

Uncertainty and Variability in Point Cloud Surface Data[†]

Mark Pauly

Niloy J. Mitra

Leonidas J. Guibas

Computer Graphics Laboratory, Stanford University, Stanford CA 94305, USA

Abstract

We present a framework for analyzing shape uncertainty and variability in point-sampled geometry. Our representation is mainly targeted towards discrete surface data stemming from 3D acquisition devices, where a finite number of possibly noisy samples provides only incomplete information about the underlying surface. We capture this uncertainty by introducing a statistical representation that quantifies for each point in space the likelihood that a surface fitting the data passes through that point. This likelihood map is constructed by aggregating local linear extrapolators computed from weighted least squares fits. The quality of fit of these extrapolators is combined into a corresponding confidence map that measures the quality of local tangent estimates. We present an analysis of the effect of noise on these maps, show how to efficiently compute them, and extend the basic definition to a scale-space formulation. Various applications of our framework are discussed, including an adaptive re-sampling method, an algorithm for reconstructing surfaces in the presence of noise, and a technique for robustly merging a set of scans into a single point-based representation.

Categories and Subject Descriptors (according to ACM CCS): I.3.5. [Computer Graphics]: Computational Geometry and Object Modeling

1. Introduction

Digital 3D geometry has become ubiquitous in science and will soon be on par with traditional multi-media data types such as sound, images, and video. Mechanical engineering, architecture, entertainment, and bio-medicine are just a few application fields that make extensive use of digital 3D shape information. In these areas, 3D acquisition devices have become a prime source for the creation of 3D geometric data. 3D scanners typically produce an unstructured cloud of samples points, where each point is a discrete sample of certain shape attributes such as 3D position, surface normal, color, or material properties. This raw data needs to be processed in various forms, e.g., to extract high level information about the scanned object, modify its shape or appearance, or create renditions that are meaningful to the user. Most of these geometry processing algorithms are based on continuous surface representations such as triangle meshes or collections of spline-patches, which are typically computed from the

given point cloud data using some surface reconstruction algorithm. Having obtained such a distinct surface model, all subsequent processing then directly operates on this representation without any reference to the origin of the data. This is suitable for applications where convincing renditions of the 3D shapes are the primary goal, such as movies or games. However, single reconstructed surfaces are by no means unique or inherent in the acquired data, since any discrete sampling provides only incomplete information about the underlying object. This shape uncertainty is further increased by measurement noise, which cannot be avoided in any physical acquisition process.

Our goal is to capture this variability and uncertainty in point-sampled surfaces. To this end we propose a new approach to surface modeling with real-world data. Instead of reconstructing a single surface, we look at the distribution of all surfaces that are plausible for a given sample set. We present a statistical representation that takes the measurement and sampling process into account, allowing a more thorough analysis of point cloud surface data. Apart from computer graphics and geometric modeling, this approach can be also beneficial in applied sciences such as engineering or bio-medicine. Whenever certain information needs to

[†] The authors wish to acknowledge support from NSF grants CARGO-0138456 and ITR-0205671, ARO grant DAAD19-03-1-0331, as well as from a Stanford Graduate Fellowship.

be extracted from the acquired data, e.g., the volume of a mechanical part, or the existence of a tumor in a CT-scan, users rely on some confidence or accuracy estimate. Thus current approaches based on a single extracted surface are not suitable for such scenarios.

To capture shape uncertainty in a discrete sample set, we locally estimate the likelihood of a surface passing through a certain point in space. These estimates are computed by propagating weighted shape extrapolators derived from least squares fits of local point neighborhoods. Additional to this *likelihood map*, we also compute a *confidence map* that quantifies the confidence of the individual shape extrapolators. Since confidence is linked to sampling density, we can use this map to guide up- and down-sampling operations on the point cloud data set.

We classify shape uncertainty into two different categories. First, the discrete sampling provides spatial information only at a finite number of points. This introduces uncertainty, since the course of the surface in between sample points is unknown and needs to be inferred from the sample set. Second, physical measurements are always corrupted by noise. Thus the measured position of a sample point cannot be treated as ground truth, but instead should be understood as the result of some stochastic process. Noise typically depends only on the physical properties of the scanner, the acquired object and the measurement environment, while the uncertainty due to discretization occurs even for noise-free data and is closely related to the sampling density. We will demonstrate how both forms of uncertainty can be integrated naturally into the likelihood and confidence maps.

Given the definition of these maps, we will show how this representation can be used to analyze discrete surface data to answer questions related to sampling and discretization. Our method allows us to quantitatively compare surfaces with respect to their quality of fit and extract the most likely surface according to certain boundary constraints. We also show how multiple sample sets of the same object can be combined into a single weighted point cloud exhibiting less uncertainty than each individual scan.

