Abstract

Visual perception is an essential sense for robots whose task is to learn about the environment. A crucial role is played by the representation of the world that the system must build and maintain, and by the algorithms adopted to compute the representation itself. In the recent years range vision and surface-based representations have been increasingly used for robotic applications. A surface-based representation of a scene is typically built from a range image by segmenting the image into a collection of surface patches, whose position, orientation and shape in 3-D space is estimated. We present a system which computes such a representation efficiently from range images. We consider patches belonging to a qualitative shape catalogue suggested by differential geometry as the main features of our representation. The surface is segmented and the patches classified by estimating the sign of mean and gaussian curvature at each nonsingular point. The parameters of the local differential structure of the surface are also estimated (augmented Darboux frame). We discuss shape distortion effects introduced by gaussian smoothing, modelled according to the diffusion paradigm. We propose a shape-preserving boundary condition for the diffusion equation. We precompute depth and orientation discontinuities maps and use them to restrict the diffusion process to non-discontinuity points, thus avoiding the creation of spurious curved regions around discontinuity contours. Some experimental results and quantitative tests are also given.

1. Introduction

Visual perception is an essential sense for robots whose task is to learn about its environment. In the past years *range vision* has been increasingly used for robotic applications involving acquisition of information about the world by a robot system, as for instance in intelligent autonomous navigation (Crowley [1], Thorpe [2]). A crucial role is played by the representation of the world that the system must build and maintain. Such a representation must contain the information necessary for the robot to perform its task; moreover, it must be computed in a time acceptable for the given application. *Surface-based representations* ([1], [14], [18]) seem to offer attractive properties for this aim. Surfaces are large and stable features; they apply to curved objects; there are generally fewer surface patches than edges in an image; finally, information about surfaces is directly contained in range images. In order to build a representation, significant surface features must be extracted from raw range data and a proper description devised. The input range image is typically segmented into a collection of *homogeneous surface patches*, whose position, orientation and shape in 3-D space is estimated.

There are several problems, however, in computing such a representation. Second-order derivatives must be estimated from the range image in order to obtain the curvatures. This is a noise-sensitive operation. Range images are quantized, typically over 255 levels. Quantization noise proves worse than the distortions introduced by most state-of-the-art range scanners (Naidu and Fisher [13]) in corrupting curvature estimates. Noise smoothing is therefore necessary, and gaussian smoothing is a commonly adopted technique. Unfortunately, gaussian smoothing introduces several side-effects. The first is the general decreasing of the absolute value of curvatures all over the image. This is due to the averaging nature of gaussian smoothing (Lowe [4]). A second side-effect of gaussian smoothing is that the sign of H and K can be changed near the boundary. The effect of this distortion increases with the standard deviation of the gaussian kernel. Finally, orientation discontinuities contours, e.g. the visible edges of a cube, are smoothed into spurious curved patches.

We present a system (sketched in Figure 1) which computes a qualitative surface-based description of a scene from a range image. Smoothing is performed by solving a diffusion equation with an efficient numerical scheme. The sign of the mean and gaussian curvature (H and K respectively) are estimated subsequently at each nonsingularity point. The patches we consider in this paper belong to a *qualitative shape catalogue* suggested by differential geometry, in which each class is characterized completely by the sign of the mean and gaussian curvatures. In actual fact, the system computes an exhaustive local representation $\mathcal{D}(P) = (P, d_1, d_2, N, \kappa_1, \kappa_2)$ known as the *augmented Darboux frame* (Ferrie *et al.* [3]), where d_1 and d_2 are the principal directions, N is the normal to the surface and κ_1, κ_2 are the principal curvatures at P. However, qualitative shape estimates are more reliable than quantitative surface structure estimates. The final output is a collection of surface patches, for each of which a qualitative shape class and estimates of the differential structure are given. An adjacency graph expresses the structure of the scene as observed by the sensor. This versatile representation embeds both *intrinsic* patch properties like shape class, and *relational* patch properties like adjacency. It seems appropriate for several robotic tasks, including location, recognition (Fisher [14]), intelligent navigation (Crowley [1], Thorpe [2]) and planning visual strategies (Bajcsy [16], Sakane [17]).

