

Finite Elements on Point Based Surfaces

U. Clarenz¹, M. Rumpf¹, A. Telea²

¹ Institut für Mathematik, Duisburg University, Germany
{clarenz|rumpf}@math-uni-duisburg.de

² Department of Mathematics and Computer Science, Eindhoven University of Technology, the Netherlands
alex@win.tue.nl

Abstract

We present a framework for processing point-based surfaces via partial differential equations (PDEs). Our framework efficiently and effectively brings well-known PDE-based processing techniques to the field of point-based surfaces. Our method is based on the construction of local tangent planes and a local Delaunay triangulation of adjacent points projected onto this plane. The definition of tangent spaces relies on moment-based computation with proven scaling and stability properties. Once local couplings are obtained, we are able to easily assemble PDE-specific mass and stiffness matrices and solve corresponding linear systems by standard iterative solvers. We demonstrate our framework by different classes of PDE-based surface processing applications, such as texture synthesis and processing, geometric fairing, and segmentation.

1. Introduction

Surface processing tools and techniques are widespread in computer graphics, animation, medical imaging, computer aided modeling, and computer vision. Many surface processing operations can be described via partial differential equations (PDEs). Using PDEs to implement surface processing has a long history and several advantages, as compared to other more algorithmic surface processing techniques. First, PDEs describe concisely and naturally a large spectrum of transformations, such as deformations, smoothing, or denoising. Secondly, PDE-based approaches come with a solid mathematical basis that provides quantitative and qualitative results about the way they alter a given surface. Finally, many efficient and exact methods for PDE discretization and solving are readily available. Among the latter, the by far most used approach is the combination of finite element (FE) discretization and iterative numerical methods, which naturally matches the triangular mesh models ubiquitous in computer graphics.

Recently, point based representations have been proposed as an alternative to triangles for 3D surfaces, with a number of advantages. No 'mesh', or connectivity information has to be stored explicitly. This allows a simple and compact representation, ideal for fast rendering and editing. When combined with advanced rendering techniques such as splat-

ting [ZPvBG01, Pau03], point based surfaces outperform triangle meshes in terms of rendering quality and data storage flexibility.

Processing point-based surfaces via PDEs should add the modeling power of PDE representation to the flexibility of the point based model. However, defining and solving PDEs using finite elements on point based surfaces is not straightforward. The main problem here is that point-based surfaces are mesh-less, so there is no direct way to define the finite elements underlying the PDE space. We shall not consider the option of building a global mesh from the point set [HDD*92, LP01] here, as it undermines the fundamental philosophy and advantages of point based models. Moreover, global point cloud triangulations and PDEs on triangle meshes respectively have been already extensively treated.

Instead, we propose an alternative approach for finite element based PDEs on point surfaces. We proceed by constructing a number of local FE matrices that represent the surface properties over small point neighborhoods. These matrices are next assembled in a single matrix which allows PDE discretization and solving on the complete surface. We illustrate our approach by a number of well-known surface processing PDEs including anisotropic diffusion, texture synthesis, and surface fairing.

We first review the basics of point based representations

and point cloud triangulations (Sec. 2), some basic PDE problems on surfaces (Sec. 3), and the general finite element approach on triangular surfaces (Sec. 4). Next, we detail the difficulties of treating PDEs on point clouds (Sec. 5). Section 6 presents our construction of tangent spaces and local meshes. We use this basis to build finite elements on point surfaces, by assembling local and global FE matrices (Section 7). Using such matrices, we solve several PDEs on point surfaces, leading to segmentation and texture processing (Sec. 8.1), texture synthesis (Sec. 8.2), and surface fairing (Sec. 8.4). Section 10 concludes the paper.

2. Related Work

In the last years, a large number of papers related to point-based surfaces has emerged. We briefly outline those related to our work, without attempting a complete overview.

Point set methods have two main components: approximating a usually smooth surface from the point set, followed by rendering the approximation. Approximating surfaces from points can be done by many techniques. These mainly differ in the assumptions laid on the point set and the smoothness model. A quite simple approximation replaces the point set with a triangulated surface model, or triangle mesh. Several efficient triangulation methods for point clouds exist [HDD*92, GKS00, Boi84, LP01]. In many cases, such techniques can be seen as producing piecewise C^1 approximations of the point cloud. Although efficient and reasonably simple to implement, such techniques may produce surfaces lacking the desired smoothness, as described in [ABCO*01, AA03]. To alleviate this, smoothing operations can be applied to the obtained triangle mesh, such as iterative Laplacian smoothing [Tau95], curvature flow fairing [DMSB99], or discrete variational fairing [Kob97]. Alternatively, increased smoothness can be obtained by using higher local approximations, such as piecewise polynomials [XWH*03] or radial basis functions [Pau03, ZPvBG01].

Rendering point surfaces follows the surface approximation assumptions. Different rendering primitives may be used, ranging from simple flat shaded planar discs, such as used by the QSplat system [RL00], up to elliptically weighted Gaussian splats [ZPvBG01] and differential points [VK01]. More complex primitives encode more information on the vicinity of a rendered point, such as surface geometry, at additional rendering expense. Simple primitives render faster but may have limited quality, especially for nonuniformly sampled surfaces.

3. PDEs on Surfaces

We start by defining basic notions of differential operators on surfaces. To this aim, we will use the concept of tangential gradients on embedded surfaces \mathcal{M} in \mathbb{R}^3 . The tangential

gradient $\nabla_{\mathcal{M}}u$ for a function u defined in a neighborhood of \mathcal{M} , is given by

$$\nabla_{\mathcal{M}}u = \nabla_{\mathbb{R}^3}u - (n \cdot \nabla_{\mathbb{R}^3}u)n,$$

where $n : \mathcal{M} \rightarrow S^2$ is the normal mapping. The gradient in the ambient space is thus projected onto the surface's tangential space. The result coincides with the classical geometric gradient $\nabla_{\mathcal{M}}u$ for embedded surfaces. In the following, we denote the scalar product of two vectors $a, b \in \mathbb{R}^m$ by $a \cdot b$. The components of $\nabla_{\mathcal{M}}u$ are denoted by $\nabla_i u$.

