cse 252c Fall 2004 Project Report: A Model of Perpendicular Texture for Determining Surface Geometry

Steven Scher

December 2, 2004

Steven Scher
SteveScher@alumni.princeton.edu

Abstract

Three-dimensional objects distributed over a surface occlude each other in a manner dependent upon the angle (sigma) between the surface normal and the viewer. Leung and Malik note that occlusion allows the top of the objects to be seen, but not the bottom. For the case where the objects are differently colored as a function of height, the color content of the scene will thus depend upon sigma. A model is described that predicts the fraction of color expected to be present in each portion of the image, for a given value of sigma. This model can be compared to the actual color fractions in an image of unknown sigma, in order to furnish an error function, so that the sigma can be found using an optimization scheme. A set of synthetic images of cylinders on a plane are generated, and sigma is found to a mean accuracy of two degrees.

1 Project Description

Solution of the Shape from Texture problem infers the local surface orientation from local feature vectors extracted from the image. Much work on this problem focuses on repeating 2-D patterns, painted onto a surface. For example, 2-D texture elements (e.g. ovals) found in an image may be warped to match a template texture element (e.g. circles), and the warp used to recover the local surface normal.

A different approach assumes a parameterized distribution of 3-D texture elements over a surface, where the texture elements vary in color with height [Leung and Malik, On Perpendicular Texture, 1997]. For example, a field of flowers is modeled as a field of cylinders whose top is yellow, and bottom is green. This model is used to find the probability of a pixel being a given color, as a function of the slant (the angle between the surface tangent and the viewing

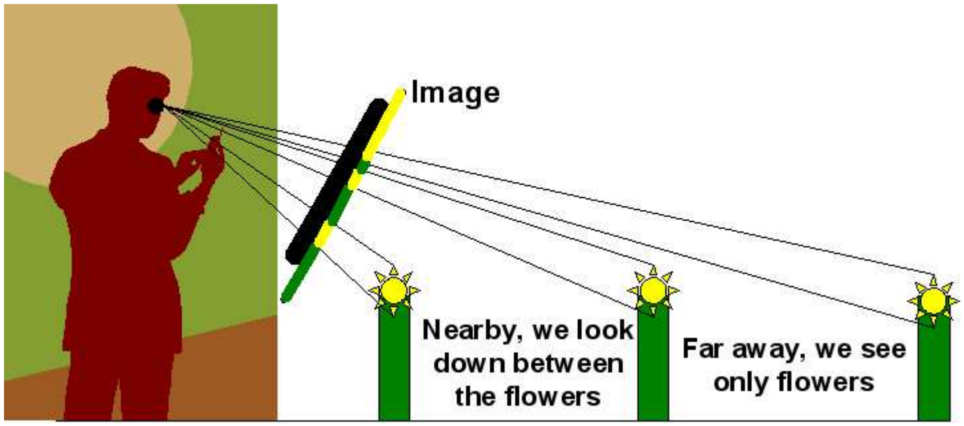

Figure 1: Far-off flowers stems are more likely to be occluded by other flowers.

direction) and of the parameters of the distribution. Significantly, the effects of perspective are incorporated into the calculation.

The observed proportion of a color in a region is then used to find the slant of the surface, if the parameters of the distribution of elements are known. If these parameters are not known, they can be found by minimizing a cost function that matches the observed frequencies of the image to those calculated by the parameters.

This project implemented the technique described above with synthetic images akin to those used in Leung and Malik. The camera was pitched down toward the ground at an unknown slant angle. Each image was separated into horizontal bands, and the fraction of each band occupied by flower, stem, and background was measured. A cost function was formed measuring the dissimilarity between the observed flower fractions in each band, and those predicted by a candidate slant angle. The unknown slant angle was found by minimizing the cost function, and was determined to an accuracy of a few degrees.

