

MESHLESS GEOMETRIC SUBDIVISION

By

Carsten Moenning

Facundo Mémoli

Guillermo Sapiro


Nira Dyn

and

Neil A. Dodgson

IMA Preprint Series # 2001

(October 2004)



INSTITUTE FOR MATHEMATICS AND ITS APPLICATIONS

UNIVERSITY OF MINNESOTA
514 Vincent Hall
206 Church Street S.E.
Minneapolis, Minnesota 55455-0436

Phone: 612/624-6066 Fax: 612/626-7370

URL: <http://www.ima.umn.edu>

Meshless Geometric Subdivision

CARSTEN MOENNING

University of Cambridge

FACUNDO MÉMOLI and GUILLERMO SAPIRO

University of Minnesota

NIRA DYN

Tel Aviv University

and

NEIL A. DODGSON

University of Cambridge

Point-based surface processing has developed into an attractive alternative to mesh-based processing techniques for a number of geometric modeling applications. By working with point clouds directly, any processing is based on the given raw data and its underlying geometry rather than any arbitrary intermediate representations and generally artificial connectivity relations. We extend this principle into the area of subdivision surfaces by introducing the notion of meshless, or point cloud, geometric subdivision. Our meshless subdivision approach replaces the role of mesh connectivity with intrinsic point proximity thereby avoiding a number of limitations of widely-used mesh-based surface subdivision schemes. Apart from introducing this idea of meshless subdivision, we put forward a first intrinsic meshless subdivision scheme and present a new method for the computation of intrinsic means on euclidean manifolds.

Categories and Subject Descriptors: I.3.5 [Computer Graphics]: Computational Geometry and Object Modelling—*Curve, surface, solid, and object representations*

General Terms: Algorithms

Additional Key Words and Phrases: Subdivision, point-based surface processing, Fast Marching

This work was partially supported by the Office of Naval Research and the National Science Foundation. CM performed this work whilst visiting the University of Minnesota at Minneapolis. Authors' current contact details: CM, Computer Lab., University of Cambridge, 15 JJ Thomson Avenue, Cambridge CB3 0FD, UK, email: cm230@cl.cam.ac.uk; FM, Electrical and Computer Engineering Department, University of Minnesota, Minneapolis, MN 55455, USA, email: memoli@ece.umn.edu; GS, Electrical and Computer Engineering Department, University of Minnesota, Minneapolis, MN 55455, USA, email: guille@ece.umn.edu; ND, School of Mathematical Sciences, Tel Aviv University, Tel Aviv 69978, Israel, email: niradyn@post.tau.ac.il; NAD, Computer Lab., University of Cambridge, 15 JJ Thomson Avenue, Cambridge CB3 0FD, UK, email: nad@cl.cam.ac.uk.

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 20YY ACM 0730-0301/20YY/0100-0001 \$5.00

1. INTRODUCTION

Mesh subdivision has developed into a powerful and widely-used tool for the free-form design, editing and representation of smooth surfaces. Subdivision schemes recursively apply a local subdivision operator to a coarse base mesh thereby producing a sequence of refined meshes which quickly converges to a smooth limit surface. The advantages of mesh subdivision include guaranteed global surface smoothness whilst supporting local feature control, the ability to handle surfaces of arbitrary topology and being efficiently and simple to apply once a base mesh is available [Dyn and Levin 2002; Zorin and Schröder 2000].

Unfortunately, when dealing with point-sampled geometry, mesh subdivision requires frequently costly and generally non-geometric surface reconstruction, see, e.g., [Amenta et al. 1998; Amenta et al. 2001; Bernadini et al. 1999; Boissonnat and Cazals 2000; Curless and Levoy 1996; Edelsbrunner and Mücke 1994], often followed by mesh simplification ([Gotsman et al. 2002] and references therein), parameterisation [Floater and Hormann 2004] and remeshing ([Alliez et al. 2003] and references therein), all pre-processing steps to obtain a base mesh. During the reconstruction step, any measurement noise, misalignment of scans, etc. may translate into topological artifacts in the form of erroneous connectivity and genus [Wood et al. 2004]. This hinders subsequent mesh simplification and remeshing and thus mesh subdivision processing. Also, in the case of extremely high-dimensional manifolds by samples [Tenenbaum et al. 2000], mesh pre-processing breaks down at the surface reconstruction step and it needs to be worked directly with the point cloud.

This paper advocates the use of meshless, or point cloud, geometric subdivision. We propose to avoid the consideration of mesh connectivity graphs and the associated pre-processing steps and instead to work with the point-sampled geometry directly using intrinsic subdivision rules. In this paper, we show the conceptual viability of this notion of meshless geometric subdivision by introducing a first intrinsic meshless subdivision scheme using geodesic centroids of intrinsic point neighbourhoods. We put forward a new method for the computation of these geodesic means, which by itself is of interest in other areas such as intrinsic statistical shape analysis [Fletcher et al. 2003] and variational theory [Jost 1994; Kendall 1990].

Some of the related geometric operations may be approximated in an Euclidean context when working with large sampling densities, regular meshes and very local subdivision rules. Working with the point-sampled geometry directly and intrinsically, however, avoids the need for non-geometric pre-processing steps and special rules dealing with irregular mesh connectivity. Unlike the non-geometric nature of Euclidean-based (pre-)processing, meshless intrinsic subdivision inherently captures the non-linear intrinsic structure of the object geometry and the principle applies equally well to three and higher-dimensional surfaces. Although it is possible to obtain a geodesic mesh representation of a point-sampled surface [Peyré and Cohen 2003; Sifri et al. 2003], to continue performing subdivision truly intrinsically, this mesh would have to be modified repeatedly during each iteration to re-enforce its intrinsic nature and more than one mesh would be required to be able to compute correct geodesic distances.

Our algorithm operates intrinsically throughout without the need for prior surface

reconstruction with the help of the geodesic distance mapping algorithm for point clouds presented in Mémoli and Sapiro [2003]. Following an overview of related work in Section 2 and some mathematical preliminaries in Section 3, we therefore summarise the main aspects of this technique in Section 4. Section 5 presents our intrinsic meshless subdivision algorithm. Section 6 gives experimental results and comments on implementational aspects. In the concluding Section 7, we briefly remark on our ideas and some preliminary results for the theoretical analysis of intrinsic meshless subdivision schemes.

2. RELATED WORK

We start with a summary of the ideas underpinning mesh-based subdivision of surfaces in \mathbb{R}^3 . The overview is motivational in nature, for a thorough formal treatment of mesh subdivision, see Dyn and Levin [2002]. The section is concluded with remarks on recent progress in meshless, point-based surface processing related to our work.

Following the notation of Dyn and Levin [2002], surface subdivision schemes consist of a subdivision operator S recursively applied to control nets $N_l = N(V^l, E^l, F^l)$ of arbitrary topology, with $l \in \mathbb{Z}_0$ denoting the subdivision level, V representing a set of control vertices in \mathbb{R}^3 and E and F describing the topological relations in the form of edges and faces respectively. The iterative application of this scheme generates a sequence $N_{l+1} = SN_l$. More specifically, starting with a coarse base net N_0 , at each iteration, new control vertices are inserted and connected according to the scheme's refinement rule and re-positioned following the operator's geometric averaging rule. Both the refinement and the geometric averaging rule give the position of control vertices in N_{l+1} in the form of weighted averages of topologically neighbouring vertices in N_l . The careful choice of these rules in relation to the control vertex valency, i.e. the number of edges emanating from the vertex, guarantees the convergence of the scheme, in each component and in the uniform norm, to a limit surface of provable continuity.

Not every existing mesh subdivision operator allows for this simple distinction between a topological refinement and a geometric averaging rule. However, those that do allow for this kind of distinction, include the most widely-used schemes. For example, the Loop [1987] subdivision scheme for triangular control nets may be cast in this form. In the case of this scheme, the refinement rule adds points related to each edge in the triangulation using face splitting. The averaging rule of points, which depends on the vertices and edges of the non-refined triangulation, then yields the final positions of the vertices in the refined triangulation.

