

Point Set Surfaces

Marc Alexa
TU Darmstadt

Johannes Behr
ZGDV Darmstadt

Daniel Cohen-Or
Tel Aviv University

Shachar Fleishman
Tel Aviv University

David Levin
Tel Aviv University

Claudio T. Silva
AT&T Labs

Abstract

We advocate the use of point sets to represent shapes. We provide a definition of a smooth manifold surface from a set of points close to the original surface. The definition is based on local maps from differential geometry, which are approximated by the method of moving least squares (MLS). We present tools to increase or decrease the density of the points, thus, allowing an adjustment of the spacing among the points to control the fidelity of the representation.

To display the point set surface, we introduce a novel point rendering technique. The idea is to evaluate the local maps according to the image resolution. This results in high quality shading effects and smooth silhouettes at interactive frame rates.

CR Categories: I.3.3 [Computer Graphics]: Picture/Image Generation—Display algorithms; I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Curve, surface, solid, and object representations;

Keywords: surface representation and reconstruction, moving least squares, point sample rendering, 3D acquisition

1 Introduction

Point sets are receiving a growing amount of attention as a representation of models in computer graphics. One reason for this is the emergence of affordable and accurate scanning devices generating a dense point set, which is an initial representation of the physical model [28]. Another reason is that highly detailed surfaces require a large number of small primitives, which contribute to less than a pixel when displayed, so that points become an effective display primitive [33, 36].

A point-based representation should be as small as possible while conveying the shape, in the sense that the point set is neither noisy nor redundant. It is important to have tools which adequately adjust the density of points so that a smooth surface can be well-reconstructed. Figure 1 shows a point set with varying density. Amenta et al [1] have investigated the problem from a topological point of view, that is, the number of points needed to guarantee a topologically equivalent reconstruction of a smooth surface. Our approach is motivated by differential geometry and aims at minimizing the geometric error of the approximation. This is done by locally approximating the surface with polynomials using moving least squares (MLS).



Figure 1: A point set representing a statue of an angel. The density of points and, thus, the accuracy of the shape representation are changing (intentionally) along the vertical direction.

We understand the generation of points on the surface of a shape as a sampling process. The number of points is adjusted by either up-sampling or down-sampling the representation. Given a data set of points $P = \{p_i\}$ (possibly acquired by a 3D scanning device), we define a smooth surface S_P (MLS surface) based on the input points (the definition of the surface is given in Section 3). We suggest replacing the points P defining S_P with a reduced set $R = \{r_i\}$ defining an MLS surface S_R which approximates S_P . This general paradigm is illustrated in 2D in Figure 2: Points P , depicted in purple, define a curve S_P (also in purple). S_P is resampled with points $r_i \in S_P$ (red points). This typically lighter point set called the *representation* points now defines the red curve S_R which approximates S_P .

The technique that defines and resamples S_P provides the fol-

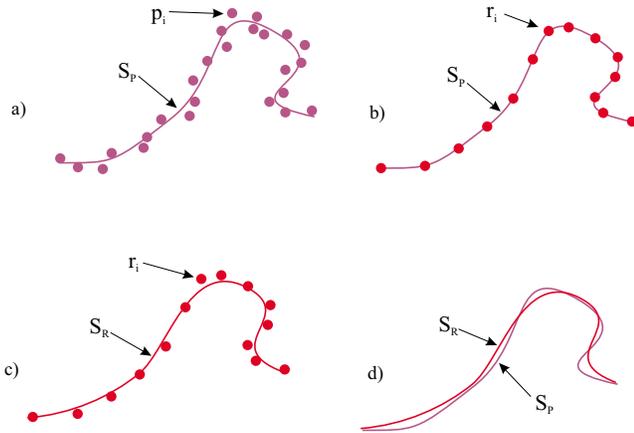


Figure 2: An illustration of the paradigm: The possibly noisy or redundant point set (purple points) defines a manifold (purple curve). This manifold is sampled with (red) representation points. The representation points define a different manifold (red curve). The spacing of representation points depends on the desired accuracy of the approximation.

lowing important property:

Smooth manifold surface: The surface defined by the point set is guaranteed to be a 2-manifold and C^∞ smooth, given that the points are sufficiently close to the surface being represented.

In the context of surface consolidation (see Section 2), our approach has two important features:

Noise reduction: Filtering imperfect data and generating a thin point set, in the sense that the points do not deviate from the surface being represented.

Redundancy reduction: The point set is down-sampled by removing redundant information introduced by oversampling the surface.

Finally, we present a rendering method that approximates S_R from the local polynomial approximations offering:

High quality: Since S_R is a smooth surface, proper resampling leads to smooth silhouettes and normals resulting in superior rendering quality at interactive frame rates.

Single step procedure: Resampling respects screen space resolution and guarantees sufficient sampling, i.e. no holes have to be filled in a postprocessing step.

2 Related work

Consolidation

Recent technological and algorithmic advances have improved the process of automatic acquisition of 3D models. Acquiring the geometry of an object starts with data acquisition, usually performed with a range scanner. This raw data contains errors (e.g., line-of-sight error [16, 37]) mainly due to noise intrinsic to the sensor used and its interaction with the real-world object being acquired. For a non-trivial object, it is necessary to perform multiple scans, each in its own coordinate system, and to register the scans [5]. In general, areas of the objects are likely to be covered by several samples

from scans performed from different positions. One can think of the output of the registration as a *thick* point set.

