

Prediction Intervals for Surface Growing Range Segmentation

James V. Miller *

Electrical, Computer and Systems Engineering
Rensselaer Polytechnic Institute
Troy, NY 12180-3590
<http://www.cs.rpi.edu/~millerj/>

Charles V. Stewart

Department of Computer Science
Rensselaer Polytechnic Institute
Troy, NY 12180-3590
<http://www.cs.rpi.edu/~stewart/>

Abstract

The surface growing framework presented by Besl and Jain [2] has served as the basis for many range segmentation techniques. It has been augmented with alternative fitting techniques [17], model selection criteria [11, 15], and solid modelling components [6]. All of these surface growing approaches, however, require global thresholds. Range scenes typically cannot satisfy the global threshold assumption since it requires data noise characteristics to be constant throughout the scene. Furthermore, these approaches can only be applied to range scenes where large seed regions can be isolated. As scene complexity increases, the number of surfaces, discontinuities, and outliers increase, hindering the identification of large seed regions.

We present statistical criteria based on multivariate regression to replace the traditional decision criteria used in surface growing. We use local estimates and their uncertainties to construct criteria which capture the uncertainty associated with extrapolating estimated fits. Our criteria allow small robust seed regions to grow to large surface patches without the use of global thresholds. To make the best use of these criteria, we restrict the surface expansion process to very localized extrapolations. This increases the sensitivity to discontinuities and allows regions to refine their estimates and uncertainties as they expand. Our approach has a small number of parameters which are either statistical thresholds or cardinality measures, i.e. we do not use thresholds defined by specific range distances or orientation angles.

1 Introduction

Surface growing range segmentation [2, 4, 7, 17, 11, 15, 6] is a local-to-global approach to surface reconstruction. Seed regions — initially isolated from discontinuities — are expanded until the boundaries of the surfaces are detected. The initial regions must be selected such that they avoid straddling discontinuities and avoid including the outliers range sensors introduce along discontinuities. If either of these constraints is violated, the reconstruction

may contain surfaces which bridge across the discontinuities (under-segmentation).

The growing process involves repeatedly searching for data which is consistent with the current region's fit. A location consistency test is used to identify potential inliers. Then, an orientation consistency test checks local surface normal estimates against the current fit. The former serves to collect inliers, identify step discontinuities and avoid outliers while the latter serves to identify crease discontinuities.

The benefit of surface growing is that it leverages off regions of high quality data while searching for discontinuities. Its strengths become increasingly important as range sensors are used in more unstructured environments — where the number of surfaces, discontinuities, and outliers increase. Existing regularization (see [3] for a review) and image region growing [5, 1] techniques, demand a high degree of local data integrity in image plane neighborhoods, faltering when presented with the numerous discontinuities and outliers from unstructured environments. Even robust estimators, designed for immunity to outliers, are thwarted by the intrinsic structure in the outlying samples from across discontinuities [16].

Current surface growing techniques, however, have two shortcomings. The first is a dependence on global thresholds. For instance, points are identified as inliers if their residual distances are within 3σ range units of the current fit. Unfortunately, the noise in a range image varies with depth, orientation, and surface reflectance (see Figure 1). Therefore, a global scale parameter may not be available. Second, current surface growing techniques rely on large seed regions. But as scene complexity increases, it becomes increasingly difficult to isolate large seed regions.

Here, we extend multivariate regression analysis techniques, developing criteria to control surface growth. Our criteria are based on local estimates and their uncertainties. This allows us to expand very small seed regions — even though the estimates from these small regions may be very uncertain. Our criteria are best used when surface extrapolation is localized, allowing estimates and uncertainties to be refined during the expansion process.

Our location compatibility test is based on “single point

*The authors would like to thank the National Science Foundation for funding this research under grants IRI-9217195 and IRI-9408700.

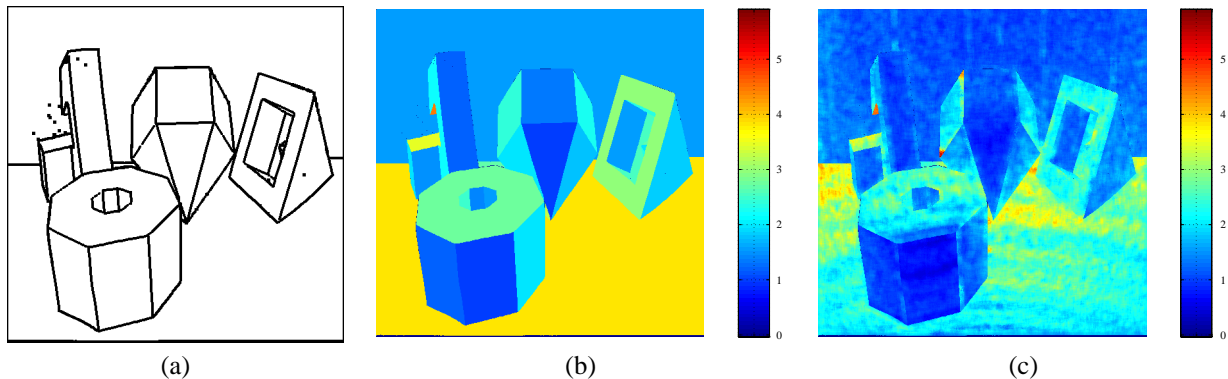


Figure 1: (a) Ground truth for a Perceptron range image. Data courtesy of the University of South Florida [10]. (b) Robust scale estimates for each ground truth surface. (c) Robust scale estimates for local regions on each ground truth surface.

prediction intervals” (SPPI). The local fit and scale estimates are used to define a region about the current fit to search for candidate data. These intervals adjust to the current region size, estimates, and confidence; allowing small fits to recover from poor initial estimates while still protecting a region from incorporating outliers or samples from across large step height discontinuities.

