MUSE: Robust Surface Fitting using Unbiased Scale Estimates

James V. Miller*
Electrical, Computer and Systems Engineering
Rensselaer Polytechnic Institute
Troy, NY 12180-3590
millerj@cs.rpi.edu
http://www.cs.rpi.edu/~millerj

Charles V. Stewart
Department of Computer Science
Rensselaer Polytechnic Institute
Troy, NY 12180-3590
stewart@cs.rpi.edu
http://www.cs.rpi.edu/~stewart

Abstract

Despite many successful applications of robust statistics, they have yet to be completely adapted to many computer vision problems. Range reconstruction, particularly in unstructured environments, requires a robust estimator that not only tolerates a large outlier percentage but also tolerates several discontinuities, extracting multiple surfaces in an image region.

Observing that random outliers and/or points from across discontinuities increase a hypothesized fit's scale estimate (standard deviation of the noise), our new operator, called MUSE (Minimum Unbiased Scale Estimator), evaluates a hypothesized fit over potential inlier sets via an objective function of unbiased scale estimates. MUSE extracts the single best fit from the data by minimizing its objective function over a set of hypothesized fits and can sequentially extract multiple surfaces from an image region. We show MUSE to be effective on synthetic data modelling small scale discontinuities and in preliminary experiments on complicated range data.

1 Introduction

A statistical method or estimator is considered robust if it can continue to operate successfully as the data diverges from the estimator's target distribution. Applying this notion to regression analysis, a robust fitting procedure should be able to extract a good fit despite a portion of the data being corrupted with random outliers. Robust operators are ideal for many computer vision applications since gross errors can be introduced by both sensors and low-level processing algorithms. In fact, robust techniques have been very successful in many computer vision applications [8, 15]; however, the greatest success is achieved when the goal is to identify or extract a single signal corrupted with random outliers.

Unfortunately, many computer vision applications do not fit this model of singly isolated objects corrupted with outliers. For instance, in range reconstruction, robust operators are used to estimate surface parameters in small image windows. Each window has the potential to contain data from multiple surfaces. These surfaces may be overlapping in the case of pseudo-transparency, for example chain link fences; and one or all of the surfaces may contain less than 50% of the region's data. Furthermore, whenever multiple surfaces are in a region, range sensors are likely to introduce random outliers aligned with the discontinuities. Reconstruction accuracy requires that all surfaces in a region be extracted, the random outliers ignored, and the extracted fits not bridge across the discontinuities (particularly in reverse engineering where accuracy is needed for small scale discontinuities, i.e. step heights $\leq 10\sigma$).

Figures 3(a) and 6(b) illustrate the problems with range data. In Figure 3(a), two surfaces of equal size are separated with a step height of 8σ and corrupted with 10% outliers. Thus, each surface contains less than 50% of the data. While the discontinuity and segmentation seem readily apparent, this scenario remains problematic for robust estimators. Figure 6(b) shows actual range data of an inbox, a few cardboard boxes, and a small flower pot. This data contains pseudo transparent regions (background surface visible through the grates of the inbox) and a large number of outliers (which tend to align with the discontinuities).

Standard robust regression techniques [4, 12] simply are not designed to address these issues. They are designed to extract a single surface whose data is corrupted with outliers, not to robustly extract multiple surfaces. While these robust regression techniques can be applied in a sequential fashion, extracting a second surface from a region once the inliers to the first surface have been removed, etc., these techniques are biased at the outset by the coherent structure of data from multiple surfaces, preferring to extract a surface that bridges across the discontinuity [13]. This is particularly evident whenever all the surfaces contain less than 50% of the data or when the size of the discontinuity drops below 10σ.

Even the robust fitting techniques developed by the vision community are hampered by the complexities of

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

CVPR'96. IEEE Computer Society Conference on Computer Vision and Pattern Recognition.

range scenes. The fixed-band techniques of Hough transforms [5], RANSAC [3], and Roth and Levine [11] require accurate *a priori* noise estimates. Our initial experiments with range data indicate noise varies with depth, image position, and surface material, making *a priori* noise estimation problematic. Furthermore, these techniques have difficulty with small scale discontinuities [13]. Stewart's MINPRAN operator [14] tolerates the large number of outliers in range images and identifies regions composed completely of outliers; but MINPRAN's assumptions about the outlier distribution are not sufficient for extracting multiple surfaces. Finally, the surface growing techniques of Mirza and Boyer [10], Leonardis *et al.* [7], and Darrell *et al.* [1] are designed specifically to extract multiple surfaces. However, these techniques rely on good seed fits, which are usually obtained using robust estimators, and therefore susceptible to the problems discussed above.

