University of Pennsylvania
ScholarlyCommons

Technical Reports (CIS)

Department of Computer & Information Science

January 1997

Fish Scales: Representing Fuzzy Manifolds

Radim Šára
University of Pennsylvania

Ruzena Bajcsy
University of Pennsylvania

Follow this and additional works at: https://repository.upenn.edu/cis_reports

Recommended Citation

Radim Šára and Ruzena Bajcsy, "Fish Scales: Representing Fuzzy Manifolds", . January 1997.

University of Pennsylvania Department of Computer and Information Science Technical Report No. MS-CIS-97-19.

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/cis_reports/80
For more information, please contact repository@pobox.upenn.edu.

Fish Scales: Representing Fuzzy Manifolds

Abstract

We address the problem of automatically reconstructing m -manifolds of unknown topology from unorganized points in metric p -spaces obtained from a noisy measurement process. The point set is first approximated by a collection of oriented primitive fuzzy sets over a range of resolutions. Hierarchical multiresolution representation is then computed based on the relation of relative containment defined on the collection. Finally, manifold structure is recovered by establishing connectivity between these primitives based on proximity, compatibility of position and orientation, and local topological constraints. The method has been successfully applied to the problem of surface reconstruction from polynocular-stereo data with many outliers.

Comments

University of Pennsylvania Department of Computer and Information Science Technical Report No. MS-CIS-97-19.

Czech Technical University
Center for Machine Perception

Fish-Scales: Representing Fuzzy Manifolds

Radim Šára Ruzena Bajcsy
University of Pennsylvania
GRASP Laboratory
3401 Walnut St., Philadelphia, PA, U.S.A.
radim@grip.cis.upenn.edu

Published in: *Proc. Int. Conf. Computer Vision, ICCV'98*. pp. 811-817. Bombay, India, January 1998

This publication can be obtained via anonymous ftp from
`ftp://cmp.felk.cvut.cz/pub/cvl/articles/sara/iccv98fin.ps.gz`

Czech Technical University, Faculty of Electrical Engineering
Department of Control Engineering, Center for Machine Perception,
121 35 Prague 2, Karlovo náměstí 13, Czech Republic
FAX +420 2 24357385, phone +420 2 24357458, <http://cmp.felk.cvut.cz>

Fish-Scales: Representing Fuzzy Manifolds

Radim Šára Ruzena Bajcsy
University of Pennsylvania
GRASP Laboratory
3401 Walnut St., Philadelphia, PA, U.S.A.
radim@grip.cis.upenn.edu

Abstract

We address the problem of automatically reconstructing m -manifolds of unknown topology from unorganized points in metric p -spaces obtained from a noisy measurement process. The point set is first approximated by a collection of oriented primitive fuzzy sets over a range of resolutions. Hierarchical multiresolution representation is then computed based on the relation of relative containment defined on the collection. Finally, manifold structure is recovered by establishing connectivity between these primitives based on proximity, compatibility of position and orientation, and local topological constraints. The method has been successfully applied to the problem of surface reconstruction from polynocular-stereo data with many outliers.

1 Introduction

In this paper we are not concerned with arbitrary unorganized point sets, but only with those obtained by a real measurement process.

Stereo reconstructions are rarely interpreted as surfaces because computer stereo vision typically produces a large fraction of outliers and ill-behaved noise. When many cameras are used, the problem becomes significant as in the dataset shown (Fig. 1). If it were not for the outliers, the dataset contains many more measurements than are necessary for an interpretation.

In this paper we propose a method of interpreting noisy measurements as manifolds. The problem is as follows: there is an m -manifold in the world one observes, which may be a curve ($m = 1$), a visual surface ($m = 2$), etc. One takes a finite number of local measurements whose dimension is $p > m$ and uses them to create a representation with manifold structure. Of course, one does not know if a manifold is observed; the only information is present in the measurements and in the a priori assumption that the perceived object is locally homeomorphic to an open m -ball or a half-ball and can be covered by a finite number of subsets.

To our knowledge, the majority of current bottom-up (non-parametric) methods for recovering curves or surfaces from visual data are based on the above assumption. *What is rarely acknowledged in full consequence, though, is that measurements are always corrupted by noise.*¹ The uncertainty not only affects the local properties of the reconstructed entity but also the global structure one wants to

Fig. 1: A top view of a cloud of isolated points in $\mathbb{R}^3$ as obtained from a four-camera stereo. The human face is looking towards the bottom edge of the page and the cameras were located in front of it. The matching artifacts fill the space around but are not compatible with it.

recover: if we triangulate a set of points and a slightly perturbed set of the same points, the triangulations may significantly differ, even non-locally. The goal then is to suppress the influence of uncertainty on the recovered structure.