Our orientation consistency test is based on the distribution of estimated regression surfaces. The “normal prediction interval” (NPI) compares a local surface normal estimate against the current fit. This criterion takes into account the uncertainty in the two normal estimates, avoiding extrapolation into areas which have high confidence but incompatible normal estimates.

Finally, we add a criterion to the surface growing process which tests a set of points as a group against the current estimated fit. The “simultaneous prediction intervals” (SPI) improve discontinuity detection by identifying when the set of candidate data favors one side (or halfspace) of the estimated fit. Figure 2(b) illustrates such a case, where points from the upper half of a step discontinuity fall within the limits imposed by the location or SPPI test. By testing the candidate data as a set, this situation — and hence small step discontinuities — can be detected.

Basing our criteria on the multivariate regression model has a number of benefits. Aside from using only local estimates, the model allows for any surface order approximation (although we do not address the model selection problem here). Additionally, the multivariate regression model allows for various dependent-independent variable combinations, allowing the same criteria to be used for both surface and space curve reconstruction.

We begin with a review of multivariate regression, present condensed derivations of the three criteria, and finally discuss the alterations to the basic surface growing model necessary to make the best use of these criteria.

2 Multivariate Regression

The multivariate regression model provides great flexibility for working with various surface orders (planar, quadratic, cubic) and ranks (surfaces, space curves), provides great inference power, and is a fairly good match to physical range sensors (i.e. sensor measurements have independent and dependent components). When concentrating on surfaces modelled as one dependent variable and two or more independent variables, regression surfaces are of the form¹

$$\hat{y} = \hat{B}'x \quad (1)$$

where for linear surfaces, $x(q \times 1)$ is a vector

$$x = [1 \quad x_1 \quad x_2]', \quad (2)$$

and for quadratic surfaces

$$x = [1 \quad x_1 \quad x_2 \quad x_1^2 \quad x_1x_2 \quad x_2^2]', \quad (3)$$

etc. However, the general multivariate regression model [13, Chapter 6] is

$$Y = XB + U \quad (4)$$

where $Y(n \times p)$ is a matrix of n observations of the p -dimensional dependent variables, $X(n \times q)$ is a matrix of n observations of the $(q - 1)$ -dimensional independent variables (the first column of X is a vector of 1's), $B(q \times p)$ is the matrix of regression parameters, and $U(n \times p)$ is a matrix of independent, identically distributed random disturbances (typically with mean 0 and covariance matrix Σ).²

The maximum likelihood estimates of B and Σ ,³

$$\hat{B} = (X'X)^{-1}X'Y, \quad \hat{\Sigma} = \frac{1}{n}Y'(I_n - X(X'X)^{-1}X')Y, \quad (5)$$

¹Throughout the text, we use the nomenclature from multivariate regression. The resulting criteria can then be easily mapped to a surface reconstruction setting.

²Real sensors are not iid (Figure 1); however, locally Σ varies slowly, so in a local region Σ is approximately constant.

³The maximum likelihood estimates require nonrandom independent data. If the independent data are randomly distributed, the maximum likelihood estimates are interpreted as “conditional on X ”.

yield estimated residuals $\hat{U} = Y - X\hat{B}$. If U is a normal data matrix from $N_p(0, \Sigma)$ (i.e. each row of U is an independent p -dimensional Gaussian random variable with mean zero and constant covariance), then

1. \hat{B} is an unbiased estimate of B ,
2. $E[\hat{U}] = 0$,
3. \hat{B} and \hat{U} are multivariate normal,
4. \hat{B} and \hat{U} are statistically independent, and \hat{B} and $\hat{\Sigma}$ are statistically independent,
5. $\hat{\Sigma}$'s distribution is Wishart, a matrix generalization of chi-squared, specifically $n\hat{\Sigma} \sim W_p(\Sigma, n - q)$, and
6. the covariance between elements of \hat{B} is $COV(\hat{\beta}_{ij}, \hat{\beta}_{kl}) = \sigma_{jl}g_{ik}$, where $G = (X'X)^{-1}$.

3 SPPI: Single Point Prediction Intervals

Given a region of data and an estimated fit, our first goal is to define a volume about the fit that restricts the search for candidate data to add to the fit. The search volume should contain the current fit's data and reflect the uncertainties in the fit estimates. Thus, when a surface is small, containing little data, the uncertainties in the estimates are high and the search volume should be fairly large (in comparison to the region size). This allows for small seed regions to recover from bad initial fit estimates. On the other hand, as surfaces grow larger, the fit uncertainties reduce, and the search volume should be a tighter bound about the fit. In this manner, outliers and samples from nearby surfaces can be avoided during the surface growing process, resulting in large surface segments which yield more accurate surface estimates.