A common approach is to generate a triangulated surface model over the thick point set. There are several efficient triangulation techniques, such as [1, 2, 4, 6, 12]. One of the shortcomings of this approach is that the triangulated model is likely to be rough, containing bumps and other kinds of undesirable features, such as holes and tunnels, and be non-manifold. Further processing of the triangulated models, such as smoothing [41, 11] or manifold conversion [15], becomes necessary. The prominent difficulty is that the point set might not actually interpolate a smooth surface. We call *consolidation* the process of “massaging” the point set into a surface. Some techniques, such as Hoppe et al [18], Curless and Levoy [9], and Wheeler et al [45] consolidate their sampled data by using an implicit representation based on a distance function defined on a volumetric grid. In [18], the distances are taken as the signed distance to a locally defined tangent plan. This technique needs further processing [19, 17] to generate a smooth surface. Curless and Levoy [9] use the structure of the range scans and essentially scan convert each range surface into the volume, properly weighting the multiply scanned areas. Their technique is robust to noise and is able to take relative confidence of the samples into account. The work of Wheeler et al [45] computes the signed distance to a consensus surface defined by weighted averaging of the different scans. One of the nice properties of the volumetric approach is that it is possible to prove under certain conditions that the output is a least-square fit of the input points (see [9]).

The volumetric sign-distance techniques described above are related to a new field in computer graphics called *Volume Graphics* pioneered by Kaufman and colleagues [20, 44, 40] which aims to accurately define how to deal with volumetric data directly, and answer questions related to the proper way to convert between surface and volume representations.

It is also possible to consolidate the point set by performing weighted averaging directly on the data points. In [43], model triangulation is performed first, then averaging is performed in areas which overlap. In [39], the data points are first averaged, weighted by a confidence in each measurement, and then triangulated.

Another approach to define surfaces from the data points is to perform some type of surface fitting [13], such as fitting a polynomial [25] or an algebraic surface [34] to the data. In general, it is necessary to know the intrinsic topology of the data and (sometimes) have a parametrization before surface fitting can be applied. Since this is a non trivial task, Krishnamurthy and Levoy [22] have proposed a semi-automatic technique for fitting smooth surfaces to dense polygon meshes created by Curless and Levoy [9]. Another form of surface fitting algorithms couples some form of high-level model recognition with a fitting process [35].

The process of sampling (or resampling) surfaces has been studied in different settings. For instance, surface simplification algorithms [7] sample surfaces in different ways to optimize rendering performance. Related to our work are algorithms which use particle systems for sampling surfaces. Turk [42] proposes a technique for computing level of details of triangular surfaces by first randomly spreading points on a triangular surface, then optimizing their positions by letting each point repel their neighbors. He uses an approximation of surface curvature to weight the number of points which should be placed in a given area of the surface. A related approach is to use physically-based particle systems to sample an implicit surface [46, 10]. Crossno and Angel [8] describe a system for sampling isosurfaces, where they use the curvature to automatically modulate the repulsive forces.

Lee [24] uses a moving-least squares approach to the reconstruction of curves from unorganized and noisy points. He proposes a solution for reconstructing two and three-dimensional curves by thinning the point sets. Although his approach resembles the one

used here (and based on theory developed in [27]), his *projection* procedure is different, and requires several iterations to converge to a clean point set (i.e., it is not actually a projection, but more of a converging smoothing step).

Point sample rendering

Following the pioneering work of Levoy and Whitted [29], several researchers have recently proposed using “points” as the basic rendering primitive, instead of traditional rendering primitives, such as triangulated models. One of the main reasons for this trend is that in complex models the triangle size is decreasing to pixel resolution. This is particularly true for real-world objects acquired as “textured” point clouds [31].

Grossman and Dally [14] presented techniques for converting geometric models into point-sampled data sets, and algorithms for efficiently rendering the point sets. Their technique addresses several fundamental issues, including the sampling rate of conversion from triangles to points, and several rendering issues, including handling “gaps” in the rendered images and efficient visibility data structures. The Surfels technique of Pfister et al [33] builds and improves on this earlier work. They present alternative techniques for the sampling of the triangle mesh, including visibility testing, texture filtering, and shading.

Rusinkiewicz and Levoy [36] introduce a technique which uses a hierarchy of spheres of different radii to model a high-resolution model. Their technique uses small spheres to model the vertices at the highest resolution, and a set of bounding spheres to model intermediate levels. Together with each spherical sample, they also save other associated data, such as normals. Their system is capable of time-critical rendering, as it adapts the depth of tree traversal to the available time for rendering a given frame.

All the above techniques account for local illumination. Schaffler and Jensen [38] propose to compute global illumination effects directly on point-sampled geometry by a ray tracing technique. The actual intersection point is computed, based on a local approximation of the surface, assuming a uniform sampling of the surface.

Point-based rendering suffers from the limited resolution of the fixed number of sample points representing the model. At some distance, the screen space resolution is high relative to the point samples, which causes undersampling. Tackling this problem by interpolating the surface points in image space is limited. A better approach is to resample the surface during rendering at the desired resolution in object-space, guaranteeing that sampling density is sufficient with respect to the image space resolution.

3 Defining the surface - projecting

Our approach relies on the idea that the given point set implicitly defines a surface. We build upon recent work by Levin [27]. The main idea is the definition of a projection procedure, which projects any point near the point set onto the surface. Then, the MLS surface is defined as the points projecting onto themselves. In the following, we explain Levin’s projection procedure [27] and, then, how to efficiently compute it.

3.1 The projection procedure

Let points $p_i \in \mathbb{R}^3, i \in \{1, \dots, N\}$, be sampled from a surface S (possibly with a measurement noise). The goal is to project a point $r \in \mathbb{R}^3$ near S onto a two-dimensional surface S_P that approximates the p_i . The following procedure is motivated by differential geometry, namely that the surface can be locally approximated by a function.

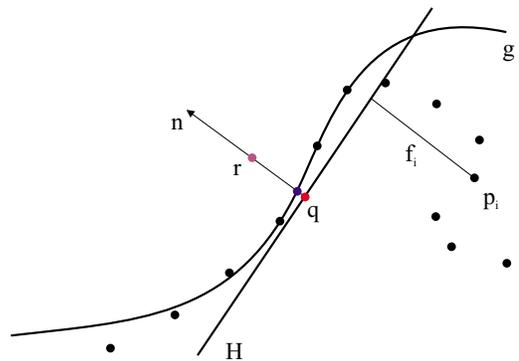


Figure 3: The MLS projection procedure: First, a local reference domain H for the purple point r is generated. The projection of r onto H defines its origin q (the red point). Then, a local polynomial approximation g to the heights f_i of points p_i over H is computed. In both cases, the weight for each of the p_i is a function of the distance to q (the red point). The projection of r onto g (the blue point) is the result of the MLS projection procedure.