The two goals for our operator are an ability to extract surfaces containing less than 50% of the data and an ability to detect small scale discontinuities. First, our objective function rates a hypothesized fit based on a series of unbiased robust scale estimates (variance in fit residuals)[1]. Each scale estimate s_k is based on the fit's k smallest absolute residuals, and a fit's representative scale estimate is the minimum of its s_k values. The unbiased scale estimates for a set of hypothesized fits are compared, and the fit with the smallest unbiased scale estimate is extracted.

Minimizing unbiased scale estimates within and between hypothesized fits produces robustness to outliers and large scale discontinuities. Outliers' characteristically large residuals produce large scale estimates. By minimizing s_k over a given fit, the fit's representative scale estimate is not dominated by outliers. Comparing unbiased scale estimates between fits produces robustness to large scale discontinuities (step heights $> 8\sigma$) since a bridging fit's representative scale estimate tends to overestimate the true scale.

Second, to address small scale discontinuities, our operator can compare unbiased scale estimates for hypothesized fits defined over different subsets of a region[2]. Here, scale estimates are localized to a sub-region defined by the hypothesized fit, creating the potential for the data considered to be from a single surface. Such a fit will have a scale estimate smaller than that of a bridging fit, even for small magnitude discontinuities.

The remainder of this paper details our new estimator, concentrating on the two features outlined above, and showing how to embed the estimator in a sequential surface extraction algorithm. The presentation emphasizes development of the robust technique, and as such, the presented extraction algorithm is limited to extracting multiple planar approximations in each window. A complete range segmentation algorithm and higher order surface extraction await future development.

As a final preliminary note, we were inspired to develop our new estimator by a recent, related technique called ALKS [6] (communicated to us by Peter Meer). Although ALKS is limited in its ability to handle extreme outliers, and is therefore not directly comparable to our new estimator, it did open our eyes to the possibility of developing this estimator. We will discuss this further in Section 3.

2 Least Median of Squares

We begin with a review of Least Median of Squares (LMS) [12] since MUSE builds on LMS's computational technique and its ability to robustly estimate the variance in the data. LMS has been very successful when applied to a lone signal corrupted with outliers but is less effective when presented with multiple surfaces and fails completely if a signal contains fewer than 50% of the data.

LMS searches a space of hypothesized fits using an objective function based on the median squared residual

$$\min_{\hat{\theta}} \ \underset{i}{\text{median}} \ r^2_{i,\hat{\theta}},$$

where $r_{i,\hat{\theta}}$ is the residual of the i^{th} data point measured relative to the hypothesized fit $\hat{\theta}$. To search efficiently, LMS usually employs random sampling [12][Chapter 5].

An important aspect of LMS is its ability to simultaneously generate a robust estimate of the variance in the inliers. LMS's scale estimate is defined as

$$\hat{\sigma}_{\hat{\theta}} = C\sqrt{\underset{i}{\text{median}} \ r^2_{i,\hat{\theta}}}, \tag{1}$$

where C is a constant chosen to make $\hat{\sigma}$ unbiased for a particular target distribution. If the inlying data is Gaussian distributed, then $C = 1.4826(1 + 5/(N-p))$ [12], where N is the number of data points, the factor 1.4826 is the inverse of the expected value of the median residual from a standardized (unit variance) Gaussian distribution, and $(1 + 5/(N-p))$ corrects for small point sets.

3 MUSE objective function

Traditionally, LMS is presented as minimizing the median squared residual. However, LMS can be equivalently viewed as minimizing the unbiased scale estimate in equation 1. From either viewpoint, LMS requires a surface to contain at least 50% of the region's data. We can move past this restriction – extracting surfaces with fewer than 50% of the data and estimating accurate fits in the presence of multiple surfaces – by building upon LMS's scale estimate and how this scale estimate is used.

LMS's scale estimate is unbiased when all the data is from the target distribution (i.e. Gaussian inliers). A small number of outliers increase the scale estimate slightly since

[1] The estimates are unbiased for a target distribution — such as measurements from a single surface. The estimates are not necessarily unbiased for mixture distributions — such as measurements at a discontinuity.

[2] Our scale estimates are unbiased regardless of point set size.

the median squared residual increases. Outliers numbering in excess of 50% increase the scale estimate dramatically since the median squared residual is an outlier. However, a meaningful scale estimate can still be constructed using a rank order statistic different from the median. This maintains an objective function based on estimated scale; but without *a priori* knowledge of the inlier percentage, we need a means to compare scale estimates across hypothesized inlier percentages.