We start with an estimated fit \hat{B} and consider a datapoint (x_f, y_f) that was not used in calculating the estimate \hat{B} . (Using the generality of equation 4, the dependent variables are $y_f(p \times 1)$ and the independent variables are $x_f(q \times 1)$ where the first element of x_f is a 1). Assuming Gaussian errors, y_f is distributed as $y_f \sim N_p(B'x_f, \Sigma)$ and y_f 's residual to the fit \hat{B} can be shown to be distributed as

$$y_f - \hat{B}'x_f \sim N_p(0, (1 + x_f'(X'X)^{-1}x_f)\Sigma), \quad (6)$$

with the increase in covariance caused by the uncertainty in the estimate \hat{B} . Since this residual is based on \hat{B} , and \hat{B} is independent of $\hat{\Sigma}$, a quadratic form involving $y_f - \hat{B}'x_f$ and $\hat{\Sigma}$ can be composed, resulting in the statistic⁴

$$\frac{n - q}{n} \frac{(y_f - \hat{B}'x_f)' \hat{\Sigma}^{-1} (y_f - \hat{B}'x_f)}{1 + x_f'(X'X)^{-1}x_f} \sim T_{(p, n - q)}^2. \quad (7)$$

This statistic is a Hotelling T^2 random variable, which is related to the more familiar F random variable. For the case

⁴Composing quadratic forms between Gaussian random variables and independent Wishart random variables is a standard means of creating a univariate statistic with a distribution known as the Hotelling T^2 . See [13, pages 73–79].

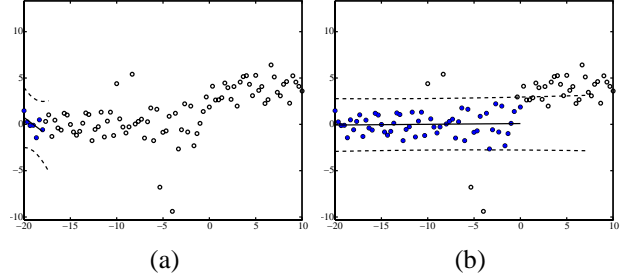


Figure 2: A step height of 4σ and a data density of 3 samples/unit. (a) Seed region with SPPI hyperbola. (b) SPPI filtering allows surface to grow and avoid outliers. Step discontinuity is detected by SPI (Section 5).

of univariate dependent variables ($p = 1$ as is typically the case in surface reconstruction), equation 7 is equivalent to

$$\frac{n - q}{n} \frac{(y_f - \hat{B}'x_f)^2}{(1 + x_f'(X'X)^{-1}x_f) \hat{\sigma}^2} \sim F_{(1, n - q)} \sim t_{(n - q)}^2. \quad (8)$$

Placing a bound on the statistic in equation 7, i.e. $Pr[T_{(p, n - q)}^2 < T_{(p, n - q); \alpha}^2] = 1 - \alpha$, the single point prediction interval (SPPI) for the value of the dependent variables $y_f(p \times 1)$ measured at $x_f(q \times 1)$ are the zeros of the implicit function

$$\frac{n - q}{n} \frac{(y_f - \hat{B}'x_f)' \hat{\Sigma}^{-1} (y_f - \hat{B}'x_f)}{1 + x_f'(X'X)^{-1}x_f} = T_{(p, n - q); \alpha}^2. \quad (9)$$

Converting equation 9 to an inequality creates a criterion which tests whether a point (x_f, y_f) — not used in estimating \hat{B} — is close enough to the estimated fit \hat{B} to potentially be considered one of its inliers. The residual distance is normalized by the estimated covariance and scaled by an amount reflecting the uncertainty in $\hat{\Sigma}$ and \hat{B} and scaled by the normalized distance from x_f to the mean of X (the independent variable values on which \hat{B} and $\hat{\Sigma}$ are based). Thus, SPPI intervals increase the further a fit is extrapolated.

All samples whose independent variables lie within the extrapolated region and whose dependent variables lie within their respective SPPI's can be considered candidate data (see Figure 2). Note, equation 9 is multivariate and can define different search shapes for different combinations of independent and dependent variables (Figure 3).

4 NPI: Normal Prediction Intervals

The normals comparison framework is based on regression surfaces with a single dependent variable and two independent variables.⁵ We test the compatibility of the estimated normals for surfaces \hat{B}_1 and \hat{B}_2 at a given point $\vec{p} = [x_1 x_2]'$, which requires extrapolating the surfaces' normal functions and uncertainties to the point \vec{p} .

⁵A similar criterion could be constructed for local tangent directions for the case of reconstructing space curves.




Figure 3: Traces of SPPI (equation 9) for $p + q - 1 = 3$ with (a) $p = 1$ (plane), (b) $p = 2$ (space curve), and (c) $p = 3$ (point cloud). A comparison of the three prediction interval regions for the same variance structure is shown in (d).

The normal function for a fit \hat{B} maps a point \vec{p} to the normal vector to the surface \hat{B} at \vec{p} , i.e. $\hat{n}(\hat{B}, \vec{p}) = \hat{n}(\hat{B}, [x_1, x_2]')$. We define the normal function by writing our regression surface parametrically

$$S(x_1, x_2) = [x_1 \ x_2 \ [1 \ x_1 \ x_2 \ \dots] \hat{B}]', \quad (10)$$

taking the partials of S with respect to x_1 and x_2 , and finally taking the cross product of these partials

$$\hat{n}\left(\hat{B}, \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}\right) = \begin{bmatrix} -\hat{B}' \frac{\partial x}{\partial x_1} \\ -\hat{B}' \frac{\partial x}{\partial x_2} \\ 1 \end{bmatrix}. \quad (11)$$

Note the normal function does not have unit length. As defined in equation 11, the normal function is Gaussian distributed on the $y = 1$ plane since it is a linear mapping of \hat{B} which is Gaussian distributed (see Section 2 and [13, Theorem 3.2.1, page 62]). Normalizing the normal function would result in a non-linear mapping of the distribution onto the unit sphere, resulting in non-Gaussian random variables (see Figure 4).