1. **Reference domain:** Find a local reference domain (plane) for r (see Figure 3). The local plane $H = \{x | \langle n, x \rangle - D = 0, x \in \mathbb{R}^3, n \in \mathbb{R}^3, \|n\| = 1\}$ is computed so as to minimize a local weighted sum of square distances of the points p_i to the plane. The weights attached to p_i are defined as the function of the distance of p_i to the projection of r on the plane H , rather than the distance to r . Assume q is the projection of r onto H , then H is found by minimizing

$$\sum_{i=1}^N (\langle n, p_i \rangle - D)^2 \theta(\|p_i - q\|) \quad (1)$$

where θ is a smooth, radial, monotone decreasing function, which is positive on the whole space. Since θ decreases as the distance of the points increases, the plane approximates the tangent plane to S near the point r . The local reference domain is then given by an orthonormal coordinate system on H so that q is the origin of this system.

2. **Local map:** The reference domain for r is used to compute a local bivariate polynomial approximation to the surface in a neighborhood of r (see Figure 3). Let q_i be the projection of p_i onto H , and f_i the height of p_i over H , i.e. $f_i = n \cdot (p_i - q)$. Compute the coefficients of a polynomial approximation g so that the weighted least squares error

$$\sum_{i=1}^N (g(x_i, y_i) - f_i)^2 \theta(\|p_i - q\|) \quad (2)$$

is minimized. Here (x_i, y_i) is the representation of q_i in a local coordinate system in H . Note that, again, the distances used for the weight function are from the projection of r onto H . The projection of r onto S_P is defined by the polynomial value at the origin, i.e. $q + g(0, 0)n$.

Levin proves that the surface defined as the points that project onto themselves is a two-dimensional manifold [27]. Further, a general analysis of moving least squares [26] leads to the conjecture that the resulting surface is infinitely smooth as long as $\theta \in C^\infty$ (provided that θ has the above-mentioned properties).

The approximation of single points is mainly dictated by the radial weight function θ . The weight function suggested in [27] is a



Figure 4: The effect of different values for parameter h . A point set representing an Aphrodite statue defines an MLS surface. The left side shows an MLS surface resulting from a small value and reveals a surface structure resulting from the wood carving. The right side shows a larger value for h , smoothing out small features or noise.

Gaussian

$$\theta(d) = e^{-\frac{d^2}{h^2}} \quad (3)$$

where h is a fixed parameter reflecting the anticipated spacing between neighboring points. By changing h the surface can be tuned to smooth out features of size $< h$ in S . More specifically, a small value for h causes the Gaussian to decay faster and the approximation is more local. Conversely, large values for h result in a more global approximation, smoothing out sharp or oscillatory features of the surface. Figure 4 illustrates the effect of different h values.

3.2 Computing the planes and polynomials

We explain how to efficiently compute the projection and what values should be chosen for the polynomial degree and h . Furthermore, we discuss trade-offs between accuracy and speed.

Step 1 of the projection procedure is a non-linear optimization problem. By setting $q = r + tn$ for some $t \in \mathbb{R}$, (1) can be rewritten as:

$$\sum_{i=1}^N \langle n, p_i - r - tn \rangle^2 \theta(\|p_i - r - tn\|) \quad (4)$$

Usually, the function will have more than one local minimum. Intuitively, the plane should be close to r , which means t should be small. Thus, we want to choose the local minimum with smallest t . For minimizing (4), we have to use some iterative scheme, which descends towards the next local minimum. Currently, a standard iterative solver is employed, together with a good heuristic choice of the initial values $\{n, t\}$, which ensure that the minimization converges to a local minimum with small t .

We find this initial value by setting $t = 0$ and solving (4):

$$\sum_{i=1}^N \langle n, p_i - r \rangle^2 \theta_i, \quad \theta_i = \theta(\|p_i - r\|) \quad (5)$$

This is a quadratic function in n and can be solved by setting its gradient

$$\sum_{i=1}^N 2\theta_i \langle n, p_i - r \rangle (p_i - r) \quad (6)$$

to zero.

The resulting initial value is refined using Powell iteration. Note that the global minimum of (4) is approached for $t = \infty$. To avoid this solution, (4) is normalized with the sum of weights θ .

The second step of the projection procedure is solved in the same way as the initial value problem for the first step: Once the plane H is computed, the weights $\theta(\|p_i - q\|)$ are known. The gradient of (2) over the unknown coefficients of the polynomial leads to a system of linear equations of size equal to the number of coefficients, e.g. 10 for a third degree polynomial.

Through experimentation, we found that high degree polynomials are likely to oscillate. Polynomials of degree 3 to 4 have proven to be very useful as they produce good fits of the neighborhood, do not oscillate, and are quickly computed.

The most time-consuming step in computing the projection of a point r is collecting the coefficients from each of the p_i . Implemented naively, this process takes $O(N)$ time, where N is the number of points. We compute these terms using a hierarchical method inspired by solutions to the N-body problem [3]. The basic observation is, that a cluster of points far from r can be combined into one point. To exploit this idea, an Octree is filled with the p_i . Leaf nodes contain the p_i ; inner nodes contain information about the number of points in the subtree and their centroid. Then, terms are collected from the nodes of the Octree. If the node's dimension is much smaller than its distance to r , the centroid is used for computing the coefficients; otherwise the subtree is traversed. In addition, we neglect nodes, for which the distance to r is larger than a predefined constant.

A simple way to trade accuracy for speed is to assume that the plane H passes through the point to be projected. This assumption is reasonable for input points, which are expected to be close to the surface they define (e.g. input that has been smoothed). This simplification saves the cost of the iterative minimization scheme.