Our new operator is based on this intuition. For any hypothesized fit, our operator calculates an unbiased estimate of the scale from the fit's k smallest (magnitude) residuals, for all possible values of k, $1 \leq k \leq N-p$ (a fit defined by randomly selecting p points has $N-p$ residuals). The smallest scale estimate over all possible k is the representative value for the hypothesized fit. This value is used in comparing different hypothesized fits, the optimum fit having smallest scale estimate. Inliers are then identified using the optimum fit's scale estimate, $\hat{\sigma}$, i.e. those points with $|r| \leq 2.5\hat{\sigma}$.

3.1 Unbiased scale estimate, s_k

Constructing an unbiased scale estimate from the k smallest absolute residuals is the key to MUSE. We use a residual normalization method similar to LMS but compute a series of scale estimates and combine them to produce our scale estimate s_k.

Given a residual density $\phi(r;\sigma)$ — ϕ is the density of random variable r having zero mean and standard deviation σ — r can be written as a scaled random variable u from a standardized distribution ($u \sim \phi(u;1)$ and $r = \sigma u$). Given N points from $\phi(r;\sigma)$, the k^{th} ordered absolute residual $r_{k:N}$ can be written in terms of the k^{th} ordered absolute residual from a standardized distribution, $r_{k:N} = \sigma u_{k:N}$. This leads to the scale estimate

$$s_k = \frac{r_{k:N}}{E\left[u_{k:N}\right]} \tag{2}$$

where $E\left[u_{k:N}\right]$ can be computed directly from $\phi(u;1)$[3].

There are many good estimates for $E\left[u_{k:N}\right]$. For instance, if the density of signed standardized residuals $\phi(u;1)$ is symmetric, then the expected values of ordered absolute standardized residuals can be approximated by

$$E\left[u_{k:N}\right] \approx \Phi^{-1}\left(0.5 * \left(1 + \tfrac{k}{N+1}\right)\right), \tag{3}$$

where $\Phi(u,1)$ is the cumulative distribution of the signed standardized residuals.

Equation 2 converts each residual to a hypothesized fit into a scale estimate. To create a scale statistic summarizing the first k residuals, we could compute a weighted linear combination of $s_{k'}$ values for $k' \leq k$, where the weights are derived from the covariance between the $s_{k'}$'s. However, ordering the absolute residuals creates highly correlated s_k variables. Therefore combining $s_{k'}$ variables via any weighted linear combination does not produce an estimate marketedly better (in terms of variance in the estimate) than simply using s_k as the scale estimate for the first k absolute residuals.[4]

3.2 MUSE at a discontinuity

The effectiveness of minimizing unbiased scale estimates between hypothesized fits when the data is composed of multiple surfaces is illustrated in Figure 1. We consider a step discontinuity between two surfaces of equal size, having unit variance Gaussian inliers, and corrupted with 10% outliers. Using analysis tools developed in [13], we calculate the expected absolute residuals for two hypothesized fits – one a correct fit to one side of the discontinuity, the other a fit bridging the discontinuity – and calculate the minimum s_k for each fit using equation 2.

For intermediate and large step heights, the correct fit's minimum s_k is smaller than the bridging fit's minimum s_k, implying our new estimator should select a correct fit even though each surface contains less than 50% of the data. However, MUSE still prefers a bridging fit at small step heights because it can find a large amount of data packed tightly about the bridging fit. This issue is addressed by the extension to MUSE presented in Section 4.

Finally, the correct fit's minimum s_k is nearly a constant overestimate with respect to step height. It is a constant since the minimum s_k never directly depends on the 55% of the data not from the correct fit. It is an overestimate because this 55% indirectly affects the inlying data's scale estimate by shifting $E[u_{k:N}]$ for the inliers.

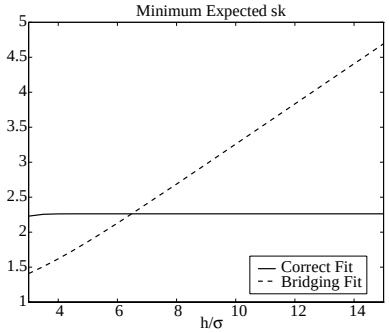

Figure 1: The minimum of the expected values of s_k at a fit bridging a step discontinuity and at a correct fit to one side of the discontinuity. $\sigma = 1$ is the true scale.

3.3 Variance in s_k

Given data from the target distribution, the N scale estimates from equation 2 are all unbiased; however, their variability differs dramatically. Figure 2(a) shows the variance

[3] To compute $E\left[u_{k:N}\right]$, convert $\phi(u;1)$ to a density for $|u|$, form the order statistic density $f(u_{k:N})$ [2], and integrate to calculate the expected values.

