A Benchmark for Surface Reconstruction

MATTHEW BERGER

Air Force Research Laboratory, Information Directorate

JOSHUA A. LEVINE

Clemson University

LUIS GUSTAVO NONATO

Universidade de São Paulo

GABRIEL TAUBIN

Brown University

and

CLAUDIO T. SILVA

Polytechnic Institute of New York University

We present a benchmark for the evaluation and comparison of algorithms which reconstruct a surface from point cloud data. Although a substantial amount of effort has been dedicated to the problem of surface reconstruction, a comprehensive means of evaluating this class of algorithms is noticeably absent. We propose a simple pipeline for measuring surface reconstruction algorithms, consisting of three main phases: surface modeling, sampling, and evaluation. We use implicit surfaces for modeling shapes which are capable of representing details of varying size and sharp features. From these implicit surfaces, we produce point clouds by synthetically generating range scans which resemble realistic scan data produced by an optical triangulation scanner. We validate our synthetic sampling scheme by comparing against scan data produced by a commercial optical laser scanner, where we scan a 3D-printed version of the original surface. Last, we perform evaluation by comparing the output reconstructed surface to a dense uniformly distributed sampling of the implicit surface. We decompose our benchmark into two distinct sets of experiments. The first set of experiments measures reconstruction against point clouds of complex shapes sampled under a wide variety of conditions. Although these experiments are quite useful for comparison, they lack a fine-grain analysis. To complement this, the second set of experiments measures specific properties of surface reconstruction, in terms of sampling characteristics and surface features. Together, these experiments depict a detailed examination of the state of surface reconstruction algorithms.

1. INTRODUCTION

Over the past two decades there has been an immense amount of effort dedicated to the problem of surface reconstruction. The problem of surface reconstruction may be formulated as follows: given a sampling of points measured on a surface, recover the *original* surface from which these points came. The faithful representation of real-world objects has a long history in computer graphics and other fields such as cultural heritage [Levoy et al. 2000; Funkhouser et al. 2011] and urban simulation [Frueh et al. 2005].

The generality of the problem has given rise to a wide variety of surface reconstruction algorithms. The algorithms primarily differ by the type of input point data and output reconstructed surface. The input may be a single depth image, a registered point cloud, or a registered point cloud equipped with normals. Moreover, the modality of the point data plays a major role in reconstruction, where various modalities from the 3D vision literature include
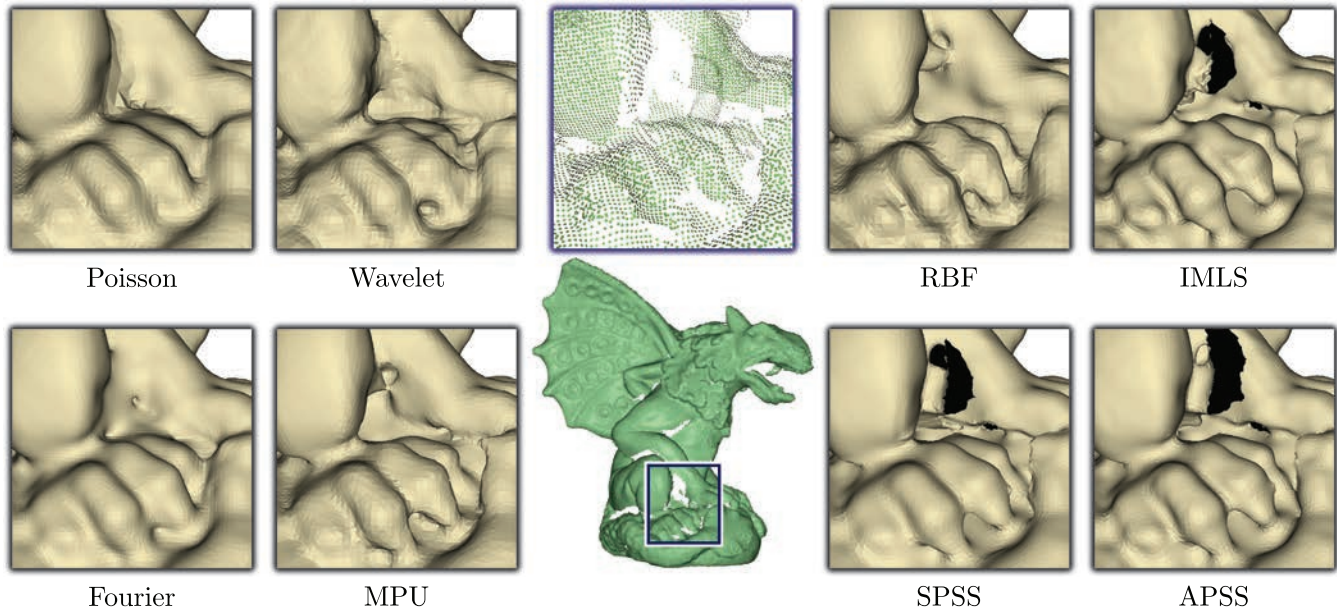
20