Actual timings for the projection procedure depend heavily on the feature size h . On a standard Pentium PC the points of the bunny were projected at a rate of 1500-3500 points per second. Smaller values for h lead to faster projection since the neighborhoods and, thus, the number of points taken into account are smaller. The code has not been optimized to be memory efficient and no storage analysis had been performed as the available main memory of 512MB was sufficient in all tests.

4 Generating the representation point set

A given point set might have erroneous point locations (i.e. is noisy), may contain too many points (i.e. is redundant) or not enough points (i.e. is undersampled). The problem of noise is handled by projecting the points onto the MLS surface they define. The result of the projection procedure is a thin point set. Redundancy is avoided by decimating the point set, taking care that it persists to be a good approximation of the MLS surface defined by the original point set. In the case of undersampling, the input point set needs to be up-sampled. In the following sections, we show techniques to remove and add points.

Down-sampling

Given a point set, the decimation process repeatedly removes the point that contributes the smallest amount of information to the shape. The contribution of a point to the shape or the error of the sampling is dictated by the definition of the shape. If the point set is reconstructed by means of a triangulation, criteria from the specific triangulation algorithm should control the resampling. Criteria include the distance of points on the surface [18], curvature [12], or distance from the medial axis of the shape [1]. For a direct display of the point set on a screen homogeneous distribution of the points

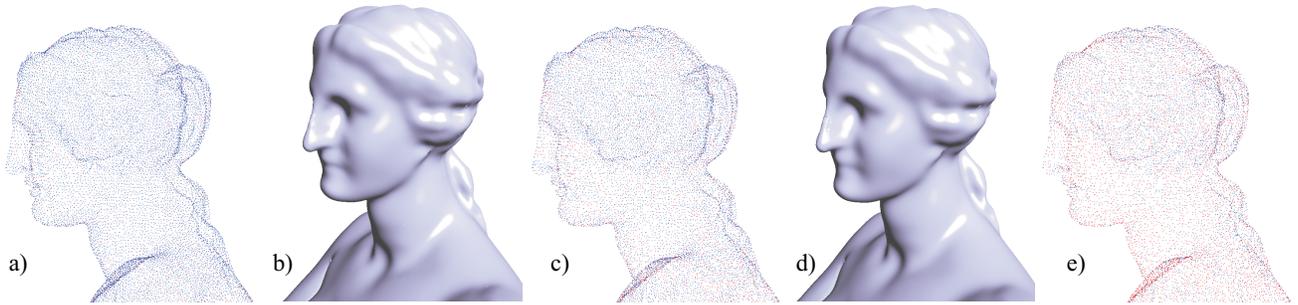


Figure 5: The point set representing an Aphrodite statue is projected onto a smooth MLS-surface. After removing redundant points, a set of 37K points represents the statue (a). The corresponding rendering is shown in (b). The point set is decimated using point removal. An intermediate stage of the reduction process is shown in (c). Note that the points are color-coded with respect to their importance. Blue points do not contribute much to the shape and might be removed; red points are important for the definition of the shape. The final point set in (e) contains 20K points. The corresponding rendering is depicted in (d) and is visually close to the one in (b).

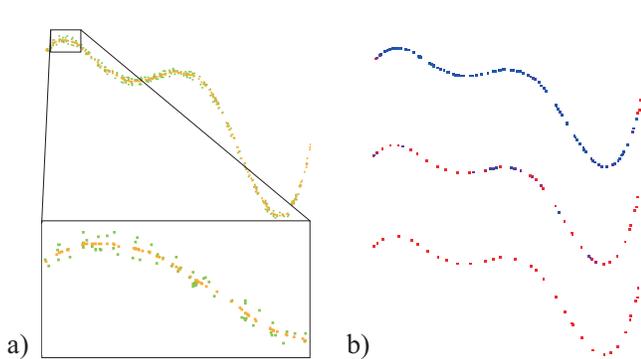


Figure 6: Noisy input points (green points) are projected onto their smooth MLS curve (orange points). The figures in (a) show the point sets and a close-up view. The decimation process is shown in (b). Points are color-coded as in Figure 5.

over the surface is required [14, 33]. Here, we derive a criterion motivated by our definition of the surface.

The contribution of a projected point p_i to the surface S_P can be estimated by comparing S_P with $S_{P-\{p_i\}}$. Computing the Hausdorff distance between both surfaces is expensive and not suitable for an iterative removal of points of a large point set. Instead, we approximate the contribution of p_i by its distance from its projection onto the surface $S_{P-\{p_i\}}$. Thus, we estimate the difference of S_P and $S_{P-\{p_i\}}$ by projecting p_i onto $S_{P-\{p_i\}}$ (projecting p_i under the assumption it was not part of P).

The contribution values of all points are inserted into a priority queue. At each step of the decimation process, the point with the smallest error is removed from the point set and from the priority queue. After the removal of a point, the error values of nearby points have to be recalculated since they might have been affected by the removal. This process is repeated until the desired number of points is reached or the contributions of all points exceed some prespecified bound.

Figure 6 illustrates our decimation process applied on the set of red points depicted in (a). First, the red points are projected on the MLS surface to yield the blue points. A close-up view over a part of the points shows the relation between the input (red) points and the projected points. In (b), we show three snapshots of the decimation process, where points are colored according to their error value; blue represents low error and red represents high error. Note that in the sparsest set, all of the points have a high error, that is, their