[4] Lee *et al.* [6] use a similar scale estimate as part of their reconstruction technique, but not as the objective function and do not compensate for the estimate's variability nor for the bias introduced during minimization.

of each s_k for $N = 100$ unit variance Gaussian points[5]. Note for small k, the standard deviation in s_k approaches σ, i.e. the scale estimate s_1 varies as much as any given residual.

The large variance in the first few s_k's affects the minimum s_k for the hypothesized fit. In Figure 2(b), we plot $E[\min s|k]$, the expected value[6] of s_k given that residual k produced the minimum scale estimate over all N residuals. The large variance in the first few s_k's can cause the minimum s_k for a fit to be a severe underestimate of the scale.

However, the variance in s_k and $E[\min s|k]$ both "converge" relatively quickly: while the standard deviation in s_1 is approximately 10 times the standard deviation in s_{90}, the standard deviation in s_{10} is only 3 times that of s_{90}. Therefore, we can increase the stability of our scale estimates – while still tolerating a large outlier percentage and still extract surfaces with less than 50% of the data – by simply ignoring the scale estimates from the first 10-15% of the sorted absolute residuals[7].

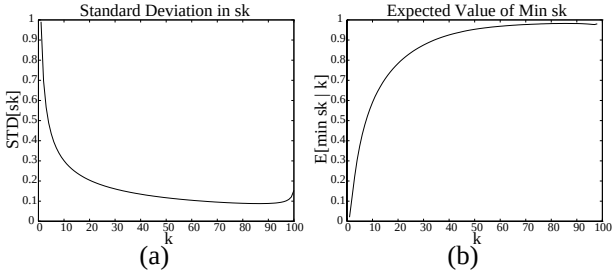

(a) (b)

Figure 2: (a) Standard deviation in s_k. (b) Empirical plot of the expected minimum value s_k given absolute residual k produces the minimum scale estimate.

3.4 Correcting for bias due to minimization

Selecting the minimum scale estimate as the representative value for a hypothesized fit biases the estimate low (in Figure 2(b), the expected minimums are all below the true scale $\sigma = 1$). In order to compare minimum scale estimates between hypothesized fits, we must correct for this bias.

An unbiased minimum scale estimate for a hypothesized fit is constructed by normalizing the minimum s_k. The normalization factor is the expected value of the minimum scale estimate from the standardized distribution given that the minimum occurs at absolute residual k, i.e.

$$\hat{\sigma} = \frac{\min\limits_k s_k}{E\left[\min v_{k'} \,\middle|\, k' = \operatorname{argmin}\limits_k s_k\right]}, \tag{4}$$

where v_k represents the k^{th} scale estimate from a standardized distribution. $E[\min v|k]$ can be calculated in $O(N^3)$ time using a dynamic programming algorithm (see [9]) and stored in a table for MUSE.

[5]The variance of s_k can be calculated from the variance in $r_{k:N}$ since they are linearly related. See David [2] for the density of $r_{k:N}$.

[6]Section 3.4 discusses the calculation of $E[\min s|k]$.

[7]In our implementation, the user sets this percentage. In the remainder of the presentation, we represent the number of points "ignored" by $k_0 - 1$.

3.5 Extracting multiple surfaces

To process a complete range image, we partition the image into small regions or windows, allowing MUSE to sequentially extract multiple surfaces from each window. The steps in processing each window are:

① Determine the number of random samples needed to extract the next surface.

② For each random sample fit:

- ❶ Calculate the sorted absolute residuals, $r_{k:N}$.
- ❷ For $k \geq k_0$, calculate $s_k = r_{k:N}/E[u_{k:N}]$.
- ❸ Assign $s^*_{\hat{\theta}}$ to $\min_{k \geq k_0} s_k$.
- ❹ Convert $s^*_{\hat{\theta}}$ to an unbiased estimate $\hat{\sigma}_{\hat{\theta}} = \min\limits_{k \geq k_0} s_k / E\left[\min v_{k'} \,\middle|\, k' = \operatorname{argmin}\limits_{k \geq k_0} s_k\right]$.

③ Select the random sample fit $\hat{\theta}$ with the smallest $\hat{\sigma}_{\hat{\theta}}$.

④ Remove all data points within $\mu\hat{\sigma}_{\hat{\theta}}$ of the optimum fit.

⑤ If the window contains another surface, goto ①.