For a vector field $v : \mathcal{M} \rightarrow \mathbb{R}^3$ of components $v = (v_1, v_2, v_3)$, we define its divergence using the components of the tangential gradient, i.e.,

$$\operatorname{div}_{\mathcal{M}}v := \nabla_i v_i$$

Here and in the following, we use the Einstein summation convention. In this notation, the Laplace operator on surfaces is given by:

$$\Delta_{\mathcal{M}}u = \nabla_i \nabla_i u.$$

For surfaces in \mathbb{R}^3 , curvature may be expressed by the shape operator S which is – using tangential gradients – given by the 3×3 -matrix $S = D_{\mathcal{M}}n = \nabla_{\mathcal{M}}n$. This matrix operates as a symmetric endomorphism on the tangent spaces and may be diagonalized by the principal curvatures κ_1 and κ_2 . The classical mean curvature is then given by $h = \operatorname{tr} \nabla_{\mathcal{M}}n = \kappa_1 + \kappa_2$ and we have the well known identity $\Delta_{\mathcal{M}}x = -hn$ where x is the position vector of the surface.

We can now formulate, on surfaces, the same type of problems as on Euclidean domains. A basic problem type is the boundary value problem:

For a subset $\Omega \subset \mathcal{M}$, a diffusion tensor A , and a right hand side f we ask for a function $u : \bar{\Omega} \rightarrow \mathbb{R}$ which solves

$$-\operatorname{div}_{\mathcal{M}}(A \nabla_{\mathcal{M}}u) = f \tag{1}$$

on \mathcal{M} and reaches values $u = u^{\partial}$ on the boundary $\partial\Omega$ of \mathcal{M} .

An example application for such an elliptic problem is the inpainting of a locally destroyed coloring of the surface. As a second problem type we consider reaction diffusion problems on \mathcal{M} :

Find a function $u : \mathbb{R}_0^+ \times \mathcal{M} \rightarrow \mathbb{R}$ with $u(0, \cdot) = u_0$ for some function u_0 , such that

$$\partial_t u - \operatorname{div}_{\mathcal{M}}(A \nabla_{\mathcal{M}}u) = f, \tag{2}$$

where A is the diffusion tensor and f the source term.

Here u can be a scalar or a vector valued quantity, and $A = A[u]$ and $f = f[u]$ may depend nonlinearly on u . An example of such a problem is segmentation via diffusion of a “marker color” which stops at the surface's feature lines. Another example is the smoothing of a gray scale surface texture, which leads to a scalar diffusion problem, where A is the nonlinear anisotropic diffusivity. In addition, we consider reaction diffusion systems for texture synthesis, such

as introduced by Turk [Tur91]. Given several components, or species, f encodes the coupling of the species' concentrations via a particular reaction. Reaction diffusion systems are a simple and effective way for synthesizing repetitive textures on surfaces. As the next application, if we consider $u = x$, where x is the position vector of the surface itself, we obtain a curvature motion problem. Indeed, if $A = 1$ and $f = 0$, it is well-known that Eqn. 2 is equivalent to mean curvature motion, i. e. $\partial_t x = -h(x)n(x)$. Using a nonlinear diffusivity, we obtain a feature preserving fairing method. Different PDEs can be modelled similarly, if desired.

4. Reviewing Finite Elements on Triangular Surfaces

Before we consider solving these PDEs on point based surfaces, we briefly discuss the by now classical finite element discretization scheme, frequently used for the numerical treatment of PDEs on discrete triangular surfaces. Let \mathcal{M}_h denote an approximating triangulation of \mathcal{M} , where h indicates the corresponding grid size. We define the space of piecewise affine, discrete functions \mathcal{V}_h and consider the fully discrete weak formulation of the elliptic problem (1) in \mathcal{V}_h . We ask for a discrete function $U \in \mathcal{V}_h$ such that

$$\int_{\mathcal{M}_h} A \nabla_{\mathcal{M}_h} U \cdot \nabla_{\mathcal{M}_h} \Phi = \int_{\mathcal{M}_h} f \Phi \quad (3)$$

for all discrete test functions $\Phi \in \mathcal{V}_h$. Functions U in \mathcal{V}_h can be represented by nodal vectors \bar{U} in \mathbb{R}^n , where $\bar{U} = (U_j)_{j=1, \dots, n}$ is a vector with components U_j . Let $\{\Phi_i\}_{i=1, \dots, n}$ be the usual nodal basis of \mathcal{V}_h with $\Phi_i(x_j) = \delta_{ij}$ for all vertices x_j of the grid \mathcal{M}_h . We can express a discrete function U in terms of its nodal values $U_j = U(x_j)$ and get $U = U_j \Phi_j$. The discrete problem can be expressed in matrix vector notation by

$$L\bar{U} = M\bar{F},$$

where the mass matrix M and the stiffness matrix L are given by

$$M = \left(\int_{\mathcal{M}_h} \Phi_i \Phi_j \right)_{i,j},$$

$$L = \left(\int_{\mathcal{M}_h} A \nabla \Phi_i \cdot \nabla \Phi_j \right)_{i,j},$$

and $\bar{F} = (f(x_i))_i$ is the vector of nodal values $f(x_i)$ at vertices of the triangulation x_i . Hence, the discrete elliptic operator in matrix form turns out to be $M^{-1}L$. In case of the parabolic problem (2), we consider a time discretization with time step τ and have to find a sequence $\{U^k\}_{k=1, \dots, n} \subset V^h$ of approximations to the continuous solution ($U^k(\cdot) \approx u(\tau k, \cdot)$) such that

