Partitioning Range Images Using Curvature and Scale

André Lejeune and Frank P. Ferrie

McGill Research Center for Intelligent Machines, McGill University,
3480 Université, Montréal, Québec, Canada, H3A 2A7
e-mail: andre@lightning.mcrcim.mcgill.edu, ferrie@lightning.mcrcim.mcgill.edu

Abstract

We present a method for partitioning a set of surface estimates obtained with a laser range finding system into subsets corresponding to parts of an object. Our strategy uses two complementary representations for surfaces: one that describes local structures in terms of differential properties (e.g. edges, lines, contours) and the other that represents the surface as a collection of smooth patches at different scales. By enforcing a consistent interpretation between these two representations, it is possible to derive a partitioning algorithm that is both efficient and robust.

1 Motivation

A fundamental problem in the interpretation of visual forms is that of *parts decomposition*, the partitioning of the surfaces of an object into regions corresponding to each part [4]. Traditionally two approaches have been considered. So-called edge-based approaches work by first identifying features corresponding to part boundaries and then interpolating to form smooth contours. Region-based approaches, by contrast, first identify clusters or regions that are continuous in specified features and then seek to merge smaller groupings until some continuity criteria are violated. Each strategy taken separately is often insufficient for lack of robustness, non-uniqueness or complexity.

We argue that each process is, in fact, complementary to each other as the detected edges provide cues to validate region boundaries and these same region boundaries specify how the edges must be interpolated [3]. An object is defined to be a solid mass composed of convex parts without holes of any kind. Consequently in our model of a part, an edge can correspond to a discontinuity in depth, (C^0), a concave discontinuity in orientation, (C^1), or a negative local minimum in normal curvature [2].

A region is defined as a subset of the surface of an object that is void of discontinuities and extrema of curvature. As such, the set of regions that comprise the surface implicitly define the loci of potential boundary points. Regions are represented as parametric surface patches computed at different resolutions to account for the appearance of features at different scales [1]. We refer to the resulting representation as a curvature scale space.

The relation of this scale space to the interpolation problem can be explained as follows. A single region implicitly defines a closed contour. When two regions are merged, the resulting region defines yet another contour. The set of all possible surface contours can thus be determined by enumerating all possible merges that form spatially continuous patches. By construction, the resulting space must contain the true surface boundaries. The computational task is to search this space for the correct interpretation.

Given that the edge and region descriptions must both correspond to the same object, the search problem can be precisely formulated as that of finding the smallest set of covers, i.e. unions of regions, that is most consistent with the edge-based description.

2 Curvature scale-space

A curvature scale space was developed to reduce the dimension of the search space by organizing data in a precisely structured fashion. It has the property of smoothing regions of convex curvature to reduce noise while preserving the positions of features corresponding to part boundaries.

Starting from the definition of a scale-space that can be found in Kimia [6], it is possible to derive such a smoothing process. Given a 2-D array of data $x[m][n]$, the diffused signal after k iterations $x_k[m][n]$ is given by

$$\begin{aligned}
a &= \tfrac{1}{2}(x_{k-1}[m][n-1] + x_{k-1}[m][n+1])\\
b &= \tfrac{1}{2}(x_{k-1}[m-1][n] + x_{k-1}[m+1][n])\\
c &= \tfrac{1}{2}(x_{k-1}[m-1][n-1] + x_{k-1}[m+1][n+1])\\
d &= \tfrac{1}{2}(x_{k-1}[m+1][n-1] + x_{k-1}[m-1][n+1])
\end{aligned}$$

$$e = \min(a, b, c, d)$$

$$x_k[m][n] = \begin{cases} e & \text{if } x_{k-1}[m][n] < e \\ & \text{(convex surface)} \\ x_{k-1}[m][n] & \text{if } x_{k-1}[m][n] \geq e \\ & \text{(other types)} \end{cases} .$$

The effect of this diffusion process is shown in the image of a doll presented in Figure 1. The images present the contours of convex regions found in the range image.

