

# A survey of point-based techniques in computer graphics

Leif Kobbelt\*, Mario Botsch

*Computer Graphics Group, RWTH Aachen University, 52056 Aachen, Germany*

---

## Abstract

In recent years point-based geometry has gained increasing attention as an alternative surface representation, both for efficient rendering and for flexible geometry processing of highly complex 3D-models. Point-sampled objects do neither have to store nor to maintain globally consistent topological information. Therefore they are more flexible compared to triangle meshes when it comes to handling highly complex or dynamically changing shapes. In this paper, we make an attempt to give an overview of the various point-based methods that have been proposed over the last years. In particular we review and evaluate different shape representations, geometric algorithms, and rendering methods, which use points as a universal graphics primitive.

© 2004 Elsevier Ltd. All rights reserved.

*Keywords:* Point-sampled geometry; Point-based graphics; GPU processing

---

## 1. Introduction

Triangle meshes are still the most common surface representation in many computer graphics applications. Because of their simplicity and flexibility, they replace traditional CAD surface representations, like NURBS surfaces, in many areas where processing performance matters.

The reason for this is that triangle meshes are significantly more flexible, since surfaces of any shape and topology can be represented by a single mesh without the need to satisfy complicated inter-patch smoothness conditions. The simplicity of the triangle primitive allows for easier and more efficient geometry generation and geometry processing algorithms. This is most evident for interactive graphics, where the highly optimized (and specialized) graphics hardware is able to process several millions of triangles per second.

Obviously, since the triangle primitive is mathematically much simpler compared to a NURBS patch, more of them have to be used to obtain the same approximation quality. However, if a smooth surface is to be represented by a triangle mesh (a piecewise linear surface), the approximation order is quadratic, i.e., halving the edge lengths reduces the error by a factor of 4 which means the number of triangles is inversely proportional to the approximation error. Hence, even with the weaker asymptotic behavior, a good approximation (for the typical precision requirements in graphics applications) can be achieved with a moderately fine mesh whose vertex density and distribution adapt to the surface curvature, i.e., to the shape complexity.

Although being much more flexible than NURBS, triangle meshes can also have restrictions and disadvantages in some special applications. Most algorithms working on triangle meshes require topologically consistent two-manifold surfaces. As a consequence, manifold-extraction or topology cleanup steps are necessary for mesh generation methods [1,2]. Maintaining the

---

\*Corresponding author.

*E-mail addresses:* [kobbelt@infomatik.rwth-aachen.de](mailto:kobbelt@infomatik.rwth-aachen.de) (L. Kobbelt), [botsch@cs.rwth-aachen.de](mailto:botsch@cs.rwth-aachen.de) (M. Botsch).

topological consistency throughout the mesh processing pipeline makes these algorithms sometimes significantly more complicated, even if the eventual target application is, e.g., visualization where a consistent topology is actually not necessary. Dynamic mesh connectivity [3,4] is one example where frequent topology changes occur because the mesh is locally restructured in order to avoid too much stretching after extreme deformations or to provide more degrees of freedom in certain regions of interest. In cases like this, point-based representations would allow for more flexibility when a globally consistent surface topology is not necessarily required (e.g., for rendering). Hence, the consequent next step towards simpler and more flexible geometric primitives seems to be the use of unstructured clouds of points or disks where we do not need to store and maintain globally consistent connectivity information.

From a rendering point of view, triangle meshes may not be suitable as meshes are becoming more and more complex while the typical screen resolution does not grow as fast. The steadily increasing performance of CPUs and graphics hardware, the cheap memory, and the wide availability of range-scanning devices result in the acquisition and generation of massive highly detailed geometry data—models consisting of several millions of triangles are common nowadays [5]. When the number of triangles exceeds the number of pixels on the screen, rendering such massive data sets leads to triangles whose projected area is less than one pixel. In this situation, the traditional incremental rasterization methods become inefficient because of the expensive triangle setup. Hence, points seem to qualify as a better suited rendering primitive for such highly complex models.

In this survey paper we will first discuss several point-based geometry representations (Section 2) and show special application scenarios where geometry processing algorithms can greatly benefit from the flexibility offered by point-sampled geometries (Section 3). In Section 4 we will finally review various point-based rendering methods. More details on point-based geometry processing, point-based rendering, or point-based computer graphics in general can be found in [6–9].

## 2. Point-based representations

A point-based geometry representation can be considered a sampling of a continuous surface, resulting in 3D positions  $\mathbf{p}_i$ , optionally with associated normal vectors  $\mathbf{n}_i$  or auxiliary surface properties like, e.g., colors or other material properties.

### 2.1. Neighborhoods and normals

If normal vectors are not given, they can be estimated by analyzing the local neighborhood of each sample

point. Because there is no connectivity information available, these local neighborhoods are usually constructed using either Euclidean neighborhoods or  $k$  nearest neighbors. In the first case, all samples within an  $\varepsilon$ -ball around a query point are defined to be its neighbors. This naïve method is not suited for irregularly sampled models, since either too many or too few neighbors may be found within the  $\varepsilon$ -ball. The neighborhood estimate will also get unreliable as soon as the local feature size [1] becomes smaller than  $\varepsilon$ , e.g., if two separate surface sheets are located spatially close to each other.

In contrast, the  $k$  nearest neighbors provide a naturally adaptive neighborhood relation in such situations. If the points satisfy certain sampling criteria, like adaptation to the local feature size, then the neighborhood estimate is guaranteed to be reliable [1,10]. Moreover, the  $k$  nearest neighborhood structure has linear complexity in the number of sample points which is similar to polygon meshes.

Both Euclidean neighborhood and  $k$  nearest neighborhood can be computed efficiently by using some hierarchical space partitioning technique. In contrast to triangle meshes, where (consistent) neighborhood information is represented explicitly, the local neighborhoods are usually dynamically re-computed in the case of point-sampled geometries. However, in certain applications a caching of local neighborhood information may also be suitable [11].

Let  $\mathbf{p}_0$  be a sample point and  $\{\mathbf{p}_1, \dots, \mathbf{p}_k\}$  its  $k$  nearest neighbors. The covariance matrix

$$C := \sum_{i=0}^k (\mathbf{p}_i - \bar{\mathbf{p}})(\mathbf{p}_i - \bar{\mathbf{p}})^T \in \mathbb{R}^{3 \times 3},$$

with  $\bar{\mathbf{p}} := \sum_{i=0}^k \mathbf{p}_i / (k+1)$ , is symmetric and positive semi-definite. The eigenvector corresponding to the smallest eigenvalue gives an estimate for the normal direction. Since this determines the normal vector up to its sign only, a consistent orientation over all sample points has to be constructed by a propagation along a minimum spanning tree [12].

### 2.2. Purely point-based representations

From an approximation point of view, a point cloud is a piecewise constant surface approximant. Hence, the resulting approximation power is linear, i.e., with an average spacing  $h$  between the samples  $\mathbf{p}_i$ , the approximation error with respect to each coordinate function is of the order  $O(h)$  [13] (cf. Fig. 1). For the approximation of a geometric shape this means that the approximation error is determined by the *spacing* between the samples. As a consequence, the sampling has to be adjusted to the surface area, implying a dense sampling even in flat surface regions. In fact, the number of samples is

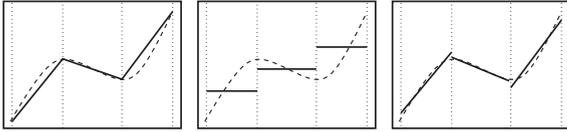


Fig. 1. Comparison of the different shape approximations: piecewise linear  $C^0$  polygons (left), piecewise constant  $C^{-1}$  points (center), and piecewise linear  $C^{-1}$  splats (right). Splats provide the same approximation power as triangle meshes, but due to the  $C^{-1}$  continuity they offer the same flexibility as point clouds.

inversely proportional to the *squared* approximation error.