The boundary distortion problem is approached with a curvature-preserving boundary treatment for diffusion smoothing, in the form of an *adaptive-leakage boundary condition* enforced at depth and orientation discontinuities. This method has some advantages over related methods, in particular Cai's fixed small-leakage boundary treatment [5]. In order to avoid the creation of noisy patches due to discontinuity smoothing, we precompute depth and orientation discontinuities maps and use them to restrict the diffusion process to non-discontinuity points, thus avoiding the creation of spurious curved regions around discontinuity contours. We have tested our system in order to estimate quantitatively its performance. We sketch in this paper a few results concerning accuracy in H estimation and classification. Similar experiments on accuracy have been reported by Flynn and Jain [15].

Figure 1: Architecture of the system.

2. Discontinuity detection

Depth and orientation discontinuity maps are computed first. The raw range data are first initially smoothed using a nonlinear, conservative filter to remove spikes of noise.

2.1. Depth discontinuities

At each image point P on a surface $(x, y, h(x, y))$, the absolute value of the directional derivative

$$d(P, \mathbf{n}) = \frac{\partial h}{\partial \mathbf{n}}(P)$$

is estimated in a 3x3 neighbourhood and for four directions of $\mathbf{n}$. If this quantity is large enough P is labelled as a depth discontinuity point. At present, discontinuities are detected by comparing $\left|\frac{\partial h}{\partial \mathbf{n}}\right|$ with a user-supplied threshold τ_d, tuned manually to an optimal value (usually between 10 and 20 pixels). All subsequent computations are performed on world-coordinate values.

2.2. Orientation discontinuities

Orientation discontinuities are detected as loci of discontinuity points for the tangent plane. Given the parametrization $\mathbf{s}(x, y) = (x, y, h(x, y))$ of surface S and a point $P \in S$, the tangent plane T at P is the plane through P and perpendicular to the normal

$$\mathbf{N}(P) = \frac{(-h_x, -h_y, 1)}{(1 + h_x^2 + h_y^2)^{1/2}}$$

where the subscripts indicate partial differentiation. Least squares are used to estimate $\mathbf{N}(P)$ in a 3x3 local environment of P, so that the estimated gradient (h_x, h_y) at P is given by

$$h_x = \frac{\sum_{k \in I(P)} \Delta x_k \Delta h(\Delta x_k)}{\sum_{k \in I(P)} {\Delta x_k}^2} \quad h_y = \frac{\sum_{k \in I(P)} \Delta y_k \Delta h(\Delta y_k)}{\sum_{k \in I(P)} {\Delta y_k}^2}$$

where $I(P)$ is a local neighbourhood of P, Δx_k and Δy_k are the incremental steps in the x and y directions and $\Delta h(\Delta q)$ is the increment of h over the increment Δq. In the continuous case, a discontinuity point for T corresponds to a discontinuity in the normal field $\mathbf{N}$. In the discrete case, an orientation discontinuity is a point P such that at least one point P_1 in a local neighbourhood of P satisfies

$$|\mathbf{N}(P) \cdot \mathbf{N}(P_1)| < \tau_o$$

where τ_o is a user-defined threshold. In our experiments, usual values of τ_o were between 0.75 and 0.9.

3. Diffusion smoothing

Diffusion smoothing provides an elegant mathematical framework for gaussian smoothing and scale-space analysis (Witkin [7], Lindeberg [8]). Intuitively, diffusion smoothing regards a surface $S = (x, y, h(x, y))$, expressed here through a parametrization on the cartesian plane, as the initial configuration of a heat distribution $u(x, y, t)$ at $t = 0$, which evolves according to the diffusion equation

$$\frac{\partial u}{\partial t} = b\nabla^2 u \tag{1}$$

with initial value $u(x, y, 0) = h(x, y)$.

