

Estimating Model Parameters and Boundaries By Minimizing a Joint, Robust Objective Function

Charles V. Stewart Kishore Bubna Amitha Perera

Department of Computer Science, Rensselaer Polytechnic Institute Troy, New York 12180-3590
 {stewart,bubnak,perera}@cs.rpi.edu

Abstract

Many problems in computer vision require estimation of both model parameters and boundaries, which limits the usefulness of standard estimation techniques from statistics. Example problems include surface reconstruction from range data, estimation of parametric motion models, fitting circular or elliptic arcs to edgel data, and many others. This paper introduces a new estimation technique, called the “Domain Bounding M-Estimator”, which is a generalization of ordinary M-estimators combining error measures on model parameters and boundaries in a joint, robust objective function. Minimization of the objective function given a rough initialization yields simultaneous estimates of parameters and boundaries. The DBM-Estimator has been applied to estimating line segments, surfaces, and the symmetry transformation between two edgel chains. It is unaffected by outliers and prevents boundary estimates from crossing even small magnitude discontinuities.

1 Introduction

Estimation techniques from the statistics literature are not well-suited for many computer vision problems. Standard definitions of least-squares and maximum likelihood estimators give procedures for estimating the parameters of lines, circles, ellipses, planes, quadratic surfaces and many other function models from a data set, and these methods may be made robust [18, 20] and unbiased [4]. Missing from these estimation techniques and indeed from the function models themselves is any notion of the domain or subset of the data over which the models apply. This is problematic because function models in computer vision are usually applicable over limited domains, which generally are not known in advance. For example, a surface in the world has a limited spatial extent and therefore a surface model estimated from range data must cover a limited image region; the same observation holds about the image description of a moving object. In a very different context, opposing outlines of a symmetric object should map onto each other using an estimated symmetry transform [5, 7], but parts of the object and of the nearby background may not be symmetric and therefore should not be mapped by the transform. In these and many other cases, accurate estimation of model or transform parameters requires accurate estimation of the domain, the boundaries, or the subset of the data over which the model




Figure 1: Example illustrating the DBM-Estimator in fitting a line segment to synthetic 2D points near a small step discontinuity: the upper frame shows a skewed initial fit; the middle frame shows the DBM-Estimator recovering from the skewed fit and growing the segment to the left but **not** across the discontinuity on the right (in fact it pulls away slightly); the lower frame shows the converged fit and boundary estimate.

applies. Inclusion of points outside the proper domain can produce biased parameter estimates, while inclusion of too few points can lead to inaccurate parameter estimates and redundant models.

This problem, referred to here as the “domain estimation problem”, has been addressed under many different guises in many different areas of computer vision. Existing techniques, which work with varying degrees of effectiveness [12], include many different segmentation methods [1, 9], snake and balloon models [6, 16], robust estimators [18, 20, 21], and mixture models [14, 22]. The approach introduced in this paper is quite different and more general than any of these. It attacks the limitations of standard estimation techniques directly by incorporating model parameters and boundaries into a single, robust objective function. The objective function is a generalization of standard M-estimator objective functions, and depends on two distance measurements for each point: the usual fitting error or residual distance and the distance of the point from the current domain. The latter will be called the “domain distance”.

The new technique, called the Domain Bounding M-Estimator (“DBM-Estimator”) has several important properties. These can be contrasted against existing techniques for the domain estimation problem.

- Boundary estimation is entirely driven by the optimiza-

tion process, placing it on par with parameter estimation and making heuristic growth techniques unnecessary. The optimization process may cause the surface boundary to grow in some directions and shrink in others, with growth rates determined by the objective function’s gradient. For 2D domains, this growth is similar to that of a balloon model [6] with no need to simulate an artificial force!

- The DBM-Estimator inherits local robustness properties of standard M-estimators while avoiding their susceptibility to leverage points, which causes their poor global robustness (breakdown point).
- Domain growth is determined by the cumulative effect of points in small regions near the domain boundary. This means the objective function does not need to be tuned to the tail of the error distribution to ensure proper domain growth [19], allowing the DBM-Estimator to distinguish between noise and small scale discontinuities.

An example illustrating the behavior of the DBM-Estimator is shown in Fig. 1.