A point-based representation that is mainly targeting at rendering is proposed by Grossman and Dally [14]. Instead of a completely unorganized point cloud, they use a set of depth images that are orthogonally sampled from a given input geometry. Similar to image-based approaches, this representation is also constructed from several views of an input object, but it differs in that each pixel is a surface sample containing geometric position and (view-independent) surface color.

If the input geometry contains a certain amount of (tolerable) noise, an exact sampling may not be necessary. Instead, Kalaiah and Varshney [15] proposed a non-deterministic statistical point-based representation, where a hierarchical PCA analysis partitions the geometry and its attributes (normals and colors) into a set of local Gaussian probability distributions. Reconstruction and rendering can then be performed by random sampling these probability distribution functions.

For a compact representation, the quantization precision to be used for storing the sample coordinates also has to be taken into account. Since the average sample spacing and the approximation error are of the same order, the quantization error should also be about the same magnitude in order to minimize redundancy (cf. Fig. 2) [16]. If we sample with a resolution of  $n \approx h^{-1}$ , we need  $O(n^2)$  samples to cover a given surface. Using a quantization precision of  $O(n)$  then leads to a total memory consumption of  $O(n^2 \log(n))$  to store the coordinates of all sample points.

This memory cost can be reduced significantly by a hierarchical point-based representation. We start by using a regular 3D binary voxel grid to represent a point-based model, such that for each non-empty grid cell its center position is used as sample point. This sampling respects the proportionality between sample spacing and quantization precision, but it has complexity  $O(n^3)$ , since empty cells have to be stored as well. Recursively sub-sampling the regular grid and applying zero-tree pruning results in an adaptive octree structure. This hierarchical organization of the full cells yields a very compact point-based LoD-representation that

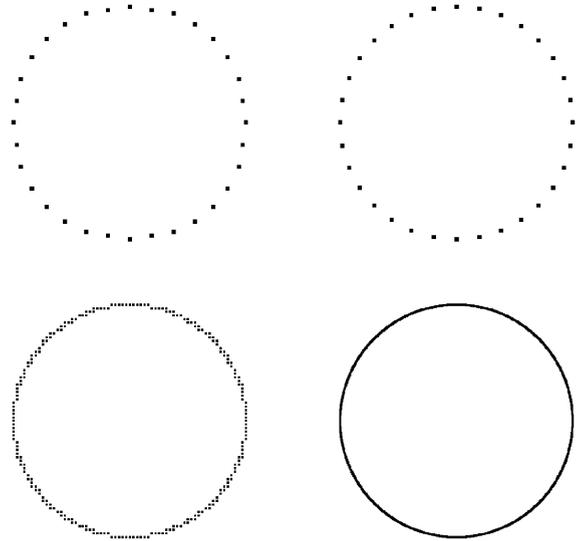


Fig. 2. Point-sampling of a circle with different quantization levels (left: 5 bit, right: 10 bit) and different sampling densities (top:  $2\pi/32$ , bottom:  $2\pi/1024$ ). In the top row the approximation error between the *continuous* circle and the *discrete* point sets is dominated by the distance between samples while in the bottom row the error is dominated by the quantization. Top left and bottom right are good samples since quantization error and sampling density are of the same order thus minimizing redundancy.

needs about  $8/3$  bit per sample point. By entropy encoding, this amount is further reduced down to about 1–1.5 bit per point [16]. Since the bit rate is independent of the sampling density  $h$ , the total memory costs in fact reduce to  $O(n^2)$ .

However, the insufficient object-space approximation power of purely point-based representations usually requires a very dense sampling. Therefore, *splats* seem to be a better compromise between the simplicity of the geometric primitive and the number of primitives that have to be used, as shown in the next section.

### 2.3. Surface splats

Surface splats have first been proposed for rendering purposes by Zwicker et al. [17]. In order to bridge the gaps between neighboring point samples, points  $\mathbf{p}_i$  are associated with a normal vector  $\mathbf{n}_i$  and a radius  $r_i$ , turning them into object-space circular disks.

A locally optimal adaptation to the curvature of the underlying surface is provided by *elliptical* splats, that are defined by two tangential axes  $\mathbf{u}_i$  and  $\mathbf{v}_i$  and their respective radii. Optimal local approximation is achieved if the two axes are aligned to the principal curvature directions of the underlying surface and the

radii are inversely proportional to the corresponding minimum and maximum curvatures. The tangent vectors can be scaled according to the corresponding ellipse radii such that an arbitrary point  $\mathbf{q}$  in the plane spanned by  $\{\mathbf{p}_i, \mathbf{u}_i, \mathbf{v}_i\}$  lies in the interior of the splat if it satisfies the condition

$$(\mathbf{u}_i^T (\mathbf{q} - \mathbf{p}_i))^2 + (\mathbf{v}_i^T (\mathbf{q} - \mathbf{p}_i))^2 \leq 1.$$

Differential geometry tells us that locally a proper ellipse is the best linear approximant to a smooth surface. In this sense splat-based representations are superior to triangle meshes. Since splats are piecewise linear surface primitives, they provide the same quadratic approximation order as triangle meshes (cf. Fig. 1). Just like for triangle meshes, the sampling density can therefore adapt to the surface curvature, such that highly detailed regions are sampled with a higher density, while flat surface regions are sampled more sparsely. Because of the higher approximation power, the quantization precision for splat coordinates has to be  $O(n^2)$ ,  $n$  again denoting the sampling density. Since splats do not have to join continuously (like triangles in a mesh), but are  $C^{-1}$  continuous instead, they still provide the same topological flexibility as pure point clouds (cf. Figs. 5, 13).

Representing sharp features, like edges or corners in technical data sets, is a well-studied problem for triangle meshes. Because the surface is no longer differentiable, the approximation power breaks down to linear order. Additionally, alias artifacts are introduced by insufficient sampling, which cannot be removed by increasing the sampling density [18]. In order to remove these artifacts and reduce normal noise, the sampling has to be aligned to the principal curvature directions [19]. If surface splats are to represent sharp features, all splats that sample these features have to be clipped against one (*edges*) or two (*corners*) clipping

lines that are specified in their local tangent frames (cf. Fig. 3) [20].

#### 2.4. Moving least-squares surfaces

The *moving least-squares* (MLS) surfaces of [21,22] provide an approximating or interpolating surface for a given set of point samples by local higher order polynomials and have first been applied to point-based methods by Alexa et al. [23,24]. An MLS surface  $\mathcal{S}$  is defined in terms of the MLS projection operator  $\Psi: \mathcal{B} \rightarrow \mathbb{R}^3$ , that projects points from a vicinity  $\mathcal{B}$  of the MLS surface onto the surface itself. Hence, the surface  $\mathcal{S}$  is defined by all fix points of this operator that project onto themselves:

$$\mathcal{S} := \{\mathbf{x} \in \mathcal{B} : \Psi(\mathbf{x}) = \mathbf{x}\} = \text{range}(\Psi).$$

The domain  $\mathcal{B}$  of this projection operator is the neighborhood of the input points  $\mathbf{p}_i$  and can be defined, e.g., by a union of balls centered at the  $\mathbf{p}_i$ :

$$\mathcal{B} = \bigcup_i \{\mathbf{x} \in \mathbb{R}^3 \mid \|\mathbf{x} - \mathbf{p}_i\| < r_{\mathcal{B}}\}.$$

The actual computation of the projection  $\mathbf{r} \mapsto \Psi(\mathbf{r})$  is split into three steps:

- (1) Find a local *reference plane*  $H_{\mathbf{r}} = \{\mathbf{x} \in \mathbb{R}^3 : \mathbf{n}^T \mathbf{x} = \mathbf{n}^T \mathbf{q}\}$  by minimizing the non-linear energy functional

$$E_{MLS}(\mathbf{q}, \mathbf{n}) := \sum_i (\mathbf{n}^T \mathbf{p}_i - \mathbf{n}^T \mathbf{q})^2 \theta(\|\mathbf{p}_i - \mathbf{q}\|)$$

for any  $\mathbf{n}, \mathbf{q} \in \mathbb{R}^3$  with  $\mathbf{n} = \mathbf{n}(\mathbf{q}) = (\mathbf{r} - \mathbf{q}) / \|\mathbf{r} - \mathbf{q}\|$ , where  $\theta: \mathbb{R} \rightarrow \mathbb{R}$  denotes a smooth, positive, and monotonically decreasing weight function.