Our framework is general in the sense that we do not assume any additional information on the distribution of shapes. Rather we impose this distribution by accumulating local shape extrapolators. It should be noted that additional context knowledge can greatly reduce the uncertainty and variability in acquired data. For example, if a scanned mechanical part is known to be of a certain type, the search space can typically be described with a few parameters. Our method does not exploit such specific information explicitly. It should be understood as a tool for analyzing shape variability when no prior on the distribution of shapes is given.

2. Related Work

Point-based surface representations have recently become popular in computer graphics. Earlier work introduced point primitives for rendering [LW85] and has inspired a significant amount of work in that direction, e.g., [RL00, ZPvG01, KV01, BWK02].

Curve and surface reconstruction has been an active research field in geometric modeling. Given a point cloud as input these methods typically extract a triangle mesh, e.g., [HDD*94, ABK98], or an implicit representation, e.g., [CBC*01, ZOF01]. More recently, various methods have been presented to directly approximate surfaces from point cloud data [Lev03, ABCO*01, AA03]. Point-sampled surfaces have also been used for geometric processing [PG01], surface re-sampling [WH94, PGK02], shape and appearance modeling [ZPKG02, AD03, PKKG03], and feature extraction [PKG03]. Our work is based on these prior efforts and we use various tools and concepts from the above papers to define our statistical shape modeling framework. We will comment on these techniques in subsequent sections.

Kalaiah and Varshney [KV03] introduced a new representation that uses statistical methods for compression and stochastic rendering of point cloud data sets. They use hierarchical PCA to compactly encode point attributes, such as position, normal, and color. Grigoryan and Rheingans [GR02] presented a point rendering method for visualizing stochastic variations in medical data using uncertainty data provided with the point samples. Schneider analyzed shape uncertainty from a more abstract point of view in [Sch01]. He identifies various sources for shape uncertainty and stresses the importance of additional context information to reduce the uncertainty.

Our work is probably most closely related to *tensor voting*, a formalism introduced by Medioni and co-workers that is based on tensor calculus (see [MLT00] for an overview). This approach is similar to ours in that it tries to describe shape information by combining local estimates using Gaussian influence functions. The scope of their work is quite different though, since they mainly concentrate on robustly classifying and reconstructing features in discrete data sets.

3. Likelihood and Confidence

In this section we introduce our statistical framework for modeling shape variability and uncertainty. Our goal is to explore the space of all continuous surfaces that are compatible with a given point set. We first consider the case of noise-free data sets, where by compatible we mean interpolating. The analysis in the presence of noise will then be given in Section 4.

Assume that M_P is the set of all continuous $(d - 1)$ -dimensional surfaces interpolating a given point cloud $P =$



Figure 1: A given point cloud (left image) could be a sample from any of an infinite number of surfaces. In the middle, four such surfaces are shown, where the gray value indicates the prior, i.e., likelihood that the point cloud is sampled from this surface. The accumulated likelihood for all surfaces as computed with our method is shown on the right.

$\{\mathbf{p}_1, \dots, \mathbf{p}_N | \mathbf{p}_i \in \mathbb{R}^d\}$, i.e., $\mathbf{p}_i \in S$ for all $\mathbf{p}_i \in P$ and all $S \in M_P$. To analyze the distribution of surfaces in M_P , we define a function $F_P : \mathbb{R}^d \rightarrow \mathbb{R}_+$ that quantifies for each point $\mathbf{x} \in \mathbb{R}^d$ the likelihood that a surface S interpolating P passes through \mathbf{x} . Conceptually we can define a likelihood map F_P as

$$F_P(\mathbf{x}) = \int_{S \in M_P} \chi_S(\mathbf{x}) p(S) dS, \quad (1)$$

where $\chi_S(\mathbf{x})$ is the characteristic function of S , i.e.,

$$\chi_S(\mathbf{x}) = \begin{cases} 1 & \mathbf{x} \in S \\ 0 & \mathbf{x} \notin S \end{cases}$$

and $p(S)$ is a weight function that specifies a prior on the distribution of surfaces $S \in M_P$ (see Figure 1). For certain applications it might be possible to explicitly define such a prior and describe the set M_P with a finite number of parameters. In general, however, Equation 1 is intractable, since M_P is infinite-dimensional and the prior is not known. We thus follow a constructive approach for defining the likelihood map F_P . This means that by defining F_P , we implicitly specify M_P and p , i.e., impose a prior on the distribution of surfaces interpolating P .

