Robust Fitting Using Mean Shift: Applications in Computer Vision

D. Suter and H. Wang

Abstract. Much of computer vision and image analysis involves the extraction of "meaningful" information from images using concepts akin to regression and model fitting. Applications include: robot vision, automated surveillance (civil and military) and inspection, biomedical image analysis, video coding, human-machine interface, visualization, historical film restoration etc.

However, problems in computer vision often have characteristics that are distinct from those usually addressed by the statistical community. These include: **Pseudo-outliers:** In a given image, there are usually several populations of data. Some parts may correspond to one object in a scene and other parts will correspond to other, rather unrelated, objects. When attempting to fit a model to this data, one must consider all populations as outliers to other populations - the term pseudo-outlier has been coined for this situation. Thus it will rarely happen that a given population achieves the critical size of 50% of the total population and, therefore, techniques that have been touted for their high breakdown point (e.g., Least Median of Squares) are no longer reliable candidates, being limited to a 50% breakdown point.

Computer vision researchers have developed their own techniques that perform in a robust fashion. These include RANSAC, ALKS, RESC and MUSE. In this paper new robust procedures are introduced and applied to two problems in computer vision: range image fitting and segmentation, and image motion estimation. The performance is shown, empirically, to be superior to existing techniques and effective even when as little as 5-10% of the data actually belongs to any one structure.

1. Introduction

Computer Vision involves the extraction of "meaningful" information from images. Subareas include: robot vision, automated surveillance (civil and military) and inspection. Moreover, the techniques involved in computer vision also find their

Received by the editors XXXXX.
1991 *Mathematics Subject Classification.* Primary 68T45; Secondary 62F35.
Key words and phrases. Robust Statistics, Model Fitting, Computer Vision, Breakdown-point.
This work was completed with the support of the Australian Research Council.

way into a diverse range of other areas, such as biomedical image analysis, video coding, human-machine interface, visualization, historical film restoration etc.

Many of the tasks to be carried out, for the purposes of computer vision, can be cast as forms of statistical estimation and fitting. This has lead to a strong community interest in statistical methods - particularly robust statistical methods. Indeed, robust statistical methods in computer vision have a history, at least two decades old, using or adapting standard methods from the statistical community: for example, M-Estimators [1] and Least Median of Squares [2] [3]. At least one novel technique, RANSAC [4] was developed by vision researchers in the early days (and is still a widely used technique).

However, problems in computer vision often have characteristics that are distinct from those addressed by the statistical community. These include:

- **Pseudo-outliers.** In a given image, there are usually several populations of data. Some parts may correspond to one object in a scene and other parts will correspond to other, rather unrelated, objects. When attempting to fit a model to this data, one must consider all populations as outliers to other populations - the term pseudo-outlier has been coined [5]. Thus it will rarely happen that a given population achieves the critical size of 50% of the total population and, therefore, techniques that have been touted for their high breakdown point (e.g., Least Median of Squares) are no longer reliable candidates from this point of view.
- **Unknown sizes of populations - and unknown location.** Computer vision requires fully automated analysis in, generally, rather unstructured environments. Thus, the sizes and locations of the populations involved, will fluctuate greatly. Moreover, there is no "human in the loop" to select regions of the image dominated by a single population, or to adjust various thresholds. In contrast, statistical problems studied in most other areas usually have a single dominant population plus some percentage of outliers (typically mis-recordings - not the pseudo-outliers mentioned above). Typically a human expert is there to assess the results (and, if necessary, crop the data, adjust thresholds, try another technique etc.).
- **Large data sizes.** Modern digital cameras exist with around 4 million pixels per image. Image sequences, typically at up to 50 frames per second, contain many images. Thus, computer vision researchers typically work with data sets in the tens of thousands of elements, *at least*, and data sets in the 10^6 and 10^9 range are not uncommon.
- **Emphasis on fast calculation.** Most tasks in computer vision must be performed "on-the-fly". Offline analysis that takes seconds, let alone minutes or hours, is usually a luxury afforded by relatively few applications.