Our assumption is that the dimension p of the measurements must be greater than the dimension of the manifold m . Since the measurements are corrupted by noise, our task must include simultaneous filtration and ‘projection’ on an m -dimensional structure. We propose that it is sufficient to locally filter and project the measurements by grouping them into local geometric primitives that sit in p -space and are expected to ‘cover’ the manifold the same way scales cover most fish. The smallest approximation error occurs when the subspace of the primitives eigenspace corresponding to m largest eigenvalues is parallel to the tangent (hyper-)plane of the manifold. The covering will be more compact if the primitives are compatible at their overlap. In this paper we show how to evaluate this overlap and how to build the global structure from local primitives.

After this introductory section, we describe in Section 2 a new m -manifold representation which will later be narrowed to surfaces in 3-D. Some results on the dataset shown in Fig. 1 are presented in Section 3. Finally, in Section 4 we discuss our approach and propose future research topics.

¹A problem specific to 3-D vision is that retinal images are mere projections of the world and the fully-dimensional structure must be reconstructed, which also introduces various *artifacts*.

1.1 Related Work

There are roughly two distinct approaches to the problem of surface reconstruction: (1) *top-down*: the connectivity information (the topological type of the surface) is a priori known, and (2) *bottom-up*: the connectivity is recovered along with geometry. The former problem has been extensively treated in the literature [12, 28, 30, 14]. Recently, a modification for closed compact surfaces has been reported in [6, 23, 20]. For a review of these methods, see [2]. See also [22, 26, 19] for parametric models recovery.

Except for a restricted class of tasks, the topological type of the observed surface is not known in advance. In the recent work of Caselles et al., the topological structure is allowed to evolve during the recovery from an initial simple model (homeomorphic to a sphere) [3]. The models are restricted to closed surfaces. Many of the top-down methods require collecting all the raw measurements prior to recovering the model and do not allow for (1) *efficient* incorporation of new *local* measurements to improve either the model accuracy or its descriptiveness (there were some attempts to do this in triangulations [32, 5]), or for (2) removing redundant primitives to increase the ratio of descriptiveness over the description length of the final representation (some attempts are reported in [1, 27]).

The bottom-up approach—known as reconstruction from unorganized points—is much more difficult and has not been treated as often. Tonnesen and Szeliski proposed modeling objects (surfaces) as collections of (isolated) *oriented particles* [29], which was a generalization of the particles of Reeves [24]. They attempted neither to recover the model from real range data nor to establish the connectivity among the particles. This was later done by Fua, who recovered the models from unorganized points without the connectivity [9] and also including the connectivity [11, 10]. He used 2-D Delaunay triangulation over a projection of the centers of the local primitives. This is only possible when the surface can be mapped on a plane by a one-to-one projection. Oriented particles are generally well suited for modeling elastic surfaces or surfaces under deformation [21]. Our work has been inspired by the work of Fua, but it differs in establishing the proper invariant connectivity among the local geometric primitives which, in turn, can help reject some inconsistent measurements.

Hoppe et al. recover local orientation *at each data-point* and then recover the signed distance function zero set by the means of a contouring algorithm [16] (see also [15, 4]). The distance is measured from the estimated *oriented* tangent plane and its zero set is a manifold by its definition. The (initially unknown) absolute orientation is propagated from points of known orientation in Euclidean Minimum Spanning Tree. The algorithm complexity is $O(n \log n)$, where n is the number of primary measurements. Recently, this work has been greatly elaborated upon [7] but the authors do not deal with the problem of outliers or uncertainty in the measurements. This approach fails when the surface is not orientable: consider just a Möbius strip which is a perfectly visually observable entity and there is no reason why it should not be recoverable from visual data.

Edelsbrunner et al. introduced the formal notion of the family of α -shapes of a finite point set in $\mathbb{R}^3$ [8]. It is a generalization of the convex hull of a point set and represents the “shape” of the set.

1.2 The Bottom-Up Approach

In the presence of outliers in raw data or when the data describe just portions of objects, one wants to recover the manifold in a bottom-up process since it is much easier to deal with unreliable and/or incomplete information locally. Also for fusing information from many partial reconstructions, some intermediate representation is preferred, one that is more compact and more descriptive than isolated points. If we accept the bottom-up paradigm, we face two relatively independent subproblems, both of which are discussed in this paper:

1. How to group isolated points to meaningful, compact, and descriptive semi-local geometric primitives. This is a local *clustering problem*.
2. How to establish connectivity between those primitives and create a (pseudo-)manifold structure. This is a global *data interpretation problem* and we want to separate it from the first step as much as possible since it depends on the purpose of data interpretation.