The closed-form solution of this problem is the gaussian convolution

$$u(x, y) = \frac{1}{4\pi bt} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} h(\xi, \eta) exp[-\frac{(\eta - x)^2 + (\xi - y)^2}{4bt}] d\xi d\eta$$

where the relation

$$\sigma = \sqrt{2bt} \tag{2}$$

links the time t to the standard deviation σ of gaussian smoothing. Given the same initial values, therefore, the solution of equation (1) at time t_k is equivalent to the result of the convolution of the initial surface with a gaussian of variance $\sqrt{2bt_k}$.

Equation 1 is solved numerically using Feng's implicit scheme described in [5]:

$$\frac{1}{\tau}(u_{p,q}^{k+\frac{1}{2}} - u_{p,q}^{k}) - \frac{b}{h^2}(u_{p+1,q}^{k+\frac{1}{2}} - 2u_{p,q}^{k+\frac{1}{2}} + u_{p-1,q}^{k+\frac{1}{2}}) = 0 \tag{3}$$

$$\frac{1}{\tau}(u_{p,q}^{k+1} - u_{p,q}^{k+\frac{1}{2}}) - \frac{b}{h^2}(u_{p,q+1}^{k+1} - 2u_{p,q}^{k+1} + u_{p,q-1}^{k+1}) = 0 \tag{4}$$

$$u_{p,q}^{0} = f_{p,q} \quad p, q = 0, \ldots, M \quad k = 0, \ldots, \infty \tag{5}$$

where τ is the time unit, h the spatial grid unit, M the image size (see Cai [5] for an extended discussion of this scheme), $p, q \in [0, M)$ and $k \in [0, \infty)$.

In matrix form, we have

$$AU_q^{k+\frac{1}{2}} = U_q^k \quad AU_p^{k+1} = U_p^{k+\frac{1}{2}}$$

where A is a tridiagonal and diagonal-dominant matrix, U_r^t is the array of u values for constant r (r row or column index as appropriate).

There are several practical advantages which make the diffusion approach attractive. Cai [5] has shown that the scheme (3, 4) is unconditionally stable and faster than repeated averaging. Moreover, the scheme can be adopted with a variable time step to obtain a scale space representation. The finite scale space produced is denser than that of gaussian smoothing (Cai [5]). We take advantage of such features in our implementation: the image is smoothed at a desired spatial scale σ by calculating the corresponding $t = t(\sigma)$ and solving system (3, 4) only once. Finally, the diffusion equation form allows an elegant and coherent boundary treatment on the surface being smoothed, as discussed in Section 3.

4. Computing the curvature images

The task of this module is to estimate the sign of the gaussian and mean curvatures at each non-discontinuity point of the surface. The principal curvatures at each point P of the smoothed surface are commonly estimated from the coefficients of a regular patch fitted locally to a neighbourhhod of P (Ferrie *et al.* [3], Monga *et al.* [10]). We have adopted an efficient, mixed approach. The idea is to estimate the curvatures from a local spline approximation of the surface, using *cubic B-splines*

$$\Omega_3(x) = \begin{cases} 0 & if\,|x| \geq 2 \\ \frac{1}{2}|x|^3 - x^2 + \frac{2}{3} & if\,|x| \leq 1 \\ -\frac{1}{6}|x|^3 + x^2 - 2\,|x| + \frac{4}{3} & if\,|x| \in (1,2) \end{cases} \tag{6}$$

Instead of fitting coefficients explicitly - a computationally expensive operation to be performed at each non-discontinuity point of the surface - we compute the discrete convolution

$$u'_{p,q} = \sum_{i=-1}^{m} \sum_{j=-1}^{m} u_{i,j}\Omega_3(p-i)\Omega_3(q-j) \tag{7}$$

which, when iterated, will converge to a C^2 continuous surface which preserves the concavity-convexity of the original surface (Ahlberg *et al.* [11], Reinsch [12]). Moreover, for each iteration, the computation

K	H	shape class
0	0	plane
0	+	negative cylindrical
0	−	positive cylindrical
+	+	negative elliptic
+	−	positive elliptic
−	any	hyperbolic

Table 1: Surface patches classification scheme.

can be done in parallel within 3x3 windows. The gaussian and mean curvature K and H are then given by the relations (Do Carmo [9])

$$K = \frac{h_{xx}h_{yy} - h_{xy}^2}{(1 + h_x^2 + h_y^2)^2} \tag{8}$$

$$2H = \frac{(1 + h_x^2)h_{yy} - 2h_x h_y h_{xy} + (1 + h_y^2)h_{xx}}{(1 + h_x^2 + h_y^2)^{3/2}} \tag{9}$$

where the subscripts indicate partial differentiation and we have assumed as usual a surface parametrization $\mathbf{s}(x, y) = (x, y, h(x, y))$.