- (2) Find a local bivariate *polynomial approximation*  $g: H_{\mathbf{r}} \rightarrow \mathbb{R}^3$  by a weighted least-squares fit to the points  $\mathbf{p}_i$  in the neighborhood of  $\mathbf{q}$ . Let  $\mathbf{q}_i$  be the orthogonal projection of  $\mathbf{p}_i$  onto  $H_{\mathbf{r}}$ ,  $(x_i, y_i)$  its local 2D coordinates, and  $f_i := \|\mathbf{q}_i - \mathbf{p}_i\|$  its height over

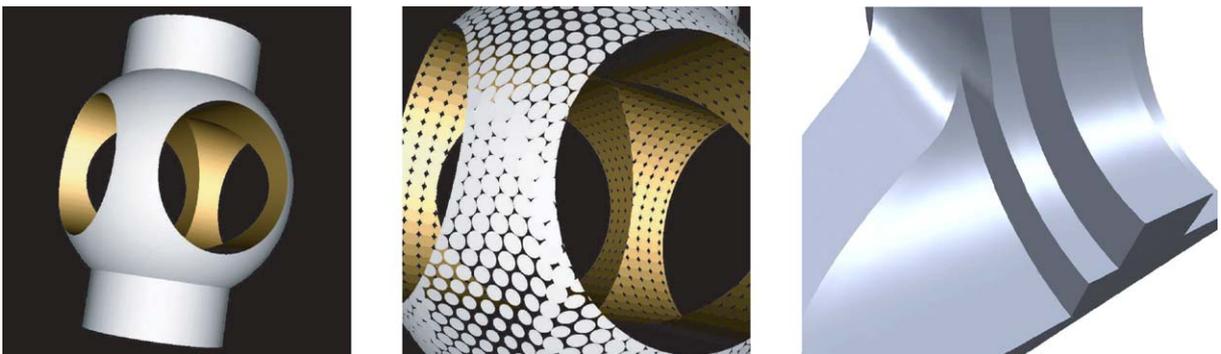


Fig. 3. In order to represent sharp features in technical data sets, each splat can be clipped against one or two clipping lines. The left example was generated using CSG operations, the right example shows a splat-based approximation of the well-known fan disk model.

$H_r$ . Then the error

$$\sum_i (g(x_i, y_i) - f_i)^2 \theta(\|p_i - q\|)$$

is to be minimized.

- (3) The projection of  $r$  is finally defined by  $\Psi(r) := q + g(0, 0)n$ .

The typical choice for the decreasing weight function  $\theta$  is a Gaussian  $\theta(d) = e^{-d^2/h^2}$ , resulting in  $C^\infty$  continuous MLS surfaces. The parameter  $h$  of the Gaussian corresponds to the globally estimated sample spacing  $h$  and can be used to control the degree of smoothing. For irregularly sampled point sets this parameter could be chosen locally instead, as proposed by Pauly et al. [25]. The approximation power of MLS surfaces is conjectured to be  $O(m + 1)$ , with  $m$  being the degree of the local polynomial approximant  $g$ .

The drawback of MLS surfaces is the computationally involved non-linear optimization problem for finding the reference frame  $H_r$ . A slightly different but considerably simpler projection approach is proposed by Alexa and Adamson [26], where they also take a correct normal computation into account in order to properly define implicit surfaces from point cloud data [27,28]. Their projection procedure repeatedly projects a given point  $x$  onto local reference planes that are defined by a weighted average position

$$a(x) := \frac{\sum_i \theta(\|x - p_i\|) p_i}{\sum_i \theta(\|x - p_i\|)}$$

and a normal computed by either a weighted least-squares fit

$$n(x) := \operatorname{argmin}_{\|n\|=1} \sum_i \|n^T(x - p_i)\|^2 \theta(\|x - p_i\|)$$

or by a weighted averaging of input normals  $n_i$

$$n(x) := \frac{\sum_i \theta(\|x - p_i\|) n_i}{\|\sum_i \theta(\|x - p_i\|) n_i\|}.$$

The iterative projection procedure is then proposed as follows (cf. Fig. 4):

- (1) Initialize  $x' \leftarrow a(x)$ .
- (2) Compute  $n \leftarrow n(x')$  and  $a \leftarrow a(x')$ .
- (3) If  $\|n^T(a - x')\| < \epsilon$  return  $x'$ .
- (4) Else project  $x' \leftarrow x' + nn^T(a - x')$  and go back to (2).

If this iteration converges, it yields a point  $x$  on the MLS surface. However, this simplified projection procedure does not result in an orthogonal projection, but it can be enhanced to do so [26].

In contrast to the algorithmic construction of MLS surfaces using the projection operator, Amenta and Kil [29] give an explicit definition of MLS surfaces in terms of critical points of the energy function  $E_{MLS}$  along lines determined by a vector field. In this work they also discuss stability issues of the traditional projection operator for points not being sufficiently close to the surface and propose an alternative projection technique.

MLS surfaces can be used to define a smooth surface from a set of points, but they are also a versatile tool to generate additional sample points on a point-sampled surface, e.g., for up- or down-sampling a model, for low-pass filtering it, or for mapping points back onto the initial surface after local restructuring.

MLS surfaces are a valuable and efficient tool for point-based geometry processing. However, rendering them, e.g. by sufficient up-sampling, is quite involved and does not map to graphics hardware (see Section 4). Hence, in most interactive applications, usually a combination of surface splats and MLS surfaces is used.

### 3. Point-based geometry processing

Several concepts and algorithms known from triangle meshes have been transferred to point-based geometry representations by simply replacing the manifold neighborhood relation implied by the triangle connectivity with the more general  $k$  nearest neighborhood relation

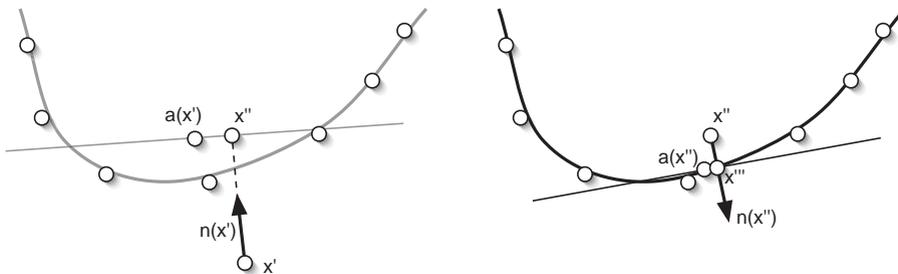


Fig. 4. Simplified MLS projection procedure. In each step the current approximation  $x'$  is updated by projecting it orthogonally onto the reference plane given by  $a(x')$  and  $n(x')$ .