In many clustering methods, one has to select a scale (resolution) at which the clustering should be done or employ global constraints like minimum description length [31]. Since selecting the scale requires a task-dependent knowledge and we want to postpone the interpretation to a later stage, we shall address the clustering problem by the means of simple and fast *multiresolution* clustering.

The reconstruction task is then accomplished in three steps:

1. *Preprocessing*. Grouping of isolated points according to an a priori *local* geometric model to geometric primitives that are instances of this model.
2. *Verification and refinement*. By recurring back to (some of) the images the orientation and position of the primitives may be refined.
3. *Interpretation at a given scale*.
 - (a) Establishing the containment order of the primitives and removing those that are redundant at the given resolution.
 - (b) Selecting the manifold order.
 - (c) Establishing connectivity between the primitives based on proximity, *mutual compatibility*, and global topological constraints.

In the preprocessing step local problems (like removing outliers) are resolved and redundancy in the dataset is suppressed at given scale. This step is repeated at multiple resolutions. Multiple such collections are fused in a simple union operation.

In the interpretation step, primitives redundant at given scale are first removed and the collection of primitives is then interpreted by imposing additional non-local constraints as manifold structure (of selected order) with boundary.

The next section introduces a surface representation that allows us to deal with the reconstruction task as we outlined it here.

2 Discrete Fuzzy Manifolds

The fuzzy-set framework provides us with tools suitable for all partial problems and allows us to generalize. We will start with simple concepts that will be used in a more rigorous task formulation.

2.1 Fuzzy Sets and Fuzzy Relations

We build upon the standard notion of fuzzy set. For an introduction to this topic we refer the reader to [33]. Here we use the following notation.

A fuzzy subset $\mathcal{S}$ of $\mathbb{R}^p$ is defined by its membership function $\mu(\mathbf{x}|\mathcal{S}) : \mathbb{R}^p \rightarrow [0, 1]$, where the value of $\mu(\mathbf{x}|\mathcal{S})$ represents the degree of membership of a point $\mathbf{x} \in \mathbb{R}^p$ in $\mathcal{S}$. The interval $[0, 1]$ is a special case of a membership set. For shortness, we shall often write $\mu(\mathcal{S})$ instead of $\mu(\mathbf{x}|\mathcal{S})$. In addition to standard set-theoretic operations of complement, intersection, and union and the algebraic operations of product and sum, we require $\mu(\mathbf{x}|\mathcal{S})$ to be integrable and we define the norm

$$\|\mu(\mathcal{S})\| = \int_{\mathbb{R}^p} \mu(\mathbf{x}|\mathcal{S}) d\mathbf{x} \leq \infty. \quad (1)$$

We will also fuzzify n -ary relations as follows: Let $\mathbf{P}$ be a product set of n sets, $\mathbf{P} = \mathbf{E}_1 \times \mathbf{E}_2 \times \dots \times \mathbf{E}_n$ and M its membership set, a fuzzy n -ary relation is a fuzzy subset of $\mathbf{P}$ taking its values in M .

2.2 Primitive Fuzzy Sets

We are not interested in isolated *primitive sets* $\mathcal{S}$, rather in a finite *collection* $C = \{\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_n\}$. We shall assume that the primitive sets are finite and non-singular

$$0 < \|\mu(\mathcal{S})\| < \infty \quad (2)$$

and that they are local models to some more global structure. We will be looking for relations between the primitive sets that generalize the local properties of the members of C to the properties of the collection C in the large.

Since our primitive sets are intended as local manifold models, some geometric properties are required, namely, the primitives must be oriented to guarantee the existence of subspace which they can uniquely be projected to and which is supposed to be parallel to the manifold tangent (hyper-)plane. The knowledge of the orientation sign is not required, though.

2.3 Primitive Fuzzy Relations

Based on the standard fuzzy-set operations we define two fuzzy relations:

1. relative intersection

$$\omega(\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_N) = \frac{\|\mu(\bigcap_{i=1}^N \mathcal{S}_i)\|}{\|\mu(\bigcup_{i=1}^N \mathcal{S}_i)\|}, \quad (3)$$

2. relative inclusion (containment)

$$\iota(\mathcal{S}_1, \mathcal{S}_2) = \frac{\|\mu(\mathcal{S}_1 \cap \mathcal{S}_2)\|}{\|\mu(\mathcal{S}_1)\|}. \quad (4)$$

Fig. 2: Hyperellipsoids that are perpendicular (a) or parallel but non-collinear (b) intersect much less than those that are (nearly or fully) collinear (c). The degree of intersection is measured by ω . Hyperellipsoids that are embedded in larger ones may be redundant at a certain scale (d). The order is defined by ι .

