

Reverse Engineering of Pipe Layouts and 3D Point Set Damage Models

David Gallup and Thomas C. Henderson

UUCS-04-020

School of Computing
University of Utah
Salt Lake City, UT 84112 USA

27 December 2004

Abstract

This paper focuses on obtaining pipe layout information from depth images, and on the development of damage models for 3D shapes. Techniques are given for recognizing various pipes and pipe features. Methods for generating compact geometric descriptions of pipes and pipe features are discussed. These methods and techniques are analyzed for error, and error measurements are given. 3D shape damage models are proposed based on a separable Gaussian kernel model. Experiments on synthetic and real data were performed and results discussed.

1 Introduction

Surface curvature analysis is a powerful method for identifying pipes and pipe features in range images. Straight pipes, for example, are equivalent to cylinders which can be effectively identified by their surface curvature. Surface curvature for depth images can be calculated using Monge patch formulas. Once the depth image has been segmented into pipes and pipe features, additional work must be done to give a compact geometric description of these objects. For example, rather than describe a straight pipe as a region of surface points exhibiting cylindrical surface curvature, the pipe could be described as a line segment with a radius.

Surface curvature can be calculated using the Monge patch formulas [1]. (Also see [2, 3].) Let the Monge patch, or depth image, be $h(u, v)$. The surface normals over the image can be computed as:

$$n(x, y) = \frac{1}{\sqrt{1 + \frac{\partial z}{\partial x}^2 + \frac{\partial z}{\partial y}^2}} \left(-\frac{\partial z}{\partial x}, -\frac{\partial z}{\partial y}, 1 \right)$$

The Gaussian curvature K over the image is:

$$K(x, y) = \frac{\frac{\partial^2 z}{\partial x^2} \cdot \frac{\partial^2 z}{\partial y^2} - \left(\frac{\partial^2 z}{\partial x \partial y} \right)^2}{\left(1 + \left(\frac{\partial z}{\partial x} \right)^2 + \left(\frac{\partial z}{\partial y} \right)^2 \right)^2}$$

and the mean curvature H over the image is:

$$H(x, y) = \frac{\frac{\partial^2 z}{\partial x^2} + \frac{\partial^2 z}{\partial y^2} + \frac{\partial^2 z}{\partial x^2} \cdot \left(\frac{\partial z}{\partial y} \right)^2 + \frac{\partial^2 z}{\partial y^2} \cdot \left(\frac{\partial z}{\partial x} \right)^2 - 2 \frac{\partial z}{\partial x} \cdot \frac{\partial z}{\partial y} \cdot \frac{\partial^2 z}{\partial x \partial y}}{2 \left(1 + \left(\frac{\partial z}{\partial x} \right)^2 + \left(\frac{\partial z}{\partial y} \right)^2 \right)^{\frac{3}{2}}}$$

The principle curvatures k_1 and k_2 can then be calculated using K and H :

$$k_1 = H + (H^2 - K)^{\frac{1}{2}}$$

$$k_2 = H - (H^2 - K)^{\frac{1}{2}}$$

Once surface curvature has been calculated, cylinder surface regions can be identified as regions where $K = 0$ and $H < 0$. The next step is to fit cylinders to these regions.

We propose a method of simplifying the problem of fitting a cylinder to a surface. Having already computed the principle curvatures and directions, we observe that one principle curvature must be zero, and its corresponding direction is the orientation of the cylinder. The reciprocal of the non-zero curvature is the negative radius of the cylinder. Cylindrical surface regions can now be segmented by orientation, radius, and position. Furthermore, we observe that a surface normal scaled by its negative radius can be added to its surface point to produce a new point along the axis of the cylinder. In other words, the set of $P' = P + \frac{1}{k}N$ is a set of points along the cylindrical axis. Now the problem of fitting a cylinder to the surface reduces to the problem of fitting a line to the axis points.

2 Data

We have created some synthetic data (see Figures 1 through 4). The figures show the given range image, followed by the surface recovered from the image.




Figure 1: Synthetic Range Image 1.




Figure 2: Synthetic Range Image 2.



Figure 3: Synthetic Range Image 3.




Figure 4: Synthetic Range Image 4.