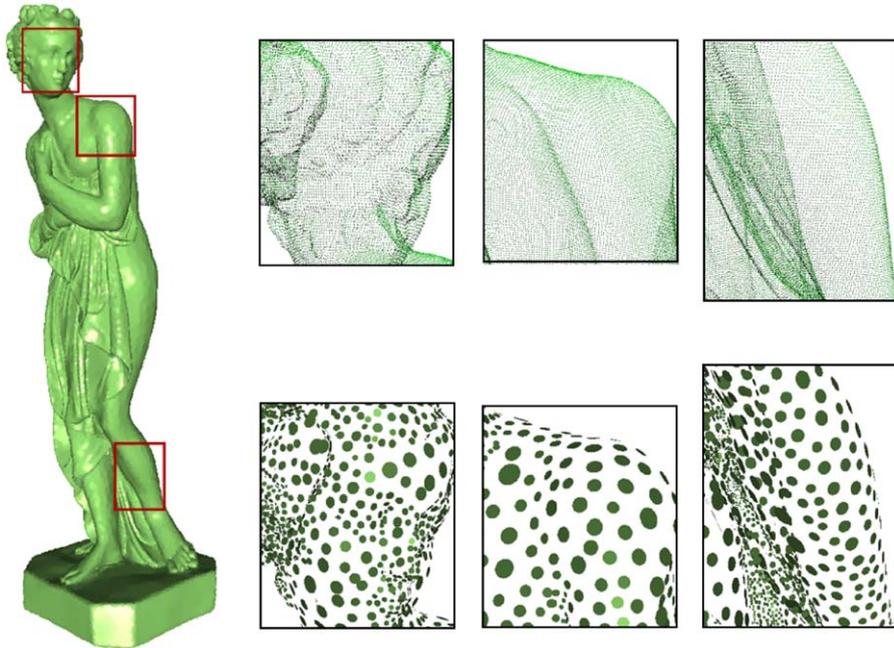


Fig. 5. Optimized sub-sampling of the Iphigenie (top, 350k points) using 30k circular splats (bottom). While respecting a prescribed error tolerance, a global optimization yields a very regular splat distribution.

for point clouds. However, several recently developed techniques are especially designed for point-sampled geometries and actually exploit the additional flexibility offered by them. Since no connectivity information or topological consistency has to be maintained, local re-sampling or restructuring is much easier than it would be using triangle meshes.

Point-based models are usually acquired by range scanning or image-based reconstruction methods. In both cases, the point samples contain a certain amount of noise due to physical measurement errors, which is to be reduced by low-pass filtering techniques. Using a decomposition into patches and local height-field approximations, Pauly and Gross [30] transferred Fourier-based spectral methods to point-sampled geometries. Their method allows for mathematically well-defined frequency analysis and filtering. While the frequencies could also be used to find feature regions of the geometry, a more robust feature detection method was presented by Pauly et al. [31]. Local covariance analyses on multiple scales robustly find feature regions that can be connected to feature lines based on a minimum spanning tree approach.

Scanned data usually contains more artifacts than just measurement noise. Occlusions in concave areas or specular surface materials may lead to holes or insufficient sampling, requiring some kind of hole-filling mechanism. Difficult material reflectances or textured models can also lead to outliers that have to be detected

and removed before further processing. These two and several other intuitive tools for post-processing scanned data have recently been presented by Weyrich et al. [32].

Because real-world objects are captured with as much detail as possible, the resulting models are usually highly complex and have to be simplified in order to be of suitable complexity for further processing. Pauly et al. [25] transferred several mesh decimation methods—vertex clustering, incremental decimation using error quadrics, and re-meshing by particle simulation—to point-sampled surfaces. These methods show the same behavior as their mesh counterparts, i.e., greedy incremental decimation gives good results, but the point distribution can greatly be equalized by several point repulsion iterations. However, both methods do not allow to prescribe an exact global error for the decimation, and they are not taking the full splat geometry into account, but rather rely on the proper distribution of splat *centers*.

In contrast, Wu and Kobbelt [33] propose a simplification method that is especially designed for splat-based surface approximations. Taking the full piecewise linear geometry of surface splats into account, their method provides an exact global error bound as well as a high-quality splat distribution (cf. Fig. 5). In a first step, a patch, respectively a splat, is grown from each input sample, such that it does not violate the error tolerance. From this set of candidate splats, a subset that fully covers the input geometry is chosen in a greedy manner.



Fig. 6. PointShop3D offers various manipulations of point-sampled models, like texture painting (left), low-distortion parameterization and displacement mapping (center), and carving (right).

Finally, a global optimization of the patch selection improves the splat distribution, similar to a repulsion-based particle simulation. Because all candidate splats have been constructed to stay within the error tolerance, the coverage tests can be done by simple set operations. This technique exploits both the improved local approximations by elliptical splats as well as the added flexibility that comes from not having to maintain a globally consistent topology, in order to obtain a considerably better approximation compared to triangle meshes (cf. Fig. 12, left, and Fig. 13).

The biggest advantage of point-based representations over polygonal meshes is that they can easily be restructured without the need to take care of manifold conditions. Hence, applications requiring frequent geometry re-sampling will benefit most from point-based methods. One prominent example is PointShop3D of Zwicker et al. [34], which is a Photoshop-like tool for manipulating point-sampled models (cf. Fig. 6). Since the point samples can represent surface positions as well as material properties or textures, these values can easily be modified on a per-point basis. Tools like sculpting and carving change the point positions, texture smoothing or painting modifies the point colors. In order to be able to paint arbitrary fine details onto the model or to fill up gaps after displacing points, local re-sampling is heavily used. Another very recent approach for painting point-sampled surfaces has been proposed by Adams et al. [35]. Targeting at realistic painting, their method provides haptic feedback and physically simulated paint transfer. Again, this method allows the user to paint arbitrarily fine detail on the surface by locally up-sampling the affected regions on the fly in order to provide a sufficient sample resolution.

The simultaneous point-sampling of geometry, normals, and appearance attributes greatly simplifies these painting methods. In contrast, for triangle meshes there is no such tight coupling between geometry and appearance, requiring the use of a texture atlas for painting the mesh. To enable the painting of fine details,

this texture atlas has to be refined dynamically [36], which is considerably more complicated as the natural point-sampling approach.

Another challenging topic is shape deformation. Among the many different approaches to model free-form shapes, volumetric freeform deformation seems to be best suited for the interactive modification of point-sampled geometries. These deformation techniques first generate a displacement function  $d: \mathbb{R}^3 \rightarrow \mathbb{R}^3$ . Then every surface sample  $\mathbf{p}_i$  (e.g. a mesh vertex or a splat center) is shifted by  $\mathbf{p}_i \mapsto d(\mathbf{p}_i)$ .

When applying an extreme deformation to a triangle mesh, certain triangles exhibit strong stretching, leading to numerically and visually undesirable triangles that have to be removed by local restructuring [3,4,37]. For point-based models, the problem is even worse, since the model can tear apart under too much stretch. However, the splats exhibiting too much stretching are easily detected by looking at the Jacobian of the displacement function  $d$ . Splitting those splats into two, followed by a local tangential relaxation (using particle repulsion and MLS projections), solves the dynamic restructuring problem for point clouds and allows for extreme deformations of point-based models (cf. Fig. 7) [20]. In a similar sense, oriented particles have been used to represent deformable (parametric or implicit) surface models that can easily be stretched, split, or joined, while a repulsion-based particle simulation achieves a uniform and sufficiently dense surface sampling throughout the modeling process [38,39].

Constructive Solid Geometry (CSG) is a very common technique for building complex models by boolean combinations of simpler ones. Since this requires inside/outside tests w.r.t. the given solids, an implicit representation of the models is best-suited [40]. Such a volumetric surface representation can easily be derived from a point-sampled surface by observing the (direction of the) shift vector  $\mathbf{p}_i - \Psi(\mathbf{p}_i)$  induced by the MLS projection operator relative to the oriented normal vector.



Fig. 7. Shape modeling with point-sampled models. Frequent local re-sampling accounts for stretching and allows for extreme deformations (left). A combination of free-form deformation, CSG operations, and displacement mapping was used to create the mug (right).