The relative intersection expresses the degree of mutual overlap of $\mathcal{S}_1, \dots, \mathcal{S}_N$. The relative inclusion expresses the degree to which $\mathcal{S}_1 \subset \mathcal{S}_2$. For an intuitive example see Fig. 2. Since $\|\mu(\bigcap_{i=1}^N \mathcal{S}_i)\| \leq \|\mu(\bigcup_{i=1}^N \mathcal{S}_i)\|$, $\|\mu(\mathcal{S}_1 \cap \mathcal{S}_2)\| \leq \|\mu(\mathcal{S}_1)\|$, and $\|\mu(\mathcal{S}_i)\| > 0$, the values of ω and ι are both confined to the interval of $(0, 1]$.

2.4 Fish-Scales

From now on, we shall assume that the membership functions have this particular form

$$\mu(\mathbf{x}|\mathcal{S}) = f\left(\frac{1}{2}(\mathbf{x} - \mathbf{x}_0)^T \mathbf{S}^{-1}(\mathbf{x} - \mathbf{x}_0)\right), \quad (5)$$

where $\mathbf{S}$ is a $p \times p$ symmetric and positive-definite *structure matrix*, the vector $\mathbf{x}_0$ is the *center*, and f is a monotone non-increasing scalar *influence function* such that:

$$0 < f(t) \leq 1, \quad f(0) = 1, \quad \text{and} \quad \int_0^\infty f(t)^k dt < \infty, \quad k = 1, 2. \quad (6)$$

We shall call such sets *fish-scales* and they will be used as local fuzzy geometric primitives. The eigenvectors and eigenvalues of the structure matrix define the geometric properties of the fish-scale: its eigenvector corresponding to the extremal eigenvalue defines the fish-scale orientation and the ratios of eigenvalues define how flat, thin, etc., the fish-scale is. If λ_{max} is the maximum eigenvalue among $\lambda_i, i = 1, \dots, p$ of $\mathbf{S}$, we define the *scale rank* ρ

$$\rho = \sum_{i=1}^p \frac{\lambda_i}{\lambda_{max}}, \quad \rho \in (0, p]. \quad (7)$$

We may then classify 3-D fish-scales ($p = 3$) as follows:

type	rank	eigenvalues
spherical	3	$\lambda_1 = \lambda_2 = \lambda_3 > 0$
flat	2	$\lambda_1 \rightarrow 0, \lambda_2 = \lambda_3 \gg 0$
linear	1	$\lambda_1 = \lambda_2 \rightarrow 0, \lambda_3 \gg 0$

Based on the rank, our reconstruction may be selective when the dataset contains measurements originating from manifolds of different order.

Fig. 3: The collection (a) of 0-simplices (vertices) and 1-simplices (edges) is not a complex (some edges intersect in between their endpoints). If we remove some of the edges we transform it to a complex (b). The complex can be further transformed to a pseudo-manifold (c) by removing some more edges. This process can be guided by the values of ω associated with the edges.

2.5 The Collection of Simplices Generated by ω

At this point we will need some concepts from standard topology. For the lack of space here, we refer the reader to standard books, e.g. [17, 13]. Let us just give our definition of pseudo-manifold: A pseudo-manifold with boundary is the carrier of an m -complex in which every $(m-1)$ -simplex is a face of at most two m -simplices (we will refer to this as the pseudo-manifold constraint). The neighborhood of every point in a pseudo-manifold is homeomorphic to an open ball or half-ball, except for a finite number of singular points that are homeomorphic to k open balls (or half-balls) identified to the point.

The relation ω generates a collection of associated simplices as follows: the primitive sets centers $\{(i)\}$ may be thought of as 0-simplices, the labeled edges $\{(i, j) \mid \omega(i, j) > 0\}$ as 1-simplices, the triangles $\{(i, j, k) \mid \omega(i, j) \wedge \omega(j, k) \wedge \omega(i, k) > 0\}$ as 2-simplices, etc.² The collection C_ω of all these simplices do not necessarily form a complex, since some of their intersections need not be simplices from the collection. We may remove some simplices from C_ω to obtain a complex, see Fig. 3. Moreover, we may remove simplices violating the pseudo-manifold constraint to get a pseudo-manifold $P_\omega \subseteq C_\omega$. If $\omega > 0$ there is at least one pseudo-manifold that can be formed from C_ω : it is the convex hull of C_ω . Let the weight of P_ω be the sum of the weights of all its 1-simplices that are incident on at least one 2-simplex. If (2) holds and no two weights $\omega(i, j)$ and $\omega(k, l)$, $(i, j) \neq (k, l)$ are equal, there is a unique pseudo-manifold $P_\omega^* \subseteq C_\omega$ that maximizes the weight over all possible pseudo-manifolds P_ω that can be formed from C_ω .

