geometry. This example depicts a distance of 3 pixels.

The fitting of surfaces to 3D points is beyond the scope of this lecture, but it is possible easy to envision fitting surfaces of different types (e.g. planes, cones, splines) to clouds of 3D points defined by segmented regions in a range image. Models can then be finalized from the fitted surfaces.