Fig. 1. Here we have synthetically sampled the Gargoyle model, and ran eight separate reconstruction algorithms on this point cloud. Note the differences between the algorithms on the claw, where some algorithms oversmooth the data, while others result in spurious holes being produced. Our benchmark aims to generate such imperfect point cloud data and study these various forms of error. The Gargoyle model is courtesy VClab, ISTI - CNR.

optical laser scanners, structured lighting, structure from motion, and photometric stereo.

The form of output can be broken down into two main components: surface representation and the dependency on the input data. The surface representation may be a parametric surface, an implicit surface, or a triangulated surface mesh. The dependency on the input data can range from interpolating all of the input data, interpolating a subset of the input, or approximating the input.

The focus of this work is on the evaluation and comparison of surface reconstruction algorithms which take as input a registered point cloud equipped with normals and output a triangulated surface mesh that approximates the input data. More specifically, we focus on input data acquired through triangulation-based scanning, where normals are absent and must be computed from the points themselves. This class of input is extremely broad, and quite common in point cloud data due to the rising ubiquity of triangulation-based scanners such as optical laser scanners. This class of output is flexible for surface reconstruction, as triangle meshes are capable of representing surfaces of arbitrary detail, while the approximation requirement gives freedom in handling point clouds that contain large imperfections.

Despite the vast amount of work in this class of algorithms, to date there has been an insufficient means of evaluation. These algorithms are typically run on acquired scan data. However, there does not exist a computational representation of the surface from which the scanned points were measured. Hence, it is not possible to compare the reconstructed surface to the original surface, and it is quite common for such approaches to instead provide a visual comparison. Quantitative measures are typically done using synthetically generated data, but existing quantitative evaluation approaches contain several shortcomings, such as the representation of the reference shape and the sampling model.

Our benchmark for surface reconstruction rectifies these deficiencies in evaluation, providing the following contributions.

—*Realistic data*. We use a collection of both simple and complex shapes, where an implicit surface is used as the computational representation. We then synthetically scan the implicit surface to provide a collection of point clouds, where our scanning simulation is validated against real data.
—*Accuracy*. By using implicit surfaces we have a precise means of performing evaluation, in both positional and differential measures. We use particle systems to uniformly sample both the implicit surface and the reconstructed surface mesh, minimizing any potential bias of measure from the corresponding triangulation.
—*Comprehensiveness*. The set of experiments demonstrates a broad range of behavior across surface reconstruction algorithms.

Part of the difficulty in establishing a comprehensive set of experiments is the large variability in point clouds. In triangulation-based scanning, a surface may be sampled under a wide variety of conditions, producing point clouds containing such characteristics as noise, outliers, nonuniform sampling, and missing data. This variability is further enhanced when scan data is processed to produce an oriented point cloud, where registration and normal orientation must be performed. Considering all of these factors, it is difficult to determine the effectiveness of a surface reconstruction algorithm with respect to an arbitrary point cloud; see Figure 1 for an illustration. In light of this, we divide our experiments into two sets, one giving a macro view of reconstruction, and the other giving a micro view.