3.1. Likelihood Map

To determine $F_P(\mathbf{x})$ for a certain \mathbf{x} , we accumulate local fitting estimates $F_i(\mathbf{x})$ from each \mathbf{p}_i . $F_i(\mathbf{x})$ measures the likelihood of a linear extrapolation from \mathbf{p}_i to \mathbf{x} , given the spatial distribution of the sample points in P . This likelihood is derived from a weighted sum of squared distances from the points in P . Let $\mathbf{q}_i(\mathbf{x}) = (\mathbf{x} - \mathbf{p}_i) / \|\mathbf{x} - \mathbf{p}_i\|$ the normalized direction vector from \mathbf{p}_i to \mathbf{x} , ϕ_i a monotonically decreasing weight function, c_i a normalization constant, and

$\mathbf{p}_{ij} = \mathbf{p}_i - \mathbf{p}_j$. $F_i(\mathbf{x})$ can then be computed as

$$\begin{aligned} F_i(\mathbf{x}) &= \frac{1}{c_i} \sum_{j=1}^N (\mathbf{p}_{ij}^T \mathbf{q}_i(\mathbf{x}))^2 \phi_i(\|\mathbf{p}_{ij}\|) \\ &= \frac{1}{c_i} \sum_{j=1}^N \mathbf{q}_i(\mathbf{x})^T \mathbf{p}_{ij} \mathbf{p}_{ij}^T \mathbf{q}_i(\mathbf{x}) \phi_i(\|\mathbf{p}_{ij}\|) \\ &= \frac{1}{c_i} \mathbf{q}_i(\mathbf{x})^T \left(\sum_{j=1}^N \mathbf{p}_{ij} \mathbf{p}_{ij}^T \phi_i(\|\mathbf{p}_{ij}\|) \right) \mathbf{q}_i(\mathbf{x}) \\ &= \frac{1}{c_i} \mathbf{q}_i(\mathbf{x})^T \mathbf{C}_i \mathbf{q}_i(\mathbf{x}), \end{aligned} \quad (2)$$

where

$$\mathbf{C}_i = \sum_{j=1}^N \mathbf{p}_{ij} \mathbf{p}_{ij}^T \phi_i(\|\mathbf{p}_{ij}\|). \quad (3)$$

The normalization constant c_i can be computed efficiently as

$$c_i = \int_{S^d} \mathbf{q}^T \mathbf{C}_i \mathbf{q} d\mathbf{q} = \frac{\pi}{2} \sum_{j=1}^d \lambda_j^l = \frac{\pi}{2} \text{tr}(\mathbf{C}_i), \quad (4)$$

where S^d is the d -dimensional sphere of directions, \mathbf{q} is a unit direction vector, λ_j^l is the l -th eigenvalue of \mathbf{C}_i , and $\text{tr}(\mathbf{C}_i)$ denotes the trace of \mathbf{C}_i . Since \mathbf{C}_i is symmetric and positive semi-definite, all eigenvalues are non-negative and the corresponding eigenvectors \mathbf{v}_i^l span an orthonormal frame. The quadratic form defined by Equation 2 defines an ellipsoid with principal axes $\lambda_j^l \mathbf{v}_i^l$ that describes the distribution of points in the neighborhood of \mathbf{p}_i . The likelihood F_i is then simply the weighted least squares error of the $(d-1)$ -dimensional sub-space that is orthogonal to the line defined by $\mathbf{q}_i(\mathbf{x})$ (see Figure 2).

Each $F_i(\mathbf{x})$ measures the likelihood that a surface passes through \mathbf{x} from the point of view of \mathbf{p}_i . Combining these local estimates then yields an expression for the likelihood map F_P :

$$F_P(\mathbf{x}) = \sum_{i=1}^N F_i(\mathbf{x}) \phi_i(\|\mathbf{x} - \mathbf{p}_i\|). \quad (5)$$

Fitting estimates closer to \mathbf{x} will be assigned a higher weight than those that are far away from the point of interest. Effectively, we make the assumption that the influence of a point \mathbf{p}_i on the course of the surface diminishes with increasing distance to \mathbf{p}_i . To capture this behavior we use a radial Gaussian influence function ϕ_i with standard deviation σ_i both in Equation 2 and Equation 5. Figure 3 shows an example of a 2D likelihood map.

Note that a direct computation of the F_i needs order $O(N)$ operations. Thus a single evaluation of Equation 5 requires order $O(N^2)$ computation. However, the matrix \mathbf{C}_i is constant as it only depends on P , not on \mathbf{x} . Thus we can precompute all the \mathbf{C}_i 's and use the quadratic form of Equation 2 to evaluate F_i in constant time. Since the Gaussian weight function drops rapidly with distance, the computa-




Figure 2: Construction of the likelihood map. Left: Fitting estimates are computed from weighted least squares fits, right: Combining individual fitting estimates yields the global likelihood map.