Point-based surface processing advocates the use of points as editing and display primitive thereby avoiding the need for a mesh connectivity graph and the manifold reconstruction step. Recent progress in this field has been sufficiently substantial to realise powerful point-based editing, free-form and multiresolution modelling [Alexa et al. 2003; Linsen 2001; Pauly et al. 2002; Pauly et al. 2003; Zwicker et al. 2002] and visualisation [Grossman and Dally 1998; Kalaiah and Varshney 2001; Levoy and Whitted 1985; Pfister et al. 2000; Rusinkiewicz and Levoy 2000; Zwicker et al. 2001; 2002] alternatives to mesh-based processing methods. With numerous applications in medical imaging, reverse engineering, cultural heritage

preservation, natural science experimentation, geoscience, numerical simulations, etc. acquiring point sets of significant density from three or higher-dimensional manifold surfaces, it is particularly attractive to avoid the generally costly, error-prone and non-geometric mesh-related pre-processing steps and to work with point-sampled geometry directly.

In this paper, we propose to replace the role of mesh connectivity in subdivision with intrinsic point neighbourhood information and formulate a set of meshless refinement and geometric averaging rules in the form of weighted geodesic centroids of these local neighbourhoods.

To the best of our knowledge, only Fleishman et al. [2003] and Guennebaud et al. [2004] touch upon the notion of point cloud subdivision. In Fleishman et al. [2003], the authors generate progressive levels-of-detail of point clouds by transferring the mesh-based idea of subdivision displacement maps to the point cloud case. They devise a point cloud simplification method for the generation of a base point set and present both a projection and a local, uniform upsampling operator with the help of local surface reconstruction using Moving Least Squares [Alexa et al. 2003]. The authors successfully mimic the principle of mesh-based subdivision displacement mapping for surfaces in point cloud form but do not consider the idea of meshless subdivision.

Similarly, Guennebaud et al. [2004]¹ are concerned with the upsampling of “surfel sets”, i.e. input points equipped with normal, local sampling density (surfel radius) and texture information, for magnified point rendering by splatting. Their interpolatory method requires that the underlying point cloud is regularly uniformly distributed and noise-free and that features such as crease lines have been detected and adequately sampled in a pre-processing step. Their technique is restricted in applicability to surfaces in \mathbb{R}^3 and intended for the operation on top of a splatting algorithm such as [Zwicker et al. 2001; 2002] providing the surfel information. As is typically the case for independently determined neighbourhoods such as the authors’ polygonal fan neighbourhoods, their (extrinsic) proximity and refinement operators are affected by overlapping neighbour relations and the need for refinement rules varying with the number of neighbours of the point under consideration. This interpolatory mesh subdivision-inspired method is used for the locally smooth upsampling of surfel sets as part of a dynamic splatting algorithm, the more general notion of meshless subdivision is not considered.

This paper is concerned with the notion of meshless surface subdivision and makes no assumptions on the availability of normal or local density information or the regularity of the input point cloud. An intrinsic neighbourhood concept is used which is based on a partitioning of the surface and avoids the shortcomings of point neighbourhoods determined independently from each other. The approach is not restricted to surfaces in \mathbb{R}^3 but extends to higher dimensions. Inspired by recent work on geodesic curve subdivision [Wallner and Dyn 2003; 2004; 2005], it offers a theoretical basis for global convergence and smoothness analysis.

In the following, we present our intrinsic framework for meshless subdivision. We start with notation and the definition of some key notions used throughout the

¹This work was produced at about the same time as ours, see [Moenning et al. 2004] for an abbreviated publication of the ideas we now put forward in detail.

paper. This is followed by the discussion of the intrinsic distance mapping algorithm of Mémoli and Sapiro [2003] and the presentation of the framework itself.

3. PRELIMINARIES

Let M represent a differentiable (smooth), compact and connected Riemannian submanifold in \mathbb{R}^m , $m \geq 3$. By *intrinsic* processing, we mean processing directly on this submanifold rather than in its embedding space. M is represented by a finite point set, or (unstructured) point cloud, $P = \{p_1, p_2, \dots, p_n\} \subset M$. The Riemannian metric on M at point $x \in M$ is a smoothly varying inner product $\langle \cdot, \cdot \rangle$ on the tangent space $T_x M$. The norm of a vector v in $T_x M$ is given by $\|v\| = \langle v, v \rangle^{\frac{1}{2}}$. We endow M with the metric inherited from \mathbb{R}^m , hence $\langle v, w \rangle$ will coincide with the usual inner product for vectors v and w in \mathbb{R}^m . Consider a (sectionally) smooth curve $\gamma : [a, b] \in \mathbb{R} \rightarrow M$ parameterised by t . The length of $\gamma(t)$ follows from integrating the norm of its tangent vectors, $\dot{\gamma}(t)$, along the curve, i.e.

$$L(\gamma) = \int_a^b \|\dot{\gamma}(t)\| dt$$

The curve γ is called the (minimising) *geodesic* from a point x to a point y in M , if it represents the minimum length-curve among all the curves on M joining x and y . Since we are assuming M to be compact, it is geodesically complete and there exists at least one such curve on M . This curve may not be unique. The length of the geodesic between x and y gives the *intrinsic*, or *geodesic*, *distance*, $d_M(x, y)$, between the points, i.e.

$$d_M(x, y) = \inf_{\gamma} \{L(\gamma)\},$$

with $\gamma(a) = x$ and $\gamma(b) = y$. The function giving the intrinsic distance from a point $x \in M$ to every point in M , $d_M(x, \cdot)$, is called the *intrinsic distance function*, or *intrinsic distance map*, of x . Unique geodesics between two points x and y on M may thus be computed from $d_M(x, \cdot)$ by backtracking from y towards x in the direction of the (negative) gradient of $d_M(x, \cdot)$.

The *extrinsic distance* between points x and y on M , $d(x, y)$, is computed in the metric of the embedding space. Since we are concerned with manifold surfaces in \mathbb{R}^m , the extrinsic distance is Euclidean and given by the length of the Euclidean line segment between x and y in the ambient space. Apart from its end points, this line segment generally does not lie on the manifold. More detail on the above notions may be found in, for example, Chavel [1997].

By *intrinsic*, or *geodesic*, *centroid*, we understand the mean of a local neighbourhood of points on M computed in terms of intrinsic distances between the points. This is to be distinguished from the *extrinsic centroid* of the subset, computed using Euclidean distances in the ambient space and subsequently projected onto the manifold surface.

Our meshless subdivision scheme makes extensive use of the geodesic Voronoi diagram concept: Define the bisector $BS(p_i, p_j)$ of $p_i, p_j \in P$, $p_i \neq p_j$, as geodesically equidistant loci with respect to p_i, p_j , i.e. $BS(p_i, p_j) = \{q \in M \mid d_M(p_i, q) = d_M(p_j, q)\}$. Let the dominance region of p_i , $D(p_i, p_j)$, denote the region of M containing p_i bounded by $BS(p_i, p_j)$. The Voronoi region of p_i with respect to point

set P is given by $R(p_i, P) = \bigcap_{p_j \in P, p_j \neq p_i} D(p_i, p_j)$ and consists of all points for which the geodesic distance to p_i is smaller than the geodesic distance to any other $p_j \in P$. The boundary shared by a pair of Voronoi regions is called a Voronoi edge. Voronoi edges meet at Voronoi vertices. The *geodesic Voronoi diagram* of P is defined as

$$VD(P) = \bigcup_{p_i \in P} \partial R(p_i, P),$$

where $\partial R(p_i, P)$ denotes the boundary of $R(p_i, P)$. Kunze et al. [1997] and Leibon and Letscher [2000] give further details on the notion of geodesic Voronoi diagrams.

4. INTRINSIC DISTANCE MAPPING ACROSS POINT CLOUDS

To compute intrinsic distance maps, and from them (minimal) geodesics, across point clouds, we use the extension of the original Fast Marching method [Helmsen et al. 1996; Sethian 1996; Tsitsiklis 1995] to surfaces in point cloud form introduced by Mémoli and Sapiro [2003]. In the following, we summarise the main aspects of this technique. For full details, see Mémoli and Sapiro [2001; 2003].