⑥ Refine scale estimates for the region's fits using points within window with small enough residuals,

$$\hat{\sigma}_{\hat{\theta}} = \operatorname*{STD}_{i \text{ in region}} (r_{i,\hat{\theta}} \mid |r_{i,\hat{\theta}}| \leq \mu\hat{\sigma}_{\hat{\theta}}).$$

⑦ Assign each region data point to the surface to which it has the smallest normalized residual if $r/\hat{\sigma}_{\hat{\theta}} \leq \mu$.

⑧ If a surface has enough inliers, refine fit using least squares on identified inliers.

Determining the number of random samples in step ① depends on the amount of data remaining, the expected number of surfaces remaining in the data, and the expected number of outliers (see [14]). Overall, assuming the normalizing factors are stored in look-up tables, the algorithm requires $O(N \log N)$ time to analyze each hypothesized fit. This time is comparable to that of LMS.

Steps ⑤ and ⑧ hinge on distinguishing between data that is a random collection of outliers and data that has an underlying surface structure. We are pursuing two techniques to address this need. The first leverages off the strengths of Stewart's MINPRAN operator [14]. The second searches the data for reliable statistical asymmetries.

3.6 Examples

Figure 3 shows four steps of the above algorithm operating on $2D$ data. The scene contains a small step discontinuity of 8σ, 10% random outliers, with each surface having less than 50% of the data. Multiple surfaces are successfully extracted.

Section 3.2 predicted MUSE will not always fit correct surfaces at small magnitude discontinuities. Figure 4 shows a dataset, similar to the previous one, but where the first surface extracted is a bridging fit incorporating almost all of the inliers to both surfaces.

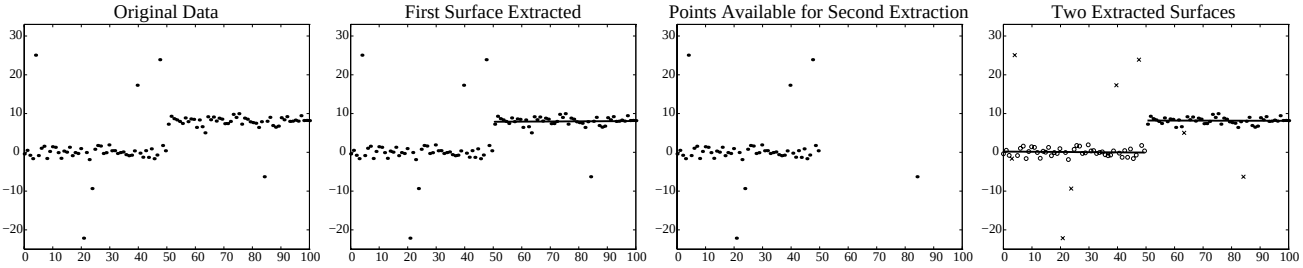

Figure 3: (a) Multiple surface range data with a step height of 8σ and 10% outliers. (b) First surface extracted. (c) Points remaining after the potential inliers to the first surface are removed. (d) Final reconstruction. Two surfaces and their inliers have been identified. "x"'s mark the points not assigned to any surface and hence are outliers.

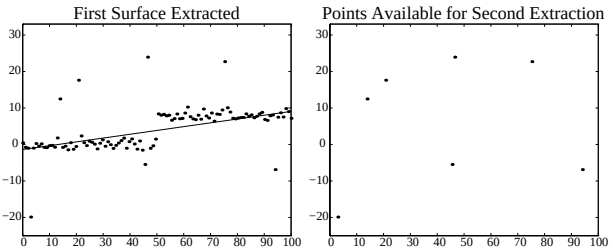

Figure 4: MUSE can still produce bridging fits to small step height discontinuities. Left: Bridging fit to a step height of 8σ. Right: Points remaining after potential inliers to the bridging fit have been removed.

4 Interior point selection

MUSE successfully tolerates large outlier percentages and can extract surfaces containing less than 50% of the data. However, MUSE still prefers a bridging fit at small discontinuities (Figure 4) because it overestimates the scale at a correct fit (Section 3.2). Equation 2 assumes all N data points are from a single surface, normalizing the inliers at a correct fit by too small a factor whenever there are multiple surfaces (for inliers numbering $N' < N$, $E[u_{k:N}] < E[u_{k:N'}]$).

Ideally, we should base our scale estimate on only the inliers to a hypothesized fit, eliminating the residuals from across discontinuities from contention for producing the minimum s_k and in calculating $E[u_{k:N'}]$. While such a localization of the inliers is impossible, we can approximate this behavior by considering only the residuals "interior to" the random sample (i.e. those points whose independent variables are inside the convex hull of the random sample's independent variables). If the random sample points are all from the same surface, then it is likely the interior data is from that same surface (pseudo transparency aside). Thus, the value N used in equation 2 is the number of interior samples and the overestimate due to multiple surfaces is avoided.