$$\int_{\mathcal{M}_h} \frac{U^{k+1} - U^k}{\tau} \Phi + \int_{\mathcal{M}_h} A[U^k] \nabla_{\mathcal{M}_h} U^{k+1} \cdot \nabla_{\mathcal{M}_h} \Phi - f[U^k] \Phi = 0 \quad (4)$$

for all $\Phi \in V^h$. Note that $\mathcal{M}_h = \mathcal{M}_h(k\tau)$ if the surface itself is evolving as e.g. in case of surface fairing applications. We obtain, for each time step of our problem, the system of linear equations

$$(M + \tau L[U^k])\bar{U}^{k+1} = M(\bar{U}^k + \tau \bar{F}[U^k]),$$

where the stiffness matrix depends on the discrete solution, i.e. $L[U] := (\int_{\mathcal{M}_h} A[U] \nabla \Phi_i \cdot \nabla \Phi_j)_{i,j}$.

Algorithmically, the integral expressions in (4) are split up into a sum over local contributions on triangles. The matrices and right hand side vector are computed as follows. We initialize $L = 0$ and next do a traversal of all triangles $T \in \mathcal{M}_h$. On each T with nodes P^0, P^1, P^2 , a corresponding local matrix $(l_{ij}(T))_{i,j}$ is computed first, corresponding to all pairings of local nodal basis functions, and next added to the matching locations in the global matrix L . For every pair i, j we update $L_{\alpha(i), \alpha(j)} = L_{\alpha(i), \alpha(j)} + l_{ij}(T)$. Here $\alpha(i)$ is the global index of the node with local index i . We proceed similarly with the mass matrix M .

5. Differences for Point Based Surfaces

One faces several difficulties when aiming to transfer the PDE discretization approach outlined above to point cloud surfaces. Conceptually, such surfaces are not described in terms of a two dimensional set in \mathbb{R}^3 . In particular, there is no global mesh available and, as outlined in Sec. 1, it would conflict with the general paradigm of point based modeling to replace the usually huge unstructured point set by a standard mesh. In particular, one would stop working on the actual, usually noisy, arbitrarily sampled data, along with their statistical properties, and completely replace them by a mesh having other properties. The standard method for handling point surfaces is to extract a local approximate tangent space on the point cloud [Pau03, ABCO*01, XWH*03, LP01]. This tangent space is generally used just for computing point normals used e.g. in shading. We will use this idea to obtain proper discrete counterparts of the differential operators $\text{div}_{\mathcal{M}}$ and $\nabla_{\mathcal{M}}$ and a metric for the discrete integration over the surface. Hence, we proceed by constructing a local tangent space and consider the local projection of the point set onto this tangent space. We are then able to define stable coupling quantities between neighbor points, using a *strictly local* Delaunay meshing. The mentioned coupling quantities to be defined will turn out to be suitable discretizations of the off-diagonal entries in the global stiffness matrix we aim to recover. The local tangent spaces of different points usually do not coincide, which induces a loss of symmetry in our matrix. To remove this problem, we finalize the matrix construction by applying a suitable symmetrization. Finally, the diagonal entries of the stiffness matrix can be defined based on a requested invariance property. Indeed, the continuous differential operator $\text{div}_{\mathcal{M}}(A \nabla_{\mathcal{M}} \cdot)$ applied to a constant function u should vanish. Hence, we require that $L\bar{U} = 0$ for $\bar{U} = (1, \dots, 1)$. We proceed similarly for the

mass matrix and the right hand side of the considered PDE. The complete approach is detailed in the following sections, leading to a stable and consistent approximation scheme for general PDEs.

6. Tangent Spaces and Local Meshes

We proceed by defining, for every point x in the considered cloud, a *tangent space*. This space attempts to approximate the points in a small neighborhood N of x . The size of N should be chosen such that a) it contains enough points for stable computation of a tangent space and b) the radius of N is smaller than the feature size we want to be visible in the approximation. For a), N can be efficiently computed using the k nearest neighbors of x , for given k . For b), N can be defined as the ball $B_\varepsilon(x)$ of given radius ε centered at x . In practice, combinations of the two criteria give the best results. We prescribe a minimum number of neighbors k_{min} , to enforce the first requirement. If the k_{min}^{th} closest neighbor of x is closer than the prescribed minimal feature size ε , we consider all additional nearest neighbors in $B_\varepsilon(x)$. Next, we use the *zero and first order moments* of N . Moments have several proven properties that allow us to robustly compute the tangent spaces as well as to distinguish between smooth and non-smooth surface parts, both as a function of the ball size ε . Robustness is clearly needed in the tangent plane computation. Distinguishing smooth from non-smooth surface areas is needed for our surface segmentation (Sec. 8.1) and fairing (Sec. 8.4) applications. In this section, we give the moment definitions and properties relevant for the tangent plane computation. Next, we discuss the concrete moment-based implementation of the tangent planes.

For a continuous surface \mathcal{M} , the zero moment is given by the local barycenter of \mathcal{M} with respect to a Euclidean ball $B_\varepsilon(x)$ centered at x :

$$M_\varepsilon^0(x) = M_\varepsilon^0 := \int_{B_\varepsilon \cap \mathcal{M}} x \, dx. \quad (5)$$