Fast Marching generally represents a very efficient technique for the solution of front propagation problems which can be formulated as boundary value partial differential equations. Take the simple case of a front propagating across a 3D Euclidean domain with speed, or weight, function $F(v)$ away from a source (boundary) u , with u and v representing 3-tuples in \mathbb{R}^3 . We are interested here in the arrival time, or offset distance, $d(u, v)$, of the front at grid point v . The relationship between the magnitude of the distance map's gradient, $|\nabla d(u, v)|$, and the given weight $F(v)$ at v can be expressed as the following boundary value formulation

$$|\nabla d(u, v)| = F(v),$$

with boundary condition $d(u, u) = 0$. The problem of determining a Euclidean weighted distance map has therefore been transformed into the problem of solving a particular type of static Hamilton-Jacobi partial differential equation, the nonlinear Eikonal equation. For $F(v) > 0$, this type of equation can be solved efficiently, in computationally optimal time, using conventional Cartesian numerics, see, [Helmsen et al. 1996; Sethian 1996; Tsitsiklis 1995].

Mémoli and Sapiro [2003] extend the applicability of this Fast Marching idea to the case of *general co-dimension* submanifolds in point cloud form in *three or higher dimensions*. For simplicity, consider the constant radius r -offset Ω_P^r , i.e. the union of Euclidean balls with radius r centred at points $p_i \in P$ (Figure 1)

$$\Omega_P^r := \bigcup_{i=1}^n B(p_i, r) = \{x \in \mathbb{R}^m : d(p_i, x) \leq r\}$$

To approximate the weighted intrinsic distance map, $d_M(q, \cdot)$, originating from a source point $q \in M$ on M , Mémoli and Sapiro [2003] suggest computing the Euclidean distance map in Ω_P^r , denoted $d_{\Omega_P^r}$. That is

$$|\nabla_M d_M(p, \cdot)| = F(p), \tag{1}$$

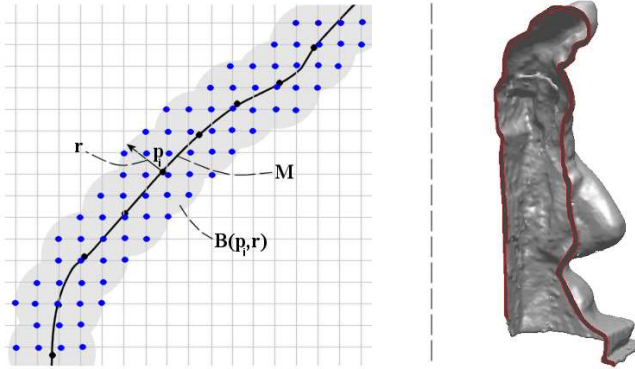


Fig. 1. Intrinsic distance mapping using Fast Marching for point clouds operates in an (not necessarily constant radius) offset band consisting of the union of balls $B(p_i, r)$ centred at (black) points p_i of the surface M (left). Only those (blue) grid points falling inside the offset band are considered during processing. A cross-sectional view of a constant radius offset band for the Michelangelo Youthful data set is shown on the right.

for $p \in M$ and with boundary condition $d_M(q, q) = 0$ is approximated by

$$|\nabla d_{\Omega_P^r}(p, \cdot)| = \tilde{F}(p), \quad (2)$$

for $p \in \Omega_P^r$ and boundary condition $d_{\Omega_P^r}(q, q) = 0$. \tilde{F} represents the (smooth) extension of the propagation speed F on M into Ω_P^r ; ∇_M denotes the intrinsic gradient operator. The problem of computing an intrinsic weighted distance map is therefore transformed into the problem of computing an Euclidean, or extrinsic, weighted distance map in the offset band Ω_P^r around the surface, i.e. in an Euclidean manifold with boundary.

Mémoli and Sapiro [2003] prove uniform (probabilistic) convergence between these two distance maps, and geodesics computed from them, for *both noise-free and noisy* (provided noise is bounded from above by r), randomly-sampled point clouds and thus show that the approximation error between the intrinsic and extrinsic distance maps is of the same theoretical order as that of the conventional Fast Marching algorithm. Fast Marching can therefore be used to approximate the solution to (2) in a *computationally optimal* manner and *without the need for any prior surface reconstruction* by only slightly modifying the conventional Cartesian Fast Marching technique to deal with bounded spaces. This is achieved by simply restricting the grid points visited by the conventional Fast Marching algorithm to those located in Ω_P^r . By performing the computations within this offset band, this method is relatively robust in the presence of noisy point samples, especially when compared to graph-based distance mapping algorithms such as [Giesen and Wagner 2003; Tenenbaum et al. 2000] in which case the geodesics pass through the noisy samples rather than an union of Euclidean balls centred at the input points.

The complexity of this algorithm is $O(N \log N)$, where N represents the number of grid points located in the (narrow) offset band Ω_P^r . Memory efficiency is achieved by storing these grid points only. Note in this context that subject to the bounds given in Mémoli and Sapiro [2003], r will generally be small and does not have to be

constant but may vary with each p_i . Implementational details are given in Section 6.2.

This Fast Marching technique underpins our geodesic computation algorithm presented as part of our intrinsic meshless subdivision framework in the following section.

5. INTRINSIC MESHLESS SUBDIVISION

Subdivision schemes incorporate refinement and geometric averaging rules in the form of weighted averages of local neighbourhoods. Mesh-based subdivision schemes are derived in a parametric setting ignoring the geometric embedding of the points in space. As a result, they are formulated in terms of local mesh connectivity rather than object geometry. We determine local neighbourhood relations from intrinsic point proximity information instead.

We start this section with the discussion of a suitable intrinsic neighbourhood concept. This is followed by the presentation of our meshless subdivision scheme and a new method for the computation of geodesic centroids, which is at the heart of this scheme.

5.1 Intrinsic point proximity information

Depending on the acquisition technique used, point-sampled geometry may feature (locally) non-uniform point distributions. In this setting, naive neighbourhood concepts such as simple ball or k nearest neighbourhoods tend to be skewed, i.e. the neighbours $q_j \in P_l$ are frequently no longer distributed spherically, i.e. all around, the point $p \in P_l$ under consideration [Floater and Reimers 2001; Linsen 2001].

To allow for local sampling non-uniformities, Linsen’s [2001] enhanced k nearest neighbourhood, which enforces a maximum angle between successive neighbours around p , or Floater and Reimer’s [2001] local Delaunay neighbourhood may be used. Their algorithm projects a local ball neighbourhood of p and p itself into their least squares plane and computes the planar Delaunay triangulation of the projected points. The points’ neighbour relations in this triangulation are then taken as the neighbour relations of p on the manifold. We suggest to collect local proximity information by considering the set of (Voronoi) neighbours, N_p , of p in the geodesic Voronoi diagram of P_l , $VD(P_l)$, instead, i.e. an *intrinsic* “natural neighbourhood” [Sibson 1980],

$$N_p = \{q : p \text{ and } q \text{ are neighbours in } VD(P_l)\},$$

for $p, q \in P_l$, $p \neq q$ (Figure 2). Due to its intrinsic nature and unlike (Euclidean) ball neighbourhoods, points from disjoint parts of a surface are prevented from being assigned to the same natural neighbourhood. The need for determining local ball radii and the problem of locally ill-defined least squares fits in the case of k nearest neighbourhoods do not arise. Since intrinsic natural neighbourhoods are derived from a partitioning of the surface, the problem of overlapping neighbour relations is avoided as well. This problem is typically encountered when using point neighbourhoods computed independently from each other such as k nearest neighbours and its variations [Guennebaud et al. 2004; Linsen 2001] and occurs




Fig. 2. As illustrated here for the planar case, the intrinsic Voronoi diagrams (solid lines) of various moderately irregularly distributed point sets provide point proximity relations in the form of *intrinsic* natural neighbourhoods (blue points) spherically distributed around the (red) point under consideration.

irrespective of the regularity of the point set distribution. Intrinsic natural neighbourhoods are, however, less straightforward to compute. We discuss our approach in detail in Section 6.

We use intrinsic natural neighbourhood information for the computation of local geodesic centroids as part of our meshless subdivision scheme presented next.

5.2 An intrinsic meshless subdivision scheme

Within our meshless geometric subdivision framework, geodesic centroids of the intrinsic neighbourhoods discussed in the previous section are used to define meshless subdivision rules. More specifically, we propose the following set of rules to be applied at each iteration:

Refinement rule: For each neighbour $q_j \in N_p$, consider the union of intrinsic neighbours of $p, q_j \in P_l, N_{pq_j}$. Upsample P_l by inserting the weighted geodesic centroid, $c(N_{pq_j}) \in P_{l+1}$, of N_{pq_j} .

Geometric averaging rule: Replace $p \in P_l$ by the weighted geodesic centroid, $c(N_p) \in P_{l+1}$, of its intrinsic Voronoi neighbourhood N_p .