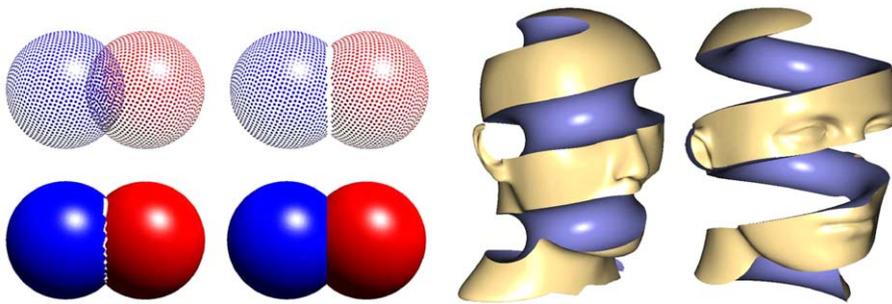


Fig. 8. The left part shows the different steps for point-based CSG operations: classification (top row) and intersection curve sampling (bottom row). Examples of difference and intersection of a head and a helix are shown on the right.

Since the intersection of two smooth objects may result in sharp feature curves, these have to be sampled properly in order to avoid aliasing artifacts. Pauly et al. [20] proposed a simple technique to snap nearby sample points to the sharp features which is based on a local Newton-type iteration method. The sharp features are then represented by clipped splats (cf. Fig. 8, left). Another very efficient method for interactive CSG computations was proposed by Adams and Dutré [41] (cf. Fig. 8, right). In a following paper they also implemented their approach on the GPU in order to exploit hardware acceleration for CSG [42].

Solving PDEs on triangle meshes is a very important topic and has been widely used for mesh smoothing, mesh editing, deformable models, and many other geometry processing tasks. When transferring these methods to point-sampled models, the missing geodesic neighborhood information again has to be replaced by the spatial  $k$  nearest neighbors. These can be used to define an approximate tangent space on the point models such that local Delaunay triangulations define the coupling between neighboring points. This allows for the straightforward discretization of differential operators and the derivation of a corresponding stiffness

matrix. This approach has been proposed by Clarenz et al. [43], where they present PDE-based segmentation, texture impainting, texture synthesis, and geometric smoothing as applications of their framework.

Even physically based modeling and animation of volumetric point-sampled models has been proposed very recently [44]. Since both the surface and the volume of the model are represented by point samples, local restructuring is easily possible, allowing for arbitrarily large distortions from the original shape. Exploiting the flexibility of point-based methods even complicated topology changes, like splashing water, can be simulated.

#### 4. Point-based rendering

The final stage for interactive geometry processing applications is the efficient rendering of point-sampled geometries. Points have first been proposed as universal rendering primitives by Levoy and Whitted [45]. Instead of deriving a rendering algorithm for each geometry representation, they propose to subdivide each representation into a sufficiently dense set of sample points.

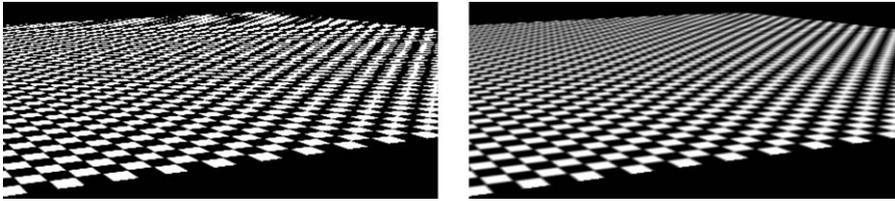


Fig. 9. A checkerboard texture rendered with surface splats: without (left) and with EWA filtering (right).

The rendering of these points can then be performed by one specially tuned technique.

Since we want to generate continuous (i.e. hole-free) images by rendering a discrete set of surface samples, methods for closing the holes and gaps in-between the samples have to be found. This can be done by image-space reconstruction techniques [14,46] or by object-space re-sampling. The techniques from the latter category dynamically adjust the sampling rate, such that the density of projected points meets the pixel resolution. Since this depends on the current viewing parameters, the re-sampling has to be done dynamically for each frame. Examples are dynamic sampling of procedural geometries [47], the randomized  $z$ -buffer [48], and the rendering of MLS surfaces [23,24,49].

In contrast to this, surface splatting [17] renders splats, i.e., object-space disks or ellipses, instead of points only. In this case, the mutual overlap of splats in object-space guarantees a hole-free rendering in image-space. On the other hand, the naïve rendering of interpenetrating splats results in shading discontinuities (cf. Fig. 12). Therefore, Zwicker et al. [17] proposed a high quality anisotropic anti-aliasing method that resembles the anisotropic EWA texture filtering of Heckbert [50]. Each splat is assigned a radially symmetric Gaussian filter kernel, such that a continuous surface signal in object-space is reconstructed by a respectively weighted averaging of splat data (e.g. colors, normals). Combining these object-space reconstruction kernels with a band-limiting image-space filter results in their high quality *EWA splatting* framework (cf. Fig. 9). If both the object-space and the image-space filters are Gaussians, and if the projection is locally (i.e., per splat) approximated by an affine mapping, then these two filters can be combined into one single Gaussian, enabling a quite efficient implementation. Nevertheless, being a purely software-based implementation, this approach is computationally too expensive for displaying highly complex models, since it achieves a splat rate of about 1M splats/s on current hardware.

Botsch et al. [16] therefore proposed the use of hierarchical representations for more efficient rendering. Their octree-based point clouds (see Section 2) can be displayed by rendering either one point or one splat for each non-empty leaf cell. They showed that model-view and projection transformations of a leaf cell center  $\mathbf{p}$  can

be implemented by an incremental summation during a depth-first traversal of the octree, followed by a final de-homogenization step. Since points sharing a common octree path also share terms of these sums, it turns out that each point transformation requires just about 4 scalar additions and 2 divisions in total. Combining this with pre-computed tables for lighting and splat footprints results in a splat rate of 5M splats/s. However, this implementation is still completely software-based and puts a high load on the CPU, blocking it from other geometry processing tasks.

In the last years, the increasing efficiency and programmability of modern graphic cards [51,52] triggered the development of hardware-based splatting methods. Targeting at the efficient visualization of the models acquired during the Digital Michelangelo project [5], Rusinkiewicz and Levoy [53] proposed a hierarchical rendering method based on a pre-computed bounding-sphere tree structure.

As this recursive tree traversal is not efficient enough to keep current GPUs busy, Dachsbacher et al. [54] proposed a sequentialization of the LoD tree structure that corresponds to a breadth-first re-ordering. Rendering the model at a certain level of detail then requires to process a continuous sub-segment of this linear point list directly by the GPU (cf. Fig. 10). Transferring the point list into video memory once and implementing the LoD selection in the vertex shaders of the GPU results in a low CPU load and an impressive splat rate of more than 50M splats/s. The drawback of this approach is that the rendering primitives used for each splat are un-filtered image space squares only, hence the quality could be improved considerably by using Gaussian filtering and elliptical splat primitives.