These rather peculiar circumstances have lead computer vision researchers to develop their own techniques that perform in a robust fashion (perhaps "empirically-robust" should be used, as few have formal proved robust properties, though many trace their heritage to techniques that do have such proved properties). These include ALKS [6], RESC [7] and MUSE [8]. Results presented in those papers *suggest* that ALKS can be tolerant to at least 65% outliers, RESC to about 80% outliers, and MUSE to 55% outliers. However, it has to be admitted that a complete solution, addressing all of the above problems, is far from being achieved. Indeed, none of the techniques, with present hardware limitations, are really "real-time" when applied to the most demanding tasks. None have been *proved* to reliably tolerate high percentages of outliers and, indeed, we have found with our experiments that RESC and ALKS, although clearly better than Least Median of Squares, in this respect, are not always reliable. For recent surveys, see [9] and [10].

The purpose of this paper is to highlight recent research carried out by the authors, in developing techniques with robust behaviour, to address various problems in computer vision - specifically, range image fitting and segmentation, and image motion estimation.

2. Basic Idea and Results

Because of the presence of multiple structures in the image, we need approaches that are robust to (pseudo)-outliers in the sense of having a high breakdown point. Established techniques, in use, that meet this criteria are based upon random sampling: e.g., Least Median of Squares, Least Trimmed Squares, and RANSAC. Random sampling techniques aim to explore the search space of possible solutions well enough to have at least one candidate which is determined solely by inliers (to a single structure in the data). However, since one doesn't have an "oracle" to tell us which of the candidates are unpolluted by outliers; we require some form of model/fit scoring. In Least Median of Squares, this is obviously the median of the residuals. In RANSAC, it is the number of data residuals inside a certain bound. Of course, each form of model scoring has potential weaknesses. In Least Median of Squares, the median of the residuals of the entire data set (with respect to the candidate model) will obviously not be a good measure if the inliers to that model contain less than 50% of the total data. Generalising to the k'th other statistic (rather than the Median) is one way out of this dilemma but now one either has to know in advance what value of k to use, or one has to attempt some form adaptation e.g., ALKS (which will perhaps be costly and limited in reliability). Even still, it is overly optimistic to expect a single statistic (the k-th order residual, for Least Median of Squares and ALKS; or the number of inliers within a certain bound, for RANSAC) to be an entirely reliable/informative measure of the quality of a solution.

This observation has lead the authors to seek alternative ways of scoring candidate models, so that greater robustness may be achieved. An early attempt [11] employed possible symmetry in the data set as one such statistic: though somewhat limited in versatility, such an approach definitely restores robustness in situations where standard Least Median of Squares will break down. Our more recent estimators seek to use more information from the residual distribution. In particular, we have used Kernel Density estimation and Mean Shift Techniques [12] to formulate model/fit scores that lead to empirically observed higher breakdown points than all existing methods.

2.1. Mean Shift and Kernel Density Estimation in Model Scoring

Consider a candidate model fit (as part of a random selection and trial fitting as in the usual Least Median of Squares algorithm). We seek to approximate the pdf of the residuals and find the mode of this distribution.

Although other kernels could be employed, we use the Epanechnikov kernel in its 1-D form:

$$K(x) = \begin{cases} \frac{3}{4}(1-x^2) & x^2 < 1 \\ 0 & otherwise \end{cases} \tag{2.1}$$

The kernel density estimator is then:

$$\hat{f}(x) = \frac{1}{nh}\sum_{i=1}^{n} K(\frac{x - X_i}{h}) \tag{2.2}$$

for a data set of n residuals X_i and using “bandwidth” h.

The mean shift is calculated by [12]:

$$\hat{f}(x) = M_h(x) = \frac{1}{n_x}\sum_{X_i \in S_h}(X_i - x) \tag{2.3}$$

where S_h is an interval of half-width h centred on x - the “mean shift window”.

Simple calculations show that iterating the above will converge to a local maximum of the estimated pdf. Because of the limited extent of the mean shift window (in the particular case of the kernel we choose, there is an inherent limit to the spatial influence of a data point, coming from the finite support of the Kernel - although a window could be imposed on other kernels not having finite support), this process is reasonably insensitive to outlier residuals. Comaniciu and co-workers have recently popularised this mean-shift technique for various applications in computer vision (e.g., clustering [13]). Here we use the method as a way of of scoring candidate fits.