This use of weighted centroids in the refinement and geometric averaging rules is reminiscent of both classical subdivision schemes [Zorin and Schröder 2000] (Section 2) and the “repeated averaging” approach towards the generation of subdivision surfaces ([Lane and Riesenfeld 1980; Oswald and Schröder 2003] and references therein). These subdivision rules are incorporated into our meshless subdivision algorithm as summarised in Algorithm 1.

In its initialisation phase, the algorithm buckets the input point set P_l in a Cartesian grid, subsequently used to support weighted distance mapping and geodesic centroid computation. Initialisation is completed with the computation of the (discrete) geodesic Voronoi diagram of P_l , $VD(P_l)$. Main processing loops over all points in P_l and proceeds with the upsampling of P_l to P_{l+1} using joint neighbourhood information readily available from $VD(P_l)$ and our refinement rule. Following this refinement step, $VD(P_{l+1})$ is computed. Meshless subdivision is concluded with the application of our geometric averaging rule to each point in the refined point

<p>Input: Point cloud $P_l \in \mathbb{R}^m$. Output: Subdivided point cloud P_{l+1}.</p> <hr/> <pre> 0 *** INITIALISATION *** 1 Bucket the base point cloud P_l in a m-dimensional Cartesian grid; 2 Compute the discrete geodesic Voronoi diagram, $VD(P_l)$, of P_l. 3 4 *** MAIN PROCESSING *** 5 FOR each point $p_i \in P_l$; 6 Determine the Voronoi neighbourhood N_{p_i} from $VD(P_l)$; 7 FOR each neighbour $q_j \in N_{p_i}$; 8 Determine the joint Voronoi neighbourhood $N_{p_i q_j}$ from $VD(P_l)$; 9 Compute the weighted geodesic centroid $c(N_{p_i q_j})$; 10 (<i>Refinement rule</i>) Upsample P_l to P_{l+1} by inserting $c(N_{p_i q_j})$; 11 ENDFOR 12 ENDFOR 13 Compute $VD(P_{l+1})$; 14 FOR each point $p_i \in P_{l+1}$; 15 Determine the Voronoi neighbourhood N_{p_i} from $VD(P_{l+1})$; 16 Compute the weighted geodesic centroid $c(N_{p_i})$; 17 (<i>Geometric averaging rule</i>) Replace p_i in P_{l+1} with $c(N_{p_i})$; 18 ENDFOR </pre>
--

Alg. 1: One iteration of meshless subdivision in pseudocode.

set P_{l+1} yielding the final, subdivided point set.

The applicability of this subdivision algorithm does not depend on a preceding simplification step. Potential algorithm applications, however, include the case in which it may be desirable to simplify an excessively dense point cloud P to a suitable base point set P_0 in the expectation that recursive subdivision of P_0 results in a smoother, more regular and more compact approximation of the underlying surface than given by P . Since our subdivision algorithm requires the computation of geodesic centroids across the base point set P_0 , for these centroids to be well-defined, any simplification of P needs to be performed subject to a minimum point density in P_0 . Due to the method's simple control of a guaranteed point density, its purely intrinsic operation, its close relationship with the natural neighbourhood concept of Section 5.1 and its efficient implementability using the geodesic distance mapping method of Mémoli and Sapiro [2003] (Section 4), we use the algorithm of Moenning and Dodgson [2004] for the simplification of an input point cloud P to a base point set P_0 still sufficiently dense to support the computation of geodesic centroids. As another result of this pre-processing step, the (discrete) geodesic Voronoi diagram of P_0 becomes available [Moenning and Dodgson 2004] so that it does not need to be computed in our subdivision algorithm's initialisation phase and the natural neighbours of a point $p_i \in P_0$ are readily known.

By performing the averaging intrinsically, our meshless subdivision rules raise the questions of how to compute geodesic centroids on manifolds and how to determine a suitable neighbour weighting scheme. In this paper, we guide the choice of weights by experimental results (Section 6) rather than theoretical evidence for the scheme's convergence towards a smooth limit surface. Future work will consider formal proofs




Fig. 3. The unweighted centroid of a (blue) subset of this set of points is expected to be located on or near the underlying surface. Due to the use of intrinsic distances, this is the case when computing the geodesic centroid (red). By contrast, in the case of the Euclidean averaging of the (blue) points, the resulting centroid (grey) is located away from the underlying surface. This effect gets more pronounced when increasing the size of the subset (from left to right). Note that for *geodesically close* neighbourhoods and those kinds of neighbourhoods only, the *orthogonally projected* ($\Pi_M : \Omega_M^r \rightarrow M$) Euclidean average, i.e. the extrinsic mean, generally provides a good (first) estimate of the geodesic centroid (leftmost).

of the scheme’s convergence to a limit surface and, consequently, any light such proofs may throw on the optimal choice of weights.

As regards the computation of geodesic centroids, Buss and Fillmore [2001] present an algorithm for the computation of geodesic averages on spheres. We generalise the underlying, earlier idea ([Karcher 1977] and references therein) of minimising a least squares expression in geodesic distances in the following section.

5.3 Geodesic centroid computation

The benefit of performing the averaging intrinsically is that it ensures that subdivision generates smoother, denser representations which remain geometrically close to the surface. This is not guaranteed to be the case when considering Euclidean instead of geodesic centroids. For the simple example illustrated in Figure 3, Euclidean averaging ignores the non-linear, intrinsic geometry of the object and moves the centroid away from the surface. By contrast, since the computation of the geodesic centroid is based on intrinsic rather than Euclidean distances, it is inherently geometry-sensitive and falls onto the surface in each case.

The weighted geodesic centroid of a set of n points is defined as the point $g \in M$ which minimises the weighted sum of squared intrinsic distances to each point

$$J(g) := \frac{1}{2} \sum_{k=1}^n w_k d_M^2(g, p_k),$$

where w_1, \dots, w_n represent the point weights, with $0 \leq w_i \leq 1$, $\sum_{i=1}^n w_i = 1$. In general, $\operatorname{argmin}_g J(\cdot)$ may not exist or may not be a single point. However, if p_1, \dots, p_n are all contained in a sufficiently small open geodesic ball B_M on M , a unique solution, g_{B_M} of $J(\cdot)$, which happens to lie in B_M [Karcher 1977], is guaranteed. The property we are alluding to here is (geodesic) convexity, i.e. for any $p_i, p_j \in B_M$, the minimal geodesic from p_i to p_j is unique in M and contained in B_M .

In the Euclidean case, direct differentiation of $J(\cdot)$ yields the minimiser $g =$

$\sum_{k=1}^n w_k p_k$. This simple result does not extend to the general case considered here but we can prove that any minimiser must satisfy:

$$V(g) := \sum_{k=1}^n w_k \nabla_M \frac{1}{2} d_M^2(g, p_k) = 0.$$

Then, starting from a *good* initial guess g_0 , we can track the minimiser g using back propagation with velocity field $V(\cdot)$. This is due to the fact that if $g_0 \in B_M$ and B_M as above, then $-V(x)$ points towards g_{B_M} , for $x \in B_M$ [Karcher 1977].