2 Model

The scene is modeled as a set of identical cylinders of height H and radius R, randomly distributed according to a Poisson process with intensity λ (the expected number of cylinders centered in a unit area is λ). The color of each cylinder from the bottom to height y differs from the color of the cylinders from y to the top.

Because the slant angle to the ground nearby differs from the angle to the ground far away, the stems of the far-off flowers are more likely to be occluded than the nearby stems, as in Figure 1.

The scene is viewed through a camera placed above the scene, with perspective projection. An example image is seen in Figure 2.

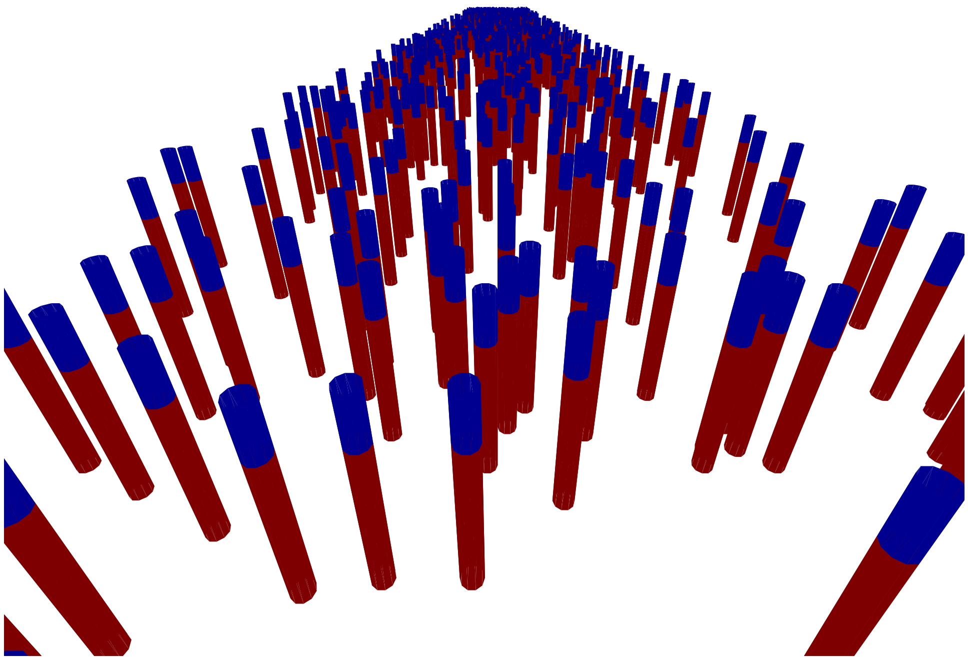

Figure 2: Synthetic Image

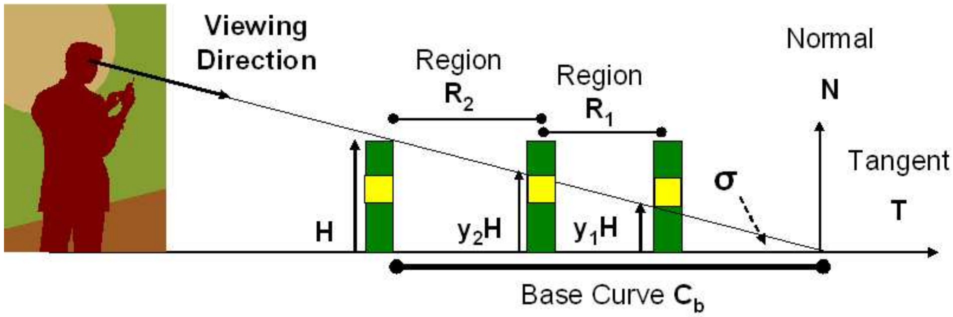

Figure 3: slant angle-dependent occlusion determines the likelihood of seeing the flower or the stem

3 Probability of a Given Color

Consider a cylinder painted the color of a stem everywhere, but painted the color of a flower from y_1 to y_2, as in Figure 3. The pixel along the line of sight will be flower-colored if (1) there is a flower to be seen, in $Region_1$, and (2) there is no flower to occlude it, in $Region_2$. Since we assumed a Poisson distribution, the probability that there is no flower in $Region_2$ is $exp(-\lambda * Area_2)$ and the probability that there does exist a flower in $Region_1$ is $1 - exp(-\lambda * Area_1)$.