The first set of experiments samples a small number of complex shapes under a large variety of scanner settings. A point cloud for a given shape provides us with a number of evaluation measures. We aggregate each measure over all point clouds to generate an *error distribution* for a given shape. This serves two purposes. First, it

provides us with an objective means to compare algorithms by looking at their performance over a distribution, rather than a single point cloud which may bias a certain class of algorithms. Second, it illustrates how effective an algorithm is at reconstructing a single shape, given that the shape may be sampled in an unbounded number of ways.

The second set of experiments complements the first, by measuring algorithmic performance in the presence of specific sampling and shape properties. For these experiments we wish to generate scans that contain specific properties such as different levels of sparsity, missing data, and noise. However, complex shapes are inappropriate to use since their complexity makes it difficult to controllably elicit these properties. Hence we use a set of simple shapes, some strictly smooth and others containing sharp features, each sampled in ways to highlight specific properties for examining surface reconstruction algorithms.

Lastly, we have made our dataset and benchmark code available to the public (at: http://reconbench.org). We expect our experiments to benefit the surface reconstruction research community in two ways. The first set of experiments may be used to obtain an immediate comparison across reconstruction algorithms, while the second set of experiments should prove useful to observe specific algorithmic behavior. Combined, our benchmark provides comprehensive insight into this class of surface reconstruction algorithms.

2. RELATED WORK

Surface reconstruction. Broadly speaking, we may classify surface reconstruction algorithms by their expected input and the type of output they produce.

One class of algorithms takes as input an unoriented point cloud and produces an *interpolating surface* in the form of a triangulation that uses a subset of the input points as vertices. Often these "connect-the-dots" algorithms are filtration-based techniques; they first build a triangulation with more elements than needed, and then prune away triangles not near the surface. By using the Delaunay triangulation coupled with modeling the point cloud as an ε-sample [Amenta and Bern 1999], many of these algorithms come with provable guarantees regarding the quality of the reconstruction. Extensive research efforts have been devoted to this model, producing the Cocone [Amenta et al. 2002] and Power Crust [Amenta et al. 2001] algorithms. Many other extensions have been compiled in a recent survey [Cazals and Giesen 2006] and monograph [Dey 2007].

Restricting the reconstruction to have vertices only on the input point cloud can be limiting when the data is nonuniform, incomplete, or noisy. Algorithms that build *approximating surfaces* give a flexible alternative in these situations. Here the output is often the triangulation of an iso-surface of a best-fit implicit function of the input. Many of these algorithms [Hoppe et al. 1992; Boissonnat and Cazals 2002] compute a distance field by estimating the tangent plane at every point and computing closest distances using these tangent planes. The method of VRIP [Curless and Levoy 1996] takes advantage of the range scans acquired through laser triangulation to construct a volumetric signed distance field that merges the scans in a least-squares sense.

Surface approximation from point sets with oriented normals has gained recent attention. Approaches range from computing an indicator function [Kazhdan 2005; Kazhdan et al. 2006; Alliez et al. 2007; Manson et al. 2008], to locally fitting functions and moving least-squares methods [Alexa et al. 2003; Ohtake et al. 2003, 2005b; Fleishman et al. 2005]. These approaches are well-equipped to handle various imperfections in the data, and comprise an interesting class of algorithms to study for comparison and evaluation.

However, normal estimation in the presence of imperfect data remains a difficult problem [Mitra and Nguyen 2003; Dey et al. 2005], and a thorough study of normal estimation is beyond the scope of our benchmark.