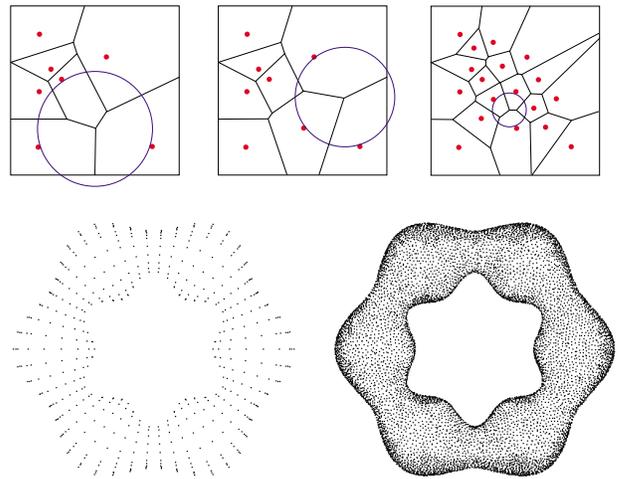


Figure 7: The up-sampling process: Points are added at vertices of the Voronoi diagram. In each step, the vertex with the largest empty circle is chosen. The process is repeated until the radius of the largest circle is smaller than a specified bound. The wavy torus originally consisting of 800 points has been up-sampled to 20K points.

removal will cause a large error. As the decimation proceeds, fewer points remain and their importance grows and the error associated with them is larger. Figure 5 shows the decimation process in 3D with corresponding renderings of the point sets.

Up-sampling

In some cases, the density of the point set might not be sufficient for the intended usage (e.g. direct point rendering or piecewise reconstructions). To alleviate this problem, we try to decrease the spacing among points. Additional points should be placed (and then projected to the MLS surface) where the spacing among points is larger than a specified bound.

The basic idea of our approach is to compute Voronoi diagrams on the MLS surface and add points at vertices of this diagram. Note that the vertices of the Voronoi diagram are exactly those points on the surface with maximum distance to several of the existing points. This idea is related to Lloyd’s method [30], i.e. techniques using

Voronoi diagrams to achieve a certain distribution of points [32].

However, computing the Voronoi diagram on the MLS surface is excessive and local approximations are used instead. More specifically, our technique works as follows: In each step, one of the existing points is selected randomly. A local linear approximation is built and nearby points are projected onto this plane. The Voronoi diagram of these points is computed. Each Voronoi vertex is the center of a circle that touches three or more of the points without including any point. The circle with largest radius is chosen and its center is projected to the MLS surface. The process is repeated iteratively until the radius of the largest circle is less than a user-specified threshold. (see Figure 7). At the end of the process, the density of points is locally near-uniform on the surface. Figure 7 shows the original sparse point set containing 800 points, and the same object after adding 20K points over its MLS surface.

5 Rendering

The challenge of our interactive point rendering approach is to use the representation points and (when necessary) create new points by sampling the implicitly defined surface at a resolution sufficient to conform to the screen space resolution.

Usually, the representation points are not sufficient to render the object in screen space. In some regions, it is not necessary to render all points as they are occluded, backfacing, or have higher density than needed. However, typically, points are not dense enough to be projected directly as a single pixel and more points need to be generated by interpolation in object space.

Culling and view dependency

The structure of our rendering system is similar in spirit to QSPat [36]. The input points are arranged into a bounding sphere hierarchy. For each node, we store a position, a radius, a normal coverage, and optionally a color. The leaf nodes additionally store the orientation of the support plane and the coefficients of the associated polynomial. The hierarchy is used to cull the nodes with the view frustum and to apply a hierarchical backface culling [23]. Note that culling is important for our approach since the cost of rendering the leaf nodes (evaluating the polynomials) is relatively high compared to simpler primitives. Moreover, if the traversal reaches a node with an extent that projects to a size of less than a pixel, this node is simply projected to the frame-buffer without traversing its subtree. When the traversal reaches a leaf node and the extent of its bounding sphere projects to more than one pixel in screen space, additional points have to be generated.

Sampling additional points

The basic idea is to generate a grid of points sufficient to cover the extent of a leaf node. However, projecting the grid points using the method described in Section 3 is not fast enough; thus, we sample the polynomials associated with the representation points. Note that the gradient of the polynomials also yields pointwise normals. Yet, this requires the point set and the associated polynomials to be near-uniform on the surface and it might be necessary to first process a given point set with the up-sampling methods presented in Section 4. This way, we ensure that the local, non-conforming (i.e. overlapping or intersecting) polynomials are a good approximation to the surface inside a patch $[-h, h]^2$ around a point and, thus, the resulting image shows a smooth surface. However, most dense point sets can be readily displayed with the approach. presented here. For example, Figure 9 shows several renderings of the original Stanford Bunny data.

It is critical to properly define the extent of a polynomial patch on the supporting plane, such that neighboring patches are guaranteed

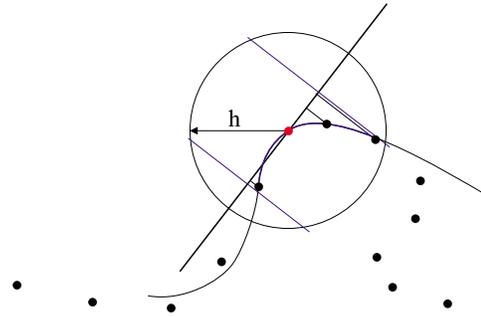


Figure 8: The patch size of a polynomial: Points inside a ball of radius h around the red point are projected onto the support plane of the red point. The patch size is defined as the bounding box (in local coordinates) of the projections. Note that using a disk of radius h or a square patch of $[-h, h]^2$ would lead to unpleasant effects in some cases, as the polynomial might leave the ball of radius h .

to overlap (to avoid holes) but do not overlap more than necessary. Since no inter-point connectivity information is available, it is unclear which points are immediate neighbors of a given point on the surface. To compute the extent of a polynomial patch on the support plane all points inside a ball of radius h are collected. These points are projected to the support plane. The extent is defined as the bounding rectangle of these projections aligned to the local coordinate system of the support plane. Since the spacing of points is expected to be less than h , patches of neighboring points are guaranteed to overlap. Note that using a constant extent (e.g. a disk of radius h on the support plane) can lead to errors, as the polynomial might leave the ball of radius h , in which a good approximation of the point set is expected. Figure 8 illustrates the computation of the patch sizes.