The basic notion is that the robust estimate should produce a strong peak in the pdf of the residuals for that fit, and that the value of the residual corresponding to that peak should be small (ideally zero, of course). We have experimented with several ways to encapsulate such notions and have found maximization of the following objective function performs well:

$$MPDE = \frac{\left(\sum_{X_i \in W_c} \hat{f}\,(X_i)\right)^{\alpha}}{exp\,(|X_c|)} \tag{2.4}$$

where the mean shift procedure is used to find the mode of the residual density $\hat{f}$ and to limit mean estimation to using only “inliers” X_i (with center value X_c)

within the mean-shift window W_c centred on the mode of the pdf. This method is similar to the Residual Concensus (RESC) method [7]. Essentially that method estimates the pdf by using a histogram (whose bin size is chosen by compressing a large histogram of residuals using a heuristic procedure so that 12% of the residuals are found in the first bin). The RESC criterion is then:

$$RESC = \sum_{i}^{m} \frac{{h_i}^{\alpha}}{|r_i|^{\beta}} \tag{2.5}$$

where h_i is the histogram value and r_i the residual of the i'th bin. The value m is chosen by another heuristic (4.4% of h_{max}), so as to exclude outliers. (Note: in [7] $\alpha = 1.3$ and $\beta = 1$.) In essence, the scores 2.4 and 2.5 differ in how one restricts attention to likely inliers (we use the mean shift window) and how one models the pdf (we use kernel density estimation *and* mean shift mode seeking). These differences lead to significant improvements in robustness.

2.2. Examples

In this section we show that approaches based upon these procedures can tolerate up to 90% or so of outliers (including pseudo-outliers) and outperform previous computer vision methods (MUSE, RESC, ALKS, RANSAC) and more widely known methods (Least Median of Squares and Least Trimmed Squares) in that regard. We demonstrate this with experiments using synthetic data and with "real life" data in the areas of: line and ellipse finding, range image segmentation and image motion estimation.

2.2.1. LINES AND ELLIPSE FINDING Various edge detection routines can be employed to try to find edges in a scene. Typically, these routines produce very noisy output (gaps in edges, isolated points etc. - see Figures 1 2) If one is seeking structures that have particular parametric forms (e.g., lines and ellipses) one can use regression type frameworks to find the structures of interest amongst the edge detector output.

Thus this group of problems is similar to the standard regression problem: although we acknowledge that one should consider using geometric distance, rather than residuals, in the minimization [14]. However, here, we avoid the non-linear theory and resultant approximations of geometric fitting, and apply our procedure to the residuals produced by substituting the data points into the defining parametric form of an ellipse or line. Note: as opposed to standard settings for regression, we have *multiple* structures. To find *every line* we could find *any* line, first, then find the inliers to that line, remove them from the pool of data and re-apply the procedure to find the next line. Here we just demonstrate the first stage - finding any valid structure (though in segmentation, section 3, we do find all of the structures sequentially.)

In computer vision, probably the most commonly used technique, for such problems, is the Hough Transform. Essentially, this transform discretizes the parameter space and each data point "votes" for the parameters it is consistent with.

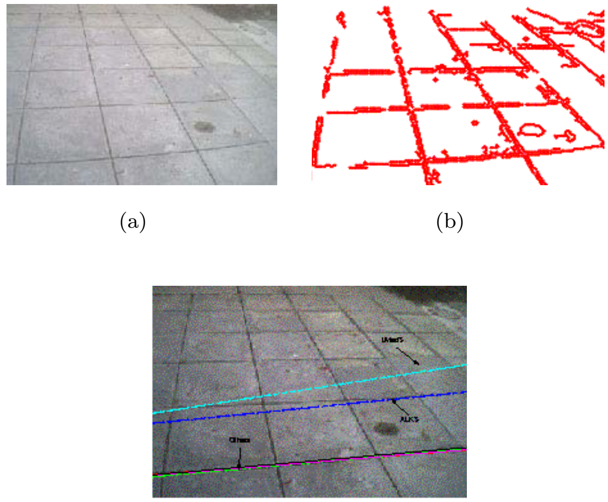

(a) (b)