Despite its advantages, the DBM-Estimator does have limitations. First, since it uses a gradient search technique, its stopping point is a local minimum, making initialization important. Usually, however, it is sufficient to avoid seed regions that substantially overlap a boundary. Second, though working better than previous techniques, the DBM-Estimator is not capable of localizing extremely small magnitude discontinuities (e.g. step heights of 2.0 standard deviations).

The DBM-Estimator is illustrated here on three different problems: line segment fitting for 2D data (Fig. 1); estimating the parameters and boundaries of a planar surface in a range image; and estimating the planar projective transform characterizing the symmetric relation between (parts of) two edgel chains. Application of the DBM-Estimator to other problems should be relatively straightforward.

2 M-Estimators

This section briefly reviews standard M-estimators [8] to set the context for derivation of the DBM-Estimator. Consider N data points \mathbf{x}_i , model parameters \mathbf{a} , which are to be estimated, and a function $r(\mathbf{x}_i; \mathbf{a})$ mapping a point and a parameter vector to a residual. In the simple case of data points $\mathbf{x}_i = (x_i, z_i)^T$ and linear regression, $r(\mathbf{x}_i; \mathbf{a}) = z_i - (1, x_i)^T \mathbf{a}$. The scale normalized residual is $u_i = r(\mathbf{x}_i; \mathbf{a})/\sigma_i$, where σ_i is the scale or noise standard deviation which may need to be estimated. The M-estimate of the parameters is

$$\hat{\mathbf{a}} = \underset{\mathbf{a}}{\operatorname{argmin}} \sum_i \rho(u_i), \quad (1)$$

where $\rho(u)$ is the “robust loss function.” Among the many possible forms of $\rho(\cdot)$, only those reaching an asymptotic




Figure 2: M-estimator ρ function (left) and weight function (right).




Figure 3: Illustration of a domain over which a surface model might apply in a range image. Both the interior $I(D)$ and boundary regions $\Omega(D)$ are shown.

value for large u are at all effective near discontinuities [21]. One of these, the Tukey biweight [2] (Fig. 2), is adapted below to the DBM-Estimator. Its equation is

$$\rho_T(u) = \begin{cases} \frac{B^2}{6} \left[1 - \left(1 - \frac{u^2}{B^2} \right)^3 \right] & \text{if } |u| \leq B, \\ \frac{B^2}{6} & \text{if } |u| > B. \end{cases} \quad (2)$$

Most often, equation 1 is solved by iteratively reweighted least squares (IRLS) [11] using a weight function entirely determined by ρ : $w(u) = \psi(u)/u = \rho'(u)/u$. Figure 2 shows the weight function computed from (2).

3 General Formulation

The DBM-Estimator is formulated by defining the domain and domain distance, by defining the objective function, and by specifying an iterative search technique. These are considered in turn below.

3.1 The Domain and the Domain Distance

The domain D consists of an interior region $I(D)$ and a boundary region $\Omega(D)$. These will be image regions for the surface fitting application (Fig. 3), sections of arc length for the curve symmetry application, and just an interval of the x -axis for line segment fitting. The next step is to define the distance $d_D(\mathbf{x}_i, \mathbf{x}_j)$ in D between any two points, \mathbf{x}_i and \mathbf{x}_j . This may be a Euclidean distance in the image plane or the difference in parameter values for an arc length parameterization of a curve. The *domain distance* of a point \mathbf{x}_i is then defined as

$$d_i = \min_{\mathbf{x}_j \in I(D)} d(\mathbf{x}_i, \mathbf{x}_j). \quad (3)$$

Points in $I(D)$, “interior points”, have $d_i = 0$. $\Omega(D)$ is then defined as points i such that $0 < d_i < d_{max}$, for some d_{max} .

Now that both a domain distance and a residual distance (Section 2) are defined, two different types of outliers must

be defined. *Domain outliers* are points with large d_i . *Residual outliers* are points with large magnitude u_i . Some points will be both. Neither type should unduly influence the estimated model parameters or boundaries.