The first order moment is then defined as as:

$$\begin{aligned} M_\varepsilon^1(x) &:= \int_{B_\varepsilon \cap \mathcal{M}} (x - M_\varepsilon^0) \otimes (x - M_\varepsilon^0) \, dx \\ &= \int_{B_\varepsilon \cap \mathcal{M}} (x \otimes x - M_\varepsilon^0 \otimes M_\varepsilon^0) \, dx \end{aligned} \quad (6)$$

where $y \otimes z := (y_i z_j)_{i,j=1,\dots,3}$. It turns out that the first moment approximates the matrix $\Pi_{\mathcal{T}_x \mathcal{M}}$ corresponding to the projection onto the tangent space $\mathcal{T}_x \mathcal{M}$. Indeed, in smooth surface regions we have:

$$M_\varepsilon^1(x) = 2c\varepsilon^2 \Pi_{\mathcal{T}_x \mathcal{M}} + o(\varepsilon^2),$$

where c is a constant that only depends on the dimension. For a proof, we refer to [CRT04b]. This shows that the eigenvectors of the first moment define an orthonormal basis of \mathbb{R}^3 , where the surface normal belongs to the vanishing (smallest) eigenvalue. For a discussion of the non smooth case we

refer to again [CRT04b]. Let $\lambda_0 > \lambda_1 > \lambda_2$ be the eigenvalues of the 3 by 3 symmetric matrix M_ε^1 , then we consider the corresponding eigenvector e_2 as the normal on the approximate tangent plane, whereas e_1 and e_0 form a 2D coordinate system in the plane itself. Figure 1 a illustrates the above in two dimensions. Our tangent plane computation is similar to the principal component analysis used in [ABCO*01, Pau03, XWH*03]. However, as explained already, we prefer our moment-based approach as it comes with proven scaling properties with respect to the ball size ε . Concisely, ε has the role of a filter size: The tangent plane ignores features significantly smaller than ε . Once the tan-

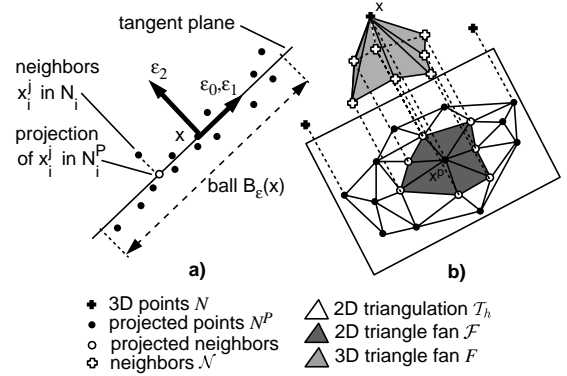


Figure 1: Tangent plane (a) and local triangulation (b)

gent plane $\{x \in \mathbb{R}^3 \mid e_2 \cdot (x - x_i) = 0\}$ is defined for a point x_i , we project all neighbors x_i^j in the neighbor set N_i onto it, yielding the projected point set $N_i^p = \{X_i^j\}_{j=0,\dots,k}$ in the 2D coordinate system (e_0, e_1) (Fig. 1 a). Here $X_i^j \in \mathbb{R}^2$ are the projections of $x_i^j \in \mathbb{R}^3$ onto the above 2D coordinate system. To simplify notation, we incorporate x_i as x_i^0 in the set of neighbors N_i^p . Next, we compute the Delaunay triangulation of the points N_i^p . This yields a triangle mesh \mathcal{T}_i in the tangent plane. The triangulation is a strictly 2D process, confined to the frame (e_0, e_1) . From the triangulation, we select the triangle fan $\mathcal{F}_i^p = \{T_i^p\}_i$ of projected triangles T_i^p around the projected seed point X_i^0 (Fig. 1 b). Finally, we define the neighbor set N_i^p of x_i as being the points x whose projections X_i^j are used by the triangles in the fan \mathcal{F}_i^p . By N_i we denote the corresponding set of 3D points before projection on the tangent plane. Note that the points N_i define a local 3D triangle fan \mathcal{F}_i of the point set whose projection on the tangent plane is exactly the triangle fan \mathcal{F}_i^p defined by the point set N_i^p .

However, the above scheme has a problem. The Delaunay triangulation may produce triangles with too small angles which, when assembled in matrices discretizing PDEs, can cause inaccuracy and instability problems when solving these PDEs (Sec. 7). These problems are well known from discretizing PDEs on ill-conforming meshes. We pre-




Figure 2: Point set (left). Three selected points with neighbor sets N_i and \mathcal{N}_i (right). Major eigenvectors along edges (bottom)

vent this as follows. If an angle smaller than a user given α_{min} , set in practice to around 25 degrees, appears in the triangulation, we remove one of its points from N^p and re-triangulate the remainder. The process is repeated until no ill-shaped triangles are created. In practice, this causes no visible slow-down. We tested a large number of noisy point sets of 30000 up to a million points. The worst case encountered contained a few tens of such triangles per point set, which were successfully removed in three re-triangulation passes. Even though these cases are rare, their removal is essential to ensure robust convergence of PDE discretization schemes.

Our tangent plane construction has several desirable properties. First, the moment-based computation is a noise-robust way to define the tangent plane. Larger neighborhoods N act as stronger noise filters. It is important to note that our approach is *not* the same as producing a smoothed mesh. Indeed, the neighbors \mathcal{N}_i of a point x_i are defined to be only the *immediate* neighbors of x_i in the Delaunay triangulation of the projected neighborhood N_i^p . As N_i increases due to increase of ϵ or k_{min} , the set \mathcal{N}_i practically stays of constant size. In practice, \mathcal{N}_i contains the average number of points in a conforming Delaunay triangulation, i.e. 4 up to 8..10 points, whereas the average size of N_i , for the point sets we worked with, ranged between 30 and 100 points. Secondly, the computation of (e_0, e_1, e_2) does *not* need to be very accurate. We use them just as a means of finding the neighbor set

\mathcal{N}_i out of the points in the point set. Even reasonably large orientation variations of e_0, e_1 cause no change in \mathcal{N}_i , as the Delaunay triangulation of $X_{i,j}^j$ uses the same closest points. We do not produce new points, but just couple the existing ones. Finally, the removal of ill-shaped triangles ensures that our stiffness matrices (Sec. 7 and further) are well conditioned, a property which is not directly enforced by classical finite elements on arbitrary triangle meshes.