Figure 3: Likelihood map in 2D, where red means high likelihood and blue indicates low likelihood for a surface passing through a specific point in space. The images on the right show the linear fitting estimate of the point centered in the black box. Top right, fitting estimate $F_i(\mathbf{x})$, bottom right, weighted fitting estimate $F_i(\mathbf{x})\phi_i(\|\mathbf{x} - \mathbf{p}_i\|)$.

tional cost can be further reduced by only considering a local neighborhood around the point of interest.

3.2. Confidence Map

Equation 5 defines the likelihood map by combining fitting estimates from all points in the point cloud. To evaluate the confidence of the fitting estimate at point \mathbf{p}_i we look at the distribution of all linear fits passing through \mathbf{p}_i . From equation 2 we can derive a confidence estimate for F_i by looking at the distribution of the eigenvalues of \mathbf{C}_i . In particular, the ratio $\bar{\lambda}_i = \lambda_i^1 / \sum_l \lambda_i^l$, where λ_i^1 is the smallest eigenvalue of \mathbf{C}_i , quantifies the quality of fit of a linear approximation at \mathbf{p}_i . A perfect fit means $\bar{\lambda}_i = 0$ and thus a high confidence in the estimate at \mathbf{p}_i . If $\bar{\lambda}_i$ reaches its maximum value of $1/d$, all directions are equally likely, indicating a low confidence at \mathbf{p}_i . We combine these individual confidence estimates into a global confidence map C_P using the same weighting scheme as in Equation 5:

$$C_P(\mathbf{x}) = \sum_{i=1}^N \bar{\lambda}_i \phi_i(\|\mathbf{x} - \mathbf{p}_i\|). \quad (6)$$

Note that the confidence map is directly related to the surface variation measure of [PGK02] and the sampling criterion proposed by Adamson and Alexa [AA03]. Figure 4 shows the confidence maps for a point cloud in 2D. Observe how the quality of the normal estimates is directly related to the confidence values.

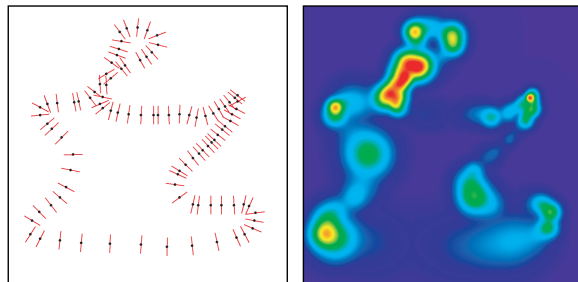


Figure 4: Normal estimates of a point set (left) and corresponding confidence map (right). Red color indicates low confidence, i.e., low preference for a specific normal direction.




Figure 5: Likelihood (middle) and confidence maps (right) for a sparsely sampled 3D data set (left).

Note that the likelihood and confidence maps only depend on relative distances between sample points and are thus invariant under similarity transforms. Since they are also defined for arbitrary dimension, 3D data sets can be processed in the same way as the above 2D examples. Figure 5 shows slices through the 3D likelihood and uncertainty maps for a point cloud in 3D.

4. Noise

Measurement noise is the second source of uncertainty that we encounter in discrete data. To define the likelihood and confidence maps in the presence of noise, we consider the point cloud P as the result of a stochastic process. We assume that each sample point \mathbf{p}_i is corrupted by zero-mean, additive noise $\xi_i \in \mathbb{R}^d$, where $g_i(\xi_i)$ denotes the probability density function of ξ_i and Ω_i the corresponding covariance matrix. Let $\mathbf{p}_i^\xi = \mathbf{p}_i + \xi_i$, $\xi_{ij} = \xi_i - \xi_j$, $\mathbf{p}_{ij}^\xi = \mathbf{p}_{ij} + \xi_{ij}$,

and $P^\xi = \{\mathbf{p}_1^\xi, \dots, \mathbf{p}_N^\xi\}$. For independently distributed ξ_i we compute the expected value of a function X_P that depends on the position of the points in P as

$$E[X_P] = \int_{\mathcal{R}} X_{P^\xi} g(\xi) d\xi,$$

where $\mathcal{R} = \mathbb{R}^d \times \dots \times \mathbb{R}^d$, $\xi = (\xi_1, \dots, \xi_N)$, and $g(\xi) = \prod_i g_i(\xi_i)$. The likelihood map for noisy data is then given as $F_P^\xi(\mathbf{x}) = E[F_P(\mathbf{x})]$ and the confidence map as $C_P^\xi(\mathbf{x}) = E[C_P(\mathbf{x})]$. To evaluate these expected values we make the assumption that the noise is small, i.e., in the range of the local sample spacing, so that the Gaussian distance weights of Equation 2 can be approximated by constants $\phi_{ij}^\xi = \phi_i \left(\sqrt{E[\|\mathbf{p}_{ij}^\xi\|^2]} \right)$. Using the fact that the samples are independent, we find that