Unfortunately, such an optimization task is an NP-complete computational problem. We have developed a sub-optimal *heuristic* algorithm that works in $O(N^2 \log N)$ time for 2-complexes. It will briefly be described in Section 2.8.

2.6 The Inclusion Graph Generated by ι

There is an oriented antisymmetric edge-labeled graph G_ι associated with ι (see (4)) whose edges are defined as follows: $\forall (S_i, S_j) \in C \times C$, $i \neq j$:

$$\iota(S_i, S_j) < \iota(S_j, S_i) \Leftrightarrow (i, j) \in G_\iota.$$

²We compute ω for edges as opposed to triangles since triangles maximizing $\omega(i, j, k)$ are not compact.

In other words, the edges (i, j) of the graph are oriented so that if S_i and S_j intersect then $\|\mu(S_i)\| > \|\mu(S_j)\|$. This means that, as long as (6) holds and no two intersecting primitive sets have the same norm, there is no oriented circuit in the graph.

The relation of relative inclusion ι is reflexive, since $\iota(S, S) = 1$, antisymmetric, since $\iota(S_1, S_2) \neq \iota(S_2, S_1)$ or $\iota(S_1, S_2) = \iota(S_2, S_1) = 0$ unless $\|\mu(S_1)\| = \|\mu(S_2)\|$, and the associated graph has no circuits, which means that ι is a fuzzy non-transitive ordinal relation [18]. The transitive closure of ι is thus the relation of *fuzzy order* which represents the hierarchical structure of C . The closure can be found as the minimum-spanning tree of G_ι , $T_\iota \subseteq G_\iota$, which is a polynomial-complexity computational problem.

We assert that T_ι captures the hierarchical structure of the collection C . If the primitive sets in C are recovered at multiple scales (resolutions), the T_ι provides a way to build a multi-resolution representation or a way to remove redundant primitive sets from C by prescribing the maximum approximation error. By selecting a threshold on ι , the T_ι can be converted to a forest, in which each root represents the primitive set that includes all other primitive sets in the component. The threshold, which, by definition of ι , is a fraction of the coarsest-level primitive set size, is to be selected. We want to pursue further research to develop this proposition.

2.7 The Fish-Scale Surface Model

From now on we shall restrict the general discussion to surfaces embedded in Euclidean p -space. Extending the results to any manifold dimension is trivial, since the key structures are defined for any dimensions.

We have selected the exponential influence function $f(t) = e^{-t}$, since the relative intersection and inclusion can be easily estimated and since the moments of the first and second orders of $\mu(x|S)$ are directly related to the primitive set center and covariance matrix.

Let S_i be the structure matrices of the primitive sets S_i , and let $\mathbf{x}_{0i}$ be their respective centers. Let $|S|$ denote the determinant of S , let $P_{12} = (S_1^{-1} + S_2^{-1})^{-1}$, and let

$$d_{12}^2 = \frac{1}{2}(\mathbf{x}_{01} - \mathbf{x}_{02})^T (S_1 + S_2)^{-1} (\mathbf{x}_{01} - \mathbf{x}_{02})$$

be the Mahalanobis distance between the fish-scale centers.

The fuzzy relations (3) and (4) for a pair of primitive sets are then estimated as follows:

1. Relative intersection $\bar{\omega}(S_1, S_2) \leq \omega(S_1, S_2)$:

$$\begin{aligned} \bar{\omega}(S_1, S_2) &= \frac{\|\mu(S_1 \cdot S_2)\|}{\|\mu(S_1 + S_2)\|} \\ &= \frac{\sqrt{|P_{12}|} e^{-d_{12}^2}}{\sqrt{|S_1|} + \sqrt{|S_2|} - \sqrt{|P_{12}|} e^{-d_{12}^2}} \stackrel{\text{def}}{=} \kappa e^{-d_{12}^2}, \end{aligned} \quad (8)$$

where the term d_{12} captures the decrease in intersection with increasing distance between the centers and the ratio κ captures the perpendicularity of the primitive sets, see Fig. 4. Note that $\bar{\omega}(S, S) = (2^{(\frac{p}{2}+1)} - 1)^{-1}$, where p is the dimension of the vector $\mathbf{x}$, whereas $\omega(S, S) = 1$. The ability of $\bar{\omega}$ to estimate ω thus decreases with increasing dimensionality of the measurement space. Nevertheless,

Fig. 4: The value of relative intersection $\bar{\omega}$ for increasing perpendicularity (angle ϕ , in degrees) of two 2-D primitive sets of different rank ρ centered at the same point.

the $\bar{\omega}$ is has similar properties as ω (it is also a resemblance relation).