5. Patch formation

The H and K sign images are combined together to give the final surface segmentation into patches belonging to several shape classes. The shape classes are illustrated in Table .

The case $K > 0$, $H = 0$ makes no geometric sense and has been omitted. Notice also that we do group the various possible saddle subclasses identified by H when $K < 0$ in one class. The reason is that they do not seem to be perceptually significant for the human vision system. The above classification allows to represent a large number of complex objects which are still a challenge for state-of-the-art recognition systems. Zero-curvature points are detected by thresholding the H, K maps. We adopt Cai's consistent curvature thresholding

$$\varepsilon_K \geq {\varepsilon_H}^2 + 2\,|H|\,\varepsilon_H \tag{10}$$

which is *consistent* in the sense that a small perturbation ξ in the principal curvatures will still lead to a correct classification of H and K (Cai [5]).

6. Quality enhancement on H, K sign images

Figure 2: Left: Typical distortion of quantized planar profile after gaussian smoothing. The original profile is bent at the border: towards the background on the right, away from the background on the left. Right: Examples of performance of adaptive boundary condition for the diffusion equation.

Small, insignificant spots in the H and K sign images are eliminated by a simple erosion-expansion technique. This is based on the idea that, in discrete mathematics, an invertible transformation T does not always satisfy $f^{-1}f = I$ (identity matrix), as it is the case in continuous mathematics. In our case, eroding the regions in the H and K sign images is aimed to make small spots disappear completely, so that nothing of them is left to be grown by the subsequent expansion.

The erosion operation is a *wall-following algorithm* which erodes region contours to a depth specified by the user (generally 1 or 2 pixels). The expansion poses the problem of which regions should be grown first. For each region contour pixel in the H and K images, a local neighbourhood is inspected to decide which label should grow first. The criterion adopted is *local maximum energy*: for each label l a local energy function $E_l = \sum_{(a,b)\in \mathbf{N}(i,j)} \delta_l(a,b)$ is computed in a neighbourhood I of pixel (i,j), where

$$\delta_l(a,b) = \begin{cases} 1 & if \;\; H(a,b) = l \\ 0 & otherwise \end{cases}$$

The label to be grown is then the one associated with the maximum E_l in I.

7. Preserving shape at boundaries

Contours of orientation discontinuities are preserved against curvature distorions introduced by gaussian smoothing by means of the discontinuity maps. Each discontinuity point is considered a boundary point for the diffusion process and the adaptive boundary condition detailed in the next section is enforced. This results in a good preservation of both discontinuities and curvature sign near patch boundaries. Cai [5] suggested a *small leakage* boundary condition for attenuating the boundary deformation effect (corruption of H, K signs) in the diffusion approach.

From this point of view Gaussian smoothing is equivalent to diffusion smoothing with “perfectly insulated” boundaries; the typical distortion arising is avoided by allowing the surface to leak into the background. Figure 2 (left) shows the smoothed data both changing inmagnitude and with an introduced curvature at the extrema. It is

Figure 3: $\frac{\partial u}{\partial x}$ changes for differently sloping boundaries (see text).