In practise, we set $g_0 = \Pi_M(\sum_{k=1}^n w_k p_k)$, where $\Pi_M : \Omega_M^r \rightarrow M$ is the orthogonal projection operator². We now show that in the light of the considerations presented above, this extrinsic mean represents either a reasonable initial condition for the back propagation or a first approximation to the intrinsic centroid. By a simple application of Lemma 17 of [Wallner and Dyn 2003], we have that

$$\left\| \sum_{k=1}^n w_k p_k - \Pi_M \left(\sum_{k=1}^n w_k p_k \right) \right\| \leq C(\text{diam}(B))^2,$$

where C is a global constant which depends on the curvatures of M . Then, let $B = B_M(x, \epsilon)$, for some $x \in M$ and $\epsilon > 0$. Since $\|p_i - x\| \leq d_M(p_i, x) \leq \epsilon$, we also have $\|\sum_{k=1}^n w_k p_k - x\| \leq \epsilon$. Therefore, since $\|g_0 - x\| \leq \|g_0 - \sum_{k=1}^n w_k p_k\| + \|\sum_{k=1}^n w_k p_k - x\|$, we obtain

$$\|g_0 - x\| \leq C\epsilon^2 + \epsilon,$$

which implies $d_M(g_0, x) \leq \epsilon(1 + H\epsilon)(1 + C\epsilon)$, for another constant H depending on global metric properties of M [Mémoli and Sapiro 2003]. We only care for a simplified bound

$$d_M(g_0, x) \leq E\epsilon.$$

Finally, let $\bar{\delta} > 0$ be the maximal $\delta > 0$ such that $B_M(x, \delta)$ is (geodesically) convex. Note that it is a fact that if $\delta \leq \frac{1}{2} \min\left(\text{inj}(M), \frac{\pi}{\sqrt{K}}\right)$, where $\text{inj}(M)$ is the injectivity radius of M and K bounds all sectional curvatures in M from above, then $B_M(x, \delta)$ is convex for any $x \in M$. See §7.6 and §7.7 in [Chavel 1997]. For such a $\bar{\delta} > 0$ and provided $\epsilon \leq \bar{\delta}/E$, and $\{p_1, \dots, p_n\} \subset B_M(x, \epsilon)$ for some $x \in M$, $g_0 \in B_M(x, \bar{\delta})$ and $-V(g_0)$ will be pointing towards g_{B_M} . Also, in case we want to use g_0 as an approximation to g_{B_M} , we have the (weak) bound $d_M(g_{B_M}, g_0) \leq (E+1)\epsilon$. Therefore, g_0 , as defined above, represents a sensible choice as the initial condition of an eventual back propagation step, or, in any case, a rough approximation to g_{B_M} with known error bound. See also Figure 3 (left). Note in particular that it is also a useful choice from the point of view of computational ease. The algorithm is summarised in Algorithm 2.

To demonstrate the applicability of this approach in the context of meshless subdivision, we first consider the case of M representing the unit sphere in the following section.

²To start from any of the points p_k represents another simple choice.

Input: Intrinsic Voronoi neighbourhood N_p of point $p \in P_l$. Weights w_i at points $q_i \in N_p$.

Output: Weighted geodesic centroid g .

```

0 *** Computation of extrinsic centroid  $g_0$  ***
1   Compute the Euclidean weighted centroid of  $N_p$ ;
2   Compute  $g_0$  by orthogonally projecting the weighted Euclidean centroid;
3
4 *** Computation of intrinsic centroid  $g$  ***
5   Compute local weighted distance maps  $d_{\Omega_p}(q_i, \cdot)$  from each neighbour  $q_i \in N_p$ 
   outwards and accumulate their squared values at the grid vertices;
6   Approximate the gradient of the accumulated distance maps using finite
   difference approximation;
7   Back propagate from  $g_0$  towards  $g$  by following the negative gradient;

```

Alg. 2: Procedure for computation of a weighted geodesic centroid in pseudocode.

6. EXPERIMENTAL RESULTS AND IMPLEMENTATIONAL DETAILS

We begin with the intrinsic meshless subdivision of a set of points sampled relatively regularly uniformly from the surface of the unit sphere in \mathbb{R}^3 . This initial restriction to spherical geometry allows for the computation of precise geodesic distances without the need for numerical techniques. This way qualitative and quantitative aspects of our operator can be presented without the influence of the particular projection and gradient descent techniques utilised when processing more complex geometry. This presentation is followed by applications of our subdivision operator to more complex geometry. The section concludes with comments on some implementational aspects.

6.1 Results and discussion

To implement the intrinsic centroid computation method, techniques for the computation of intrinsic distances between points on the surface, the projection of the starting point for the back propagation onto the surface and the computation of the back propagation itself are required. In the case of the unit sphere, these techniques are readily available and no numerical techniques are required. Intrinsic distances between points follow trigonometrically and orthogonal projection is trivial. Similarly, the exponential map and its inverse are directly available and may be used to implement the back propagation procedure. As a result, for the case of spherical geometry, our approach for geodesic centroid computation narrows down to the technique of Buss and Fillmore [2001].

We apply our intrinsic meshless subdivision operator to a base point set P_0 of 2144 points sampled relatively regularly uniformly from the unit sphere (Figure 5). The application of our subdivision operator to P_0 using the initial intrinsic proximity information from $VD(P_0)$ yields the subdivided point set P_1 of Figure 6. The result, P_2 , obtained from the application of the operator to P_1 using natural neighbourhood information from $VD(P_1)$ is shown in Figure 7. Given the relatively strong regularity of the input data, uniform weighting was used for both the refinement and the geometric averaging rule in both iterations. The results produced by our meshless subdivision operator are presented alongside the point sets produced




Fig. 4. Histograms of the (spherical) distance from each point in P_1 (left) and P_2 (right) to its closest neighbour in the respective set.

Model \ k	1	2	3	4	5	6	7	8	9	10
P_1	0.807	0.832	0.832	0.821	0.788	0.809	0.8132	0.786	0.818	0.828
P_2	0.781	0.815	0.823	0.804	0.785	0.799	0.7980	0.784	0.810	0.822

Table I. Values of the density uniformity measure $\widehat{\rho}(k)$ for P_1 and P_2 and with $k \in \{1, 2, \dots, 10\}$. The values underline the regular uniformity of the subdivided point sets generated by our algorithm.

by the application of Loop subdivision to a triangular mesh representation of P_0 . As indicated by the detail views of Figure 7, the point distributions obtained from the two operators after two iterations are qualitatively similar with the slight irregularities in the distribution of P_2 being slightly more pronounced in the case of meshless subdivision due to the use of uniform weights. There are, however, no noticeable differences in the smoothing effect of these two operators.

In order to analyse the point set distributions generated by our subdivision operator quantitatively, we compute the mean and the standard deviation of the distance from each point in the set to its closest neighbour(s) for the subdivided point sets P_1 and P_2 . For each p in the point set P_l , let $sd_k(p)$ denote the (spherical) distance from p to its k th closest neighbour. As an indicator for the uniformity of the density of point set P_l , consider $\rho(k) = \frac{\min_P sd_k(p)}{\max_P sd_k(p)}$. Since $\rho(k)$ represents an absolute measure, it may be too sensitive, therefore we compute instead $\widehat{\rho}(k) = \frac{\text{mean}(sd_k) - \text{std}(sd_k)}{\text{mean}(sd_k) + \text{std}(sd_k)}$, where mean and std stand for the mean and standard deviation of the spherical distances over the point set respectively.

The histograms of $sd_1(x)$ corresponding to the two sets of points are given in Figure 4. Note in particular in Table I that the values of $\widehat{\rho}(k)$, for $1 \leq k \leq 10$, are quite close to 1.0 therefore indicating small dispersion up to the 10th closest neighbour.

Figure 8 presents an application example dealing with more complex geometry. We apply our meshless subdivision operator to a base point set of 10088 points generated from the Michelangelo Youthful data set with the help of [Moenning and Dodgson 2004]. Experimentation revealed the base point set to be regular enough to allow for simple reciprocal distance weighting in the computations of the weighted

geodesic centroids. The flatly shaded renderings of the surfaces reconstructed from the subdivided point sets P_1 and P_2 illustrate the smoothing effect of the meshless subdivision. As indicated by the comparative illustrations in Figure 9, this approach may be used to obtain a smoother, more regular and more compact representation of an highly dense point cloud. A similar effect is shown in Figure 10. The 50% decimated versions of the rocker arm and screwdriver CAD data sets available from the Cyberware website were meshlessly subdivided twice. The smoothing effect of these iterations is again clearly apparent when comparing the surfaces reconstructed from the subdivided point sets to those reconstructed from P_0 and the non-simplified, non-subdivided data sets respectively.

The detail view of Figure 8 highlights the local clustering effect caused by overlapping neighbour relations as discussed in Section 5.1, an effect typically only encountered with point neighbourhoods which determine the neighbour relations independently of each other. Our discrete approximation of the intrinsic Voronoi diagram, however, implies discretisation error and thus individual instances of a point being assigned to the wrong Voronoi region. The generally limited number of these instances is considered preferable to the complications associated with addressing the problem of overlapping neighbour relations (see, e.g. Guennebaud et al. [2004]) when using alternative neighbourhood concepts.