The grid spacing d is computed so that a grid perpendicular to the viewing direction is sufficiently sampled in image space. If the grid is, indeed, perpendicular to the viewing direction, the sampling is also correct on the polynomial. If the grid is not perpendicular to the viewing direction, the projected area might be oversampled. However, note that a polynomial over that grid might not be parallel to the viewing direction. More precisely, if the derivative of the polynomial is less than one, sufficient sampling is guaranteed for all orientations of the support plane. If the derivative is higher, the sampling density needs to be adjusted.

Upon the view-dependent grid spacing d , the polynomials are evaluated by a forward difference approach, where the polynomial is scanned across its extent in its local u, v parametric space. The affine map transforming from support plane coordinates to world coordinates is factored into polynomial evaluation, thus, generating points in world coordinates. These points are then fed into the graphics pipeline to be projected to the screen.

Grid pyramids

The time-critical factor is the view-dependent evaluation of the points on the polynomial. Optimally, these are recomputed whenever the projected screen space size changes. To accelerate the rendering process, we store a grid pyramid with various resolutions per point. Initially, the pyramid levels are created, but no grid is actually evaluated. When a specific grid resolution is needed, the system creates and stores the level that slightly oversamples the polynomial for a specific resolution, such that small changes in the viewing position do not result in new evaluations.

To enhance the interactivity of our approach, we also allow the point size to adapt to changing viewing conditions. For example,

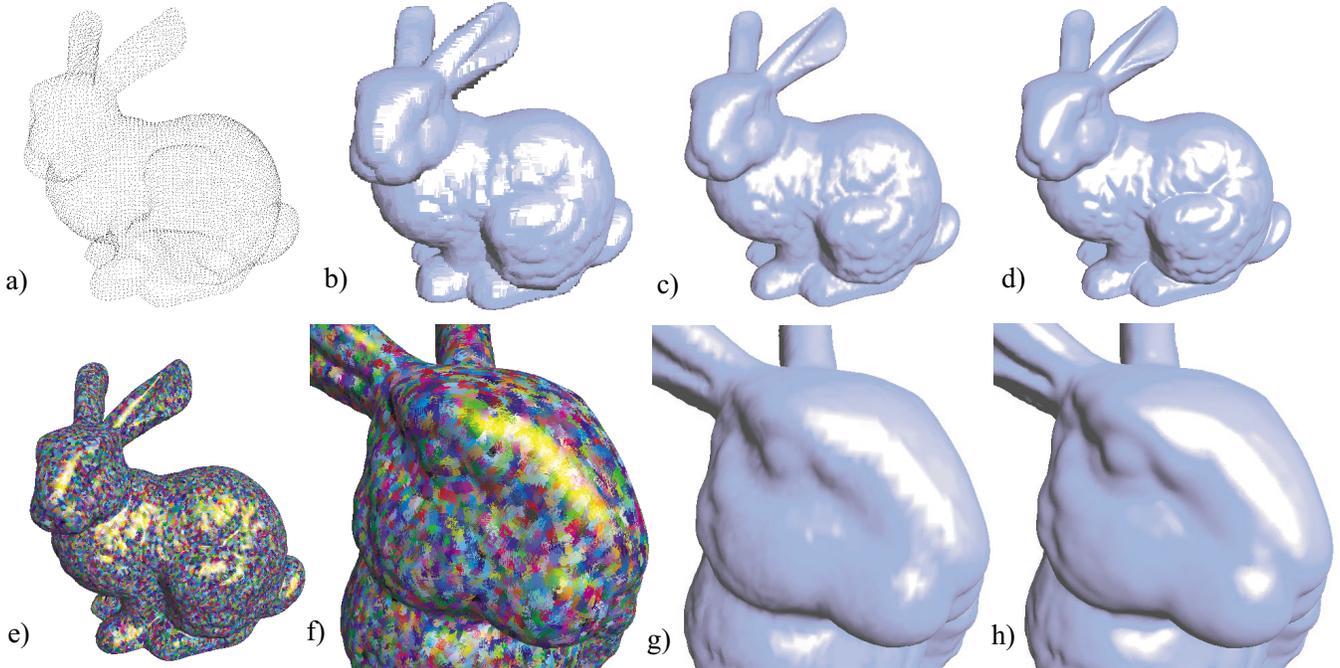


Figure 9: The Stanford Bunny: The points defining the bunny are depicted in (a) (some points are culled). Points are splatted in (b) to satisfy screen space resolution. Note the difference of a piecewise linear mesh over the points (c) and close-up in (g) to the rendering of non-conforming polynomial patches (d) and (h). The patches are color-coded in (e) and (f).

while rotating or zooming, sparser grids with large points are used to guarantee an interactive frame rate. Once the viewer stops moving, a proper grid is chosen from the pyramid.

Results

We have tested our approach on a variety of point sets. Figure 9 shows the renderings of the Stanford Bunny. In (a), the original point set is shown. Splatting (b) is not leading to good results, because the model is not sampled densely enough. The traditional Gouraud-shaded mesh in (c) and (g) is compared to our approach in (d) and (h). Note the accuracy of the highlights. The non-conforming local polynomial patches are shown color-coded in (e) and (f). An example of an environment mapping to demonstrate the normal continuity is given in Figure 10. Note that silhouettes and normals are smooth, which leads to less distortions on the boundary and in the reflections.

The frame rates we achieve are mainly dictated by the number of visible representation points (i.e. graph traversal time) and the number of pixels to be filled. All models depicted in the paper are displayed at more than 5 frames per second in a 512^2 screen window (see the accompanying video for more information). The number of representation points ranges from 1000 (for the torus) to 900K (for the angel statue). Tests are performed on a PC with GeForce2 graphics board.

6 Conclusion

In differential geometry, a smooth surface is characterized by the existence of smooth local maps at any point. In this work we use this as a framework to approximate a smooth surface defined by a set of points and we introduced new techniques to resample the surface to generate an adequate representation of the surface.