2.1 Measures

To measure the results of surface normal computation we can compare the ideal surface normals obtained during the dataset generation with the normals computed with the Monge patch formulas. Error can be measured by taking the dot product of the computed normal and the ideal normal and subtracting it from one.

Measuring segmentation and shape identification can also be done by comparison with the ideal values obtained from the dataset. The error can be measured as the percentage of incorrectly segmented or identified pixels.

The error in cylinder fitting can be measured by comparing the cylinder with the surface points used to fit the cylinder. For each point, the error is the absolute value of the distance of that point to the cylindrical axis minus the cylindrical radius.

3 Experiments

3.1 Surface Normals

Surface normals have been computed in a synthetic image as a baseline, and Figure 5 shows the error (brighter is higher error) in surface normal computation. It should be clear that there is a high degree of error at occlusion edges of objects. This is probably due to the fact

that these normals are generated erroneously using height values from neighboring objects. This problem could be eliminated by detecting discontinuities in the surface using some threshold and then ignoring these values.




Figure 5: Error in normal computation. White areas show high error.

3.2 Noise

The effectiveness of the system is affected by the amount of noise in the data. Typically this noise is modeled as additive Gaussian noise. Figures 6 and 7 show the effects of noise on shape identification and cylinder fitting.

Notice that while the performance of shape identification degrades quickly, the cylinder (line) fitting process does not. This is because most of the incorrectly identified points are barely cylindrical and nearly flat, which means their radius of curvature is near infinite. Therefore their computed cylindrical axis points fall outside the image or at least are scattered enough to have little effect on the line fitting.

3.3 Real Range Data

Figure 8 shows the intensity data from a range scanner. Corresponding range point datasets are shown in Figure 9.

Fitting cylinders to the splitting pipe data gave good results (see Figure 10). However,




Figure 6: The left image shows correctly identified pixels in white. The right image shows the computed cylindrical axis points. The noise is Gaussian noise with sigma of 1.0. Shape identification error is 6.18%. The mean and max error of the fitted cylinder are 0.0036 and 0.6529.

fitting to the cable roll data proved to be more difficult. Although the orientation of the fitted cylinder is correct, the position is not. We believe this is due to the asymmetric error in shape identification (see Figure 11).

4 Damage Models

Given the ideal nature of our pipe layout model in terms of cylinders, cones, planes, etc., it is useful to have a high-level definition of error in the shape where such error is caused by physical processes acting on actual surfaces. We therefore introduce the following kinds of damage to ideal shapes:

- dents: this type has several subtypes:
 - scratch
 - warp
 - gouge
- holes: this is anything that changes the genus of the shape.

Our goal here is to propose 3D models for these, as well as algorithms to recover them from actual range datasets.




Figure 7: The noise is Gaussian noise with sigma of 5.0. Shape identification error is 21.37%. The mean and max error of the fitted cylinder are 0.0274 and 0.7032.



Figure 8: Intensity Images from Range Scanner.

4.1 Dents

A dent is a deformation of a surface like that shown in Figure 12. A dented shape can be characterized as a regular pipe shape with some points missing. The shape of the missing point set is much like a 2-D Gaussian surface, and therefore, we model it as such.

A dent is modeled as a 2D Gaussian separable filter (see Figure 13):

$$SG(x, y | \sigma_1, \sigma_2) = \left(\frac{1}{\sqrt{2\pi}\sigma_1} \exp\left(-\frac{x^2}{2\sigma_1^2}\right) \right) \times \left(\frac{1}{\sqrt{2\pi}\sigma_2} \exp\left(-\frac{y^2}{2\sigma_2^2}\right) \right)$$

To introduce a dent into a shape, the location and orientation of the dent must be specified,




Figure 9: 3D Data Points from Range Scanner.




Figure 10: The results of shape identification and cylinder fitting for the dataset from Figure 8.

and then a set of points in the shape are generated that satisfy:

1. the points are inside the analytically defined shape, and
2. the points are not under the surface of the Gaussian patch generated with a frame attached at the specified surface point and z-axis into the shape.