3.2 The Objective Function Estimate

The DBM-Estimator objective function is

$$\mathcal{E}(\mathbf{a}, D) = \sum_i \rho(u_i, d_i), \quad (4)$$

and the DBM-Estimate of model parameters and boundaries is

$$(\hat{\mathbf{a}}, \hat{D}) = \operatorname{argmin}_{\mathbf{a}, D} \sum_i \rho(u_i, d_i), \quad (5)$$

where $\rho(u, d)$ is the generalization of the robust loss function for ordinary M-estimators. (No specific representation of the domain is assumed here, so the domain is just written as D , with \hat{D} as its estimate. Also, the dependence of u_i and d_i on \mathbf{a} and D is left out of these equations to simplify notation.) The summation is taken over all data points, although, as will be clear later, evaluating this summation will only require points in $I(D)$ and $\Omega(D)$. Defining the generalized loss function $\rho(u, d)$ is clearly the heart of the issue.

3.3 Designing $\rho(u, d)$

Intuitively, $\rho(u, d)$ should be designed to grow the domain to include residual inliers, stop growth or shrink away from discontinuities, and be unaffected by residual outliers — properties evident in the example in Fig. 1. Hand specified data illustrating the requirements for line segment fitting are shown in Fig. 4; the intuitions are the same for the other applications. Plots (a) and (b) show small scale depth and orientation discontinuities where the domain should not expand further to the right despite points in $\Omega(D)$ having fit residuals less than 3σ . (This illustrates a limitation of region growing techniques which typically include as inliers all neighboring points falling within $\pm 3\sigma$ of an extrapolated model.) In (c), the structure to the right, well outside the domain, should neither bias the line segment parameters nor force expansion of the domain. In (d), residual outliers falling in $\Omega(D)$ should not prevent domain growth from crossing occluded regions that are at most d_{max} wide.

Here are four requirements on $\rho(u, d)$ imposed by the required behavior.

1. For points in $I(D)$, ρ should behave as a standard M-estimator. Thus, $\rho(u, 0) = \rho_T(u)$.
2. $\rho(u, d)$ should be constant for $d > d_{max}$, independent of u , preventing domain outliers from influencing model parameters (Figure 4c). Thus, as d increases from 0 to d_{max} , $\rho(u, d)$ should evolve from $\rho_T(u)$ to a constant function (for fixed d).¹ Points that are initially domain

¹This has the added advantage of reducing $\partial\rho/\partial u$ (and therefore the IRLS weight function) with increasing d , so that points in $\Omega(D)$ have reduced influence on estimated fit parameters.

outliers may eventually be drawn into the domain, but only through its continued expansion.

3. For small u , $\rho(u, d)$ should decrease monotonically as d decreases, driving the domain toward the point. For larger u , such as for points on the tail of the residual distribution, $\rho(u, d)$ should decrease monotonically as d increases, causing the domain to back away from the point. The transition value of u between these two effects, a value denoted as u_0 and shown in Fig. 5(a) and (b), is important since the cumulative influence of the points in a small section of $\Omega(I)$ will determine whether D expands or shrinks locally. The exact choice of u_0 is based on statistical analysis summarized in Sec. 3.5.
4. Residual outliers within $\Omega(D)$ should not, by themselves, prevent continued domain growth. This requirement implies that for all d , $\lim_{u \rightarrow \infty} \rho(u, d) = \rho_\infty$.

The first three requirements are met by a simple function (see Figure 5a):

$$\rho(u, d) = \begin{cases} \rho_T(u) & \text{if } d \leq 0, \\ (1 - f(d))\rho_T(u) + f(d)\rho_\infty & \text{if } 0 < d < d_{max}, \\ \rho_\infty = \rho_T(u_0) & \text{if } d \geq d_{max}, \end{cases} \quad (6)$$

This joint loss function combines a standard M-estimator loss function for points in $I(D)$ (the first requirement) with a constant loss function for points outside $\Omega(D)$ (the second requirement). The function $f(d)$ controls the transition from one to the other as a point moves through $\Omega(D)$, and is piecewise quadratic over $[0, d_{max}]$ with $f(0) = 0$, $f(d_{max}) = 1$, and $f'(0) = f'(d_{max}) = 0$. All three requirements are clearly met by $\rho(u, d)$.