The limited impact of these assignment errors is illustrated by the detail view of Figure 11. The Isis data set was subdivided once with the regularity of the subdivided point set being only mildly affected by erroneous neighbourhood assignments. In contrast to the processing of the Youthful base point set, meshless subdivision of the Isis point cloud using geodesic centroids was found to yield results not significantly different from the more efficient subdivision by extrinsic centroids. We exploit this observation, due to the high density of the initial point set, by presenting the results from orthogonal projection of the uniformly weighted Euclidean centroids, i.e. without subsequent gradient descent towards their geodesic counterparts.

Table II summarises the parameter settings and point set sizes for the various application examples. Using our non-optimised implementation, the offset band computation pre-processing step and a meshless subdivision iteration took a maximum of a few hundred seconds each. Both offset band and intrinsic Voronoi diagram computation efficiency generally depend strongly on the offset band radius r and the grid spacing, our settings of which are presented in the table. Note in this context that the numerical error inherent in the intrinsic distance computations increases with r , the admissible minimum value of which necessarily increases with the grid spacing [Mémoli and Sapiro 2001]. Details on various aspects of the implementation of our algorithm are provided in the following section.

6.2 Implementational details

The algorithm was implemented in C++ (Microsoft Visual C++ 7.1) with the help of the “Blitz++ 0.7 Numerical Library” [Blitz++ 2003] on a Pentium 4 2.8GHz, 512MB Windows XP machine. In the following, we summarise relevant implementational details.

Iteration 1				
	$\Delta x (= \Delta y = \Delta z)$	r	$ P_0 $	$ P_1 $
Sphere	0.25	0.8	2144	8570
Youthful	1.0	2.2	10080	39888
Screwdriver	0.5	1.0	13577	56220
Rocker arm	1.0	1.9	20088	81442
Isis	0.1	0.2	187644	760162

Iteration 2				
	$\Delta x (= \Delta y = \Delta z)$	r	$ P_1 $	$ P_2 $
Sphere	0.25	0.75	8570	34275
Youthful	0.25	0.75	39888	208010
Screwdriver	0.25	0.6	56220	295110
Rocker arm	0.25	0.45	81442	488212

Table II. Parameter settings and point set sizes for the application examples; $\Delta x, \Delta y, \Delta z$ refer to the grid spacing in the three principal Cartesian grid directions; r represents the constant offset band radius.

6.2.1 *Offset band computation and intrinsic distance mapping.* Intrinsic distance mapping requires, firstly, the computation of the offset band Ω_P^r and, secondly, weighted Euclidean distance mapping within the offset band. We meet both of these requirements with the help of conventional Fast Marching. For a given r , offset band computation amounts to the simultaneous propagation of (circular) fronts from each input point outwards until the front’s extent equals r . The offset band consists of those grid vertices visited during the propagation. All other grid vertices of the point set’s bounding volume are discarded to minimise memory usage. When dealing with relatively large point sets, we perform this offset band computation as a pre-processing step which makes the set of valid grid vertices available for subsequent weighted distance mapping. Irregularly distributed point sets may require the offset band to be adaptive. This can be achieved by computing radii r_i using a minimum spanning tree or local principal component analysis as discussed in [Mémoli and Sapiro 2001; 2003]. Weighted distance mapping within Ω_P^r (or $\Omega_P^{r_i}$) is achieved by another application of conventional Fast Marching restricted to operate within the offset band only, i.e. the set of grid vertices returned by the preceding band fitting step.

A single min-heap as typically used for the efficient implementation of the conventional Fast Marching method [Sethian 1999] and a Cartesian grid represent the main data structures required in this context. The grid data structure is implemented in the form of a lookup table with each grid vertex mapped to, amongst others, its distance map value and min-heap index and offset band membership status. By using a lookup table, invalid grid vertices can be quickly discarded and different grid spacing in each direction is supported.

6.2.2 *Geodesic Voronoi diagrams and intrinsic natural neighbourhoods.* The natural neighbourhood information of a point $p_i \in P_l$ is available from $VD(P_l)$. We approximate this diagram using weighted geodesic distance maps, i.e. in analogy to the dropping of pebbles in still water, circular fronts move across the surface from the points of impact. The locations where wave fronts meet define the geodesic

Voronoi diagram of the points of impact. Figure 12 gives a triangular mesh-based example produced using public domain software [Peyré and Cohen 2003]. The wave propagation is discretised and simulated accurately by solving (2) using our implementation of the intrinsic distance mapping for surfaces in point cloud form discussed in Section 6.2.1. This way Voronoi edges/vertices and thus each point’s Voronoi neighbours are obtained *during front propagation* as loci of intersection between geodesic offset curves [Cohen 2001].

The initial geodesic Voronoi diagram, $VD(P_0)$, is either computed during the algorithm’s initialisation phase by propagating fronts simultaneously from all $p_i \in P_0$ outwards (Figure 12) or follows from prior intrinsic point cloud simplification as indicated in Section 5.2.

The neighbourhood information is held in the form of a lookup table mapping each p to its set of natural neighbours N_p represented by a set of indices referencing the corresponding input points. The lookup table of (valid) grid vertices, already discussed above, represents the only other main data structure used in this context and is required here, as above, for distance mapping support.

6.2.3 Geodesic centroid computation. Using the natural neighbourhood information N_p of point p , its weighted Euclidean centroid is readily available. We use the “almost” orthogonal projection operator of Alexa and Adamson [2004] to project this centroid onto M thereby obtaining g_0 (Section 5.3). Fast Marching for point clouds [Mémoli and Sapiro 2003] is then again used to compute distance maps within the offset band from each point in N_p outwards with the squared distance map values being accumulated at the grid vertices. To avoid unnecessary propagation, the extent of the distance mapping is limited to the radius of the sphere enclosing the points in N_p . A standard Runge-Kutta gradient descent procedure with multilinear interpolation is finally employed to back propagate from g_0 towards the weighted geodesic centroid g by following the (negative) gradient of the accumulated distance maps estimated with the help of the normalised central difference operator provided by Blitz++ 0.7 [Blitz++ 2003].

7. CONCLUSION

We introduced the concept of meshless, or point cloud, subdivision based on the computation of weighted geodesic centroids on manifolds represented by noise-free or noisy point clouds. We implemented and showed the applicability of this technique with the help of a new method for the computation of such centroids.

The consideration of local intrinsic point proximity instead of mesh connectivity helps to overcome some of the limitations of existing mesh subdivision schemes. For example, by working with the raw data and operating directly across the point cloud, problems associated with the complexity of the topological subdivision of high-dimensional meshes are avoided. For extremely high-dimensional data, meshless subdivision operators need to be devised, however, which upsample the point-sampled geometry more conservatively than the operator suggested in this paper. In this direction, we are considering introducing adaptive neighbourhoods based on curvature estimators such as those reported in Cazals and Pouget [2003] and Mitra and Nguyen [2003]. We leave this to future work.

As first proposed in the context of univariate spline subdivision schemes [Lane

and Riesenfeld 1980], we plan to investigate the practical and theoretical aspects of repeating the geometric averaging step several times after each refinement step. We expect to get higher smoothness with a higher number of averaging steps in each iteration of the meshless subdivision process.

Depending on the extent of any non-uniformities of the input point cloud, the experimental selection of point weights tends to be elaborate. We are currently working on the more systematic choice of weights. This issue is of course closely related to the theoretical analysis of our meshless intrinsic subdivision scheme. In this context, we intend to build upon ideas of Wallner and Dyn [Wallner and Dyn 2003; 2004; 2005], which this research was inspired by, in particular their convergence and smoothness analysis of non-linear geodesic curve subdivision by proximity to a corresponding linear extrinsic subdivision scheme.

ACKNOWLEDGMENTS

We gratefully acknowledge the permission to use the Michelangelo point sets granted by the Stanford Computer Graphics group. The Isis, 50% decimated and non-simplified CAD data sets and the Bunny point set were taken from the Cyberware and Stanford 3D Scanning Repository websites respectively. The (extended) Loop subdivision implementation was obtained from Henning Biermann's website at <http://mrl.nyu.edu/~biermann/subdivision/>. Surface reconstructions were performed with the help of Paraform's Points2Polys software available at <http://paraform.com/ppdl/>.