2. Relative inclusion $\bar{\iota}(\mathcal{S}_1, \mathcal{S}_2) \leq \iota(\mathcal{S}_1, \mathcal{S}_2)$:

$$\bar{\iota}(\mathcal{S}_1, \mathcal{S}_2) = \frac{\|\mu(\mathcal{S}_1 \cdot \mathcal{S}_2)\|}{\|\mu(\mathcal{S}_1)\|} = \frac{\sqrt{|\mathbf{P}_{12}|}}{\sqrt{|\mathbf{S}_1|}} e^{-d_{12}^2}, \quad (9)$$

which is—like $\iota(\mathcal{S}_1, \mathcal{S}_2)$ —also a non-transitive ordinal relation because of (6).

2.8 Surface Reconstruction in $\mathbb{R}^3$

To implement the algorithm described in Section 2.5 we need a local condition acting on 2-simplices (triangles) and 1-simplices (edges) that helps transform the initial collection $\mathbf{C}_\omega$ to a pseudo-manifold $\mathbf{P}_\omega$ by removing some edges (and incident triangles). We use the following:

Condition 1 *The angle between any two triangles incident on a common edge in a complex must be greater than $\frac{2}{3}\pi$.*

This is a sufficient but not necessary condition for a complex to be pseudo-manifold. Clearly, no three triangles can be incident on a single edge without violating this condition. The condition also constrains the surface curvature (which is necessary since we do not constrain the boundary curvature and/or length).

The heuristic algorithm then works as follows:

1. Compute 3-D Delaunay triangulation³ $\mathbf{T}$ over the centers of all primitive sets $\mathcal{C}$.
2. For all edges in $\mathbf{T}$ of increasing ω do the following: if the edge is incompatible with any of the triangles incident on any of its endpoints, remove it (and all incident triangles) from $\mathbf{T}$. This step transforms $\mathbf{T}$ to pseudomanifold $\mathbf{P}_\omega$.

An edge (i, j) is incompatible with a triangle (j, k, l) if both angles $\gamma_1 = \angle\{(j, l, i), (j, k, l)\}$ and $\gamma_2 = \angle\{(i, k, j), (j, k, l)\}$ are smaller than $\frac{2}{3}\pi$ (see Fig. 5).

That we require the pseudosurface $\mathbf{P}_\omega$ to be a sub-complex of Delaunay triangulation is a sufficient but not

³Tetrahedralization, in fact.

Fig. 5: Edge (i, j) is incompatible with triangle (j, k, l) if both $\gamma_1 < \frac{2}{3}\pi$ and $\gamma_2 < \frac{2}{3}\pi$.

Fig. 6: Four images that were used as input data to the polynocular stereo algorithm.

necessary condition. The second condition is also sufficient but not necessary (it is much stronger than Condition 1). If we find the necessary and sufficient conditions our reconstructions will contain less holes when the dataset is severely corrupted by noise.

The computational complexity of this procedure is $O(N^2 \log N)$, which is the complexity of 3-D Delaunay triangulation. The algorithm finds a sub-optimal solution to the optimization problem described in Section 2.5.

3 Experimental Results

We show an example of reconstruction from a quadrinocular stereo system in which all the cameras are fully calibrated. The input images (see Fig. 6) are rectified and sub-sampled to half-resolution. Stereo matching is run independently on three of the six possible image pairs. Matches are selected based on maximum Normalized Cross-Correlation using a 5×5 image window and on left-to-right, right-to-left symmetric ordering and disparity gradient constraints. Sub-pixel resolution disparity is computed in a post-processing step that uses affine distortion model and least-squares parametric estimation. The reconstructed points from all three stereo pairs are shown in Fig. 7a. (It is a close-up of the set shown in Fig. 1. No spatial clipping of the point set was done to get the results shown later.)