Our method of triangulating tangent plane projections for the k closest neighbors resembles the local triangulation proposed by Linsen et al. [LP01]. However, we consider large k values (30 up to 100). Linsen et al. use $k = 6$, which should lead to considerably less stable tangent planes. Our triangulation quality check enforces minimal angles α_{min} , whereas [LP01] does not guarantee this.

Figure 2 (left) shows the bunny model in which we chose three points on the left ear. The points, their neighborhoods N for $k = 60$, and lines to their neighbors \mathcal{N}_i are shown in red, yellow, respectively green in the detail image Fig. 2 (right). Note that, although each N_i is large (60 points), the corresponding \mathcal{N}_i has 6 neighbors. Moreover, the neighbor set \mathcal{N}_i stays the same for $k \in [20..60]$, which outlines the stability of our method.

7. Assembling the Finite Element Matrices

We have shown how to construct, for every x_i in a point set, the neighbor set \mathcal{N}_i . In this section, we show how to build the matrices needed for solving PDEs on point surfaces. Given the local 3D triangle fan $\mathcal{F}_i = \{T_l\}_l$, we define the preliminary matrix entry \tilde{L}_{ij} as

$$\tilde{L}_{ij} = \sum_l A \nabla_{T_l} \Phi_i \cdot \nabla_{T_l} \Phi_j |T_l| \quad (7)$$

where Φ_j are the affine linear basis functions on the triangles of F defined by $\Phi_k(x_i^j) = \delta_{kj}$ for all k and j . A is the application-dependent discrete diffusivity term. Here the gradient ∇_{T_l} is the gradient on the affine triangle T_l .

We could alternatively define \tilde{L}_{ij} by integrating *in the tangent plane* only, i.e., on the triangles of \mathcal{F}_i^p , instead of \mathcal{F}_i . The right choice is application dependent. If we process (e.g. denoise) the surface itself, it is incorrect to use projected quantities, as they do not take into account the spatial orientation of the points. Indeed, this would couple points in flat regions as strongly as points in curved regions. As already outlined, if we compute on the 3D triangle fan, we use the tangent planes just as a help for the triangulation and obtaining the neighbor relations \mathcal{N}^p , and perform all other computations in 3-space. When processing a fixed and very noisy point cloud, the smoothing induced by the tangent space construction may be desirable. In that case, one would replace the 3D triangle fan \mathcal{F}_i by its projection \mathcal{F}_i^p .

The quantity \tilde{L}_{ij} describes the coupling of point i with all its neighbors j , from the point of view of i . Clearly, \tilde{L}_{ji} is

not necessarily equal to \tilde{L}_{ij} , as the neighbor computations of i and j are strictly speaking independent (Sec. 6). Moreover, we must still define a point's self-coupling, i.e. the matrix's diagonal entries. To produce a complete, symmetric 'stiffness' matrix L , we now define

$$L_{ij} = \frac{1}{2}(\tilde{L}_{ij} + \tilde{L}_{ji}) \quad (8)$$

for $i \neq j$ and for the diagonal entries

$$L_{ii} = - \sum_{x_j \in \mathcal{N}(x_i)} L_{ij}. \quad (9)$$

The latter ensures - as already mentioned in Sec. 4 - the desirable property that $L(1, \dots, 1)^T = 0$. The matrix L has now the same properties as the classical stiffness matrix on a triangular mesh. However, L is *not* produced via a global triangulation, but via our local, on-the-fly triangulation.

Finally, for the mass matrix M , we consider a diagonal, lumped mass matrix, and set

$$M_{ii} = \frac{1}{3} \sum_l |T_l|. \quad (10)$$

Here, as for the stiffness matrix, we can either consider triangles in the 3D fan \mathcal{F}_i or alternatively their 2D projections in the projected fan \mathcal{F}_i^p .

8. Applications

8.1. Surface Segmentation

Nonlinear diffusion methods are well known in image processing applications [Kim97, PR99]. In these applications, one solves Equation 2, where u is the scalar gray value or vector-valued image color. Time plays the role of a scaling parameter: $u(t=0)$ is the initial image, and $\{u(t)\}_{t>0}$ is a family of progressively smoothed images. Appropriate choices for the diffusivity A and source term f yield different diffusion types. For example, $A = 1$ and $f = 0$ gives the well known heat equation with its isotropic smoothing effect. A better choice for image processing is to set A small in areas where we want to keep image details and large in areas where we desire strong smoothing. Finally, we can enforce the diffusion direction to follow the feature lines.

As a more challenging application, we consider the segmentation of regions on a surface \mathcal{M} which are bounded by sharp edges. For this, we use a diffusion process where we limit diffusion across and close to edges and have it large in smooth areas. For this, we can set $A = C_\epsilon$, where

$$C_\epsilon = G \left(\frac{\|M_\epsilon^0(x) - x\| \lambda_2(M_\epsilon^1(x))}{\epsilon \lambda_0(M_\epsilon^1(x))} \right),$$

with $G(s) = (\alpha + \beta s^2)^{-1}$ with suitably chosen $\alpha, \beta > 0$. This is the surface classifier presented in [CRT04b] which is small in the vicinity of edges on the surface and almost 1 in smooth surface regions. However, the above choice for A may not stop diffusion *completely* close to edges, which




Figure 3: Seed points (a), diffused signal u for yellow seed after 5 iterations (b), and two views of the obtained segments (c,d) for a surface having 75781 points