To render such surfaces, the surface is covered by a finite number, as small as possible, of non-conforming, overlapping, polynomial patches. We showed that the error of these approximations is bounded and dependent on the spacing among points. Thus, it is possible to provide a point set representation that conforms with a specified tolerance.

Our paradigm for representing surfaces advocates the use of a point set (without connectivity) as a representation of shapes. This representation is universal in the sense that it is used from the beginning (i.e. acquisition) to the end (i.e. rendering) of a graphical representation's life cycle. Moreover, we believe that this work is a step towards rendering with higher order polynomials. Note that we have used simple primitives (points) to achieve this goal. This admits to the current trend of integrating high quality rendering features into graphics hardware.

It would be interesting to integrate our approach with combinatorial methods such as the one of Amenta et al [1]. This would combine topological guarantees with the additional precision of higher order approximations and the possibility of smoothing out noise or small features.

Using different values for h , it is easy to generate more smooth or more detailed versions of a surface from one point set (see, for example, Figure 4). A set of different versions could be used as a smooth-to-detailed hierarchy and would allow for multiresolution modeling [21]. Of course, h is not necessarily a global parameter and could be adapted to the local feature size. Varying h has several implications and utility in handling point sets (see [24] for a nice introduction to the issues in two dimensions), such as properly accounting for differences in sampling rate and levels of noise during the acquisition process. Also, non radial functions might be necessary to properly account for sharp features in the models.

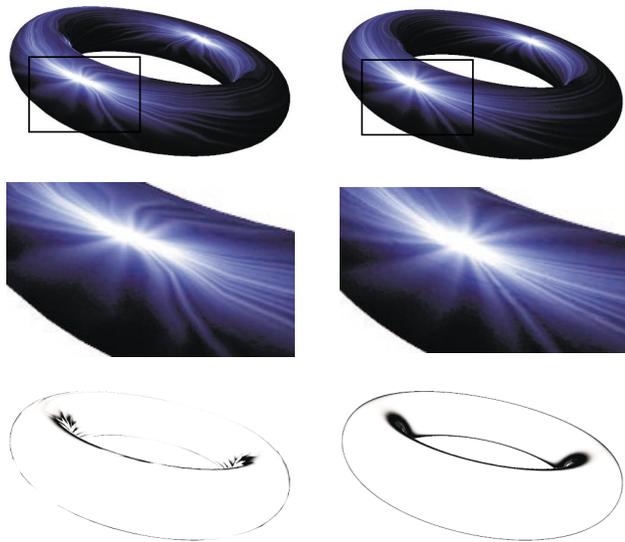


Figure 10: Comparison of mesh rendering with our technique with environment mapping. The left column shows renderings of a mesh consisting of 1000 vertices. The right column shows our technique using the vertices as input points. The environment maps emphasize the improved normal and boundary continuity.

Acknowledgements

This work was supported by a grant from the Israeli Ministry of Science, a grant from GIF (German Israeli Foundation) and a grant from the Israeli Academy of Sciences (center of excellence). The bunny model is courtesy of the Stanford Computer Graphics Laboratory. The angel statue was scanned by Peter Neugebauer using a structured light scanner and the QTSculptor system.

References

- [1] N. Amenta, M. Bern, and M. Kamysseis. A new voronoi-based surface reconstruction algorithm. *Proceedings of SIGGRAPH 98*, pages 415–422, July 1998.
- [2] C. L. Bajaj, F. Bernardini, and G. Xu. Automatic reconstruction of surfaces and scalar fields from 3d scans. *Proceedings of SIGGRAPH 95*, pages 109–118, August 1995.
- [3] J. Barnes and P. Hut. A hierarchical $O(N \log N)$ force calculation algorithm. *Nature*, 324:446–449, Dec. 1986. Technical report.
- [4] F. Bernardini, J. Mittleman, H. Rushmeier, C. Silva, and G. Taubin. The ball-pivoting algorithm for surface reconstruction. *IEEE Transactions on Visualization and Computer Graphics*, 5(4):349–359, October - December 1999.
- [5] P. Besl and N. McKay. A method for registration of 3-d shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):239–256, February 1992.
- [6] J.-D. Boissonnat. Geometric structures for three-dimensional shape representation. *ACM Transactions on Graphics*, 3(4):266–286, October 1984.
- [7] P. Cignoni, C. Montani, and R. Scopigno. A comparison of mesh simplification algorithms. *Computers & Graphics*, 22(1):37–54, February 1998.
- [8] P. Crossno and E. Angel. Isosurface extraction using particle systems. *IEEE Visualization '97*, pages 495–498, November 1997.
- [9] B. Curless and M. Levoy. A volumetric method for building complex models from range images. *Proceedings of SIGGRAPH 96*, pages 303–312, August 1996.
- [10] L. H. de Figueiredo, J. de Miranda Gomes, D. Terzopoulos, and L. Velho. Physically-based methods for polygonization of implicit surfaces. *Graphics Interface '92*, pages 250–257, May 1992.
- [11] M. Desbrun, M. Meyer, P. Schröder, and A. H. Barr. Implicit fairing of irregular meshes using diffusion and curvature flow. *Proceedings of SIGGRAPH 99*, pages 317–324, August 1999.
- [12] M. Gopi, S. Krishnan, and C. T. Silva. Surface reconstruction based on lower dimensional localized delaunay triangulation. *Computer Graphics Forum*, 19(3), August 2000.