To remove shading discontinuities and achieve a visually smoother rendering, a Gaussian filter kernel should be assigned to each splat also in the case of hardware-accelerated rendering. The difficulty is that when two splats are projected to the same pixel, only closely overlapping splats should be blended, while in other cases when the  $z$ -distance between the splats is above a certain threshold, the front-most splat should overwrite the splats behind. An implementation of this functionality requires a fuzzy depth test or a more general  $a$ -buffer [55]. Both are not available on current GPUs, therefore multiple rendering passes have to be

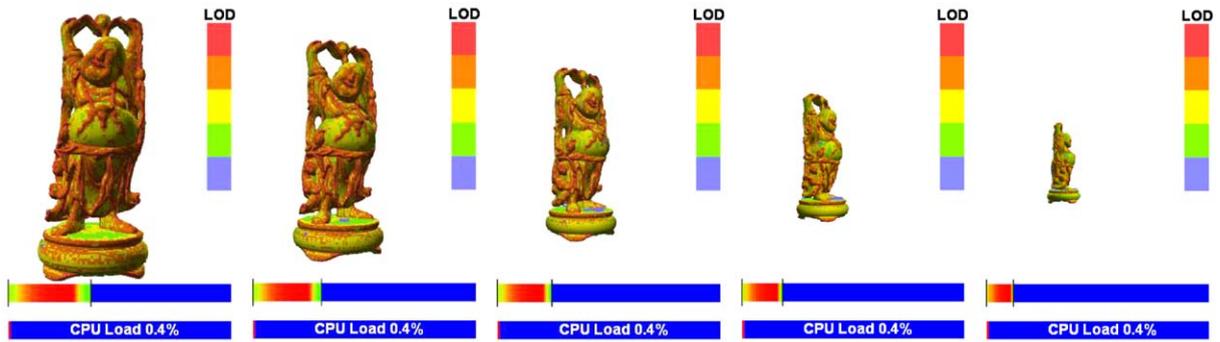


Fig. 10. A sequentialization of a LoD hierarchy results in the rendering of continuous segments of a linear point list. This, and the delegation of fine-grained LoD control to the GPU, results in a hierarchical LoD rendering with almost no CPU load.

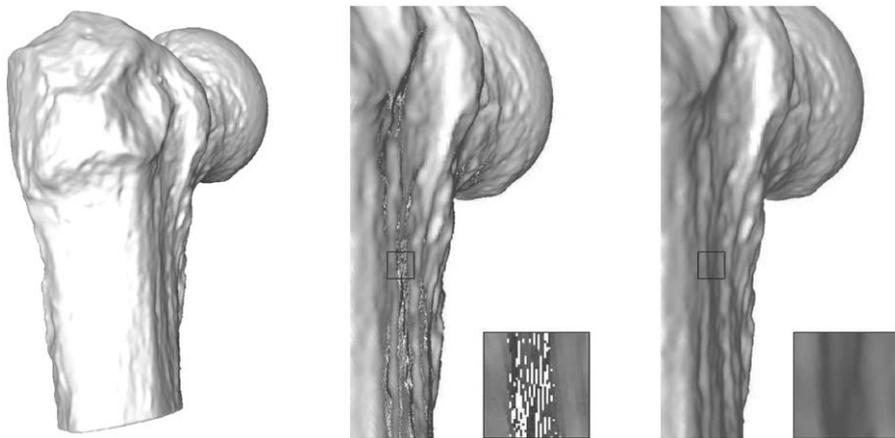


Fig. 11. Previous local affine approximations to the projective mapping caused holes for extreme viewing angles (center). These holes can be closed either by an affine approximation that correctly maps the outer splat contour (right) or by per-pixel ray-casting that exactly computes the inverse projective mapping.

used instead. The first one does so-called *visibility splatting*, i.e., it fills the  $z$ -buffer by rendering (without lighting) all objects slightly shifted away from the viewer by  $\varepsilon$ . The second pass renders all splats with Gaussian blending turned on, but it does not alter the  $z$ -buffer, thereby blending only those splats that differ by less than  $\varepsilon$  in depth [53]. This process accumulates a weighted sum of colors ( $\sum_i \alpha_i rgb_i, \sum_i \alpha_i$ ) in each RGBA pixel, such that a per-pixel division by its accumulated alpha value calculates the correct weighted average. This per-pixel normalization can efficiently be done by rasterizing a window-sized quad that is textured by the result of the second rendering pass [56,57].

Using the pixel shaders of current graphics hardware allows the rasterization of elliptical splats by rendering just one vertex per splat. Computing the projected size in a vertex shader triggers the rasterization of an image space square. A fragment shader processes each of its pixels and constructs the elliptical shape by discarding

pixels outside the ellipse. An implementation using circular object-space splats and two-pass Gaussian filtering was presented by Botsch and Kobbelt [56], achieving a splat rate of 10M splats/s.

However, this method, like most others, uses an affine approximation to the projective mapping in order to compute the splat's footprint in image-space. This simplification may lead to holes in the image, when the model is viewed from extremely flat angles and is shifted away from the main viewing axis (cf. Fig. 11). This problem was addressed first by the perspective accurate splatting of Zwicker et al. [58], using a different affine approximation to the projection, such that the outer splat contour (instead of the splat center) is correctly mapped. Although being slightly incorrect in the splat's interior, this method effectively avoids holes in the image. As the projection is still approximated affinely, the full EWA splatting framework can be integrated into this approach. Per-pixel correct projections can be



Fig. 12. Comparing different rendering techniques for a simplified model of a scanned statue: naïve splatting without filtering is comparable to flat shading (center left), Gaussian blending corresponds to Gouraud shading (center right), Phong splatting yields the same quality as Phong shading for triangle meshes (right).

achieved by performing a local ray casting in the fragment shaders, as proposed by Botsch et al. [59].

Most point-based rendering methods use a constant normal vector for lighting each individual splat, leading to results comparable to flat shading for triangle meshes (cf. Fig. 12, center left). Blending by Gaussian reconstruction kernels smooths out the shading artifacts and is hence comparable to Gouraud shading: the shading varies smoothly, but the image may appear blurry (cf. Fig. 12, center right). The use of non-constant normal fields and per-pixel lighting significantly improves the visual quality [17,60,61]. For triangle meshes, Phong shading achieves superior rendering quality by computing per-pixel lighting based on interpolated normal vectors of the triangle's vertices in object space. This concept is not directly applicable to point-based rendering since access to the normal vectors of neighboring points or splats is not possible due to the lack of explicit connectivity information. In order to achieve per-pixel normals while maintaining the simplicity of splat-based representations, Botsch et al. [59] assign a linear normal field to each splat individually. These normal fields are derived by least squares fitting to a given dense set of sample normals. Rendering such Phong-splats can easily and efficiently be implemented on the GPU and yields high visual quality comparable to Phong shading for triangle meshes (cf. Fig. 12, right).

In order to represent sharp features by point-sampled geometries, Pauly et al. [20] proposed to clip splats against clipping lines defined in their local tangent frames (cf. Fig. 3). This representation can easily be rendered by integrating a per pixel-clipping test into the

fragment shaders, as proposed by Zwicker et al. [58] and Botsch [59].

The presented techniques exploit programmable shaders of current GPUs to make hardware acceleration directly available for point-based rendering. The resulting rendering approaches are very efficient and provide a high visual quality. Since elliptical splats represent an optimal approximation to the local surface geometry and since the  $C^{-1}$  continuity of splats provides further degrees of freedom, splatting techniques lead to superior visual rendering results compared to triangle rendering with the same number of linear primitives (cf. Fig. 13). Nevertheless, since current graphics hardware is highly optimized and specialized for triangle rendering, points or splats still cannot keep up with triangles in terms of effective rendering performance.

## 5. Conclusion

Point-based representations have proven to be a valuable alternative to polygonal meshes in several special applications. In particular when highly complex or dynamic models are to be processed and a proper surface topology is not necessarily required, the simplicity of point-based representations leads to simpler and more efficient algorithms.