Next, spatial binning with parameter σ is used to group points to clusters: least-squares estimates of location $\mathbf{x}_0^0$ are found within each non-empty $3 \times 3 \times 3$ neighborhood of each bin. New location estimates $\mathbf{x}_0^1$ and covariance matrix estimates $\mathbf{S}$ are found in a spherical region of radius 1.5σ centered at $\mathbf{x}_0^0$. The parameters $(\mathbf{x}_0^1, \mathbf{S})_i$ describe the primitive sets. The result at scale $\sigma = 6.9\text{mm}$ is shown in Fig. 7b. In this experiment we did not compute the hierarchical representation based on ι as described in Section 2.6.

Reconstruction results for various scale parameters σ are shown in Fig. 8. The rms error of the reconstructions is about 0.3mm as measured for planar surfaces with random texture slanted no more than 45 degrees with respect to the cameras and at 800mm distance from the rig. More

Fig. 7: Isolated spatial points reconstructed by quadrinocular stereo (a), the primitive sets—each is rendered as an ellipsis and a short needle showing the orientation (b).

Fig. 8: Reconstructed surface, samples from the resolution spectrum (σ is the scale parameter). The spectrum was *not* computed by decimating the triangulation. The shading is simulated.

examples for synthetic data are shown in Fig. 9.

4 Conclusions

In this paper we proposed a novel model for hierarchical discrete manifold representation and discussed the corresponding reconstruction problem. We demonstrated the general method on surface reconstruction from isolated points.

The proposed p -D geometric representation is suitable for incremental reconstruction and for fusion of partial models of arbitrary topology; the reconstruction procedure can effectively suppress outliers and remove redundant measurements without using any computationally expensive robust methods, which makes it particularly suitable for fast processing of point-reconstructions from polynocular stereo.

The advantages of the proposed method are (1) significant *data reduction* compared to the original point set (about 1:40), (2) *accuracy* of less than 0.1% in range, (3) *small sensitivity* of the reconstructed structure to *moderate noise* in the data, (4) the recovered connectivity is *invariant*

Fig. 9: Reconstruction of objects of different topological type (artificial data, Gaussian noise, varying density and SNR). Two close spherical shells (a), SNR = 54dB. The connectivity may change with increased uncertainty as in (b) (the same scale parameter σ , SNR = 14dB). A Möbius strip (c). A torus (d). Input point sets are shown in the left column, the collections of fish-scales rendered as ellipsoids in the middle column, and the reconstructed pseudo-surfaces in the last column.

to Euclidean transformation group, (5) the ability to recur to the input images for *refinement prior to high-level interpretation of the data*, (6) *data fusion is easy* since it is done at intermediate interpretation level, (7) the *topological type* of the recovered entity is allowed to change with the change of resolution, (8) the information from a *large number of cameras* can be processed fast with the only bottleneck in computing power.

The data reduction and speed are very important for telepresence, where the bandwidth is a strong constraint. The accuracy is very important for medical applications. And the stability is critical to partial models fusion.

To our surprise, it turned out that we can cope with the presence of moderate highlights on the observed surface and with faint textures without affecting the precision of the reconstruction. The real-data example we gave in this paper shows just that.

We are developing a general and scalable implementation using the message-passing paradigm of parallel computing of a system that reconstructs 3-D models from images from many cameras working simultaneously (Raymond McKendall, GRASP). We also pursue error modeling and propagation that is expected to help in the clustering step and help guarantee the performance bounds (Gerda Kamberova, GRASP).

Acknowledgements

This work has been supported by the following grants: Army DAAH04-96-1-0007, DARPA N00014-92-J-1647, NSF SBR89-20230. First author has been partly supported by the Czech Grant Agency under the grant GAČR 102/95/1378.