Here, $Area_2$ is a rectangle whose width is the cylinders diameter, and whose length is proportional to y_2, the height of the top of the flower-colored region, and the tangent of the slant angle σ. $Area_1$ is similarly defined by the geometry. This gives:

$P_{Region_2 empty} = exp(-\lambda * 2R * H(1 - y_2) * tan(\sigma))$

$P_{Region_1 Flowerd} = 1 - exp(-\lambda * 2R * H(y2 - y1) * tan(\sigma))$

Thus, the probability that a given pixel will be flower-colored is dependent upon the slant angle. Averaged over a window, the probability of seeing the flowers color is approximated by the fraction of pixels having that color. Given the parameters of the cylinders and their distribution, the observed fraction of flower-colored pixels can reveal the slant angle of the ground with respect to the camera. With the cameras orientation known, the scenes topography is revealed.

4 Synthetic Images

Synthetic images were created by randomly populating a dense grid with identical cylinders, according to a Poisson Process. The cylinders have height H=20 and radius R=1, and fill a rectangular area 100 units wide and 1000 units deep. There are approximately 1000 cylinders, with λ=.01. The bottom (Stem) and top (Flower) of each cylinder are separated at y=.75*H.

The scene is viewed from a camera at the edge of the rectangular area. The camera has zero roll (the horizon is horizontal) and is pitched down so that the