- [13] A. Goshtasby and W. D. O'Neill. Surface fitting to scattered data by a sum of gaussians. *Computer Aided Geometric Design*, 10(2):143–156, April 1993.
- [14] J. P. Grossman and W. J. Dally. Point sample rendering. *Eurographics Rendering Workshop 1998*, pages 181–192, June 1998.
- [15] A. P. Guezic, G. Taubin, F. Lazarus, and W. Horn. Converting sets of polygons to manifold surfaces by cutting and stitching. *IEEE Visualization '98*, pages 383–390, October 1998.
- [16] P. Hebert, D. Laurendeau, and D. Poussart. Scene reconstruction and description: Geometric primitive extraction from multiple view scattered data. In *IEEE Computer Vision and Pattern Recognition 1993*, pages 286–292, 1993.
- [17] H. Hoppe, T. DeRose, T. Duchamp, M. Halstead, H. Jin, J. McDonald, J. Schweitzer, and W. Stuetzle. Piecewise smooth surface reconstruction. *Proceedings of SIGGRAPH 94*, pages 295–302, July 1994.
- [18] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, and W. Stuetzle. Surface reconstruction from unorganized points. *Computer Graphics (Proceedings of SIGGRAPH 92)*, 26(2):71–78, July 1992.
- [19] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, and W. Stuetzle. Mesh optimization. *Proceedings of SIGGRAPH 93*, pages 19–26, August 1993.
- [20] A. Kaufman, D. Cohen, and R. Yagel. Volume graphics. *IEEE Computer*, 26(7):51–64, July 1993.
- [21] L. P. Kobbelt, T. Bareuther, and H.-P. Seidel. Multiresolution shape deformations for meshes with dynamic vertex connectivity. *Computer Graphics Forum*, 19(3):249–259, Aug. 2000.
- [22] V. Krishnamurthy and M. Levoy. Fitting smooth surfaces to dense polygon meshes. *Proceedings of SIGGRAPH 96*, pages 313–324, August 1996.
- [23] S. Kumar, D. Manocha, W. Garrett, and M. Lin. Hierarchical back-face computation. In X. Pueyo and P. Schröder, editors, *Eurographics Rendering Workshop 1996*, pages 235–244, New York City, NY, June 1996. Eurographics, Springer Wien.
- [24] I.-K. Lee. Curve reconstruction from unorganized points. *Computer Aided Geometric Design*, 17(2):161–177, February 2000.
- [25] Z. Lei, M. M. Blane, and D. B. Cooper. 3L fitting of higher degree implicit polynomials. In *Proceedings of Third IEEE Workshop on Applications of Computer Vision*, Sarasota, Florida, USA, Dec. 1996.
- [26] D. Levin. The approximation power of moving least-squares. *Mathematics of Computation*, 67(224), 1998.
- [27] D. Levin. Mesh-independent surface interpolation. *Advances in Computational Mathematics*, in press, 2001.
- [28] M. Levoy, K. Pulli, B. Curless, S. Rusinkiewicz, D. Koller, L. Pereira, M. Ginzton, S. Anderson, J. Davis, J. Ginsberg, J. Shade, and D. Fulk. The digital michelangelo project: 3d scanning of large statues. *Proceedings of SIGGRAPH 2000*, pages 131–144, July 2000.
- [29] M. Levoy and T. Whitted. The use of points as a display primitive. Tr 85-022, University of North Carolina at Chapel Hill, 1985.
- [30] S. P. Lloyd. Least squares quantization in PCM. *IEEE Transactions on Information Theory*, 28:128–137, 1982.
- [31] D. K. McAllister, L. F. Nyland, V. Popescu, A. Lastra, and C. McCue. Real-time rendering of real-world environments. *Eurographics Rendering Workshop 1999*, June 1999.
- [32] A. Okabe, B. Boots, and K. Sugihara. *Spatial Tesselations — Concepts and Applications of Voronoi Diagrams*. Wiley, Chichester, 1992.
- [33] H. Pfister, M. Zwicker, J. van Baar, and M. Gross. Surfels: Surface elements as rendering primitives. *Proceedings of SIGGRAPH 2000*, pages 335–342, July 2000.
- [34] V. Pratt. Direct least-squares fitting of algebraic surfaces. *Computer Graphics (Proceedings of SIGGRAPH 87)*, 21(4):145–152, July 1987.
- [35] R. Ramamoorthi and J. Arvo. Creating generative models from range images. *Proceedings of SIGGRAPH 99*, pages 195–204, August 1999.
- [36] S. Rusinkiewicz and M. Levoy. Qsplat: A multiresolution point rendering system for large meshes. *Proceedings of SIGGRAPH 2000*, pages 343–352, July 2000.
- [37] M. Rutishauser, M. Stricker, and M. Trobina. Merging range images of arbitrarily shaped objects. In *IEEE Computer Vision and Pattern Recognition 1994*, pages 573–580, 1994.
- [38] G. Schaufler and H. W. Jensen. Ray tracing point sampled geometry. *Rendering Techniques 2000: 11th Eurographics Workshop on Rendering*, pages 319–328, June 2000.
- [39] M. Soucy and D. Laurendeau. Surface modeling from dynamic integration of multiple range views. In *1992 International Conference on Pattern Recognition*, pages 1:449–452, 1992.
- [40] M. Sramek and A. E. Kaufman. Alias-free voxelization of geometric objects. *IEEE Transactions on Visualization and Computer Graphics*, 5(3):251–267, July - September 1999. ISSN 1077-2626.
- [41] G. Taubin. A signal processing approach to fair surface design. *Proceedings of SIGGRAPH 95*, pages 351–358, August 1995.
- [42] G. Turk. Re-tiling polygonal surfaces. *Computer Graphics (Proceedings of SIGGRAPH 92)*, 26(2):55–64, July 1992.
- [43] G. Turk and M. Levoy. Zippered polygon meshes from range images. *Proceedings of SIGGRAPH 94*, pages 311–318, July 1994.
- [44] S. W. Wang and A. E. Kaufman. Volume sculpting. *1995 Symposium on Interactive 3D Graphics*, pages 151–156, April 1995. ISBN 0-89791-736-7.
- [45] M. Wheeler, Y. Sato, and K. Ikeuchi. Consensus surfaces for modeling 3d objects from multiple range images. In *IEEE International Conference on Computer Vision 1998*, pages 917–924, 1998.
- [46] A. P. Witkin and P. S. Heckbert. Using particles to sample and control implicit surfaces. *Proceedings of SIGGRAPH 94*, pages 269–278, July 1994.