The complete distribution for $\hat{n}(\hat{B}, [x_1 \ x_2]')$ can be determined directly from the distribution of \hat{B} (given in Section 2). Since only two of the three components of \hat{n} are random, we concentrate on the distribution of

$$\hat{n}_{\hat{B}} = \begin{bmatrix} \hat{n}_{x_1} \\ \hat{n}_{x_2} \end{bmatrix}_{\hat{B}} = \begin{bmatrix} -\hat{B}' \frac{\partial x}{\partial x_1} \\ -\hat{B}' \frac{\partial x}{\partial x_2} \end{bmatrix}. \quad (12)$$

Using the covariance of linear forms [13, page 30] and the covariance of the components of \hat{B} given in the Section 2, it can be shown that the distribution of a regression surface's normal at the point \vec{p} is given by

$$\begin{bmatrix} \hat{n}_{x_1} \\ \hat{n}_{x_2} \end{bmatrix}_{\hat{B}} \sim N_2\left(\begin{bmatrix} n_{x_1} \\ n_{x_2} \end{bmatrix}, \sigma^2 \begin{bmatrix} \frac{\partial x}{\partial x_1} & \frac{\partial x}{\partial x_2} \end{bmatrix}' G \begin{bmatrix} \frac{\partial x}{\partial x_1} & \frac{\partial x}{\partial x_2} \end{bmatrix}\right) \quad (13)$$




Figure 4: An estimated normal's distribution is Gaussian in the space of $[n_{x_1} \ n_{x_2} \ 1]'$ but is not Gaussian if normalized.

where $[n_{x_1} \ n_{x_2}]'$ is the true normal, σ is the true scale, and

$$G = (X'X)^{-1} = \begin{bmatrix} m & m\bar{X}'_1 \\ m\bar{X}_1 & X'_1 X_1 \end{bmatrix}^{-1} = \begin{bmatrix} G_{11} & G_{12} \\ G_{21} & G_{22} \end{bmatrix}. \quad (14)$$

To simplify, let $H = \begin{bmatrix} \frac{\partial x}{\partial x_1} & \frac{\partial x}{\partial x_2} \end{bmatrix}' G \begin{bmatrix} \frac{\partial x}{\partial x_1} & \frac{\partial x}{\partial x_2} \end{bmatrix}$,

$$\begin{bmatrix} \hat{n}_{x_1} \\ \hat{n}_{x_2} \end{bmatrix}_{\hat{B}} \sim N_2\left(\begin{bmatrix} n_{x_1} \\ n_{x_2} \end{bmatrix}, \sigma^2 H\right). \quad (15)$$

4.1 Normal Compatibility Criterion

Our goal is to construct an orientation consistency criterion, allowing the comparison of extrapolated normals from two surfaces \hat{B}_1 and \hat{B}_2 at a given point \vec{p} . Here, we use the distributions derived in the previous section to construct a criterion which does not require knowledge of the population parameters (σ , B_1 , or B_2). The resulting NPI criterion is suitable for all regression surface orders and can even be used to compare estimated normals at a point for two surfaces of different surface orders.

We assume the two fits \hat{B}_1 and \hat{B}_2 are derived from independent sets of samples. Under these conditions, the two estimated normals are independent; and when the two fits are derived from the same underlying surface,^{6,7}

$$(\hat{n}_{B_1} - \hat{n}_{B_2}) \sim N_2(0, \sigma^2 (H_{B_1} + H_{B_2})). \quad (16)$$

Using the linear property of Gaussian distributions,

$$\frac{a' [\sigma^2 (H_{B_1} + H_{B_2})]^{-1/2} (\hat{n}_{B_1} - \hat{n}_{B_2})}{\sqrt{a'a}} \sim N_1(0, 1) \quad (17)$$

for all vectors a [13, Corollary 3.2.1.1 & 3.2.1.3, page 62].

Furthermore, since the two fits are derived from independent sets of samples, we have two independent scale estimates, $\hat{\sigma}_{B_1}$ and $\hat{\sigma}_{B_2}$. These independent scale estimates define two independent chi-squared random variables

$$m_{B_1} \frac{\hat{\sigma}_{B_1}^2}{\sigma^2} \sim \chi_{(m_{B_1} - q_{B_1})}^2 \quad m_{B_2} \frac{\hat{\sigma}_{B_2}^2}{\sigma^2} \sim \chi_{(m_{B_2} - q_{B_2})}^2 \quad (18)$$

⁶Explicit assumption: when \hat{B}_1 and \hat{B}_2 are describing the same surface, they are also describing the same scale parameter σ .

⁷We use subscripts B_1 to refer to parameters estimated for surface B_1 , i.e. m_{B_1} is the number of points used to estimate B_1 .

where the q 's are the number of parameters in each fit. Since these chi-squared random variables are independent, their sum is a chi-squared random variable with the sum degrees of freedom [9]

$$\frac{1}{\sigma^2} (m_{B_1} \hat{\sigma}_{B_1}^2 + m_{B_2} \hat{\sigma}_{B_2}^2) \sim \chi_{(m_{B_1} + m_{B_2} - q_{B_1} - q_{B_2})}^2. \quad (19)$$