REFERENCES

- ALEXA, M. AND ADAMSON, A. 2004. On normals and projection operators for surfaces defined by point sets. In *Proc. Symp. on Point-Based Graphics*. 149–156.
- ALEXA, M., BEHR, J., COHEN-OR, D., FLEISHMAN, S., LEVIN, D., AND SILVA, C. T. 2003. Computing and rendering point set surfaces. *IEEE Trans. on Visual. and Comp. Graph.* 9, 1, 3–15.
- ALLIEZ, P., COHEN-STEINER, D., DEVILLERS, O., LÉVY, B., AND DESBRUN, M. 2003. Anisotropic polygonal remeshing. In *Proc. SIGGRAPH '03*. 485–493.
- AMENTA, N., BERN, M., AND KAMVYSSELS, M. 1998. A new voronoi-based surface reconstruction algorithm. In *Proc. SIGGRAPH '98*. 415–422.
- AMENTA, N., CHOI, S., AND KOLLURI, R. 2001. The power crust. In *Proc. 6th ACM Symp. on Solid Modeling*. 249–260.
- BERNADINI, F., MITTLEMAN, J., RUSHMEIER, H., SILVA, C., AND TAUBIN, G. 1999. The ball-pivoting algorithm for surface reconstruction. *IEEE Trans. on Visualization and Computer Graphics* 5, 349–359.
- Blitz++ 2003. *Blitz++ Numerical Library*. <http://www.oonumerics.org/blitz/>.
- BOISSONNAT, J.-D. AND CAZALS, F. 2000. Smooth surface reconstruction via natural neighbour interpolation of distance functions. In *Proc. 16th ACM Symp. on Computational Geometry*. 223–232.
- BUSS, S. R. AND FILLMORE, J. P. 2001. Spherical averages and applications to spherical splines and interpolation. *ACM Trans. on Graphics* 20, 2, 95–126.
- CAZALS, F. AND POUGET, M. 2003. Estimating differential quantities using polynomial fitting of osculating jets. In *Proc. Eurographics Symp. on Comput. Geom.* 177–187.
- CHAVEL, I. 1997. *Riemannian Geometry: a modern introduction*. Cambridge University Press, Cambridge, UK.
- COHEN, L. D. 2001. Multiple contour finding and perceptual grouping using minimal paths. *Journal of Math. Imaging and Vision* 14, 3, 225–236.
- ACM Transactions on Graphics, Vol. V, No. N, Month 20YY.

- CURLESS, B. AND LEVOY, M. 1996. A volumetric method for building complex models from range images. In *Proc. SIGGRAPH '96*. 303–312.
- DYN, N. AND LEVIN, D. 2002. Subdivision schemes in geometric modelling. *Acta Numerica* 12, 73–144.
- EDELSBRUNNER, H. AND MÜCKE, E. P. 1994. Three-dimensional alpha shapes. *ACM Trans. on Graphics* 13, 43–72.
- FLEISHMAN, S., COHEN-OR, D., ALEXA, M., AND SILVA, C. T. 2003. Progressive point set surfaces. *ACM Trans. on Comp. Graph.* 22, 4, 997–1011.
- FLETCHER, P. T., LU, C., AND JOSHI, S. 2003. Statistics of shape via principal geodesic analysis on lie groups. In *Proc. IEEE Conf. on Comp. Vision and Pattern Rec.* 95–101.
- FLOATER, M. S. AND HORMANN, K. 2004. Surface parameterization: a tutorial and survey. In *Multiresolution in Geometric Modelling, to appear*. <http://vcg.isti.cnr.it/~hormann/pub-topic.html>.
- FLOATER, M. S. AND REIMERS, M. 2001. Meshless parameterization and surface reconstruction. *Computer Aided Geometric Design* 18, 2, 77–92.
- GIESEN, J. AND WAGNER, U. 2003. Shape dimension and intrinsic metric from samples of manifolds with high co-dimension. In *Proc. 19th Symp. on Computat. Geom.* 329–337.
- GOTSMAN, C., GUMHOLD, S., AND KOBELT, L. 2002. Simplification and compression of 3d meshes. In *A. Iske, E. Quak and M. S. Floater (eds.), Tutorials on Multiresolution in Geometric Modelling*. 319–362.
- GROSSMAN, J. P. AND DALLY, W. J. 1998. Point sample rendering. In *9th Eurographics Workshop on Rendering*. 181–192.
- GUENNEBAUD, G., BARTHE, L., AND PAULIN, M. 2004. Real-time point cloud refinement. In *Proc. Eurographics Symp. on Point-Based Graphics*. 41–49.
- HELMSEN, J., PUCKETT, E. G., COLLELA, P., AND DORR, M. 1996. Two new methods for simulating photolithography development in 3d. In *Proc. SPIE Microlithography IX*. 253–261.
- JOST, J. 1994. Equilibrium maps between metric spaces. *Calc. Var.* 2, 173–204.
- KALAIHAH, A. AND VARSHNEY, A. 2001. Differential point rendering. In *12th Eurographics Workshop on Rendering*. 139–150.
- KARCHER, H. 1977. Riemannian center of mass and mollifier smoothing. *Comm. Pure Appl. Math.* 30, 5, 509–541.
- KENDALL, W. S. 1990. Probability, convexity and harmonic maps with small image i: Uniqueness and fine existence. *Proc. London Math. Soc.* 61, 2, 371–406.
- KUNZE, R., WOLTER, F.-E., AND RAUSCH, T. 1997. Geodesic voronoi diagrams on parametric surfaces. In *Proc. Comp. Graph. Int. '97*. 230–238.
- LANE, M. AND RIESENFELD, R. F. 1980. A theoretical development for the computer generation of piecewise polynomial surfaces. *IEEE Trans. Pattern Anal. and Mach. Intell.* 2, 35–46.
- LEIBON, G. AND LETSCHER, D. 2000. Delaunay triangulations and voronoi diagrams for riemannian manifolds. In *Proc. 16th ACM Symp. on Computat. Geom.* 341–349.
- LEVOY, M. AND WHITTED, T. 1985. The use of points as display primitive. Tech. Rep. 85022, UNC-Chapel Hill Computer Science Dept., Chapel Hill, North Carolina, USA. <http://graphics.stanford.edu/papers/points/point-with-scanned-figs.pdf>.
- LINSEN, L. 2001. Point cloud representation. Tech. Rep. 2001-3, Computer Science Dept., Universität Karlsruhe, Germany. <http://www.ubka.uni-karlsruhe.de/vvv/ira/2001/3/3.pdf>.
- LOOP, C. 1987. Smooth surface subdivision based on triangles. M.S. thesis, University of Utah, Dep. of Math.
- MÉMOLI, F. AND SAPIRO, G. 2001. Fast computation of weighted distance functions and geodesics on implicit hyper-surfaces. *Journal of Comput. Phys.* 173, 1, 764–795.
- MÉMOLI, F. AND SAPIRO, G. 2003. Distance functions and geodesics on point clouds. Tech. Rep. 1902, Institute for Mathematics and its Applications, University of Minnesota, Minneapolis, USA. <http://www.ima.umn.edu/preprints/dec2002/1902.pdf>.
- MITRA, N. J. AND NGUYEN, A. 2003. Estimating surface normals in noisy point cloud data. In *Proc. 19th Conf. on Comput. Geom.* 322–328.