Unfortunately, $\rho(u, d)$ from (6) does not satisfy the fourth requirement and therefore does not realize the desired behavior pictured in Figure 4(d). As $\rho(u, d)$ is currently specified, residual outliers in $\Omega(D)$, which may be caused by partial occlusions or clustered outliers, can prevent growth of D or even cause it to shrink. The problem is that the first three requirements imply $\rho(u, d)$ has different asymptotic values (for large u) for different d ; the fourth implies these asymptotes must be the same. The solution is to relax the first requirement, altering $\rho_T(u)$ to redescend to an asymptote of ρ_∞ :

$$\rho_M(u) = \begin{cases} \frac{B^2}{6} [1 - (1 - \frac{u^2}{B^2})^3] & \text{if } |u| \leq B, \\ (\frac{B^2}{6} - \rho_\infty) [1 - (1 - \frac{(C-|u|)^2}{(C-B)^2})^3] + \rho_\infty & \text{if } B < |u| \leq C, \\ \rho_\infty & \text{if } |u| > C. \end{cases} \quad (7)$$

Substituting $\rho_M(u)$ for $\rho_T(u)$ in (6) yields the current $\rho(u, d)$.

Plots of $\rho(u, d)$ and the associated weight function to be used in IRLS parameter estimation are shown in Fig. 5(b)




Figure 4: Data illustrating the desired behavior of domain growth. The plots show a cross-section of an image, with the x axis as the image coordinate and the z axis as the depth. In each case, a linear “surface” is shown, with dashed lines above and below showing approximate 3σ uncertainty bands. $\Omega(D)$ and $I(D)$ are indicated by segments of the x axis.




Figure 5: Plots illustrating $p(u, d)$ (using $d_{max} = 3$); (a) shows a version meeting the first four requirements; (b) shows the final version; and (c) shows its associated weight function. Only the positive u axis is shown. In studying these plots, the reader should note that the expected value of $|u|$ under a Gaussian distribution ≈ 0.79 . This means a substantial majority of inlier fit residuals will have $|u| < u_0$.

and (c). It is important to note that the negative weights for larger u seen in (c) have only a minor effect on estimation, which in fact is advantageous: near discontinuities they push the estimated model (very) slightly away from data corresponding to the other structure.

3.4 Minimization

Conceptually, solving (5) is straightforward. The minimization mixes IRLS steps to refine the model parameters and differential search steps on the domain. The IRLS steps are exactly as in ordinary M-estimators. The differential search steps are applied to each section of the discretized boundary $\Omega(D)$. For example, in line segment fitting, the boundary segments are the two intervals $[x_l - d_{max}, x_l)$ and $(x_r, x_r + d_{max}]$ (see Fig. 4), which are determined by x_l and x_r . Hence, the differential search steps compute changes to x_l and x_r . Considering just x_r , the partial derivative of the objective function is

$$\frac{\partial \mathcal{E}}{\partial x_r} = \sum_{\{i|x_i \in (x_r, x_r + d_{max})\}} \frac{\partial \rho(u_i, d_i)}{\partial d} \cdot (-1).$$

Only partial derivatives for points in the boundary segment are non-zero, and the (-1) is just the derivative of d_i with respect to x_r . This partial derivative is normalized by its expected value calculated assuming all boundary segment residuals are normally distributed. The step change in x_r , which may be positive or negative depending on $\frac{\partial \mathcal{E}}{\partial x_r}$, is then just proportional to this normalized derivative. Extension to two-dimensional domains is straight-forward for polygonal representations of the domain boundary. In this case, the step computed is along the segment normal (see Sec. 5).




Figure 6: The minimum u_0 required for different d_{max} such that there is a 0.99 probability of domain growth across a large interval for 1-d (upper, dotted curve) and 2-d (lower, solid curve) domains.

3.5 Parameter Tuning and Scale

The DBM-Estimator is controlled by four main parameters, B , C , u_0 and d_{max} and is influenced by estimates of scale (noise standard deviation), σ . The first two parameters can be easily fixed and ignored. B , inherited from the standard Tukey biweight, is set at 4.0 [11], and $C = 2B$. The parameters u_0 and d_{max} are interdependent.² They must be set together to determine the sensitivity of the DBM-Estimator to small magnitude discontinuities, while ensuring continued growth across a fit’s correct domain. The relationship between u_0 and d_{max} can be established numerically by finding for each d_{max} the minimum u_0 such that there is a 0.99 probability of (correct) domain growth across a large (e.g. 100 pixel) interval. These values are shown in Fig. 6 for one-dimensional domains and for two-dimensional domains, where a polygonal boundary segment is assumed to be 5 pixels wide. The choice of d_{max} therefore dictates u_0 . Rather than using a fixed d_{max} , however, the DBM-

²Recall that u_0 is defined in terms of scale normalized residuals (Sec. 2).