(c)

FIGURE 1. Line finding

Therefore, it produces an estimate of the parameters with a finite precision which is determined by the "bin" size in parameter space. In this respect, it is unlike the other methods we consider: all of the other methods produce estimates with potentially infinite precision (limited by round-off errors, and, of course, in accuracy rather than precision). Moreover, since the parameter space has to be discretized, the Hough transform is simple to apply when the parameter space is compact or known limits apply (e.g., as imposed by the finite boundaries of an image) but is less simple when the parameter space has infinite extent or the practical bounds on likely encountered perameters are less clear. For these reasons, despite the early and continuing successes of such a technique, we prefer to seek improved methods. Nonetheless, it is instructive to compare the results with the Hough transform.

In addition to the Hough transform, we experiment with various established robust fitting methods Least Median of Squares, and several techniques not generally known outside of the computer vision community: RANSAC, ALKS, and RESC. An example can be found in Figure 1. The pavement image 1(a) is processed with an edge detector 1(b). After various forms of robust line fitting are applied, we obtain the results typically as shown in 1(c). Least Median of Squares (light upper line) fails because it is not robust to some many outliers. ALKS (middle dark line) also fails. The remaining techniques (MDPE, RESC and Hough - lower dark line) all produce similar results on this example - as the discretization of the parameter space for the Hough transform is adequate, its limit of precision is not apparent.

(a)

(b)

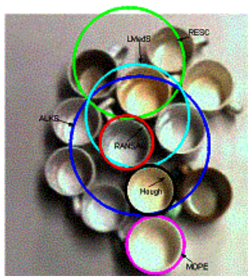

(c)

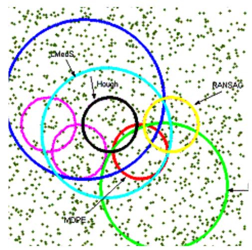

(d)

FIGURE 2. Ellipse finding

Similarly, examples of ellipse fitting can be found in Figure 2. The cups image 2(a) is processed with an edge detector 2(b). After various forms of robust line fitting are applied, we obtain the results shown in 2(c). MPDE (finds bottom cup), Hough (finds cup above it) and RANSAC (finds next cup above) all produce robust fits. Least Median of Squares, ALKS and RESC all breakdown (larger circles in the figure). As a second example, a synthetic data set was created to resemble "olympic rings" and then artificially corrupted with uniformly distributed noise samples - see Figure 2(d). In this example, the data that are "inliers" to any one ring constitute only about 5% of the data - that is, for a given fit, there are around 95% outliers to that fit. MPDE (bottom right ring), RANSAC (top right ring) and the Hough transform (top middle ring) can all find a ring. However, Least Median of Squares, ALKS and RESC again breakdown. Note: it is not significant *which structure is found*.

These and other similar experiments we have performed, shows that algorithms based upon our MPDE criterion outperform other robust techniques (Least Median of Squares, techniques based upon least k'th order statistics such as ALKS, and the RESC approach from which we drew our motivation). The technique is challenged for robustness only by RANSAC (which requires a priori knowledge of the expected sizes of the majority of the inlier residuals) and the non-regression (limited precision) Hough transform (which also requires a well-chosen bin size for the parameter discretization). We must acknowledge that almost no method has an implementation that is free of certain parameters that have to be set. In our experiments, where there are such parameters, we have tried first to use the settings advocated by the authors of those methods (if, indeed, they are available to us). We have also tried to vary those parameters and to select the best performance delivered by a method over the range of parameters investigated. We believe that this is as fair as we can be in our comparisons. The essential parameter required for our technique is the bandwidth of the kernel density estimator (although there are the inevitable minor bounds and tolerance parameters that plague code in order to guard against certain numerical limits, or to decide when to cease iteration). In the experiments reported in this section, we empirically investigated the behaviour of our approach with varying bandwidth and found that overall performance, at least over a reasonable range of bandwidth choices, was reasonably stable. In the next section we present results where we have automatically chosen the bandwidth.