- MOENNING, C. AND DODGSON, N. A. 2004. Intrinsic point cloud simplification. In *Proc. 14th GraphiCon '04*. <http://www.cl.cam.ac.uk/users/cm230/docs/pdfs/GraphiCon04.pdf>.
- MOENNING, C., MÉMOLI, F., SAPIRO, G., DYN, N., AND DODGSON, N. A. 2004. Meshless geometric subdivision. Tech. Rep. # 1977, Institute for Mathematics and its Applications, University of Minnesota, Minneapolis, USA. <http://www.ima.umn.edu/preprints/apr2004/1977.pdf>.
- OSWALD, P. AND SCHRÖDER, P. 2003. Composite primal/dual sqrt(3)-subdivision schemes. *Computer Aided Geometric Design* 20, 3, 135–164.
- PAULY, M., KEISER, R., KOBBELT, L. P., AND GROSS, M. 2003. Shape modeling with point-sampled geometry. *Proc. SIGGRAPH '03*, 641–650.
- PAULY, M., KOBBELT, L. P., AND GROSS, M. 2002. Multiresolution modeling of point-sampled geometry. Tech. Rep. 378, Computer Graphics Lab, ETH Zürich, Switzerland. <http://graphics.stanford.edu/~mapauly/Pdfs/MultiresModeling.pdf>.
- PEYRÉ, G. AND COHEN, L. 2003. Geodesic re-meshing using front propagation. In *Proc. 2nd IEEE VLSM Workshop*.
- PFISTER, H., ZWICKER, M., VAN BAAR, J., AND GROSS, M. 2000. Surfels: Surface elements as rendering primitives. In *Proc. SIGGRAPH '00*. 335–342.
- RUSINKIEWICZ, S. AND LEVOY, M. 2000. Qsplat: a multiresolution point rendering system for large meshes. In *Proc. SIGGRAPH '00*. 343–352.
- SETHIAN, J. A. 1996. Theory, algorithms, and applications of level set methods for propagating interfaces. *Acta Numerica* 5, 309–395.
- SETHIAN, J. A. 1999. *Level Set Methods and Fast Marching Methods*, 2nd ed. Cambridge University Press, Cambridge, UK.
- SIBSON, R. 1980. A vector identity for the dirichlet tessellation. In *Math. Proc. of the Cambridge Philos. Soc.* 151–155.
- SIFRI, O., SHEFFER, A., AND GOTSMAN, C. 2003. Geodesic-based surface remeshing. In *Proc. 12th Int. Meshing Roundtable*. 189–199.
- TENENBAUM, J. B., DE SILVA, V., AND LANGFORD, J. C. 2000. A global geometric framework for nonlinear dimensionality reduction. *Science*, 2319–2323.
- TSITSIKLIS, J. N. 1995. Efficient algorithms for globally optimal trajectories. *IEEE Trans. on Automatic Control* 40, 1528–1538.
- WALLNER, J. AND DYN, N. 2003. Smoothness of subdivision schemes by proximity. Tech. Rep. 112, Tech. Univ. Wien, Austria.
- WALLNER, J. AND DYN, N. 2004. Smoothness of subdivision schemes in manifolds. Tech. Rep. 125, Tech. Univ. Wien, Austria. <http://www.geometrie.tuwien.ac.at/wallner/sbd2.pdf>.
- WALLNER, J. AND DYN, N. 2005. Convergence and c1 analysis of subdivision schemes on manifolds by proximity. *Comp. Aided Geom. Design* 22. To appear. <http://www.geometrie.tuwien.ac.at/wallner/sbd1.pdf>.
- WOOD, Z., HOPPE, H., DESBRUN, M., AND SCHRÖDER, P. 2004. Removing excess topology from isosurfaces. *ACM Trans. on Graphics* 23, 2, 190–208.
- ZORIN, D. AND SCHRÖDER, P. 2000. Subdivision for modeling and animation. *SIGGRAPH '00 course notes*.
- ZWICKER, M., PAULY, M., KNOLL, M., AND GROSS, M. 2002. Pointshop3d: An interactive system for point-based surface editing. *Proc. SIGGRAPH '02*, 322–329.
- ZWICKER, M., PFISTER, H., VAN BAAR, J., AND GROSS, M. 2001. Surface splatting. *Proc. SIGGRAPH '01*, 371–378.
- ZWICKER, M., PFISTER, H., VAN BAAR, J., AND GROSS, M. 2002. Ewa splatting. *IEEE Trans. on Visualization and Computer Graphics*, 223–238.




Fig. 5. Relatively regularly uniformly distributed base point set, P_0 , of 2144 points acquired from the unit sphere (right). The triangular base mesh generated from P_0 for the support of Loop subdivision is shown on the left. A flatly shaded rendering of the reconstructed surface is shown in the centre.




Fig. 6. Results after one iteration of Loop (left) and meshless subdivision (right); $|P_1| = 8570$ points in both cases. Flatly shaded renderings of the reconstructed surfaces are shown next to each point set.




Fig. 7. Results after the second iteration of Loop (left; 34275 points) and meshless subdivision (right; 36442 points). The corresponding reconstructed surfaces are given next to each point set. The detail views indicate how the slight irregularities in the distribution of P_0 lead to slightly more pronounced local irregularities in the case of meshless subdivision due to the use of uniform weighting. These slightly more pronounced irregularities do not have any noticeable effect on the smoothness of the surface.

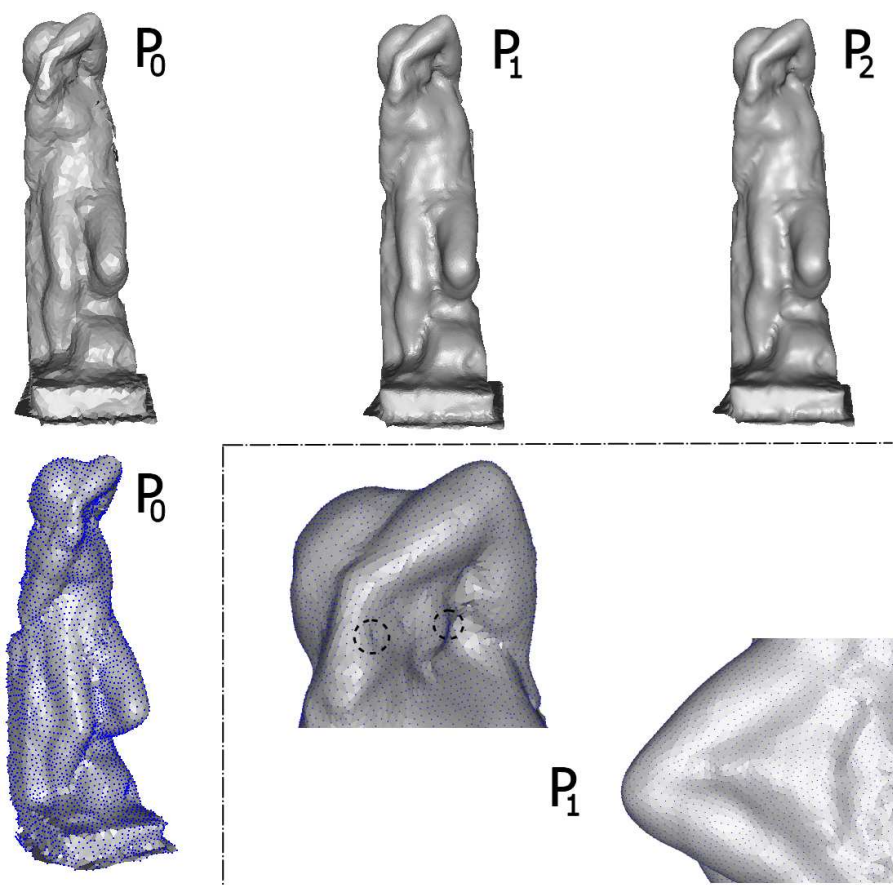


Fig. 8. Flatly shaded renderings of the surfaces reconstructed from point sets $|P_0| = 10088$ (top and bottom left), $|P_1| = 39888$ (centre and bottom right) and $|P_2| = 208010$. The smoothing effect of meshless subdivision is clearly visible. As illustrated by the detail front and side views, the subdivided point sets are distributed relatively regularly uniformly. Instances of local irregularities (encircled in black) are caused by discretisation errors during discrete intrinsic Voronoi diagram computation.




Fig. 9. Flatly shaded renderings of the surfaces reconstructed from the non-simplified Youthful point set (left) and P_2 of Figure 8 illustrating how meshless subdivision of a base point set resulting from point cloud simplification may be used to obtain a smoother, more regular and more compact representation of the original data set of 1728305 points.




Fig. 10. On the left, the flatly shaded renderings of surfaces reconstructed from the 50% decimated screwdriver (top) and rocker arm (bottom) CAD data sets (P_0) are shown. The reconstructed surfaces obtained from two iterations of meshless subdivision of these point sets are presented in the centre (P_2). The surfaces reconstructed from the non-simplified, non-subdivided CAD data sets are given on the right. The smoothing effect of meshless subdivision is again clearly visible.




Fig. 11. Point set detail views and smoothly shaded renderings of the Isis data set (left), P_0 , $|P_0| = 187644$, and the subdivided point set P_1 (right), $|P_1| = 760162$, resulting from one iteration of meshless subdivision. Due to the high density of P_0 , the difference in location between the (uniformly weighted) extrinsic and the corresponding intrinsic centroid was found to be negligible and extrinsic means were used throughout.




Fig. 12. Wave propagation for the computation of a discrete geodesic Voronoi diagram. By simultaneously propagating waves for geodesic distance mapping purposes from the (red) generator points outwards (left), an intrinsic Voronoi partitioning of the triangulated surface is obtained (right).