Figure 4: Seed points (a), diffused signal for a seed after 15 iterations (b), and obtained segments (c) for a surface having 65500 points

is what we need for segmentation. To this end, we set $A = H(C_\epsilon)$ and $f = K(u)$, where

$$H(u) = \begin{cases} 0 & ; u > \gamma \\ \alpha(u - \gamma)^q & ; u \leq \gamma \end{cases},$$

$$K(u) = \begin{cases} 0 & ; u > 1 \\ \alpha(1 - u)^q & ; u \leq 1 \end{cases}.$$

for suitably chosen $0 < q \ll 1$, $\alpha > 0$, and $\gamma > 0$. In practice, a good choice is $q = 0.5$, $\alpha = 1$, and $\gamma = 0.05$. Given that C_ϵ ranges from very small close to edges to 1 in flat

areas (Sec. 6), our choice for γ , and subsequently for H , ensures that diffusion is zero close to edges and strong in smooth surface areas. We set the initial condition $u_0 = 0$ over the whole surface \mathcal{M} , except for a small hand-picked seed area within the region to be segmented, where we set $u_0 = 1$. The diffusion process stops spreading the seed intensity and stop at the surrounding edges, due to the choice of A . Furthermore, the right hand side f serves as a contrast enhancement, which pushes u to the value 1 at any position which has a positive u value. Figures 3 and 4 illustrate the segmentation process on two different point set surfaces, where we used different colors for every region and corresponding seed. The colors' saturations correspond to the diffused signal u . In Fig. 3 we show the use of $A = C_\epsilon$, which causes a very small amount of diffusion to leak out of the segmented regions. This is visible as the light red and green tints on the model's arms that come from the respective red and green segmented regions. Exact segments can be easily obtained by e.g. upper thresholding the signal u . A better choice is shown in Fig. 4 where we use the second option $A = H(C_\epsilon)$. Here, the obtained segments are clearly separated by white areas. These areas correspond to high curvature regions, where the the function H has zero values which completely block diffusion.

Computing the diffusion-based segmentation is efficient, as the matrix A needs to be assembled just once. On a Pentium IV 1.8 MHz machine, one diffusion iteration takes 0.3 seconds for the 75781 points model in Fig. 3 and 0.57 seconds for the 121723 points model in Fig. 7.

8.2. Texture Synthesis

We describe now the use of PDEs to generate textures on point surfaces, using the reaction diffusion method presented by Turk in [Tur91]. This method uses two 'chemical species' concentration functions a and b that diffuse and react, i.e., build up or annihilate each other, on a given surface. The process is described by:

$$\frac{\partial a}{\partial t} = F_a(a, b) + D_a \Delta a \quad (11)$$

$$\frac{\partial b}{\partial t} = F_b(a, b) + D_b \Delta b \quad (12)$$

where F_a and F_b are the creation rates and D_a and D_b are the diffusivities of the species a and b respectively. The system is initialized with constant a and b values biased by a small random perturbation. After several iterations, regular patterns appear (cf. Fig. 5, top row). By using a five species system, stripe-like patterns can be generated (cf. Fig. 5, bottom row). We have used exactly the same PDEs and parameter settings as the original work by Turk [Tur91].

We discuss first the synthesis of spots and stripes textures. Figure 5 shows these types after 100, 600, and 1700 iterations for the spots (top row) and 100, 300, and 600 iterations for the stripes (bottom row). Here, we visualize the




Figure 5: Top row: spot texture synthesis. Bottom row: stripes texture synthesis

concentration (a or b) result of the reaction diffusion via a blue-to-red colormap. The patterns and the number of iterations needed are practically identical with the ones produced by [Tur91]. Next, we synthesize a zebra-like pattern, by disabling the initial random perturbation and setting the initial concentration to a given value v . We extend Turk's method by forcing the zebra pattern to follow the surface edges by relating v to the moment-based classifier value. For this, we select all points where C_ϵ is closer to its minimum value than 10 percent of its range. This delivers points on and close to surface edges. We next set v to 1 on these points and 0 on the remainder and proceed with the texture synthesis. The species start diffusing from the surface edges (red regions in Fig. 6 a) into the smooth areas (blue regions in Fig. 6 a). Figure 6 a-c shows three instants of the zebra pattern formation. On a Pentium IV 1.8 GHz machine, for the bunny dataset, the spot formation took around 30 seconds, the stripes 8 seconds, and the zebra pattern 18 seconds.




Figure 6: Aligned zebra patterns after 10 iterations (a), 100 iterations (b), and 1300 iterations (c)

8.3. Inpainting Textures

Inpainting, originally an artist's work, is the process of repairing local damages in an image, or texture, by using the

image colors outside of, and close to, the damaged area to fill in, or 'paint in', the defect itself. Several inpainting methods exist which essentially try to restore the damaged area so that various properties (statistic data, gradient information) of the valid image are extrapolated in the damaged area in a natural way [CS00, BSCB00]. We demonstrate here a simple linear inpainting method which allows repairing the color texture on a damaged region \mathcal{D} of a point set surface \mathcal{M} , by extending the texture on the boundary $\partial\mathcal{D}$ of \mathcal{D} into the interior of \mathcal{D} . For this, we consider the boundary value problem (1) with $A = \mathcal{C}_\epsilon$. This diffusivity choice prefers rather independent texture expansion on both sides of an edge. Hence, we avoid texture smearing across edges. Figure 7 shows the inpainting of a model of 121723 points. First, we created several defects by painting on the model. Such defects are shown in yellow in Fig 7 c,e. The texture was set to black in the defects area. Using the anisotropic diffusivity A (low close to creases, high in flat areas) for inpainting diminishes texture smearing close to creases: The model's black hair color and facial color are kept separate (Fig. 7 c,d). The same happens with the leg's skin color and white shoe color (Fig. 7 e,f). However, the black trouser color tends to diffuse on the left leg (Fig. 7 b), as this region is flat. If desired, this can be prevented by using an anisotropy tensor A that incorporates color gradient information.