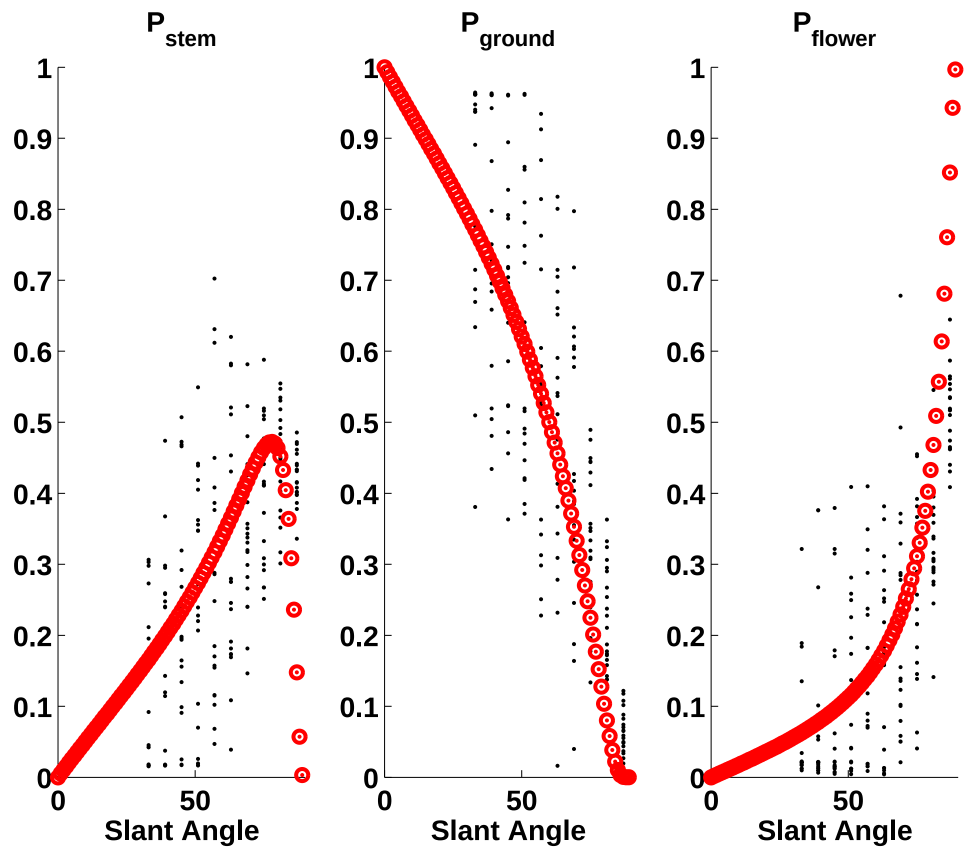

Figure 4: Probabilities based on true σ_0 provide a reasonable fit to to measured values. P_{flower} provides a better match thatn P_{stem} or P_{ground}

vector from the camera to the image center makes an angle σ_0 with the surface normal. Vectors from the camera to points above the image center make angles greater than σ_0, and those to points below the image center make angles less than σ_0. The angular field of view of the camera was set to 60 degrees.

The synthetic view is captured, and a narrow swatch is extracted, covering the full height of the image, but only the center 10% of the image width. This allowed a much smaller area to be populated with cylinders. The image is then divided into ten equally-sized horizontal bands. The colors are separated into the flower, stem, and background colors, and the fractional composition of each horizontal band is calculated: for the first band, F_{stem1}, $F_{flower1}$, $F_{ground1}$, for the second band, F_{stem2}, ...

The model provides a good description of the way the color content of the image varies with slant angle. Figure 4 shows the probabilities calculated for the true value of $sigma_0$ with the measured flower-color fractions for 10 images.

5 Cost Function

The accomplishment of Leung and Malik is the creation of a function $p_{flower}(\sigma)$ that predicts the probability of seeing a flower at a given pixel. Given a guess σ of the slant angle the ground makes with the cameras view direction, we can use P_{flower} to estimate the expected fractional area with the flowers color. The error between this estimate and the measured fraction $F_{flower1}$ creates an easily evaluated cost function.

The cost function chosen here is the sum over the 10 bands of the squared error between the fractional flower presence measured F_{flower} and the probability P_{flower} predicted by the slant angle. Since the slant angle varies from the bottom of the image to the top, P_{flower} is evaluated at a different σ for each band. With the perspective projection, the slant angle varies linearly with image height from its value σ_0 at the image center.

The cost function is minimized by a standard nonlinear search mechanism to find the best estimate of the slant angle σ_0. As seen in the example case in figure 4, the value of 58 degrees σ_0 found by the search mechanism provides a good match to the measured fraction of flower color in each band of the image.

6 Discussion

The mean error in estimating σ_0 was 2 degrees, with a 2 degree standard deviation of the estimates across multiple images formed with the same true σ_0. The result was generally robust to changes of σ_0 from 45 degrees to 70 degrees, and to changes in the cameras angular field of view between 30 degrees and 90 degrees.

However, the choice of the field of view and σ_0 must include some far-off objects (σ at least 60 degrees) in order to accurately estimate σ_0. The dense far-off object may provide the most consistent color ratios.

The image resolution could be reduced substantially without significantly affecting the ability to determine σ_0. The algorithm was used with identical images at full resolution, and with images decimated (using nearest-neighbor decimation) by factors of 4 to 64. These are shown in the figures below.

No significant loss of accuracy was apparent with decimation from the original (631x103 pixel) image up to 32-times (a 19x3 pixel image), at which point the structure of the objects is still visible. Suddenly, at 64-times decimation (a 9x1 pixel image), the algorithm is completely unable to determine the angle (errors of 30 degrees).

Since low resolution may still yield a viable estimation of surface geometry, application of the algorithm to small windows of the image may prove fruitful, allowing a more generalized surface to be discovered.