Elliptical surface splats provide the same asymptotic approximation power as triangle meshes, but achieve a locally better approximation since the axes of elliptical splats can be aligned to the principal curvature directions of the surface. Moreover, splats offer more

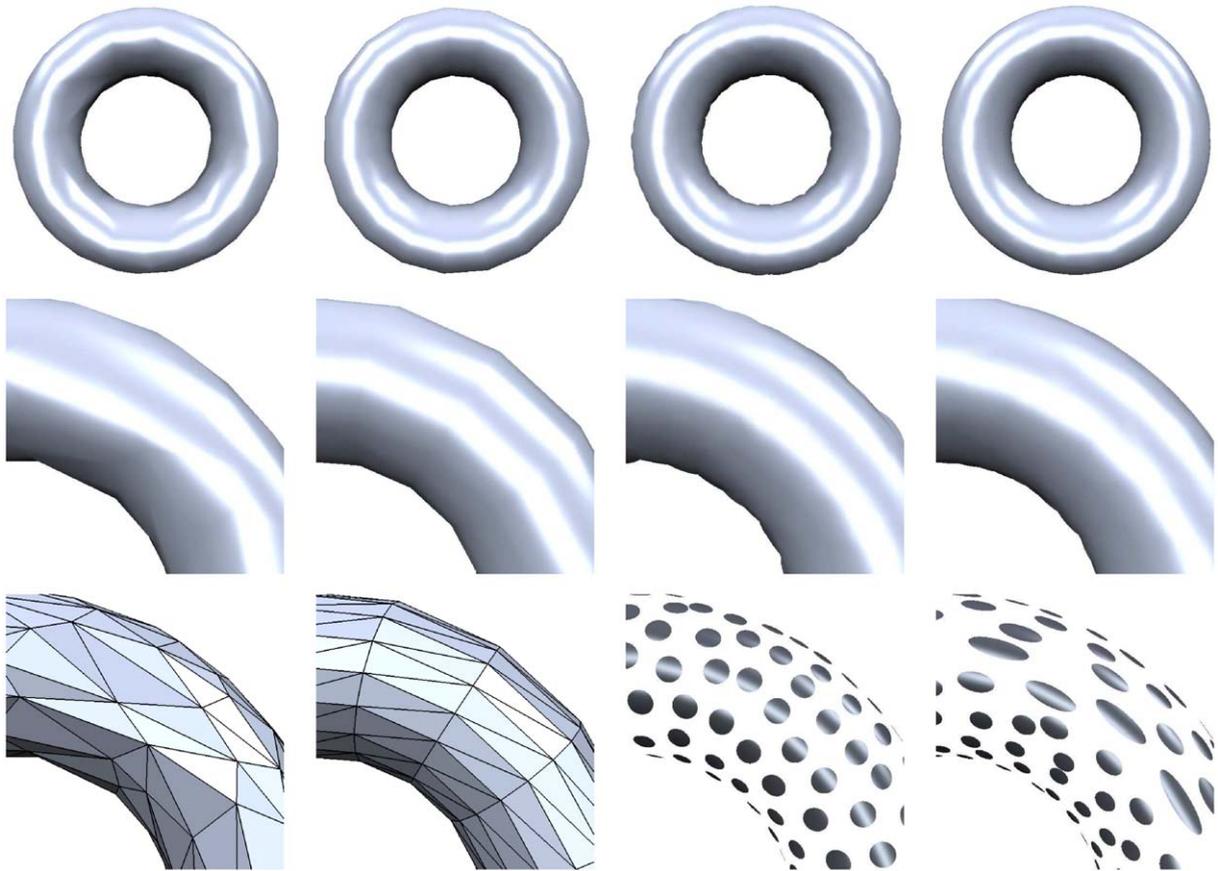


Fig. 13. Comparing different shape approximations consisting of about 730 geometry primitives: irregular triangle meshes (left), regular aligned triangle mesh (center left), circular splats (center right), elliptical splats (right). The respective approximation errors to the continuous torus are 0.64%, 0.58%, 0.20%, and 0.14%.

flexibility since they do not require  $C^0$  continuity. Hence, given a certain budget of surface primitives, splats usually yield better approximation results compared to triangles (cf. Fig. 13). Although current graphics hardware is optimized for triangle rendering, programmable shaders enable the very efficient implementation of high quality point-based rendering techniques whose performance is already coming close to the effective polygon rate.

In the future, new types of graphics hardware architectures could exploit the particular advantages of points as a general 3D graphics primitive by supporting adaptive object space sampling techniques in order to reduce the redundancy of point-based representation and the overdraw rate for screen pixels.

#### Acknowledgements

We thank Mark Pauly and Marc Alexa for helpful comments and the following people for providing us

images from their papers: Matthias Zwicker (Figs. 3, 6, 9), Mark Pauly (Figs. 7, 8), Bart Adams (Fig. 8), Carsten Dachsbacher (Fig. 10), and Jianhua Wu (Fig. 5).

#### References

- [1] Amenta N, Bern M, Kamvysselis M. A new Voronoi-based surface reconstruction algorithm. In: Proceedings of ACM SIGGRAPH 98; 1998. p. 415–21.
- [2] Bischoff S, Pavic D, Kobbelt L. Automatic restoration of polygon models. Preprint. 2004.
- [3] Welch W, Witkin A. Free-form shape design using triangulated surfaces. In: Proceedings of ACM SIGGRAPH 94; 1994. p. 247–56.
- [4] Kobbelt L, Bareuther T, Seidel H-P. Multiresolution shape deformations for meshes with dynamic vertex connectivity. In: Proceedings of Eurographics 00; 2000. p. 249–60.
- [5] Levoy M, Pulli K, Curless B, Rusinkiewicz S, Koller D, Pereira L, Ginzton M, Anderson S, Davis J, Ginsberg J, Shade J, Fulk D. The digital Michelangelo project: 3D