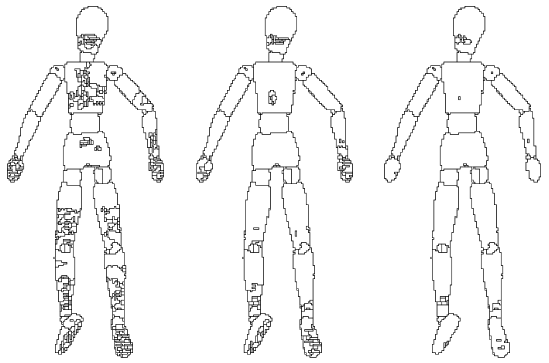

a) Region segmentation before any diffusion. The parts of the doll are over-segmented because of noise on the surface. b) After 5 iterations of diffusion, some noise still causes over-segmentation on some parts. c) After 30 iterations, each part consists mostly of a single region.

Figure 1: Curvature scale space of a doll.

3 Resolution algorithm

Once the scale space representation is built, optimizing covering of the surface so that the resulting boundaries are consistent with the detected cues becomes possible. It was found that this optimization problem maps naturally into a relaxation labelling framework [5].

A label is assigned to each surface region in the scale space representation and constraints are defined by the edges found between adjacent regions. The presence of edges at a common boundary of two regions defines a negative constraint that discourages the regions to merge. The absence of edges at a boundary defines a positive constraint that encourages the merging of two adjacent regions. By allowing the network to converge, extracted edges and regions are resolved into an appropriate parts decomposition, as shown in Figure 2.

4 Conclusion

This short paper presented an overview of a method for partitioning surface information into part-oriented descriptions of object geometry. The key idea borrowed from early work in computer vision, that is, the concurrent application of edge-based and region-based methods of image segmentation.

We presented a re-casting of this paradigm in terms of surface differential geometry and the development of a solution using constraint propagation networks (a.k.a. relaxation labelling). The results obtained demonstrate that the resulting algorithms work very well and can provide a sound basis for the generation of object models as can be found in our report [7].

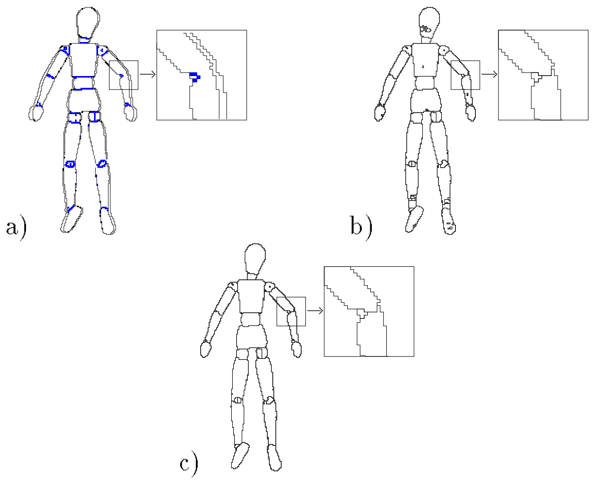

a) Detected edges in the range image of a doll after 30 iterations of diffusion. Blowup of the elbow shows enough cues to suggest segmentation of the forearm, but interpolation is necessary to obtain two closed regions. b) Region segmentation taken from the scale space. It provides the necessary contour to interpolate the cues in the elbow. Also note that there are contours in each leg that are not supported by any cues. c) After merging, the expected interpolation occurred in the elbow and regions have been properly merged in each leg.

Figure 2: Merging results of the doll

References

[1] P. J. Besl and R. C. Jain. Segmentation through variable-order surface fitting. *IEEE Tr. on PAMI*, PAMI-10(2):167–192, Mar. 1988.

[2] F. Ferrie, J. Lagarde, and P. Whaite. Recovery of volumetric object descriptions from laser rangefinder images. In *Proc. of the 1ST Eur. Conf. on CV*, Antibbes, France, April 1990.

[3] F. Ferrie, A. Lejeune, and D. Baird. Curvature, scale, and segmentation. In *SPIE App. of AI X: Mach. Vision & Rob.*, volume 1708, pages 240–250, Orlando, Florida, April 20-24 1992.

[4] D. Hoffman and W. Richards. Parts of recognition. *Cognition*, 18:65–96, 1984.

[5] R. Hummel and S. Zucker. On the foundation of relaxation labeling processes. *IEEE Tr. on PAMI*, PAMI-5(3):267–287, 1983.

[6] B. B. Kimia, A. R. Tannenbaum, and S. W. Zucker. Shapes, shocks, and deformations, I: The components of shape and the reaction-diffusion space. Technical Report LEMS 105, LEMS, Brown University, May 1992.

[7] A. Lejeune and F. P. Ferrie. Finding the parts of objects in range images. Technical Report 93-8, McRCIM, McGill University, April 1993.