$$\begin{aligned} E[\|\mathbf{p}_{ij}^\xi\|^2] &= E[(\mathbf{p}_{ij}^\xi)^T \mathbf{p}_{ij}^\xi] = E[(\mathbf{p}_{ij} + \xi_{ij})^T (\mathbf{p}_{ij} + \xi_{ij})] \\ &= E[\mathbf{p}_{ij}^T \mathbf{p}_{ij}] + E[\xi_{ij}^T \xi_{ij}] \\ &= \mathbf{p}_{ij}^T \mathbf{p}_{ij} + E[\xi_i^T \xi_i] + E[\xi_j^T \xi_j] \\ &= \mathbf{p}_{ij}^T \mathbf{p}_{ij} + \text{tr}(\Omega_i) + \text{tr}(\Omega_j), \end{aligned} \quad (7)$$

and hence

$$\phi_{ij}^\xi = \phi_i \left(\sqrt{\mathbf{p}_{ij}^T \mathbf{p}_{ij} + \text{tr}(\Omega_i) + \text{tr}(\Omega_j)} \right).$$

The mean matrix \mathbf{C}_i^ξ for the fitting estimates of Equation 2 can thus be written as

$$\begin{aligned} \mathbf{C}_i^\xi &= E \left[\sum_{j=1}^N (\mathbf{p}_{ij} + \xi_{ij}) (\mathbf{p}_{ij} + \xi_{ij})^T \phi_{ij}^\xi \right] \\ &= \sum_{j=1}^N \mathbf{p}_{ij} \mathbf{p}_{ij}^T \phi_{ij}^\xi + \sum_{j=1}^N E[\xi_{ij} \xi_{ij}^T] \phi_{ij}^\xi \\ &= \bar{\mathbf{C}}_i + \sum_{j=1}^N E[\xi_i \xi_i^T] \phi_{ij}^\xi \sum_{j=1}^N E[\xi_j \xi_j^T] \phi_{ij}^\xi \\ &= \bar{\mathbf{C}}_i + \Omega_i \sum_{j=1}^N \phi_{ij}^\xi + \sum_{j=1}^N \Omega_j \phi_{ij}^\xi, \end{aligned} \quad (8)$$

where $\bar{\mathbf{C}}_i$ is computed as in Equation 2 using ϕ_{ij}^ξ instead of $\phi_i(\|\mathbf{p}_{ij}\|)$ as distance weights. As before, the normalization constant is given as $c_i^\xi = \frac{\pi}{2} \text{tr}(\mathbf{C}_i^\xi)$. We can then write

$$F_i^\xi(\mathbf{x}) = \frac{1}{c_i^\xi} \frac{(\mathbf{p}_i^\xi - \mathbf{x}) \mathbf{C}_i^\xi (\mathbf{p}_i^\xi - \mathbf{x})}{\|(\mathbf{p}_i^\xi - \mathbf{x})\|^2}. \quad (9)$$




Figure 6: The effect of noise on the likelihood map. Left: Input point cloud, middle: Likelihood map without noise, right: Likelihood map with identically distributed Gaussian noise.

Using this expression we compute the expected likelihood map as

$$\begin{aligned} F_P^\xi(\mathbf{x}) &= \int_{\mathcal{R}} \sum_{i=1}^N \frac{1}{c_i^\xi} \mathbf{q}_i^{\xi T} \mathbf{C}_i^\xi \mathbf{q}_i^\xi g(\xi) d\xi \\ &= \sum_{i=1}^N \int_{\mathcal{R}} \frac{1}{c_i^\xi} \frac{(\mathbf{p}_i + \xi_i - \mathbf{x})^T \mathbf{C}_i^\xi (\mathbf{p}_i + \xi_i - \mathbf{x})}{\|(\mathbf{p}_i + \xi_i - \mathbf{x})\|^2} g_i(\xi_i) d\xi_i \\ &= \sum_{i=1}^N \int_{\mathcal{R}} \frac{(\mathbf{p}_i - (\mathbf{x} - \xi_i))^T \mathbf{C}_i^\xi (\mathbf{p}_i - (\mathbf{x} - \xi_i))}{c_i^\xi \|(\mathbf{p}_i - (\mathbf{x} - \xi_i))\|^2} g_i(\xi_i) d\xi_i \\ &= \sum_{i=1}^N \int_{\mathcal{R}} F_i^\xi(\mathbf{x} - \xi_i) g_i(\xi_i) d\xi_i \\ &= \sum_{i=1}^N F_i^\xi(\mathbf{x}) \otimes g_i(\mathbf{x}), \end{aligned} \quad (10)$$