- scanning of large statues. In: Proceedings of ACM SIGGRAPH 00; 2000. p. 131–44.
- [6] Pauly M. Point primitives for interactive modeling and processing of 3D geometry. PhD thesis, ETH Zürich. 2003.
- [7] Zwicker M. Continuous reconstruction, rendering, and editing of point-sampled surfaces. PhD thesis, ETH Zürich. 2003.
- [8] Alexa M, Dachsbacher C, Gross M, Pauly M, van Baar J, Zwicker M. Point-based computer graphics. In: Eurographics 03 Tutorial Notes; 2003.
- [9] Alexa M, Gross M, Pauly M, Pfister H, Stamminger M, Zwicker M. Point-based computer graphics. In: SIGGRAPH 04 Course Notes; 2004.
- [10] Andersson M, Giesen J, Pauly M, Speckmann B. Bounds on the  $k$ -neighborhood for locally uniformly sampled surfaces. In: Proceedings of Symposium on Point-Based Graphics 04; 2004. p. 167–71.
- [11] Linsen L, Prautzsch H. Fan clouds—an alternative to meshes. In: Proceedings of Dagstuhl Seminar on Theoretical Foundations of Computer Vision; 2002. p. 39–57.
- [12] Hoppe H, DeRose T, Duchamp T, McDonald J, Stuetzle W. Surface reconstruction from unorganized points. In: Proceedings of ACM SIGGRAPH 92; 1992. p. 71–8.
- [13] Davis P. Interpolation and Approximation. Dover Publications; 1975.
- [14] Grossman JP, Dally WJ. Point sample rendering. In: Proceedings of Eurographics Workshop on Rendering 98; 1998. p. 181–92.
- [15] Kalaiah A, Varshney A. Statistical point geometry. In: Proceedings of Eurographics Symposium on Geometry Processing 03; 2003. p. 107–15.
- [16] Botsch M, Wiratanaya A, Kobbelt L. Efficient high quality rendering of point sampled geometry. In: Proceedings of Eurographics Workshop on Rendering 02; 2002. p. 53–64.
- [17] Zwicker M, Pfister H, van Baar J, Gross M. Surface splatting. In: Proceedings of ACM SIGGRAPH 01; 2001. p. 371–8.
- [18] Kobbelt L, Botsch M, Schwanecke U, Seidel H-P. Feature sensitive surface extraction from volume data. In: Proceedings of ACM SIGGRAPH 01; 2001. p. 57–66.
- [19] Botsch M, Kobbelt L. Resampling feature and blend regions in polygonal meshes for surface anti-aliasing. In: Proceedings of Eurographics 01; 2001. p. 402–10.
- [20] Pauly M, Keiser R, Kobbelt L, Gross M. Shape modeling with point-sampled geometry. In: Proceedings of ACM SIGGRAPH 03; 2003. p. 641–50.
- [21] Levin D. The approximation power of moving least-squares. *Mathematics of Computation* 1998;67(224): 1517–31.
- [22] Levin D. *Geometric Modeling of Scientific Visualization*. Springer, Ch. Mesh-independent surface interpolation, 2003. p. 37–49.
- [23] Alexa M, Behr J, Cohen-Or D, Fleishman S, Levin D, Silva CT. Point set surfaces. In: Proceedings of IEEE Visualization 01; 2001. p. 21–8.
- [24] Alexa M, Behr J, Cohen-Or D, Fleishman S, Levin D, Silva CT. Computing and rendering point set surfaces. *IEEE Transactions on Visualization and Computer Graphics* 2003;9(1):3–15.
- [25] Pauly M, Gross M, Kobbelt L. Efficient simplification of point-sampled surfaces. In: Proceedings of IEEE Visualization 02; 2002. p. 163–70.
- [26] Alexa M, Adamson A. On normals and projection operators for surfaces defined by point sets. In: Proceedings of Symposium on Point-Based Graphics 04; 2004. p. 149–55.
- [27] Adamson A, Alexa M. Approximating and intersecting surfaces from points. In: Proceedings of Eurographics Symposium on Geometry Processing 03; 2003. p. 245–54.
- [28] Adamson A, Alexa M. Approximating bounded non-orientable surfaces from points. In: Proceedings of Shape Modeling International 04; 2004. p. 243–52.
- [29] Amenta N, Kil Y. Defining point-set surfaces. In: Proceedings of ACM SIGGRAPH 04; 2004. p. 264–70.
- [30] Pauly M, Gross M. Spectral processing of point-sampled geometry. In: Proceedings of ACM SIGGRAPH 01; 2001. p. 379–86.
- [31] Pauly M, Keiser R, Gross M. Multi-scale feature extraction on point-sampled surfaces. In: Proceedings of Eurographics 03; 2003. p. 281–9.
- [32] Weyrich T, Pauly M, Heinzle S, Keiser R, Scandella S, Gross M. Post-processing of scanned 3D surface data. In: Proceedings of Symposium on Point-Based Graphics 04; 2004. p. 85–94.
- [33] Wu J, Kobbelt L. Optimized sub-sampling of point sets for surface splatting. In: Proceedings of Eurographics 04; 2004. p. 643–52.
- [34] Zwicker M, Pauly M, Knoll O, Gross M. Pointshop 3D: an interactive system for point-based surface editing. In: Proceedings of ACM SIGGRAPH 02; 2002. p. 322–9.
- [35] Adams B, Wicke M, Dutré P, Gross M, Pauly M, Teschner M. Interactive 3D painting on point-sampled objects. In: Proceedings of Symposium on Point-Based Graphics 04; 2004. p. 57–66.
- [36] Carr NA, Hart JC. Painting detail. In: Proceedings of ACM SIGGRAPH 04; 2004. p. 845–52.
- [37] Lawrence J, Funkhouser T. A painting interface for interactive surface deformation. In: Proceedings of Pacific Graphics 03; 2003. p. 141–50.
- [38] Szeliski R, Tonnesen D. Surface modeling with oriented particle systems. In: Proceedings of ACM SIGGRAPH 92; 1992. p. 185–94.
- [39] Witkin A, Heckbert P. Using particles to sample and control implicit surfaces. In: Proceedings of ACM SIGGRAPH 94; 1994. p. 269–77.
- [40] Hoffmann CM. *Geometric and solid modeling: an introduction*. Morgan Kaufmann Publishers Inc.; 1989.
- [41] Adams B, Dutré P. Interactive Boolean operations on surfel-bounded solids. In: Proceedings of ACM SIGGRAPH 03; 2003. p. 651–6.
- [42] Adams B, Dutré P. Boolean operations on surfel-bounded solids using programmable graphics hardware. In: Proceedings of Symposium on Point-Based Graphics 04; 2004. p. 19–24.
- [43] Clarenz U, Rumpf M, Telea A. Finite elements on point based surfaces. In: Proceedings of Symposium on Point-Based Graphics 04; 2004. p. 201–11.
- [44] Müller M, Keiser R, Nealen A, Pauly M, Gross M, Alexa M. Point based animation of elastic plastic and melting objects. In: Proceedings of Eurographics/ACM

- SIGGRAPH Symposium on Computer Animation 04; 2004. p. 141–51.
- [45] Levoy M, Whitted T. The use of points as display primitives. Technical Report, CS Department, University of North Carolina at Chapel Hill. 1985.
- [46] Pfister H, Zwicker M, van Baar J, Gross M. Surfels: Surface elements as rendering primitives. In: Proceedings of ACM SIGGRAPH 00; 2000. p. 335–42.
- [47] Stamminger M, Drettakis G. Interactive sampling and rendering for complex and procedural geometry. In: Proceedings of Eurographics Workshop on Rendering 01; 2001. p. 151–62.
- [48] Wand M, Fischer M, Peter I, Meyer auf der Heide F, Straßer W. The randomized z-buffer algorithm: interactive rendering of highly complex scenes. In: Proceedings of ACM SIGGRAPH 01; 2001. p. 361–70.
- [49] Fleishman S, Cohen-Or D, Alexa M, Silva CT. Progressive point set surfaces. *ACM Transactions on Graphics* 2003;22(4).
- [50] Heckbert PS. Fundamentals of texture mapping and image warping. Master's thesis, University of California at Berkeley, 1989.
- [51] Lindholm E, Kilgard M, Moreton H. A user-programmable vertex engine. In: Proceedings of ACM SIGGRAPH 01; 2001. p. 149–58.
- [52] Mark WR, Ghanville RS, Akeley K, Kilgard MJ. Cg: a system for programming graphics hardware in a C-like language. In: Proceedings of ACM SIGGRAPH 03; 2003. p. 896–907.
- [53] Rusinkiewicz S, Levoy M. QSplat: a multiresolution point rendering system for large meshes. In: proceedings of ACM SIGGRAPH 00; 2000. p. 343–52.
- [54] Dachsbacher C, Vogelgsang C, Stamminger M. Sequential point trees. In: Proceedings of ACM SIGGRAPH 03; 2003. p. 657–62.
- [55] Carpenter L. The a-buffer an antialiased hidden surface method. In: Proceedings of ACM SIGGRAPH 84; 1984. p. 103–8.
- [56] Botsch M, Kobbelt L. High-quality point-based rendering on modern GPUs. In: Proceedings of Pacific Graphics 03; 2003. p. 335–43.
- [57] Guennebaud G, Pauling M. Efficient screen space approach for hard-ware accelerated surfel rendering. In: Proceedings of Vision, Modeling, and Visualization 03; 2003. p. 1–10.
- [58] Zwicker M, Räsänen J, Botsch M, Dachsbacher C, Pauly M. Perspective accurate splatting. In: Proceedings of Graphics Interface 04; 2004. p. 247–54.
- [59] Botsch M, Spornat M, Kobbelt L. Phong splatting. In: Proceedings of Symposium on Point-Based Graphics 04; 2004. p. 25–32.
- [60] Kalaiah A, Varshney A. Differential point rendering. In: Proceedings of Eurographics Workshop on Rendering Techniques 01; 2001. p. 139–50.
- [61] Kalaiah A, Varshney A. Modeling and rendering points with local geometry. *IEEE Transactions on Visualization and Computer Graphics* 2003;9(1):30–42.