Reconstruction evaluation. In the area of surface reconstruction evaluation, most of the preceding approaches employ qualitative methods when comparing to other reconstruction algorithms. This usually takes the form of a visual comparison. However, significantly less work has been devoted to obtaining *quantitative* measures. This is due to the common use of scan data, where there is no longer a computational representation of the shape. For synthetic data, the works of Kazhdan [2005], Manson et al. [2008], and Süßmuth et al. [2010] take a triangle mesh as ground truth, and randomly sample the triangles directly to obtain a point cloud. This form of sampling, however, does not reflect the type of data obtained from a scanner which is subsequently organized into a point cloud. The works of Hoppe et al. [1992] and ter Haar et al. [2005] obtain synthetic scans of a triangle mesh from ray tracing or z-buffering the mesh. These methods can produce realistic data under the assumption of clean data, but are insufficient for replicating common scan artifacts. While our approach also generates synthetic range data, it is more realistic since we simulate an optical triangulation-based scanner.

A drawback of all of these approaches is the use of a triangle mesh as ground truth. Sampling a triangle mesh, either directly or through synthetic scans, may produce "faceted scans", where multiple samples lie on a single triangle. This can be misleading for reconstruction algorithms, as the reconstruction may preserve these faceted portions. It is also problematic to use a triangle mesh as ground truth for comparing surfaces. METRO [Cignoni et al. 1998] has become quite common for comparing two triangulated surface meshes, however, for surface reconstruction we are more interested in seeing how well a reconstructed surface compares to a real shape that is smooth, not necessarily a faceted approximation. Moreover, if we are interested in comparing differential quantities, we have an ill-posed definition of surface normals when using a triangle mesh as ground truth.

Similar benchmarks. Finally, related benchmarks exist in the area of 3D stereo reconstruction—for binocular stereo [Scharstein and Szeliski 2002] and multiview stereo [Seitz et al. 2006]. Both use real-world data as input to the various acquisition methods. Multiview stereo, in particular, applies VRIP [Curless and Levoy 1996] to each range scanned surface, and uses the resulting triangle mesh as the gold standard for comparison. However, as pointed out by Kazhdan et al. [2006], VRIP is certainly not free of errors, and arguably *every* surface reconstruction algorithm will contain errors in the presence of imperfections in scanned data. Thus, having a clear understanding of these types of errors is crucial when using a reconstructed surface as a gold standard.

3. OVERVIEW

Our benchmark is broken up into three main phases: surface modeling, sampling, and evaluation. See Figure 2 for the full pipeline.

We start with an implicit surface. We model piecewise-smooth surfaces using *integrated polygonal constraints*, and approximate a triangle mesh with an implicit surface, as detailed in Section 4.

We then sample this implicit surface to obtain an oriented point cloud. We simulate the process of an optical triangulation scanner in order to produce range scans. We slightly overlap the range scans and register them through a rigid-body registration algorithm. From the registered point cloud, we then compute and orient normals
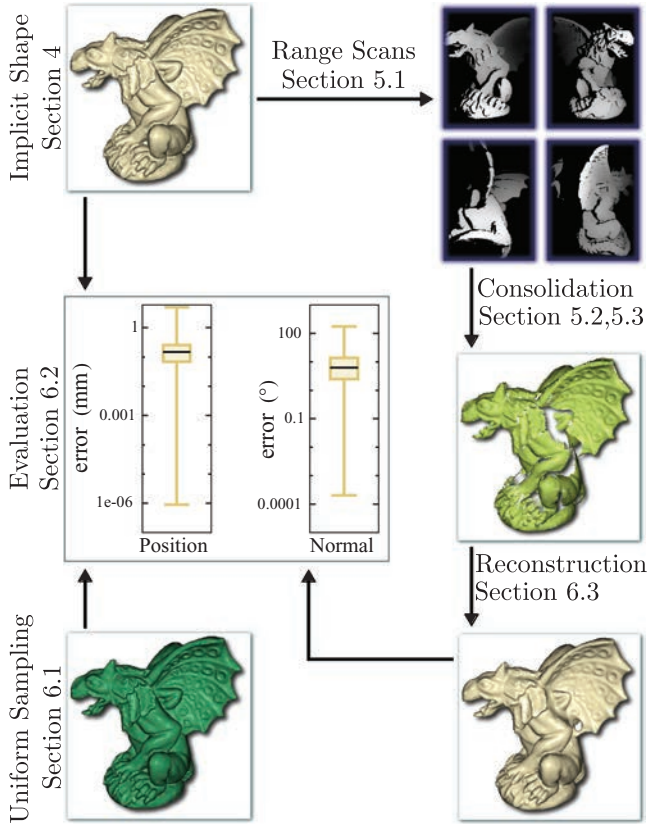
Fig. 2. Overview of our benchmark. First we create an implicit representation of a surface mesh. We then sample this implicit surface by synthetically scanning the shape to obtain individual range scans, and consolidate the scans into a single oriented point cloud via registration and normal estimation. We run a reconstruction algorithm on this oriented point cloud, and compare this output to the implicit model and a dense uniform sampling of the implicit shape to obtain quantitative results. The Gargoyle model is courtesy VClab, ISTI - CNR.