this latter effect that we eish to avoid. This is achieved by imposing at each border point the boundary condition

$$b_{in} \frac{\partial u}{\partial \mathbf{n}} = b_{out} \frac{\partial u}{\partial \mathbf{n}} \tag{11}$$

where b_{in} is the diffusion coefficient on the surface, b_{out} is the diffusion coefficient of a narrow stripe surrounding the background and $\mathbf{n}$ is the normal at the boundary in the $z = 0$ plane. Cai suggests that b_{out} should be "much smaller" than b_{in}, e.g. 10%. Condition (11) can be simplified by splitting it into two equations along the x and y axis respectively. Discretizing at the boundary pixel (p, q) at time $(k+1/2)$ for x and $(k+1)$ for y, we obtain

$$\begin{bmatrix} b_{in} & -(b_{in} - b_{out}) & b_{out} \end{bmatrix} \begin{bmatrix} u_{p-1,q}^{k+1/2} \\ u_{p,q}^{k+1/2} \\ u_{p+1,q}^{k+1/2} \end{bmatrix} = 0$$

$$\begin{bmatrix} b_{in} & -(b_{in} - b_{out}) & b_{out} \end{bmatrix} \begin{bmatrix} u_{p,q-1}^{k+1} \\ u_{p,q}^{k+1} \\ u_{p,q+1}^{k+1} \end{bmatrix} = 0$$

A fixed leakage coefficient can work, in the best case, with *one* slope type (sloping towards or away from the background) and *one* pair of directional derivative values across the boundary $(\left|\frac{\partial u}{\partial e}\right|_{in}, \left|\frac{\partial u}{\partial e}\right|_{out})$, e being x or y according to which equation of the numerical scheme is being solved. But in most cases both $sgn(\frac{\partial u}{\partial e})$ and $\left|\frac{\partial u}{\partial e}\right|$ change along the boundary, thus suggesting the introduction of an *adaptive* leakage coefficient. The sign will be different for borders sloping towards as opposed to away from the background. This is illustrated in figure 3 considering a right-hand border for the x equation of the numerical scheme. In *(a)* the signs of the discrete gradients across the boundary $(u_{i+1} - u_i)_{out}$ and $(u_i - u_{i-1})_{in}$ (assuming $\Delta x = 1$) are the same; in *(b)* they are different, but a fixed positive b_{out} would still pull the surface towards the background. Moreover, depending on the reciprocal position of surface and background, the absolute value of the internal and external gradients will change as well.

The main attractive point of boundary condition (11) is that its implementation is efficient: it requires only a small revision of matrix A for boundary pixels. A complex boundary treatment would spoil the efficiency of the numerical scheme. We have therefore tried to

% specialpsfile=puma.ps

Figure 4: Synthetic range image of a PUMA robot.

Figure 5: Discontinuity map for the PUMA robot image.

introduce a similar but more general condition to generalize (11) by introducing an *adaptive leakage coefficient* b^a_{out} given by

$$b^a_{out} = b_{in} \left(\frac{\left.\frac{\partial u}{\partial \mathbf{n}}\right|_{in}}{\left.\frac{\partial u}{\partial \mathbf{n}}\right|_{out}} \right) \qquad (12)$$

The adaptive coefficient b^a_{out} is evaluated at each border pixel and used in (11) instead of the fixed leakage b_{out}. The effect is to adapt the amount of leakage at the particular border pixel, taking into account the difference between internal and external slope before enforcing the boundary condition (11). In our experiments, detailed in [6] this condition performed better with both planar and curved surfaces (an example is given in Figure 2 right).

8. Experimental results

We present in this Section a few illustrative experiments run with a prototype implemented in C++ on a SPARC workstation under XWindows and Unix. The images used in the experiments were either synthetic or acquired by a laser striper developed by our group (Naidu and Fisher [13]). All images are 128x128 pixels, 8 bit per pixel. Note that the quantization effect has serious effects on the accuracy of curvature estimation.

8.1. Examples of segmentations

The first example (Figure 4) is a synthetic 128x128 range image of a PUMA robot model (about 170x140x100mm^3). All the surfaces are developable; the classification was therefore based on the H sign image only. Figure 5 shows the discontinuity map ($\tau_d = 7$, $\tau_o = 0.75$). The final classification is shown in Figure 6. The zero threshold for H was $\varepsilon_H = 0.009$, determined experimentally. The image was smoothed with diffusion interval $t_0 = 0$, $t_{max} = 4$ equivalent to $\sigma_{max} = 2$ for gaussian smoothing, given the diffusion coefficient of equation 1 $b = 0.5$ (notice that, with $b = 0.5$, equation 2 reduces to $\sigma = \sqrt{t}$). A few patches were lost as an effect of the combined shrinking due to the discontinuity detection and the erosion-expansion. Small patches are more vulnerable to shape distortion and therefore likely to disappear after the erosion-expansion stage.