While there is no guarantee that a given random sample will contain points from a single surface, the random sampling process (probabilistically) guarantees that at least one random sample will contain points from a single surface. Therefore, at least one hypothesized fit should benefit from interior point selection, allowing it be distinguished from bridging fits which would still overestimate the scale.

4.1 Alterations to extraction algorithm

Supporting interior point selection requires two alterations to the random sampling process. First, the points interior to the random sample (or its convex hull when fitting higher order surfaces) need to be gathered efficiently. Second, each random sample needs to contain enough interior points to generate stable s_k values ($N' \geq k_0$, Section 3.3).

4.2 Drawback to interior point selection

Interior point selection can remove the overestimate in scale due to multiple surfaces; however, the variance in the scale estimate increases since the estimate is based on a smaller number of points. The statistical efficiency of an interior point selection estimate can be increased by requiring random samples to bound larger sets of interior points. As an alternative, we choose to perform an intermediate least squares fit on the interior points identified as inliers (residuals within $\mu\hat{\sigma}_{\hat{\theta}}$), refining the fit and scale estimates before different hypothesized fits are compared.

5 MUSE versus LMS

Figure 5 compares the empirical performance of MUSE and LMS. Our performance metric is the ratio of (1) the average bias between the optimal fit extracted and the correct fit, and (2) the average bias between the least squares fit and the correct fit. A bias ratio of zero indicates the extracted fit is the correct fit and a bias ratio of 1 indicates the extracted fit is near the least squares fit (i.e. a bridging fit). Figure 5 shows the bias ratio for MUSE, the interior point selection version of MUSE, and LMS for a variety of step and crease height discontinuities. On the left, there are 10% outliers and 50% of the good points are from each surface, i.e. each surface contains less that 50% of the data. On the right, there are 10% outliers but now one of the surfaces contains 65% of the good data.

Regardless of discontinuity magnitude, MUSE outperforms LMS. While LMS selects a fit approximating the least

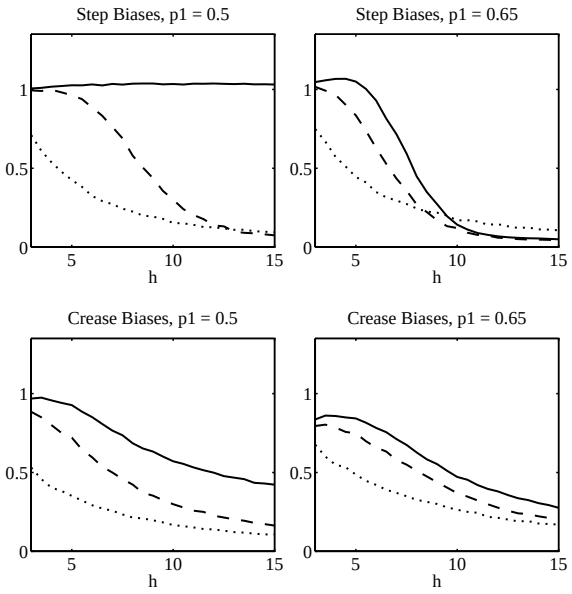

Figure 5: Empirical comparison of MUSE and LMS using a bias ratio metric. Solid curve = LMS. Dashed curve = MUSE. Dotted curve = MUSE Interior. Left: Each surface contains less than 50% of the points. Right: Largest surface contains 58.5% of the data.

squares fit when all surfaces contain less than 50% of the data, MUSE selects the correct fit with increasing probability as the discontinuity magnitude increases. For small magnitude discontinuities, the interior point selection version of MUSE shows a substantial improvement over the standard MUSE and LMS. However, the decreased statistical efficiency associated with interior point selection causes it to exhibit more bias than the standard version of MUSE at large magnitude discontinuities.

6 Discussion and Results

Figure 6 illustrates MUSE applied to a range scene composed of an stack of plastic inboxes, a few cardboard boxes, and a small flower pot. Part (a) of the figure shows an intensity image of the scene. Part (b) shows the input range data. Each range measurement is shown as a small octahedron, shaded according to its depth. Notice the range data has numerous outliers and that portions of the cardboard boxes can be seen through the grates in the inbox (pseudo-transparency). Also note the inbox has a small lip (small step discontinuity) highlighting the borders of each tray.

Figure 6(c) and (d) show the reconstructed range data. In part (c), MUSE was allowed to extract at most one surface from each 11×11 window. MUSE has eliminated the outliers in the scene, reconstructed portions of the cardboard boxes appearing "behind" the grate of the inbox, and for the most part it did *not* combine the small step height lip around the inbox trays. In part (d), MUSE was allowed to extract at most two surfaces from each window. Here, more of the grates of the inbox have been reconstructed, more of the cardboard boxes appearing behind the grates have been reconstructed, and several other regions of the range image have been filled.