Estimator changes it in a coarse-to-fine manner. Using large d_{max} initially prevents the DBM-Estimator from crossing small magnitude discontinuities, and allows it cross narrow occluded regions (gaps in the data). Using a small final d_{max} allows more precise boundary localization. We use the initial d_{max} and the results shown in Fig. 6 to fix u_0 .

Three comments about scale, σ , are important. First, robust estimation of scale within the DBM-Estimator is straightforward using standard techniques [8], although scale should be fixed once the domain is sufficiently large. Second, since both noise and modeling error contribute to the value of σ , fixing σ from outside the DBM-Estimator controls tolerance of modeling error in domain estimation. See Sec. 6 for an example of the effect of this. Third, σ , whether estimated or fixed externally, should not be used directly in calculating the scale normalized residual u_i for points in $\Omega(D)$. The calculation of u_i must instead use prediction interval techniques [19] to factor in the uncertainty of estimated model parameters. This is especially important to allow recovery from poor initial fits, as shown in Fig. 1.

3.6 Instantiating the DBM-Estimator

Applying the DBM-Estimator to a specific problem requires four issues to be addressed. (1) A mechanism is needed to provide an initial domain, and perhaps initial model parameter estimates. (2) The domain representation and domain distance metric must be specified. (3) A method for weighted least-squares model parameter estimation is needed. (4) A technique for calculating fit residuals and scale must be provided. All but the second are required for properly formulated estimation problems [15].

4 Line Segment Estimation

This and the next two sections briefly describe application of the DBM-Estimator to the problems of line segment estimation, surface fitting, and symmetry estimation.

The simple example of line segment fitting has been used to describe the DBM-Estimator throughout this paper. Experiments are presented here to characterize DBM-Estimator performance in accurate boundary location given a small starting region. Results are shown in Fig. 7 for step and crease discontinuities and for two different starting regions. Step discontinuities are characterized by scale normalized step height h/σ , while crease discontinuities are characterized by angle α , such that the two lines forming the discontinuity make an interior angle of $\pi - 2\alpha$. Fig. 1 gives an example near a step edge with $h/\sigma = 3.5$. The fit of interest is always the left half of the discontinuity. Starting regions are either far from the discontinuity or close to it. Fig. 7 shows that the starting region makes a difference for small magnitude discontinuities: starting far away allows the fit to stabilize before reaching the discontinuity. Still, however, discontinuities are localized extremely accurately by the DBM-Estimator for all but the smallest discontinu-




Figure 7: The average bias in the position of x_r , the right domain boundary of the line segment, for step (left) and crease (right) discontinuities. The left and right halves of the discontinuities were each 100 pixels wide. Datasets with $\sigma = 1.0$ noise and 10% outliers were generated from the discontinuity models. In each plot the dotted line shows the results of starting from near the discontinuity while the solid line shows the results of starting far away.

ities. This performance is far better than that of ordinary robust estimators [21], which usually fail at step heights below 6σ .

5 Bounded surface estimation

A bounded surface estimated from range data is represented by a low-order polynomial equation $z(x, y)$ and a two-dimensional region $I(D)$. This region is characterized by an enclosing polygon and one or more non-intersecting interior polygons corresponding to holes in the surface. Initialization requires only a small initial region, which may be specified manually or by finding the centers of connected regions having similar robust estimates of local surface normals. The domain polygon consists of short line segments, with differential search steps calculated in two parts: (a) the change in position along the normal to each segment is found as described in Sec. 3.4, and then (b) the change in each vertex is found by combining motion vectors from adjacent segments. Segments growing too large are split; those shrinking too small are merged. Changes in topology are detected and corrected along the polygons [17]. The overall effect is similar to that of a snake or balloon model [6, 16], without any simulation of artificial forces. Currently there is no smoothing of the DBM-Estimator boundary polygon, so it appears jagged. Example results are shown in Fig. 8.