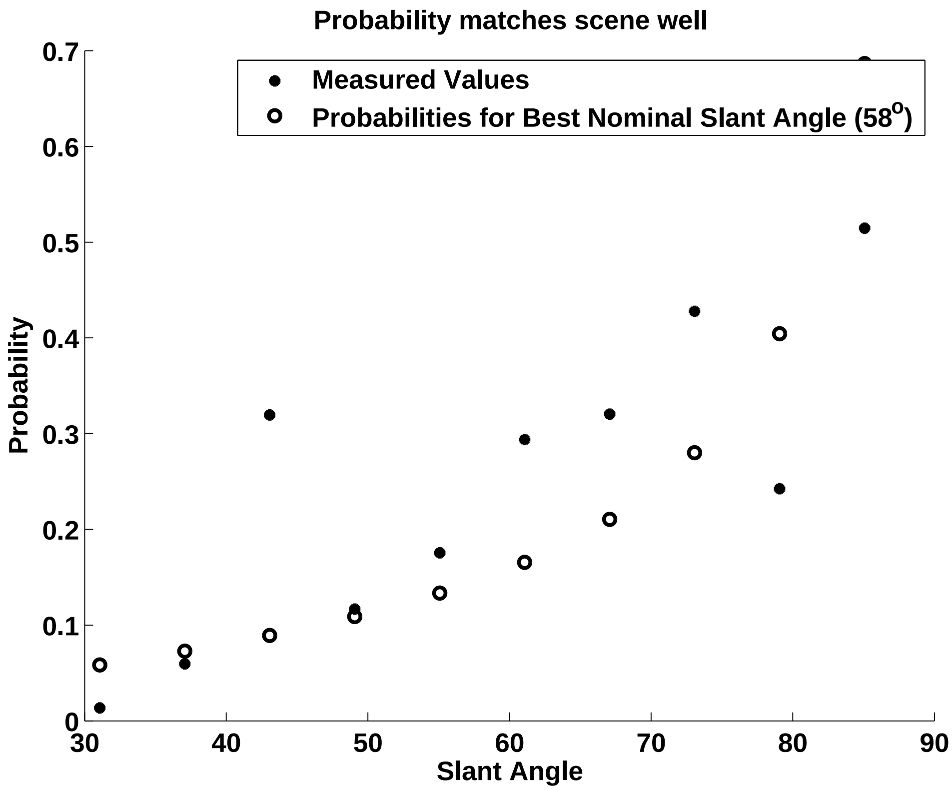

Figure 5: Estimated value of σ_0 (58 deg) is close to actual value (60 deg). Probabilities based on estimated σ_0 are close to measured values

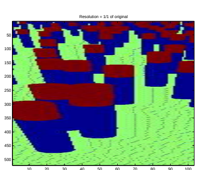

Figure 6: Full-Resolution Image

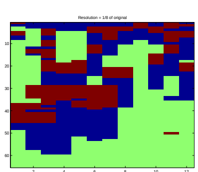

Figure 7: 1/8 Resolution Image

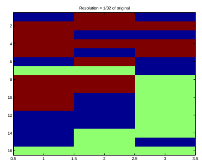

Figure 8: 1/32 Resolution Image

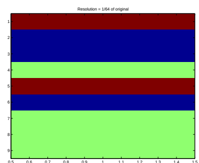

Figure 9: 1/64 Resolution Image

7 Future Work

Rather than assuming a single color for each region, each region could have its color drawn from a different distribution of colors. This would require averaging over an additional layer of probabilities.

Since the features used (e.g. yellowness of flower tops) are highly specific to the training images, new features are needed for a new environment. When moving to real scenes, allowing each region of the cylinder to be colored according to a distribution would facilitate determination of appropriate feature vectors for the cylinder regions, as these distributions could be calculated from example images of known geometry.

When higher resolution than needed is available, application of this method to many overlapping windows in an image would generate surface normal measurements on a grid across the image. This may be beneficial for more complex scenes.

8 Principal Reference

Leung and Malik

On Perpendicular Texture: Why do we see more flowers in the distance?

IEEE Conf. Computer Vision and Pattern Recognition, June 1997, Puerto Rico