Equations 16 and 19 are independent (see Section 2) and combine to form a t random variable, where for all \vec{a}

$$\frac{\frac{a'}{\sqrt{a'a}} [\sigma^2 (H_{B_1} + H_{B_2})]^{-1/2} (\hat{n}_{B_1} - \hat{n}_{B_2})}{\sqrt{\frac{1}{\sigma^2} (m_{B_1} \hat{\sigma}_{B_1}^2 + m_{B_2} \hat{\sigma}_{B_2}^2) / (m_{B_1} + m_{B_2} - q_{B_1} - q_{B_2})}} \sim t_{(m_{B_1} + m_{B_2} - q_{B_1} - q_{B_2})}. \quad (20)$$

Equation 20 represents a conversion of the multivariate problem of comparing two two-dimensional random variables (the normals to the two surfaces) to an infinite set of univariate statistics. Since this statistic must be true for all vectors \vec{a} , we only need to examine the vector \vec{a} at which the statistic attains its maximum value.⁸ It can be shown that [13, Corollary A.9.2.2, page 480]

$$\max_a \frac{(a'x)^2}{a'a} = x'x. \quad (21)$$

Therefore, we can reduce the infinite set of statistics defined in equation 20 by squaring the statistic and applying the property of the maximum just described, arriving at the single univariate statistic

$$\frac{(\hat{n}_{B_1} - \hat{n}_{B_2})' [H_{B_1} + H_{B_2}]^{-1} (\hat{n}_{B_1} - \hat{n}_{B_2})}{(m_{B_1} \hat{\sigma}_{B_1}^2 + m_{B_2} \hat{\sigma}_{B_2}^2) / (m_{B_1} + m_{B_2} - q_{B_1} - q_{B_2})} \sim t_{(m_{B_1} + m_{B_2} - q_{B_1} - q_{B_2})}^2. \quad (22)$$

Setting a confidence bound on this statistic, we obtain a criterion for comparing the estimated normal from two surfaces at a particular point.

5 SPI: Simultaneous Prediction Intervals

SPPI's define an interval or region about a current fit which is likely to contain the dependent variable y_f at a prescribed independent variable location, x_f . As a search for candidate data, SPPI's test each point individually for potential inlier status. However, this test alone tells us nothing about the candidate data set as a whole. Consider the case where all the data points in the extrapolated region fall just within the upper bound of their SPPI's. Figure 2(b) illustrates such a case, where points from the upper half of the step fall just inside of the SPPI's for the lower half of the step. Although each point passes an individual inlier test, the likelihood of all the data points falling near the upper tail of their respective distributions is collectively small. This situation should indicate that the candidate data is at least partly composed of samples from across a small step

discontinuity. Thus, if we use a test on the collection of candidate data, we can identify smaller step heights than with SPPI's alone.

There is another motivation for testing the candidate data as a set. While the candidate data are assumed to be independently distributed, their residual values relative to an estimated fit are not independent. In fact, it can be shown that the covariance between the residual⁹ r_i^* measured at x_i^* and residual r_j^* measured at x_j^* is given by

$$COV[r_i^*, r_j^*] = x_i^{*'} (X'X)^{-1} x_j^* \Sigma \quad i \neq j. \quad (23)$$

The collection of candidate points needs to be compared as a set against its "simultaneous prediction interval" (SPI). For the case of one-dimensional dependent variables, the simultaneous prediction interval reduces to a **k-dimensional** Student-t random variable for which there exist multiple approximate interval solutions [12, 8]. Here, we derive a summary statistic which is itself **univariate** yet still applicable to multivariate dependent variables.

Consider a set of k points, not used to estimate \hat{B} and $\hat{\Sigma}$, represented in the matrices $Y^*(k \times p)$ and $X^*(k \times q)$. The set of k residuals to the regression surface are given by $R^* = Y^* - X^* \hat{B}$, where each residual (row of R^*) is from a different distribution (equation 6) and the residuals' pairwise dependence is given by equation 23. Since each row of R^* has a different distribution and the rows are not independent, we cannot treat R^* as a normal data matrix [13, page 65]. However, we can operate on R^* as a single random variable by "vectorizing" R^* . The variable R^{*V} is the $(kp \times 1)$ vector formed by stacking the columns of $R^*(k \times p)$. Considering the distribution of each r_i^* and their interdependence, the distribution of R^{*V} is given by

$$R^{*V} \sim N_{kp}(0, \Sigma \otimes A), \quad (24)$$

where " \otimes " is the Kronecker matrix product [13, Appendix A] and $A = I_k + X^* (X'X)^{-1} X^{*'}.$

Since R^{*V} is a kp -dimensional normal random variable,

$$a' R^{*V} \sim N_1(0, a'(\Sigma \otimes A)a) \quad (25)$$

is univariate normal for *all* fixed vectors $a(kp \times 1)$ (see the definition of a multivariate normal [13, pages 60-62]).

Furthermore, since the regression surface's estimated covariance is Wishart, i.e. $n\hat{\Sigma} \sim W_p(\Sigma, n - q)$, then

$$\frac{nc' \hat{\Sigma} c}{c' \hat{\Sigma} c} \sim \chi_{(n-q)}^2 \quad (26)$$

is a chi-squared random variable with $n - q$ degrees of freedom for *all* fixed vectors $c(p \times 1)$ [13, page 67].

Given equation 25 is true for all vectors $a(kp \times 1)$, then it true for vectors of the form $a = c \otimes 1_k$, where c is the $(p \times 1)$ vector used in equation 26 and 1_k is a $(k \times 1)$ vector of ones.

⁹We use the "asterix" to indicate these points were not used in calculating the estimated fit.

⁸Union Intersection Principle, see [13, Chapter 5].