8.4. Fairing of Point Based Surfaces

The last application of our framework for PDEs on point based surfaces is surface fairing using anisotropic geometric diffusion. Here, surface geometrical noise is smoothed out, whereas features such as edges are preserved or possibly even enhanced [CDR00, DMSB99]. This is especially useful for point surfaces acquired via noisy scanning. Given an initial compact embedded manifold \mathcal{M}_0 in \mathbb{R}^3 , we compute a family of faired manifolds $\{\mathcal{M}(t)\}_{t \in \mathbb{R}_+^+}$, with corresponding coordinate mappings $x(t)$. The time t describes the fairing process and $x(t)$ are given by solving the system of anisotropic evolution equations:

$$\partial_t x - \operatorname{div}_{\mathcal{M}}(A \nabla_{\mathcal{M}} x) = 0 \quad (13)$$

We start with the initial condition $\mathcal{M}(0) = \mathcal{M}_0$. We define the tensor A such that we have strong diffusion along surface features and weak diffusion across them. As before, we use a moment-based classifier. When computing it, we also obtain a basis w^1, w^2 in the tangent plane $T_x \mathcal{M}$, defined by the major and medium eigenvectors of the first order moment (Sec. 6). In this basis, the tensor A is defined as

$$A = \begin{pmatrix} 1 & 0 \\ 0 & \mathcal{C}_\epsilon \end{pmatrix}$$

Since \mathcal{C}_ϵ is high in smooth regions and low close to edges and corners, Eqn. 13 smooths the surface by keeping the features. Due to the anisotropy A , we enforce a signal enhancement in the direction of the eigenvector w^1 . In the direction of w^2 , the diffusion is proportional with the classifier \mathcal{C}_ϵ , i.e.,




Figure 7: Texture inpainting after 1 step (a) and 30 steps (b). Details: defects (yellow) (c,e) and their inpainting (d,f)

strong in smooth areas and weak close to edges, which is exactly what we desire. We refer to [CRT04a], which describes this application, but without detailing the actual PDE discretization we consider here. Figure 8 shows several results, all obtained with a few tens of diffusion iterations. The important surface edges, such as the bunny's ear edges, body-hip contact line, the transversal femur cut, and the chiselled letters, are preserved. Small 'noise' details, such as the bunny's skin ripples, bone irregularities, and stone graininess, are removed.

9. Implementation

Several aspects are essential for an efficient implementation. One of the costliest computations in the whole process is the nearest neighbor search used for the classifier and tangent plane computation. (Sec. 6). We accelerate this search




Figure 8: Top row: initial surfaces. Bottom row: surfaces faired via diffusion

using the Kd and/or Bd trees provided by the ANN package [Mou], also used by the PointShop point rendering system [ZPKG02]. However, ANN's standard Kd and Bd tree implementations treat the (usually very numerous) points in the point cloud independently: searching the neighbors of every point implies, in a worst case, a full leaf-to-root search tree traversal. In many point sets, the points are not completely randomly distributed. Points geometrically close to each other come close to each other in the point vector too. We use this to accelerate the neighbor search, as follows. We do not try to return the *exact* k closest points, but k points contained within a small given radius ϵ from a given point. These points are kept in a cache. If the cache is empty, we fill it by executing the standard k closest neighbor search. If the cache is not empty, it contains search results for the *previous* point, so we retain those k' points closer than ϵ from the *current* point. The cache miss, i.e. the remaining, usually few, $k - k'$ points are found by the usual tree search. This acceleration pays off as function of k . Indeed, as k increases, neighborhoods of close points will largely overlap. For k equal to 20, 50, and 100 closest points, we got a speedup factor of 2.62, 3.92, respectively 5.46 as compared to the standard tree search. This speedup was consistently observed for point sets between 100000 and one million points. On our Pentium IV 1.8 GHz machine, for the point set and four k -closest-points values in Fig. 2, we need 0.4, 0.6, 1.05, and 1.83 seconds respectively for the nearest neighbor search, tangent plane, and classifier computations.

For the Delaunay triangulation (Sec. 6), we used the Triangle software [She96] which provides efficient checking and enforcing of various quality norms on the produced triangles, such as minimal angles. This is important for the conditioning of the assembled matrices (Sec. 7). We solve the linear systems given by the matrices using standard iterative techniques, such as conjugate gradient.

We built our system based on the QSplat rendering software [RL00] which uses a bounding sphere hierarchy to quickly and progressively render very large point sets. We perform all our moment, tangent plane, and PDE solving

computations on the finest hierarchy level, i.e. the real points themselves. If desired, the color, normal, and position results can be propagated upwards in the hierarchy, so that we immediately benefit from QSplat's efficient rendering. We could also perform our PDE computations on *coarser* hierarchy levels, e.g. to trade off accuracy for speed. Such an approach is taken in the multiresolution point set renderer described in [PKG03] to minimize the cost of local polynomial fitting.

10. Conclusions

The main aim of the presented framework is to carry over the surface processing capabilities of finite element PDE methods, well proven for mesh based surfaces, to point based surfaces. Our framework can be seen as a two-scale approach. On the fine scale, we build local point couplings by using Delaunay triangulations of point projections on local tangent planes. The local couplings define fine-scale finite elements. It is only on this scale that the actual interpretation of the data as a function is clear and straightforward. On the next scale, we consider the different tangent spaces of different points, and average the first-scale FE models of these points to obtain the 'global' stiffness matrix (Sec. 7). To interpret data as a function on the second scale, one can average the function values on first-scale local triangles and interpret them as function values on interpolated points, where point interpolation is done by averaging point interpolations from the fine scale. We use the local tangent planes solely as a means of computing the point couplings. Thus, our approach differs from other methods on point clouds, such as [Pau03, XWH*03, ABCO*01, LP01, AA03]. Let us note that, given different surface approximations, like any produced by the afore cited methods, we could easily extend our matrix assembly process to such surfaces, by reimplementing Eqns. (7) and (10) on this approximation.