for each point, producing an oriented point cloud suitable for the class of algorithms under consideration. These steps are described in more detail in Section 5.

We next run a surface reconstruction algorithm using the oriented point cloud as input. This gives us a triangle mesh, which we evaluate by comparing to the implicit surface and a dense uniformly sampled point cloud of the implicit surface. We then construct positional and normal error metrics, demonstrated in Figure 2 as individual distributions of point-to-point correspondences. This is explained in detail in Section 6.

4. SURFACE MODELING

In modeling ground-truth data, care must be taken in the surface representation, as it impacts the rest of our pipeline. Although triangulated surfaces are popular and easy to work with, we use smooth and piecewise-smooth surfaces as ground truth, as it benefits the sampling and evaluation phases as follows.

—*Sampling*. Our laser-based scanning simulator requires a surface equipped with a smooth normal field in order to best model an optical laser scanner. As the normal field of a triangulated surface is

discontinuous between triangle faces, this surface representation can adversely impact our scanning simulator.

—*Evaluation*. The surface reconstruction algorithms under consideration assume a point cloud sampled from a smooth surface, so using a smooth surface for quantitative evaluation respects an algorithm's assumptions. Moreover, a smooth normal field permits us to reliably evaluate differential quantities in the reconstruction.

The implicit surfaces that we use to model piecewise-smooth objects are defined via a novel implicit function definition from integrated smoothness constraints. This method integrates weight functions over polygons. It also allows us to define sharp features on object surfaces.

4.1 Polygonal MPU

Our implicit representation is a straightforward extension of Multilevel Partition of Unity (MPU) [Ohtake et al. 2003] applied to a triangulated surface, with the main distinction being that we integrate weight functions over polygons. We use the weight function of Shen et al. [2004], defined for a given point $\mathbf{x} \in \mathbb{R}^3$ and for an arbitrary point on a triangle t, $\mathbf{p} \in t$.

$$w(\mathbf{x}, \mathbf{p}) = \frac{1}{\left(|\mathbf{x} - \mathbf{p}|^2 + \epsilon^2\right)^2} \qquad (1)$$

Here, ϵ is a smoothing parameter. We may now integrate this weight function over the entire triangle t.

$$w(\mathbf{x}, t) = \int_{\mathbf{p} \in t} w(\mathbf{x}, \mathbf{p}) \mathrm{d}\mathbf{p} \qquad (2)$$

For evaluating Eq. (2), Shen et al. [2004] propose a method for numerical integration. However, we derive a closed-form solution for this expression. This prevents potential numerical inaccuracies caused by a quadrature scheme, which could be detrimental to having a reliable benchmark. We outline the derivation in Appendix A.

Equipped with a mechanism for integrating weights over polygons, we proceed with MPU by hierarchically fitting shape functions to a triangulated surface with triangle set $T = \{t_1, \ldots, t_n\}$. We adaptively build an octree over T, where for each cell we associate a sphere whose radius is the length of the cell's diagonal. We then gather all triangles that are contained in, or overlap the sphere, and fit a shape function to these triangles.