Thus, we can combine equations 25 and 26 (R^{*V} and $\widehat{\Sigma}$ are independent) to form a series of t statistics

$$\frac{(c \otimes 1_k)' R^{*V} / \sqrt{(c \otimes 1_k)' (\Sigma \otimes A) (c \otimes 1_k)}}{\sqrt{\frac{n}{n-q} c' \widehat{\Sigma} c / c' \Sigma c}} \sim t_{(n-q)} \quad (27)$$

parameterized by c , each with $n - q$ degrees of freedom.

Using techniques similar to those in Section 4, equation 27 can be reduced to the single univariate statistic¹⁰

$$\frac{k(n-q) \bar{r}^* \widehat{\Sigma}^{-1} \bar{r}^*}{n+k+k(\bar{X}_1^* - \bar{X}_1)' \widehat{\Sigma}_X^{-1} (\bar{X}_1^* - \bar{X}_1)} \sim t_{(n-q)}^2 \quad (28)$$

where \bar{r}^* is the mean of the new residuals, \bar{X}_1 and \bar{X}_1^* are the sample means of X_1 and X_1^* (where $X = [1_n X_1]$ and $X^* = [1_k X_1^*]$) and $\widehat{\Sigma}_X$ is the sample covariance of X_1 .

Placing a confidence bound on this statistic yields our SPI criterion which tests the k candidate data points for “group” inlier status. As in the SPPI and NPI case, the SPI criterion takes into account local fit and noise estimates and the uncertainties associated with these estimates. Additionally, the SPI criterion factors in “where” the candidate data is drawn and the dependence between the new residuals, allowing the mean residual to grow with the average degree of extrapolation (using the Mahalanobis distance in X_1). Finally, the SPI criterion is a univariate statistic suitable for segmenting both surfaces and space curves of arbitrary order.

6 Surface Growing

To use the SPPI, NPI, and SPI criteria in a surface growing algorithm requires a few adaptations to the basic surface growing model. We limit the presentation to extracting connected regions. This will allow us in the future to compare our techniques with other segmentation algorithms using the University of South Florida’s Segmentation Comparison Tool [10].

As with all the surface growing algorithms, we start with a set of seed regions. Here, we leverage off the criteria’s ability to factor in uncertainty, and seed with small surface patches extracted from a robust local operator. We chose to construct our seed regions using MUSE [14] due to its abilities to withstand outliers and extract multiple surfaces in a given region.

To use the NPI criterion, we calculate robust normal estimates for every pixel assigned to a seed region. To compensate for MUSE’s window based extraction, the initial normals are calculated by finding least-squares fits in 5×5 neighborhoods centered at each pixel, limiting the fit to only those points which fall within the SPPI limits of the pixel’s seed surface. This allows the initial normals to be based on pixels from different MUSE windows, yet still

limit the calculation to pixels that are likely to be from the same surface.

To allow for a greedy surface expansion algorithm (where we do not need to expand every seed region), we order the seed surfaces based on an estimate of the confidence in the seed fit. The confidence in a surface increases with its area and number of points, while it decreases with increasing scale. Therefore, we order the seed segments by

$$n \left| \widehat{\Sigma}_X \right| / \hat{\sigma}. \quad (29)$$

A seed region will not be expanded if its pixels have already been assigned to other surfaces.

Applying the criteria to surface growing is rather straight forward. A set of candidate points is selected and tested for location compatibility using the surface’s SPPI limits. Points already assigned to a non-seed surface are allowed to switch surfaces only if the new region provides a higher confidence estimate of the reconstructed point (i.e. the width of the SPPI interval is less for the candidate surface than for the surface it is currently assigned).¹¹ The points surviving are tested for orientation compatibility using their initial normal estimates (if available) and the current surface’s NPI limits. Finally, any points remaining are tested as a group using the SPI limits.

The main adaptation to the surface growing algorithm is in selecting the candidate points. Most surface growing techniques gather all the points connected to the current region that pass their location and orientation tests. However, we want to use the group location compatibility test (SPI) to assist in identifying step discontinuities. To use this criterion effectively, we localize the selection of candidate points. Otherwise, we could be testing a group of points from across several different boundaries. We localize the expansion by selecting a point on the boundary of the region and gathering the pixels within a prescribed radius on the image plane. The points within this radius (a 3 pixel radius works well) not already assigned to the region form the initial candidate set. To insure the region remains connected during expansion, the points surviving the SPPI and NPI tests are checked for connectivity with the current region and only the connected subset is tested against SPI. If the set of points passes the SPI test, they are added to the region, and the surface estimates are recalculated. Another candidate location is selected and process repeats until the surface can no longer expand. Another surface is then selected for expansion.

After a complete pass through the seed surfaces, the expansion process is repeated to refine discontinuity locations. Points near crease discontinuities can typically

¹⁰Detailed derivations will be made available on our web-site.

¹¹Requiring a surface to produce a higher confidence estimate is another diversion from existing techniques. Most techniques define the “better” surface at a point using the residual distance, not the SPPI widths.

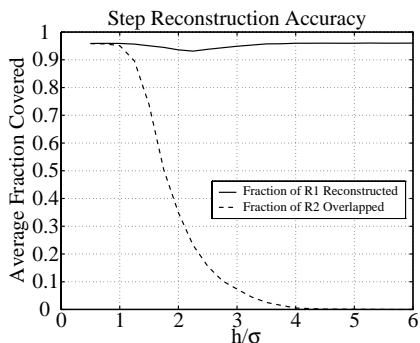


Figure 5: At $h = 4\sigma$, a reconstructed step contains two surfaces, each with 96% of the data from their side of the step and less than 0.5% of the data from across the discontinuity.

pass the SPPI test for either of the surfaces involved in the crease. Furthermore, due to the uncertainty in the initial normal calculation at a crease, these points may also pass both surfaces' NPI test. To refine crease locations, we remove the constraint that a point can only be assigned to the surface producing the smallest SPPI interval, and replace it with a constraint based on the intersection of the two surfaces (see [6]).