Running our PDEs on the same surfaces represented as triangle meshes and point sets respectively, with the same parameter settings, produced virtually identical results. Let us emphasize that we avoid building a global surface representation. Our only global object is the stiffness matrix describing the PDE to solve. Assembling this sparse global matrix allows computing the point couplings only once. If desired, however, we could *completely* avoid assembling this matrix, e.g. by using iterative methods needing only one matrix row at a time, which is computed on the fly. Such approaches are well known e.g. in the field of progressive radiosity.

Our framework can be extended in several directions. First, more types of PDEs could be solved by merely adapting the matrix assembly step. Secondly, one could use the local couplings described here to build consistent global mesh representations from point clouds. Finally, multiresolution schemes on point surfaces can be built to accelerate the PDE solving to target interactive applications.

References

- [AA03] ADAMSON A., ALEXA M.: Approximating and intersecting surfaces from points. In *Proc. Eurographics Symposium on Geometry Processing* (2003), pp. 89–97.
- [ABCO*01] ALEXA M., BEHR J., COHEN-OR D., FLEISHMAN S., LEVIN D., SILVA C.: Point set surfaces. In *Proc. IEEE Visualization* (2001), pp. 21–28.
- [Boi84] BOISSONNAT J. D.: Geometric structures for three-dimensional shape representation. *ACM Trans. Graph.* 3, 4 (1984), 266–286.
- [BSCB00] BERTALMIO M., SAPIRO G., CASELLES V., BALLESTER C.: Image inpainting. In *Proc. ACM SIGGRAPH* (2000), pp. 417–424.
- [CDR00] CLARENZ U., DIEWALD U., RUMPF M.: Nonlinear anisotropic diffusion in surface processing. *Proc. IEEE Visualization* (2000), 397–405.
- [CRT04a] CLARENZ U., RUMPF M., TELEA A.: Fairing of point based surfaces. In *Proc. Comp. Graphics Intl. (CGI)* (2004). to appear.
- [CRT04b] CLARENZ U., RUMPF M., TELEA A.: Robust feature detection and local classification for surfaces based on moment analysis. *to appear in IEEE TVCG* (2004).
- [CS00] CHAN T., SHEN J.: Mathematical models for local deterministic inpainting. In *Tech. Report CAM-00-11, Image Processing Group, UCLA* (2000).
- [DMSB99] DESBRUN M., MEYER M., SCHROEDER P., BARR A.: Implicit fairing of irregular meshes using diffusion and curvature flow. In *Proc. ACM SIGGRAPH* (1999), pp. 317–324.
- [GKS00] GOPI M., KRISHNAN S., SILVA C. T.: Surface reconstruction based on lower dimensional localized delaunay triangulation. In *Proc. Eurographics* (2000), vol. 19(3).
- [HDD*92] HOPPE H., DEROSE T., DUCHAMP T., McDONALD J., STUETZLE W.: Surface reconstruction from unorganized points. In *Proc. ACM SIGGRAPH* (1992), pp. 71–78.
- [Kim97] KIMMEL R.: Intrinsic scale space for images on surfaces: The geodesic curvature flow. *Graphical Models and Image Processing* 59(5) (1997), 365–372.
- [Kob97] KOBBELT L.: Discrete fairing. In *Proceedings of the 7th IMA Conference on the Mathematics of Surfaces* (1997), pp. 101–131.
- [LP01] LINSEN L., PRAUTZSCH H.: Global versus local triangulations. In *Proc. Eurographics (short presentations)* (2001), pp. 71–78.
- [Mou] MOUNT D. M.: Ann: A library for approximate nearest neighbor searching. www.cs.umd.edu/~mount/ANN.
- [Pau03] PAULY M.: *Point primitives for interactive modeling and processing of 3D geometry*. Dissertation, Department of Computer Science, ETH Zürich, 2003.
- [PKG03] PAULY M., KEISER R., GROSS M.: Multi-scale feature extraction on point-sampled surfaces. In *Proc. Eurographics* (2003), vol. 22 (3), pp. 121–130.
- [PR99] PREUSSER T., RUMPF M.: Anisotropic nonlinear diffusion in flow visualization. In *Proc. IEEE Visualization* (1999).
- [RL00] RUSINKIEWICZ S., LEVOY M.: QSplat: A Multiresolution Point Rendering System for Large Meshes. In *Proc. ACM SIGGRAPH* (2000), pp. 343–352.
- [She96] SHEWCHUK J. R.: Triangle: Engineering a 2d quality mesh generator and delaunay triangulator. In *1st Workshop of Applied Computational Geometry* (1996), ACM Press, pp. 124–133.
- [Tau95] TAUBIN G.: A signal processing approach to fair surface design. In *Proc. ACM SIGGRAPH* (1995), pp. 351–358.
- [Tur91] TURK G.: Generating textures on arbitrary surfaces using reaction-diffusion. *Computer Graphics (SIGGRAPH '91 Proceedings)* 25, No. 4 (1991), 289–298.
- [VK01] VARSHNEY A., KALAIHAH A.: Differential point rendering. In *Proc. 12th Eurographics Workshop on Rendering* (2001), pp. 139–150.
- [XWH*03] XIE H., WANG J., HUA J., QIN H., KAUFMAN A.: Piecewise \mathcal{J}^1 continuous surface reconstruction of noisy point cloud via local implicit quadric regression. In *Proc. IEEE Visualization* (2003), pp. 198–206.
- [ZPKG02] ZWICKER M., PAULY M., KNOLL O., GROSS M.: Pointshop 3d: an interactive system for point-based surface editing. In *Proc. ACM SIGGRAPH* (2002), pp. 322–329.
- [ZPvBG01] ZWICKER M., PFISTER H., VAN BAAR J., GROSS M.: Surface splatting. In *Proc. ACM SIGGRAPH* (2001), pp. 267–275.




Figure 9: Finite elements on point surfaces applications: Segmentation (a,b,f). Inpainting (c). Fairing (d,e). Texture synthesis (g,h,i)