6 Symmetry Transform Estimation

Given two pairs of edge sequences, $\mathbf{S} = \{\tilde{p}_i | i = 1, \dots, n\}$ and $\mathbf{S}' = \{\tilde{p}'_i | i = 1, \dots, m\}$, the problem is to find a correspondence map $C : \{1, \dots, n\} \rightarrow \{1, \dots, m\}$, a subsequence $\{\tilde{p}_{i_{min}}, \tilde{p}_{i_{max}}\}$ of \mathbf{S} , and a plane projective transform T mapping each \tilde{p}_i in the subsequence onto $\tilde{p}'_{C(i)}$ (with small error). Symmetry is enforced by requiring T to map $\tilde{p}'_{C(i)}$ onto \tilde{p}_i as well, making T at least approximately an involution. An example may be seen in Fig. 9 where the left and right halves of the outline appear symmetric over a large interval.

The vagueness of this problem description, characteristic of many domain estimation problems, may be eliminated by casting it into the DBM-Estimator formulation using the steps outlined in Sec. 3.6: (1) An initial, small subsequence




Figure 8: Domains of bounded planar surfaces estimated by the DBM-Estimator using range data taken from the South Florida dataset [12]. The top left and right show intermediate stages of growth on a genus 1 surface, before and after changes in topology. The middle shows ground truth segmentation results. The bottom shows DBM-Estimator surface interiors, each of which is estimated independently based on automatically extracted seed regions. Observe that currently the results show double boundaries at discontinuities because adjacent domain estimates are unrelated. Of special note here, the three center surfaces (15, 21, and 20) on the foreground object are particularly difficult to estimate and most published techniques merge them [12, 19].

and associated correspondence map is assumed to be provided by an algorithm such as described in [5, 7]. (2) The domain and domain distance are simply represented by an arc-length reparametrization, $s(i)$, of the sequences. This makes the domain representation exactly as in line segment estimation with the x coordinate being replaced by arc-length. From this we can explicitly write the problem formulation as

$$\operatorname{argmin}_{T, s_1, s_r} \sum_{i=1}^n [\rho(\|T\tilde{p}_i - \tilde{p}'_{C(i)}\|/\sigma, d_{s(i)}) + \rho(\|T\tilde{p}'_{C(i)} - \tilde{p}_i\|/\sigma, d_{s(i)})]. \quad (8)$$

Clearly T depends on C and vice-versa, so they are estimated by alternating matching based on estimated T and refining T based on new matches, as in [3]. (3) Weighted least-squares parameter estimation is a straightforward extension of Hartley's [10] centering and normalization algorithm. (4) As discussed above, scale in this application represents modeling error — inexactness of the symmetry — more than noise, and therefore scale is assumed to be set externally. Example results illustrating use of the DBM-Estimator in estimating T and the symmetry boundaries, with emphasis on the role of scale are shown in Fig. 9.

7 Discussion

The DBM-Estimator has been proposed to address the general problem of simultaneously and robustly estimating model parameters and model boundaries. This model may be a curve, a surface, a transformation, or an image motion, depending on the problem instance. Formulation of the DBM-Estimator is intuitive, controlled by only a few parameters, and easily adapted to numerous problems. The major issue in each problem instance is defining the domain and the domain boundary representation. Example implementations and results on three problem instances have been shown. These illustrate the breadth of possible applications of the DBM-Estimator, different possible domain boundary representations, and the ability of the DBM-Estimator to tolerate outliers and localize small scale discontinuities. The polygonal boundary representation for 2D domains has the added ability to adapt to changes in the domain boundary topology during minimization of the DBM-Estimator objective function.

Several issues are still under investigation. First we have begun experimentation with incorporating the DBM-Estimator into a mixture model formulation [14, 22]. Promising results, in terms of recognizing and localizing extremely small-scale boundaries, have been obtained in the line segment estimation problem. Second, we are working on improvements to the smoothness of the polygonal boundary representation for surface fitting. Third, we are developing a generalization that uses piecewise polynomial (spline) curves and surfaces instead of a single model for the entire domain. Together, these improvements will add to the utility of what we believe is already a widely applicable new technique.

Acknowledgements

The authors gratefully acknowledge the support of the National Science Foundation under award IRI-9408700.