7 Analysis

One of the major strengths of the prediction interval framework lies in the ability to reconstruct very small magnitude step discontinuities. The SPI facilitate localized queries along the boundary of a region, identifying when a local expansion attempts to cross a discontinuity. Figure 5 illustrates the abilities of SPI. Here, we ran repeated simulations on 30×30 datasets containing a single step discontinuity of various magnitudes. The step divided the data into one 600 pixel region and one 300 pixel region. The prediction interval region growing system was seeded with two 6×6 regions, one on each side of the discontinuity. To measure the performance, we calculated the fraction of each region reconstructed by the largest extracted segment. For very small step heights, a single bridging fit was extracted, covering 95% of the data from both regions. However, by $h = 4\sigma$, two surfaces were extracted, with the largest extracted surface covering 95% of its side of the step while covering less than 0.5% of the other side of the step.

This is a significant result since without SPI, the theoretical limit for surface growing¹² is between 5σ and 6σ . This theoretical limit is twice the location threshold which must be set to 2.5σ to 3σ in order to identify and gather the inliers to a region. Also note, the prediction interval performance is achieved without a known or globally estimated scale parameter.

Figure 6 shows a prediction interval based segmentation on a scene composed of a series of 4σ steps. For comparison we include a segmentation from Fitzgibbon *et al.*'s

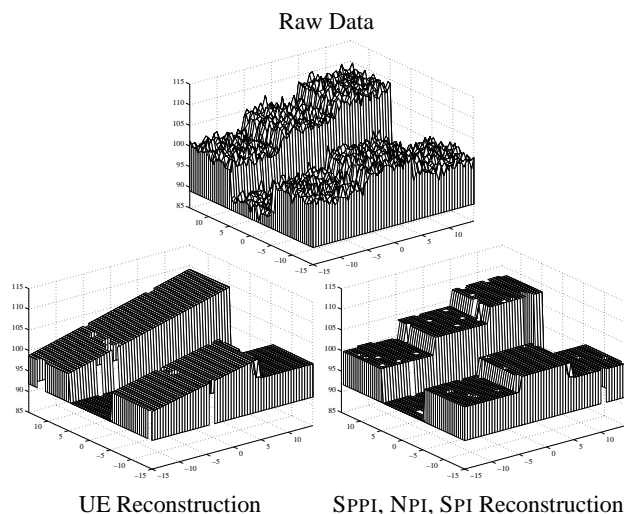


Figure 6: Small step heights ($h = 4\sigma$, data density 2 samples/unit) are bridged by the UE segmenter, however, the prediction interval segmenter handles the small steps correctly due to the SPI.

range segmenter (UE segmenter [6]). The UE segmenter was chosen for comparison since it had the best overall performance in [10] and because it is the closest in style of those discussed in [10] to the prediction interval segmenter. The parameters for the UE segmenter were set using the techniques discussed in [6] and [10]. In particular, the location threshold was set at 2.5σ . The UE segmenter results in bridging fits whereas the prediction interval segmenter properly reconstructs the steps.

8 Results

Figure 7 shows our segmentation technique on one of the Perceptron test data sets from the University of South Florida's Segmentation Comparison Project [10]. Following their framework, we limit the extraction to planar surfaces. The Perceptron data required a data smoothing stage prior to the initial normals calculation; although, we reverted back to the original data during reconstruction. For comparison, we include segmentation images from the algorithms tested in [10]. Qualitatively, our technique performs favorably, erring slightly on the side of oversegmenting. Instead of reconstructing 7 surfaces on the octagon "nut", our segmenter extracted 9 large regions. The top of the "nut" is difficult to extract since the noise varies by a factor of three across the surface (see Figure 1(c)). The oversegmentation on the front surfaces of the "nut" are caused by slight curved distortions in the data (even though the objects have planar faces). The oversegmentations, however, are locally consistent and should not pose a problem when the segments are combined to form a solid model. On the inner radius of the "nut", we extracted two of three surfaces present. Only the UB segmenter performed better on that region of the data. The crease discontinuities

¹²The limit for window based robust operators is even higher [16].