In practice we use linear functions for our shape functions, where for each cell i we associate the function $g_i(\mathbf{x}) = \mathbf{x}^{\mathrm{T}}\mathbf{n}_i + b_i$. For all triangles that belong to the sphere of cell i, $T_i \subset T$, we fit the shape function as follows.

$$\mathbf{n}_i = \frac{\sum_{t \in T_i} \mathbf{n}_t \int_{\mathbf{p} \in t} w(\mathbf{s}_i, \mathbf{p}) \mathrm{d}\mathbf{p}}{\sum_{t \in T_i} \int_{\mathbf{p} \in t} w(\mathbf{s}_i, \mathbf{p}) \mathrm{d}\mathbf{p}} \qquad (3)$$

$$b_i = -\left\langle \frac{\sum_{t \in T_i} \int_{\mathbf{p} \in t} \mathbf{p} \, w(\mathbf{s}_i, \mathbf{p}) \mathrm{d}\mathbf{p}}{\sum_{t \in T_i} \int_{\mathbf{p} \in t} w(\mathbf{s}_i, \mathbf{p}) \mathrm{d}\mathbf{p}}, \mathbf{n}_i \right\rangle \qquad (4)$$

Here, $\mathbf{n}_t$ is the triangle normal of t and $\mathbf{s}_i$ is the center of the sphere for cell i. Although one may use higher-order shape functions such as quadrics, we found the difference negligible. The main difference was that for linear functions we required a deeper octree to adequately approximate T.

The octree is built such that each cell is subdivided only if the zero set of its shape function deviates sufficiently from the sphere's triangles. If the cell's sphere is empty to start with, we increase the sphere's radius until it encompasses a sufficient number of triangles
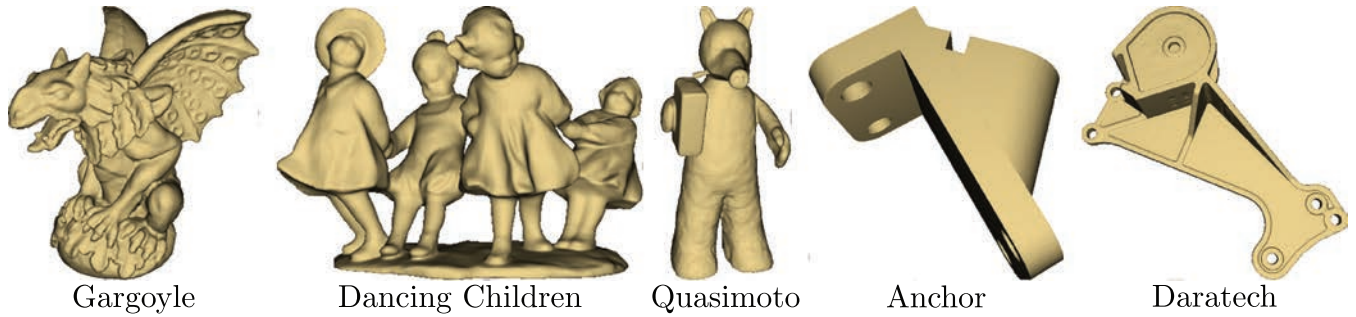
Fig. 3. Complex shapes created via our Polygonal MPU scheme. In our experiments these shapes are synthetically scanned under a wide variety of typical use case scan parameters. This class of shapes contains many interesting characteristics for scanning, such as multiple scales of detail, nontrivial topology, and sharp features. The Gargoyle model is courtesy VClab, ISTI - CNR, the Dancing Children model is courtesy AIM@SHAPE, the Anchor model is courtesy of Dey et al. [2003], and the Daratech model is courtesy of Regli and Gaines [1997] via the INRIA Gamma database.