The second example is a real image of a Renault part (Figure 7), whose size is about 190x100x80mm^3. This is a complex sculptured

Figure 6: Final segmentation for the PUMA robot image.

Figure 7: Range image of a Renault part.

object, including both developable and non-developable surfaces. The H image was experimentally set at $\varepsilon_H = 0.024$. The consistent thresholding equation (10) yields $\varepsilon_K = 0.00422$. The image was smoothed between $t_0 = 0$ and $t_{max} = 9$, equivalent to $\sigma_{max} = 3$ with $b = 0.5$.

Figure 8 gives the discontinuity map ($\tau_d = 10$, $\tau_o = 0.75$). The result of the segmentation is shown in Figure 9. Most of the significant patches appear in the final segmentation. This result is an improvement on many results achieved by analogous techniques applied to this object.

8.2. Testing the system

We ran several tests to assess quantitatively the performance of our system. A complete report will appear in a forthcoming paper. We summarize here two aspects, classification and accuracy. Similar tests were reported by Flynn and Jain [15] but were limited to accuracy.

Classification tests were run by measuring the percentage of correctly classified points in a synthetic patch against variations of the patch geometry (shape class and curvature), the number of smoothing cycles and the zero threshold ε_H. An example is given in Figure 10, which shows the upper bound of ε_H values guaranteeing a given percentage of correctly classified points on a cylindrical patch. From this we conclude that, if R_M is the maximum radius of cylindrical features in an image, then the threshold ε_H should be less than about $\frac{1}{3R_M}$. Similar graphs were generated for elliptic, planar and hyperbolic patches and contain useful suggestions about how to choose the threshold ε_H. At present, good sign classification accuracy is achieved with patches of about 14 pixel side (referring to 128x128 images smoothed with up to three smoothing cycles). The same causes make it difficult to distinguish accurately between planar patches and curved surfaces with low curvatures. Worst-case interference radii are about $R = 70$ for cylinders and $R = 140$ for spheres. The number of smoothing cycles nc is linked to both the diffusion time interval and the standard deviation of the equivalent smoothing gaussian by the relations

$$\Delta t = nc^2 \quad \sigma = \sqrt{2b\Delta t}$$

where Δt is the diffusion interval and b is the diffusion coefficient. In

Figure 8: Discontinuity map for the Renault image.

Figure 9: Final segmentation of the Renault image.

Figure 10: Upper bound of ε_H values guaranteeing given classification percentage for quantized synthetic cylinders smoothed with $nc = 2$ and $nc = 3$ against increasing radii.

our implementation b is 0.5, therefore $\sigma = nc$. Since we want to use the segmentation results for recognition purposes we are interested in values of nc preserving as many visible features as possible (i.e. not considering large, low-frequency blobs). We determined experimentally that values of nc less than 4 match this requirement in most cases. Therefore our tests were run with $nc = 1, 2, 3$ to supply us with information useful for our purposes.

Accuracy test were run by estimating H on cylindrical and spherical patches of different sizes, with and without quantization noise, and after increasing amounts of smoothing. The estimates get worse as the diffusion interval increases, as expected. For a sphere of radius $r = 19$, for instance, the deviation from the expected value is 0.0026, 0.0037 and 0.0054 for 1, 2 and 3 cycles of smoothing respectively; the correspondent percentage error in the estimated radius is 5%, 7% and 11% respectively. The estimates become more and more unreliable as the radii decrease. For radii less than 10 pixels, the estimates degrade rapidly as the number of smoothing cycles increase; still looking at the sphere case, the error in the estimated radius for $r = 10$ is already 31% for one smoothing cycle and becomes 66% and 167% for two and three smoothing cycles. The reason is that as the size of the smoothing gaussian approaches that of the surface to be smoothed, the curvature of the latter is more and more distorted.