2.2.2. OPTIC FLOW Optic flow $((u, v))$ is the apparent motion on the image plane caused by relative motion between an observer and the imaged points. If one assumes that the imaged point at position (x, y) and time t maintains a constant brightness $I(x, y, t)$, as it moves under the optic flow, then a simple differentiation reveals to first order:

$$\frac{\partial I}{\partial x}u + \frac{\partial I}{\partial y}v + \frac{\partial I}{\partial t} = 0 \tag{2.6}$$

- the so called "Optic Flow Constraint" (see, e.g., [3]). Since this constraint provides only one constraint in two unknowns, it has been solved by assuming, within a small image patch, that the flow is of a simple parametric form (e.g., locally constant, or locally affine). Such low order local polynomial approximations can be justified in the usual manner - providing the flow is locally smooth. In fact, it can be shown that if the scene contains rigid planar patches, then to a good approximation, the flow will be and 8 parameter quadratic

$$u = ax^2 + bxy + cx + dy + e; \;\; v = by^2 + axy + fx + gy + h \tag{2.7}$$

(i.e., having less degrees of freedom than the full quadratic defined over a 2-D domain.)

Within each patch, we can measure the quantities $\frac{\partial I}{\partial x}$, $\frac{\partial I}{\partial y}$, and $\frac{\partial I}{\partial t}$ so that we can solve for the parameters of the assumed flow model. However, since the measurements can contain very large errors, and since the patches may straddle different flows (moving according to different parameters), we need methods of solution that are robust to large numbers of outliers (including pseudo-outliers). Robust statistical methods such as Least Median of Squares have been demonstrated as superior to Least Squares and other competing methods in this respect [3].

Technique	Average Error
Bab-Hadiashar and Suter [3]	1.94
QMDPE ($\sigma = 2.0$, 25×25, m=30)	1.34

TABLE 1. Yosemite Valley Sequence - Optic Flow

We have applied our method, as described above but with some modifications, to the problem of optic flow estimation [15]. Firstly, we have implemented a form of automatic selection of the kernel bandwidth. Using standard techniques [16]:

$$\hat{h} = \left[\frac{243 R(K)}{35 u_2(K)^2 n} \right]^{\frac{1}{5}} s \tag{2.8}$$

where $R(K) = \int_{-1}^{1} K(\xi)^2 \, d\xi$, $u_2(K) = \int_{-1}^{1} \xi^2 K(\xi) \, d\xi$ and s is the sample standard deviation. This still requires the estimation of the sample scale in a robust fashion and, since the above is the recommended upper bound on $\hat{h}$, we also have to employ a multiplicative adjustment $c\hat{h}$ for $0 < c < 1$. Space does not permit us to go into details here, but we have investigated, with some success, both the use of order statistics for scale estimation and a rather more novel scheme of using a mean-shift like procedure to also find the pdf valley (thus providing a useful classification of inliers as those between the peak and the valley in the pdf). We have also, for speed reasons, modified the MPDE criterion to use what we call QMPDE:

$$QMPDE = \frac{\left(\hat{f}\,(X_c) \right)^{\alpha}}{exp\,(|X_c|)} \tag{2.9}$$

One sample of our results is provided in Table 1. To estimate the optic flow at a point, we centre a rectangular patch on that point and solve using all optic flow constraints from that patch. The table reports an error measure (see [3] for a definition) averaged over every result. We compare against the closest (and to this point best) competing robust method, and to other methods available in the literature. Clearly QMDPE performs well.

3. Segmentation

Segmentation is the process of dividing the image (or a sequence of images such as a movie) into spatial and/or temporal structures of interest. In principle, one can use a parametric fitting method to sequentially find structures in the data, find the inliers to that fitted model, remove the inliers and repeat by fitting another model. Such approaches, using variants of Least k'th order fitting have been shown to be reasonably promising [17] [18]. However, the implementation of such a relatively simple scheme becomes complicated by a number of practical issues that are beyond the scope of this paper and possibly of little interest to the intended audience - suffice to say that one can devise methods, based upon robust fitting schemes, the outperform traditional and other competing methods.