References

- [1] J.-D. Boissonat. Geometric structures for three-dimensional shape representation. *ACM Trans. Graphics*, 3(4):266–286, 1984.
- [2] R. M. Bolle and B. C. Vemuri. On three-dimensional surface reconstruction methods. *IEEE Trans. PAMI*, 13(1):1–13, 1991.
- [3] V. Caselles, R. Kimmel, G. Sapiro, and C. Sbert. Three dimensional object modeling via minimal surfaces. In *Proc. European Conf. Computer Vision*, pp. 97–106, 1996.
- [4] B. Curless and M. Levoy. A volumetric method for building complex models from range images. *Computer Graphics (SIGGRAPH '96)*, pp.303-312, 1996.
- [5] L. de Floriani and E. Puppo. An on-line algorithm for constrained Delaunay triangulation. *Computer Vision, Graphics, and Image Processing: Graphical Models and Image Processing*, 54(3):290–300, 1992.
- [6] H. Delingette. Simplex meshes: a general representation for 3D shape reconstruction. RR-2214, INRIA, Sophia Antipolis, 1994.
- [7] M. Eck and H. Hoppe. Automatic reconstruction of B-spline surfaces of arbitrary topological type. In *Computer Graphics (SIGGRAPH '96)*, pp. 325–334, 1996.
- [8] H. Edelsbrunner and E. P. Mücke. Three-dimensional alpha shapes. *ACM Trans. Graphics*, 13(1):43–72, 1994.
- [9] P. Fua. Reconstructing complex surfaces from multiple stereo views. Tech Note 550, AI Center, SRI International, Menlo Park, CA, 1994.
- [10] P. Fua and Y. G. Leclerc. Using 3-dimensional meshes to combine image-based and geometry-based constraints. In *Proc. European Conf. Computer Vision*, pp. 281–291, 1994.
- [11] P. Fua and P. Sander. Reconstructing surfaces from unstructured 3D points. In *Proc. ARPA Image Understanding Workshop*, 1992.
- [12] W. E. L. Grimson. A computational theory of visual surface interpolation. A.I. Memo 613, MIT, 1981.
- [13] J. L. Gross and T. W. Tucker. *Topological Graph Theory*. John Wiley & Sons, 1987.
- [14] S. Han and G. Medioni. Reconstructing free-form surfaces from sparse data. In *Proc. Int. Conf. Pattern Recognition*, pp. 100–104, 1996.
- [15] A. Hilton, A. J. Stoddart, J. Illingworth, and T. Windeatt. Reliable surface reconstruction from multiple range images. In *Proc. European Conf. Computer Vision*, pp. 117–126, 1996.
- [16] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, and W. Stuetzle. Surface reconstruction from unorganized points. In *Computer Graphics (SIGGRAPH '92)*, pp. 71–78, 1992.
- [17] K. Jänich. *Topology*. Springer-Verlag, 1984.
- [18] A. Kaufmann. *Introduction to the theory of fuzzy subsets*. Academic Press, 1975.
- [19] A. Leonardis, F. Solina, and A. Macerl. A direct recovery of superquadric models in range images using recover-and-select paradigm. In *Proc. European Conf. Computer Vision*, pp. 309–318, 1994.
- [20] R. Lin and W.-C. Lin. Recovery of 3-D closed surfaces using progressive shell model. In *Proc. Int. Conf. Pattern Recognition*, pp. 95–98, 1996.
- [21] J.-C. Lombardo and C. Puech. Oriented particles: A tool for shape memory objects modelling. In *Proc. Graphics Interface*, Quebec City, Canada, 1995.
- [22] A. P. Pentland. Perceptual organization and the representation of natural form. *Artificial Intelligence*, 28(2):293–331, 1986.
- [23] R. Poli, G. Coppini, and G. Valli. Recovery of 3D closed surfaces from sparse data. *Computer Vision, Graphics, and Image Processing: Image Understanding*, 60(1):1–25, 1994.
- [24] W. T. Reeves. Particle systems—a technique for modeling a class of fuzzy objects. *Computer Graphics (SIGGRAPH '83)*, 359–376, 1983.
- [25] J. Rossignac and P. Borrel. Multi-resolution 3D approximations for rendering complex scenes. In *Modeling in Computer Graphics*, pp. 455–465. Springer-Verlag, 1993.
- [26] F. Solina and R. Bajcsy. Recovery of parametric models from range images: The case for superquadrics with global deformations. *IEEE Trans. PAMI*, 12(2):131–147, 1990.
- [27] M. Soucy and D. Laurendeau. Multiresolution surface modeling based on hierarchical triangulation. *Computer Vision, Graphics, and Image Processing*, 63(1):1–14, 1996.
- [28] R. Szeliski. *Bayesian Modeling of Uncertainty in Low-level Vision*. Kluwer Academic Publishers, 1989.
- [29] R. Szeliski and D. Tonnesen. Surface modeling with oriented particle systems. *Computer Graphics (SIGGRAPH '92)*, 185–194, 1992.
- [30] D. Terzopoulos. Multi-level reconstruction of visual surfaces: Variational principles and finite element representations. A.I. Memo 671, MIT, 1982.
- [31] R. S. Wallace. *Finding Natural Clusters through Entropy Minimization*. Phd thesis, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, 1989.
- [32] D. F. Watson. Computing the n -dimensional Delaunay tessellation with application to Voronoi polytopes. *The Computer Journal*, 24(2):167–171, 1981.
- [33] L. A. Zadeh. Fuzzy sets. *Information and Control*, 8:338–353, 1965.