Figure 6(e) and (f) show the surfaces extracted in parts (c) and (d). There are a few bridging fits between the small lip of the inbox and tray surface and a bridging fit connects an inbox grate to the box lying behind the grate.

7 Conclusion

Current robust estimators suffer from several limitations when applied to data from multiple surfaces, such as near a depth or orientation discontinuity. No existing technique handles small magnitude discontinuities, and without knowledge of the noise in the data or the outlier distribution, no existing technique reconstructs surfaces containing fewer than half the points. Our new operator, MUSE, is designed to address both of these problems. By minimizing unbiased scale estimates within and between hypothesized fits, MUSE is robust to large outlier percentages and intermediate and large magnitude discontinuities. By localizing a hypothesized fit's scale estimates to the points interior to the random sample, MUSE can reconstruct small magnitude discontinuities.

We have demonstrated MUSE's effectiveness statistically on synthetic range data and we have shown its application to a range image of a complicated scene. In ongoing work, we will complete the theoretical analysis of MUSE, thoroughly demonstrate its application to real range data, and incorporate it into a system to reconstruct complete surfaces.

References

[1] T. Darrell, S. Sclaroff, and A. Pentland, *Segmentation by Minimal Description*, 3rd ICCV, 1990, pp. 112–116.

[2] H.A. David, *Order Statistics*, John Wiley & Sons, NY, 1970.

[3] M.A. Fischler and R.C. Bolles, *Random Sample Consensus: A paradigm for model fitting with applications to image analysis and automated cartography*, Communications of the ACM **24** (1981), 381–395.

[4] F.R. Hampel, E.M. Ronchetti, P.J. Rousseeuw, and W.A. Stahel, *Robust Statistcs: The Approach Based on Influence Functions*, John Wiley & Sons, New York, 1986.

[5] J. Illingworth and J. Kittler, *A Survey of the Hough Transform*, CVGIP **44** (1988), 87–116.

[6] K. Lee, R. Park, and P. Meer, *A Comparison of High Breakdown Point Range Image Segmentation Methods*, Submitted to Pattern Analysis and Machine Intelligence, 1995.

[7] A. Leonardis, A. Gupta, and R. Bajcsy, *Segmentation as the Search for the Best Description of the Image in Terms of Primitives*, 3rd ICCV 1990, IEEE, 1990, pp. 121–125.

[8] P. Meer, D. Mintz, D.Y. Kim, and A. Rosenfeld, *Robust Regression Methods for Computer Vision: A Review*, IJCV **6** (1991), no. 1, 59–70.

[9] J.V. Miller and C.V. Stewart, *Robust Regression Minimizing Unbiased Scale Estimates*, Tech. report, Computer Science Department, Rensselaer Polytechnic Institute, 1996.