where \otimes denotes the convolution operator. A similar derivation holds for the confidence map. To incorporate the noise model described above into our framework, we thus only need to adjust the matrix \mathbf{C}_i for each fitting estimate according to Equation 8, and apply a convolution operation to the final maps as described in Equation 10. Figure 6 shows the effect of noise on the likelihood map.

5. Filter Kernels and Scale-Space

In the most simple case, all ϕ_i have the same standard deviation σ_i so that a range query with fixed radius can be used to compute the local neighborhoods. It has been observed previously, however, that globally invariant weight functions are unsuitable for data sets with spatially varying sampling density [PGK02]. We thus use an adaptive Gaussian weight function

$$\phi_i(\mathbf{x}) = e^{-\|\mathbf{x} - \mathbf{p}_i\|^2 / \sigma_i^2}, \quad (11)$$

where the kernel radius σ_i is related to the local sampling density as $\sigma_i = \sigma \cdot \eta_i$. The variable η_i denotes the local sample spacing estimated from a k -neighborhood as described in [PKKG03], and σ is a global scale parameter. Similar to linear scale-space formulations, where a convolution with a gaussian of varying kernel width leads to a multi-scale representation of a given function f , σ can be understood as a

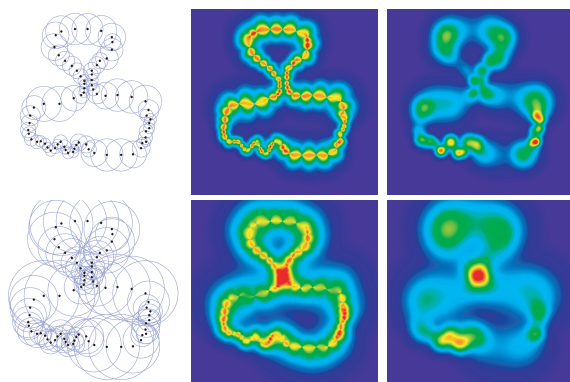


Figure 7: Likelihood and confidence at different scales. The circles in the left images show the iso-value 0.1 of the Gaussian weight functions.

scale-parameter for F_p . The effect of different choices for this scale parameter can be observed in Figure 7. This example shows that likelihood and confidence strongly depend on scale. While the narrowing part of the curve can be robustly resolved on a small scale, increasing the scale leads to increased uncertainty in that area. On the other hand, the larger kernels better handle the noisy sections of the point cloud. Thus the scale-space representation can be useful for estimating the optimal neighborhood size for point-based surface modeling (see, e.g., [MN03]). A similar approach has also been used for multi-scale feature extraction in [PKG03].

6. Results and Applications

This sections shows various applications for the shape uncertainty framework defined above. For ease of illustration, all examples are given on 2D data sets. As mentioned before, the extension to 3D is straightforward (see also Figure 5).

Visualizations of the likelihood and uncertainty maps can be of immediate use in interactive scanning applications, where the user controls a scanning process by interactively adjusting the position of the scanned object. Apart from indicating holes as in [RHHL02], this method also directs the user to regions of high uncertainty, which indicate insufficient sampling. We can also compare two given surfaces in terms of their quality of fit with respect to a given point cloud by evaluating the integral $L(S) = \frac{1}{|S|} \int_S F_p(\mathbf{x}) d\mathbf{x}$. This value can be understood as the likelihood that a point cloud P has been sampled from a surface S .

6.1. Surface Re-Sampling

One of the most fundamental geometric processing methods is surface re-sampling. The confidence map defined in Equation 4 can be used to guide re-sampling operations both




Figure 8: Adaptive surface sampling. A given continuous surface has been successively up-sampled by inserting points in regions of low confidence. The bottom row shows the confidence map of the corresponding point clouds shown in the top row.

for up- and down-sampling. It is particularly suited for iterative point removal or insertion algorithms, where it can be used to determine the importance of a point for the approximation of a particular surface. Points should be removed in regions of high confidence, while points should be inserted where the confidence is low. Similar to previous methods [PGK02, ABCO*01, Lin01], the error function, i.e., the confidence map, can be updated efficiently after such incremental operations. One advantage of our approach is that the resulting sampling distribution is not only curvature adaptive, but also concentrates more samples in regions where two distinct sheets of the surface come together. This means that subsequent point-based surface processing based on k -nearest neighbors can be robustly performed on the re-sampled data sets. Figure 8 shows an example of adaptive surface re-sampling.