References

- [1] F. Arman and J. Aggarwal. Model-based object recognition in dense range images. *Computing Surveys*, 25(1):5–43, 1993.
- [2] A. E. Beaton and J. W. Tukey. The fitting of power series, meaning polynomials, illustrated on band-spectroscopic data. *Technometrics*, 16:147–185, 1974.




Figure 9: Example of symmetry transform estimation. The left shows an approximately symmetric (with small perspective effects) orchid with two overlaid edge sequences shown as the white outline. The center and right show results of estimating symmetry starting from small corresponding subsequences taken from the lower left and right corners of the orchid. Each image displays the mapping of the interior and boundary subsequences from the right half of the outline onto the left. For the center image, σ was fixed at 2 pixels, so the DBM-Estimator domain could not grow across intervals of minor asymmetry (even though the boundaries enter these intervals). For the right image, σ was fixed at 4 pixels, so the symmetry grew much farther. The top end of the mapped subsequence shows that the boundary interval has entered a region of strong asymmetry, so growth has stopped.

- [3] P. Besl and N. McKay. A method for registration of 3-d shapes. *PAMI*, 14(2):239–256, 1992.
- [4] J. Cabrera and P. Meer. Unbiased estimation of ellipses by bootstrapping. *PAMI*, 18(7):752–756, 1996.
- [5] T. Cham and R. Cipolla. Symmetry detection through local skewed symmetries. *Image and Vision Computing*, 13(5):439–450, 1995.
- [6] L. Cohen and I. Cohen. Finite-element methods for active contour models and balloons for 2-d and 3-d images. *PAMI*, 15(11):1131–1147, 1993.
- [7] R. Curwen, C. Stewart, and J. Mundy. Recognition of plane projective symmetry. In *Proc. ICCV*, pages 1115–1122, 1998.
- [8] F. R. Hampel, P. J. Rousseeuw, E. Ronchetti, and W. A. Stahel. *Robust Statistics: The Approach Based on Influence Functions*. John Wiley & Sons, 1986.
- [9] R. Haralick and L. Shapiro. *Computer and Robot Vision*, volume 1. Addison-Wesley, 1992.
- [10] R. Hartley. In defence of the 8-point algorithm. In *Proc. ICCV*, pages 1064–1070, 1995.
- [11] P. W. Holland and R. E. Welsch. Robust regression using iteratively reweighted least-squares. *Commun. Statist.-Theor. Meth.*, A6:813–827, 1977.
- [12] A. Hoover, G. Jean-Baptiste, X. Jiang, P. Flynn, H. Bunke, D. Goldgof, K. Bowyer, D. Eggert, A. Fitzgibbon, and R. Fisher. An experimental comparison of range image segmentation algorithms. *PAMI*, 18(7):673–689, 1996.
- [13] P. J. Huber. *Robust Statistics*. John Wiley & Sons, 1981.
- [14] S. Ju, M. Black, and A. Jepson. Skin and bones: Multi-layer, locally affine, optical flow and regularization with transparency. In *Proc. CVPR*, pages 307–314, 1996.
- [15] K. Kanatani. *Statistical Optimization for Geometric Computation: Theory and Practice*. Elsevier, 1996.
- [16] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *IJCV*, 1(4):321–331, 1988.
- [17] T. McInerney and D. Terzopoulos. Topologically adaptable snakes. In *Proc. ICCV*, pages 840–845, 1995.
- [18] P. Meer, D. Mintz, A. Rosenfeld, and D. Y. Kim. Robust regression methods for computer vision: A review. *IJCV*, 6:59–70, 1991.
- [19] J. V. Miller and C. V. Stewart. Prediction intervals for surface growing range segmentation. In *Proc. CVPR*, pages 1027–1033, 1997.
- [20] P. J. Rousseeuw and A. M. Leroy. *Robust Regression and Outlier Detection*. John Wiley & Sons, 1987.
- [21] C. V. Stewart. Bias in robust estimation caused by discontinuities and multiple structures. *PAMI*, 19(8):818–833, 1997.
- [22] Y. Weiss. Smoothness in layers: Motion segmentation using nonparametric mixture estimation. In *Proc. CVPR*, pages 520–526, 1997.