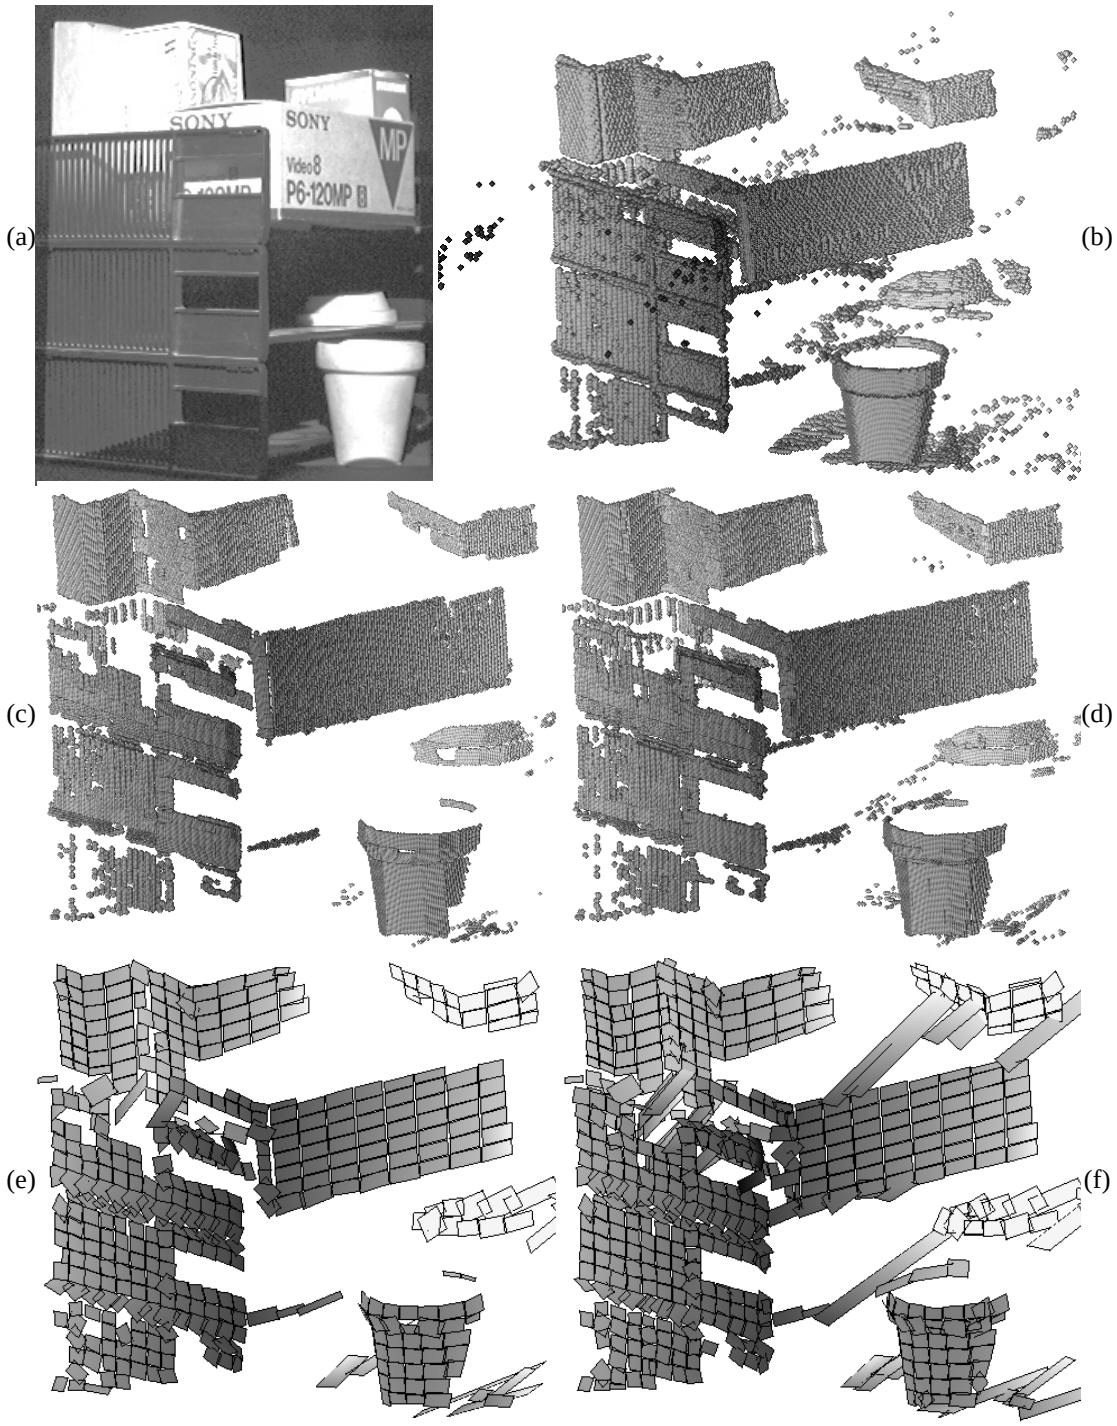

Figure 6: (a) Intensity image. (b) Range measurements. Grey level encodes depth. (c) and (d) Reconstructed range measurements when MUSE extracts (c) one and (d) two surfaces per region. (e) and (f) Planar surfaces extracted for (c) & (d).

[10] 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, pp. 366–371.

[11] G. Roth and M.D. Levine, *Extracting Geometric Primitives*, CVGIP-Image Understanding **58** (1993), no. 1, 1–22.

[12] P.J. Rousseeuw and A.M. Leroy, *Robust regression and outlier detection*, John Wiley & Sons, 1987.

[13] C.V. Stewart, *Expected Performance of Robust Estimators Near Discontinuities*, ICCV 95, 1995, pp. 969–974.

[14] ———, *MINPRAN: A New Robust Estimator for Computer Vision*, IEEE T-PAMI **17** (1995), 925–938.

[15] P.H.S. Torr, A. Zisserman, and S.J. Maybank, *Robust Detection of Degenerate Configurations for the Fundamental Matrix*, ICCV 95, IEEE, 1995, pp. 1037–1042.