Consider for example, a range image, such as one collected via a laser range-finder, or stereo camera, or structured lighting technique: ultimately, though these devices differ in their principles, they all produce an "image" which is a set of 3D points sampled from some real world scene. The challenge is then to make sense of this mass of data points - to extract meaningful structures and surfaces. We have shown that one can segment a range image, into smooth parametric surfaces [19].

Likewise, using variants of Least k'th Squares, one can segment a movie sequence purely on the basis of the motion [20]. We are working on using MDPE and QMDPE on this task. It is worthwhile noting that to completely solve the problem of segmentation, in situations where the observed objects can have varying shapes and motions, one must not only use a robust fitting technique (and solve the attendant implementation issues), but one also needs to solve the *model selection problem*. Though the cited papers show some promising investigations into that issue, robust model selection in this context remains the major (and daunting) problem.

4. Conclusion

In this paper we have adopted the philosophy that a single statistic, such as the k'th order statistic, or the number of inliers within a certain bound; is unlikely to be a sophisticated enough measure to reliably discriminate between candidate model fits in a robust procedure that has to cope with a wide range of possible outlier/pseudo-outlier populations. Instead, we have looked at characterizing the quality of a model fit by more complex measures of the residual distribution: capturing information such as how peaked around zero the residual probability density function is. To this end, we have devised procedures that, at their core, use kernel density estimation of the pdf and a mean-shift approach to located the peak of that pdf. Experiments have shown that such procedures can considerably out-perform existing robust techniques in terms of apparent breakdown point behaviour.

In concluding, we must remark on the shortcomings of the approaches we are hereby promoting. From a practical point of view, the methods are somewhat costly. Though we haven't extensively studied the issue of just how computationally efficient the approaches can become: it would be optimistic to expect that that they can compete with methods relying on much more simple measures of candidate solution quality. From a theoretical point of view, a lot remains to be studied. The reader would be justified in concluding that some aspects of our proposal are ad-hoc and that many variants can be easily dreamt up - without this paper providing any justification for the particular variants we have examined. Secondly, though we promote our schemes in terms of "breakdown point", we acknowledge a number of issues in respect of this. We have not formally defined "breakdown point"; nor, consequently, have we in any way attempted to prove attainment of a high breakdown point. In these respects, our approach is intuitive and empirical. However, we trust, despite these shortcomings, the techniques we have described will be of use to the computer vision community (and wider) as the basis of proven practical methods which can be refined, and whose theoretical underpinnings can be explored. Moreover, we must point out that, despite impressions that may be obtained by reading much of the literature, particularly that aimed more at the practitioner, more traditionally accepted techniques still have their shortcomings in similar ways. For example, though it is often cited that Least Median of Squares has a proven breakdown point of 50%, it is often overlooked that all practical implementations of Least Median of Squares are an approximate form of Least Median of Squares (and thus only have a weaker guarantee of robustness). Indeed, the robustness of practical versions of Least Median of Squares hinges on the robustness of two components (and in two different ways): the robustness of the median residual as a measure of quality of fit and the robustness of the random sampling procedure to find at least one residual distribution whose median is not greatly affected by outliers. Our procedures, like many other procedures, share the second vulnerability as we too rely on random sampling techniques. The first vulnerability is sometimes disregarded for practical versions of Least Median of Squares, because robustness is viewed as being guaranteed by virtue of the proof of robustness for the ideal Least Median of Squares. However, two comments should be made in this respect. Firstly, that proof relies on assumptions regarding the outlier distribution and it can easily be shown that clustered outliers will invalidate that proof. Secondly, there is an inherent "gap" between a proof for an ideal procedure and what one can say about an approximation to that procedure. We believe that our method of scoring the fits better protects against the vulnerabilities that structure in the outliers expose. We have presented empirical evidence to support that. The challenge is to not only continue to amass empirical evidence, but to also explore the theoretical properties of these (and similar) schemes.