More formally, let $SG(x, y | \sigma_1, \sigma_2)$ be the 2D Gaussian function, S be the shape in question, and let \bar{p} be a point on the surface of S and \bar{n} be the normal to S at \bar{p} . Then let T be the transform which moves \bar{p} to the origin and aligns $-\bar{n}$ with the z -axis. The set of points in the damaged model, D , is defined as:

$$D = S \cap \bar{q}$$

where $\bar{q} \equiv \{x, y, z\} \in \mathfrak{R} \ni z > SG(x, y | \sigma_1, \sigma_2)$.

An example is demonstrated by cross sections through a cylinder with a gouge like region removed (see Figure 14).




Figure 11: The results of shape identification and cylinder fitting for the dataset from Figure 8.




Figure 12: Example of dented surface.

Holes are modeled as negative cylindrical shapes, but will not be further considered here.

4.2 Experiments

We have developed a gradient descent model to fit the best separable Gaussian function to the particular data (using the Matlab *fminsearch* function). We provide some results here on synthetic and actual range data.

The complete parameterization of a separable Gaussian involves many parameters such as position, orientation, and magnitude. To facilitate the gradient descent model, these parameters are predetermined from the data, leaving only σ_1 and σ_2 .

Figure 15 shows the result of fitting to a dent that was generated synthetically with a separable Gaussian. As expected, the fit is quite close. In Figure 16, a separable Gaussian is fit to a box shape. The fit is reasonably close, thus demonstrating the flexibility of the separable Gaussian model. (Error in the fit is defined as the mean squared error of the data points from the separable Gaussian function surface.)




Figure 13: Dents as Separable Gaussian; symmetric dents (upper left, lower right) and linear dents (upper right and lower left).

Figure 17 shows the fit of a separable Gaussian to the dented tin can data from Figure 9. Figures 18 and 19 show the distribution of the residual error in all the surface points' distance from the recovered cylinder axis (left side of figure), and just those points used to determine the axis (right side of figure) in the pipe and cylinder datasets, respectively.

5 Conclusions and Future Work

We have proposed a framework for the specification of pipe layout dataset; this includes a method to synthesize data as well as methods to recover pipe layouts from Monge patch images. Furthermore, we have proposed a set of damage models for 3D shapes and methods for their recovery from range datasets as well. Experimental results are reasonable with tolerable error.

In the future, we intend to extend this work in two main directions:

- **Hole Models:** methods will be developed to model and recover through holes in pipe layout data.




Figure 14: Gouged Surface.

- **Physics-based Models:** the separable Gaussian dent model provides a good first approximation to damage in a surface, but we would like to explore models based on a particle set representation of a solid. Particle interaction physics will be defined (e.g., spring or potential functions), and then an external force applied and the solid shape deformed according to the model.

We believe that physics-based models offer many advantages, including potential insight into how the damage occurred. For example, it may be possible to determine the direction and magnitude of the force (e.g., this might help decide the direction from which a projectile arrived at the surface).

References

- [1] D. Forsyth and J. Ponce. *Computer Vision*. Prentice Hall, New Jersey, 2003.




Figure 15: The left image is the surface of a synthetic dented cylinder. The dent shape is Gaussian. The right image is a cross section of the dent fitted with a separable Gaussian. The fitting error is 0.0971.




Figure 16: This cylinder has a box shaped dent. The cross section shows the best fit of a separable Gaussian. The error is 2.5064.




Figure 17: A cross section of the dented tin (see Figure 9) fitted with a separable Gaussian. The error is 2.2295.




Figure 18: The distribution of error in distance from the recovered cylinder axis in the pipe dataset for (a) all surface points (left side), and (b) just those points used to determine the cylinder axis (right side). Moreover, the error mean, standard deviation and max errors are (a) 0.3721, 0.3249, 0.9938, and (b) 0.5727, 0.7456, 2.7268




Figure 19: The distribution of error in distance from the recovered cylinder axis in the cylinder dataset for (a) all surface points (left side), and just those points used to determine the cylinder axis (right side). Moreover, the error mean, standard deviation and max errors are (a) 2.5996, 2.0997, 7.7492, and (b) 0.1454, 0.2133, 1.031

- [2] M. Suk and S. M. Bhandarkar. *Three-Dimensional Object Recognition from Range Images*. Springer-Verlag, Tokyo, 1992.
- [3] E. Trucco and A. Verri. *Introductory Techniques for 3-D Computer Vision*. Prentice Hall, New Jersey, 1998.