6.2. Combining Surface Scans

Complex geometry is typically acquired using multiple overlapping scans, each covering a part of the model surface. Various algorithms have been proposed to merge a set of scans into a consistent representation, e.g., [CL96]. Typically, these methods apply some blending operator to combine sample points in regions of overlap. We propose a different method that creates a new point cloud by simply merging a set of given point clouds. However, the samples in the combined data set are enhanced by fidelity weights that are directly related to the confidence estimates obtained from each individual point cloud. This method is illustrated in Figure 9. As shown in the bottom row on the left, the reconstruction without fidelity weights exhibits severe artifacts. These are due to false normal estimates caused by the noise in the data. Increasing the radius of the reconstruction kernel can avoid these artifacts, but leads to substantial blurring




Figure 9: Combining point clouds. In the top row on the left two input data sets are shown that are corrupted by noise in different regions of the surface. The combined point cloud is shown in the third column and the weighted combined point cloud on the right (gray level indicates fidelity weight). The second and third rows show the corresponding likelihood and confidence maps, respectively. The bottom row shows three reconstructions using weighted least squares approximation as proposed in [AA03]. Left and middle: Reconstruction without fidelity weights using different reconstruction kernel widths. Right: Reconstruction with fidelity weights, using the same reconstruction kernel as in the left image.

of surface features (middle). In contrast, the reconstruction using the fidelity weights gives a stable reconstruction even for small kernel sizes that preserves salient features (right).

6.3. Surface Reconstruction

The likelihood map can also be used for surface reconstruction. We have implemented a scheme based on geodesic active contours [CKS97] that evolves an implicitly defined surface under geodesic flow defined on the likelihood map. This method tries to approximate the "most likely" surface, while at the same time ensuring certain smoothness properties of the resulting surfaces. Figure 10 shows the result of this algorithm on a noisy 2D data set. More details on geodesic active contours can be found in [CKS97].




Figure 10: Curve reconstruction on noisy data using geodesic active contours. Left: Input point cloud, middle: Corresponding likelihood map, right: Evolving curve, starting from the circle shown in the center.

7. Conclusion and Future Work

We have introduced a statistical framework for analyzing discrete surface data represented by clouds of point samples. We show that uncertainty due to both discretization and noise can be incorporated efficiently into a single representation. This representation allows us to visualize uncertainty and variability in acquired data sets, perform re-sampling and surface reconstruction operations, and merge multiple sample sets into a single point cloud.

It is important to note that the construction of the likelihood and confidence maps using weighted least squares fits defines the prior on the distribution of surfaces compatible with a given point cloud, taking only the spatial information provided by the point samples into account. This generality, while advantageous when analyzing data sets from different acquisition sources, also limits the applicability for specific applications scenarios. A significantly more accurate analysis of shape variability should be possible when considering context specific prior information about the underlying shape space. In the future we plan to extend our scheme to integrate context information into the definition of the likelihood and confidence maps. Another interesting direction for future research is the extension of our framework to analyze the variability within a family of shapes. We are also investigating a statistical classification of surface topology from point cloud data.

References

- [AA03] ADAMSON A., ALEXA M.: Approximating and intersecting surfaces from points. In *Proceedings of the Eurographics/ACM SIGGRAPH symposium on Geometry processing* (2003), Eurographics Association, pp. 230–239.
- [ABCO*01] ALEXA M., BEHR J., COHEN-OR D., FLEISHMAN S., LEVIN D., SILVA C. T.: Point set surfaces. In *Proceedings of the conference on Visualization '01* (2001), pp. 21–28.
- [ABK98] AMENTA N., BERN M., KAMVYSSELIS M.: A new Voronoi-based surface reconstruction algorithm. In *Computer Graphics* (1998), vol. 32, pp. 415–421.