In general, our tests indicate good performance for H sign estimation, but limited accuracy for absolute values. This is in accord with the results of Flynn and Jain [15], which suggest that accuracies better than 10% are difficult to obtained from quantized data even in large smoothed patches of nonzero curvature. Therefore the representation computed by the system is a good *qualitative* description of the surfaces in view, with the quantitative surface structure estimated with limited accuracy.

9. Conclusion

We have presented a technique for computing a versatile surface-based representation from range images, and its implementation. In particular, we have discussed two problems in range segmentation: preserving curvature sign at boundaries and discontinuity contours over smoothing. We believe a surface-based representation is useful for a variety of robotic tasks involving vision, such as navigation, inspection, patch recognition or verification by matching segmented patches to model patches with the same qualitative description, position estimation of parts from model-to-data patch correspondence, and robot grasping by gripper positioning on large stable patches.

Acknowledgement

This work was performed during the first author's permanence at the Department of Artificial Intelligence, University of Edinburgh. Funding was provided by the European Institute of Technology under Grant I-068.

References

[1] J. Crowley: *Using a Composite Surface Model for Perceptual Tasks*, Proc. IEEE Conference on Robotics and Automation, Raleigh (NC), 1987.

[2] C. Thorpe: *Outdoor Visual Navigation for Autonomous Robots*, Proc. 2^{nd} International Symposium on Intelligent Autonomous Systems, Amsterdam, 1989.

[3] F. Ferrie, J. Lagarde and P. Whaite: *Recovery of Volumetric Object Descriptions from Laser Rangefinder Images.* Proc. First European Conference on Computer Vision, 1990, pp. 387–396.

[4] D. Lowe: *Organization of Smooth Image Curves at Multiple Scales*, in *International Journal of Computer Vision*(**3**), 1989, pp. 119 – 130.

[5] L. D. Cai: *A "Small Leakage" Model for Diffusion Smoothing of Range Data.* Proc. 11th IJCAI, Detroit, 1989, pp.1585–1590.

[6] E. Trucco: *On Shape-Preserving Boundary Conditions for Diffusion Smoothing*, Proc IEEE Conference on Robotics and Automation, Nice, 1992.

[7] A. Witkin: *Scale-Space Filtering.* Proc. 8th IJCAI, 1983, pp.1019-1022.

[8] T. Lindeberg: *Scale-Space for Discrete Signals.* IEEE PAMI **12**, 1990, pp. 234–254.

[9] M. P. do Carmo: *Differential Geometry of Curves and Surfaces*, Prentice-Hall, 1976.

[10] O. Monga, N. Ayache and P. Sanders: *Modeling Uncertainty for Estimating Local Surface Geometry.* Proc. 7th Scandinavian Conference on Image Analysis, Aalborg, 1991, pp. 403–410.

[11] J. Ahlberg, E. Nilson and J. Walsh: *Theory of Splines and its Applications.* Academic Press, 1967.

[12] C. Reinsch: *Smoothing by Spline Functions.* Num.Math. **10**, 1967, pp.177–183.

[13] D.K. Naidu and R.B. Fisher: *A Comparative Analysis of Algorithms for Determining the Peak Position of a Stripe to Subpixel Accuracy.* Proc. British Machine Vision Conference, Glasgow, 1991.

[14] R.B. Fisher: *From Surfaces to Objects.* Wiley, 1989.

[15] P.J. Flynn and A.K. Jain: *On Reliable Curvature Estimation.* Proc. IEEE Conference on Computer Vision and Pattern Recognition, 1989, pp. 110 – 116.

[16] J. Maver and R. Bajcsy: *How to Decide from the First View Where to Look Next*, Proc. Image Understanding Workshop, Pittsburgh (PA), 1990.

[17] S. Sakane, M. Ishii and M. Kakikura: *Occlusion Avoidance of Visual Sensors Based on a Hand-Eye Motion Simulator System: HAVEN*, Advanced Robotics(**2**), 1987, pp. 149 – 165.

[18] P. Besl: *Surfaces in Range Image Understanding*, Springer-Verlag, 1987.