References

[1] M. Black and A. Rangarajan. On the unification of line processes, outlier rejection, and robust statistics with applications in early vision. *International Journal of Computer Vision*, 19(1):57–92, July 1996.

[2] P. Meer, D. Mintz, D.Y. Kim, and A. Rosenfeld. Robust regression methods in computer vision: A review. *International Journal of Computer Vision*, 6:59–70, 1991.

[3] A. Bab-Hadiashar and D. Suter. Robust optic flow computation. *International Journal of Computer Vision*, 29(1):59–77, August 1998.

[4] M. A. Fischler and R. C. Bolles. Random consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Commincations of the ACM*, 24(6):381–395, June 1981.

[5] C.V. Stewart. MINPRAN: A new robust estimator for computer vision. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 17(10):925–938, October 1995.

[6] K-M Lee, P. Meer, and R-H Park. Robust adaptive segmentation of range images. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 20(2):200–205, 1998.

[7] X. Yu, T.D. Bui, and A. Krzyzak. Robust estimation for range image segmentation and reconstruction. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 16(5):530–538, 1994.

[8] J.V. Miller and C.V. Stewart. MUSE: Robust surface fitting using unbiased scale estimates. In *Proc. Computer Vision and Pattern Recognition'96*, pages 300–306, June 1996.

[9] P. Meer, C.V. Stewart, and D. Tyler. Robust computer vision: an interdisciplinary challenge. *Computer Vision and Image Understanding*, 78:1–7, April 2000.

[10] C.V. Stewart. Robust parameter estimation in computer vision. *SIAM Reviews*, 41(3):513–537, September 1999.

[11] H. Wang and D. Suter. LTSD: A highly efficient symmetry-based robust estimator. In *Proceedings ICARCV2002*, pages 332–337, 2002.

[12] K. Fukunaga and L.D. Hostetler. The estimation of the gradient of a density function, with applications in pattern recognition. *IEEE Trans. Info. Theory*, IT-21:32–40, 1975.

[13] D. Comaniciu, V. Ramesh, and A. Del Bue. Multivariate saddle point detection for statistical clustering. In A. Heyden, G. Sparr, M. Nielsen, and P. Johansen, editors, *Proceedings of European Conf. Computer Vision (ECCV'02), Copenhagen, Denmark*, volume 3, pages 561–576, 2002.

[14] K. Kanatani. *Statistical Optimization for Geometric Computation: Theory and Practice*. Elsevier Science, Amsterdam, The Netherlands, April 1996.

[15] H. Wang and D. Suter. Variable bandwidth QMDPE and its application in robust optic flow estimation. In *Proceedings ICCV03, International Conference on Computer Vision, Nice, France*, to appear, 2003.

[16] M.P. Wand and M. Jones. *Kernel Smoothing*. Chapman and Hall, 1995.

[17] A. Bab-Hadiashar and D. Suter. Robust segmentation of visual data using ranked unbiased scale estimate. *ROBOTICA, International Journal of Information, Education and Research in Robotics and Artificial Intelligence*, 17:649–660, 1999.

[18] A. Bab-Hadiashar and D. Suter. Range and motion segmentation. In A. Bab-Hadiashar and D. Suter, editors, *Data Segmentation and Model Selection for Computer Vision*, chapter 5, pages 119–142. Springer-Verlag, 2000.

[19] H. Wang and D. Suter. MDPE: A very robust estimator for model fitting and range image segmentation. *Int. J. of Computer Vision*, page to appear, 2003.

[20] A. Bab-Hadiashar, N. Gheissari, and D. Suter. Robust model based visual data using segmentation. In R. Kasturi, D. Laurendeau, and G. Suen, editors, *Proceedings of ICPR2002*, volume 2, pages 753–757, 2002.

Acknowledgment

This work was carried out with the support of the Australian Research Council: grant A10017082. in particular, the second author was supported by that grant. This support is gratefully acknowledged.

Dept. Elect. & Comp. Syst. Eng., P.O. Box 35, Monash University, Clayton, Victoria, 3800, Australia

E-mail address: `[d.suter,hanzi.wang]eng.monash.edu.au`