- [AD03] ADAMS B., DUTRE P.: Interactive boolean operations on surfel-bounded solids. In *ACM Trans. Graph.* (2003), vol. 22, ACM Press, pp. 651–656.
- [BWK02] BOTSCH M., WIRATANAYA A., KOBBELT L.: Efficient high quality rendering of point sampled geometry. In *Proceedings of the 13th Eurographics workshop on Rendering* (2002), Eurographics Association, pp. 53–64.
- [CBC*01] CARR J. C., BEATSON R. K., CHERRIE J. B., MITCHELL T. J., FRIGHT W. R., MCCALLUM B. C., EVANS T. R.: Reconstruction and representation of 3d objects with radial basis functions. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques* (2001), ACM Press, pp. 67–76.
- [CKS97] CASELLES V., KIMMEL R., SAPIRO G.: Geodesic active contours. In *International Journal of Computer Vision* (1997), vol. 22, Kluwer Academic Publishers, pp. 61–79.
- [CL96] CURLESS B., LEVOY M.: A volumetric method for building complex models from range images. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques* (1996), ACM Press, pp. 303–312.
- [GR02] GRIGORYAN G., RHEINGANS P.: Probabilistic surfaces: Point based primitives to show surface uncertainty. In *Proceedings of the conference on Visualization '02* (2002), pp. 147–154.
- [HDD*94] HOPPE H., DE ROSE T., DUCHAMP T., HALSTEAD M., JIN H., McDONALD J., SCHWEITZER J., STUETZLE W.: Piecewise smooth surface reconstruction. In *Computer Graphics* (1994), vol. 28, pp. 295–302.
- [KV01] KALAI AH A., VARSHNEY A.: Differential point rendering. In *Proceedings of the 12th Eurographics Workshop on Rendering Techniques* (2001), Springer-Verlag, pp. 139–150.
- [KV03] KALAI AH A., VARSHNEY A.: Statistical point geometry. In *Proceedings of the Eurographics/ACM SIGGRAPH symposium on Geometry processing* (2003), Eurographics Association, pp. 107–115.
- [Lev03] LEVIN D.: Mesh-independent surface interpolation. In *Geometric Modeling for Scientific Visualization* (2003), Springer-Verlag, pp. 37–50.
- [Lin01] LINSEN L.: *Point Cloud Representation*. Tech. rep., Faculty of Computer Science, University of Karlsruhe, 2001.
- [LW85] LEVOY M., WHITTED T.: *The use of points as display primitives*. Tech. rep., The University of North Carolina at Chapel Hill, Department of Computer Science, 1985.
- [MLT00] MEDIONI G., LEE M.-S., TANG C.-K.: *A Computational Framework for Segmentation and Grouping*. Elsevier, 2000.
- [MN03] MITRA N. J., NGUYEN A.: Estimating surface normals in noisy point cloud data. In *Proceedings of the nineteenth conference on Computational geometry* (2003), ACM Press, pp. 322–328.
- [PG01] PAULY M., GROSS M.: Spectral processing of point-sampled geometry. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques* (2001), ACM Press, pp. 379–386.
- [PGK02] PAULY M., GROSS M., KOBBELT L. P.: Efficient simplification of point-sampled surfaces. In *Proceedings of the conference on Visualization '02* (2002), pp. 163–170.
- [PKG03] PAULY M., KEISER R., GROSS M.: Multi-scale feature extraction on point-sampled surfaces. In *Computer Graphics Forum* (2003), vol. 22, pp. 281–289.
- [PKKG03] PAULY M., KEISER R., KOBBELT L. P., GROSS M.: Shape modeling with point-sampled geometry. In *ACM Transactions on Graphics (TOG)* (2003), vol. 22, ACM Press, pp. 641–650.
- [RHHL02] RUSINKIEWICZ S., HALL-HOLT O., LEVOY M.: Real-time 3d model acquisition. In *Proceedings of the 29th annual conference on Computer graphics and interactive techniques* (2002), ACM Press, pp. 438–446.
- [RL00] RUSINKIEWICZ S., LEVOY M.: Qsplat: a multi-resolution point rendering system for large meshes. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques* (2000), ACM Press/Addison-Wesley Publishing Co., pp. 343–352.
- [Sch01] SCHNEIDER B.: On the uncertainty of local form of lines and surfaces. In *Cartography and Geographic Information Science* (2001), vol. 28, pp. 237–247.
- [WH94] WITKIN A. P., HECKBERT P. S.: Using particles to sample and control implicit surfaces. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques* (1994), ACM Press, pp. 269–277.
- [ZOF01] ZHAO H.-K., OSHER S., FEDKIW R.: Fast surface reconstruction using the level set method. In *Proceedings of the IEEE Workshop on Variational and Level Set Methods (VLSM'01)* (2001), IEEE Computer Society, pp. 194–202.
- [ZPKG02] ZWICKER M., PAULY M., KNOLL O., GROSS M.: Pointshop 3d: an interactive system for point-based surface editing. In *Proceedings of the 29th annual conference on Computer graphics and interactive techniques* (2002), ACM Press, pp. 322–329.
- [ZPvG01] ZWICKER M., PFISTER H., VAN BAAR J., GROSS M.: Surface splatting. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques* (2001), ACM Press, pp. 371–378.