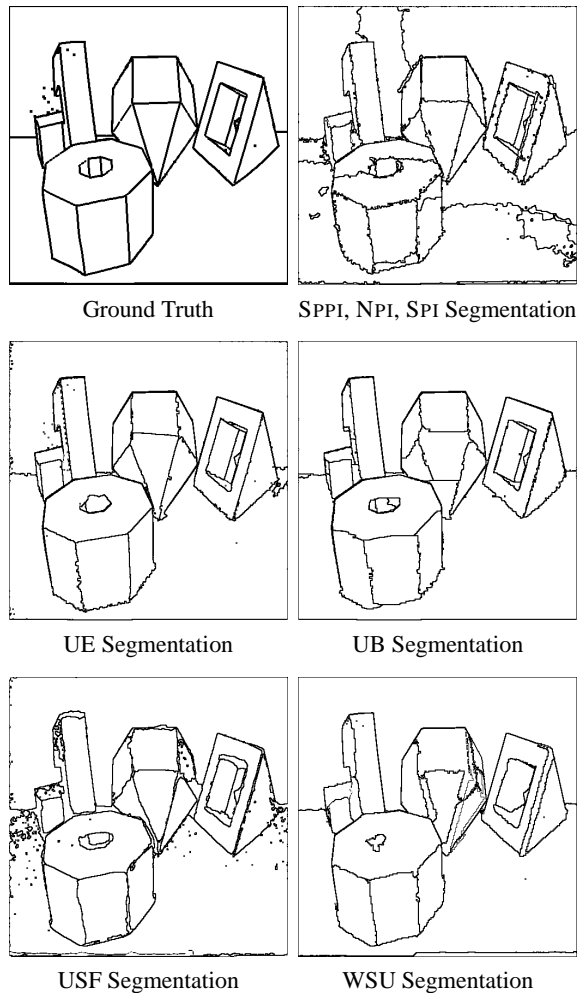


Figure 7: Segmentation comparison. Single isolated points not shown (extreme inliers/near outliers).

on the object in the center of the scene were well localized by our segmenter. Two of the seven crease discontinuities on this object were overshot by 3 to 5 pixels; however, the other creases are well localized and compare very favorably to the other segmentation algorithms. Finally, we do over-segment the background and the groundplane. We are currently investigating methods to address this problem.

9 Conclusion

Surface growing has the potential to address the issues involved with reconstructing range scenes from unstructured environments. By extending multivariate regression analysis techniques, we have developed criteria that bound the search for candidate data to add to a region's fit and that evaluate the compatibility of the candidate data to the existing fit. Our criteria differ from those currently employed by not requiring a global scale parameter, global range threshold, or a global orientation threshold. Rather, our criteria are locally adaptive, varying the range and ori-

entation thresholds for each surface based local confidence measures. Furthermore, we supply the SPI criterion which tests candidate points for group inlier compatibility, thereby increasing the sensitivity to small discontinuities.

In the future, we plan to extend the surface growing algorithm with solid modelling constraints as in [6], incorporate a model selection methodology to extract higher order surfaces, and perhaps most importantly quantitatively compare our surface growing segmentation technique against other range segmentation techniques using the University of South Florida's segmentation comparison tool [10].

References

- [1] E. Al-Hujazi and A. Sood, *Range Image Segmentation with Applications to Robot Bin-Picking using Vacuum Gripper*, IEEE T-SMC **20** (1990), no. 6, 1313–1325.
- [2] P.J. Besl and R.C. Jain, *Segmentation Through Variable-Order Surface Fitting*, T-PAMI **10** (1988), 167–192.
- [3] R. Bolle and B. Vemuri, *On Three-Dimensional Surface Reconstruction Methods*, T-PAMI **13** (1991), no. 1, 1–13.
- [4] T.E. Boult and M. Lerner, *Energy-Based Segmentation of Very Sparse Range Surfaces*, IEEE Conference on Robotics and Automation, 1990, 232–237.
- [5] O.D. Faugeras, M. Hebert, and E. Pauchon, *Segmentation of Range Data into Planar and Quadratic Patches*, IEEE CVPR, 1983, 8–13.
- [6] A.W. Fitzgibbon, D.W. Eggert, and R.B. Fisher, *High-level CAD Model Acquisition from Range Images*, Tech. report, Dept. of Artificial Intelligence, Univ. of Edinburgh, 1995.
- [7] P. Fua and P. Sander, *Reconstructing Surfaces from Unstructured 3D Points*, Proc. of the DARPA Image Understanding Workshop, 1991, 615–625.
- [8] G.J. Hahn, *Simultaneous Prediction Intervals for a Regression Model*, Technometrics **14** (1972), no. 1, 203–214.
- [9] W.L. Hays and R.L. Winkler, *Statistics: Probability, Inference, and Decision*, Holt, Rinehart, and Winston, Inc., New York, 1971.
- [10] A. Hoover, G. Jean-Baptiste, X. Jiang, P.J. Flynn, H. Bunke, D.B. Goldgof, K. Bowyer, D.W. Eggert, A. Fitzgibbon, and R.B. Fisher, *An Experimental Comparison of Range Image Segmentation Algorithms*, T-PAMI **18** (1996), 673–689.
- [11] A. Leonardis, A. Gupta, and R. Bajcsy, *Segmentation of Range Images as the Search for Geometric Parametric Models*, IJCV **14** (1995), 253–277.
- [12] G.J. Lieberman, *Prediction Regions for Several Predictions from a Single Regression Line*, Technometrics **3** (1961), no. 1, 21–27.
- [13] K.V. Mardia, J.T. Kent, and J.M. Bibby, *Multivariate Analysis*, Academic Press, New York, 1979.
- [14] J.V. Miller and C.V. Stewart, *MUSE: Robust Surface Fitting using Unbiased Scale Estimates*, IEEE CVPR, 1996, 300–306.
- [15] M.J. Mirza and K.L. Boyer, *An Information Theoretic Robust Sequential Procedure for Surface Model Order Selection in Noisy Range Data*, IEEE CVPR, 1992, 366–371.
- [16] C.V. Stewart, *Expected Performance of Robust Estimators Near Discontinuities*, ICCV, 1995, 969–974.
- [17] G. Taubin, *Estimation of Planar Curves, Surfaces, and Non-Planar Space Curves Defined by Implicit Equations with Applications to Edge and Range Image Segmentation*, T-PAMI **13** (1991), no. 11, 1115–1138.