(which we take to be six). This gives a covering of the space with overlapping spheres. We may then evaluate the implicit function at a point by blending all shape functions whose spheres contain that point.

$$f(\mathbf{x}) = \frac{\sum_i q_i(\mathbf{x}) g_i(\mathbf{x})}{\sum_i q_i(\mathbf{x})} \qquad (5)$$

Here, q_i is a quadratic b-spline function centered at $\mathbf{s_i}$.

To preserve sharp features, we follow Ohtake et al. [2003] in detecting sharp features within a leaf cell and consequently applying CSG operations for exact feature preservation. We identify sharp features using a threshold on dihedral angles. We then apply union and intersection operations on overlapping shape functions to exactly preserve the sharp feature. We support sharp feature curves and corners containing a maximum degree of four.

4.2 Benchmark Shapes

Our first set of experiments consists of shapes that contain different types of complexities; see Figure 3 for these shapes. The Gargoyle model contains details of various feature sizes, ranging from the bumps on the bottom to the ridges on its wings. The Dancing Children model is of nontrivial topology, containing tunnels of different sizes and features such as the rim of the hat on the left child and wrinkles in the cloth. The Quasimoto model is representative of a shape containing articulated parts, such as arms, legs, and head. The Anchor model contains sharp features, moderately sized tunnels, as well as a single deep concavity. Lastly, the Daratech model contains sharp features, small tunnels, as well as thin surface sheets. We note that the Gargoyle, Dancing Children, and Quasimoto models were scanned from objects and subsequently reconstructed, which has two potential consequences: every surface point is visible from some position of the scanner, and the implicit function may inherit smoothing from the original surface reconstruction algorithm.

The second set of experiments uses simple shapes; see Figure 4. The Bumpy Sphere contains smooth features at varying scales. The Spiral shape is primarily composed of a thin cylindrical feature. Lastly, the Mailbox consists of straight and curved sharp features.

5. SAMPLING

The intent of our sampling scheme is to replicate the acquisition process of a triangulation-based scanner, in order to produce realistic point clouds. To this end, sampling is composed of three intermediate stages: synthetic range scanning, registration, and orientation.

Fig. 4. Simple shapes created via our Polygonal MPU scheme. In our experiments these shapes are scanned in a precise a manner in order to replicate specific scanning difficulties, such as sparsity, missing data, and noise. The Mailbox model is courtesy of Dey et al. [2003].

(a) Uniform sampling   (b) Nonuniform sampling

(c) Noisy data   (d) Misaligned scans
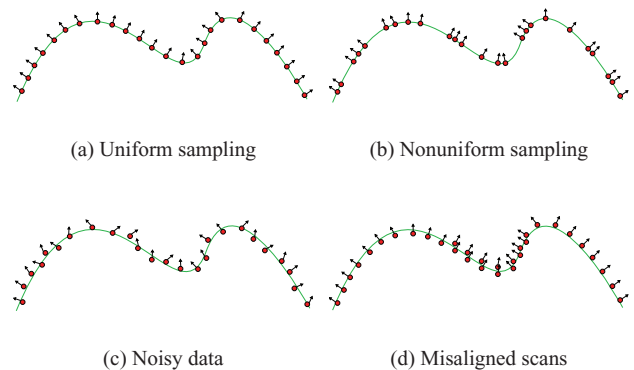
Fig. 5. Common properties of scanned data on a sampled curve. The green curve is the true curve, while the red points are the oriented points sampled from the curve.

Common properties found in scanned data are illustrated in Figure 5. See Figure 6 for an illustration of our synthetic scanner's capability in replicating such properties.

5.1 Synthetic Range Scans

We simulate the acquisition of range scans by modeling an optical laser-based triangulation scanning system. Such scanning systems suffer *random error* and *systematic error*. Random error is due to physical constraints, such as noise in the laser, variations in the reflectance due to surface materials, and nonlinear camera warping.

(a) Uniform sampling   (b) Nonuniform sampling

(c) Noisy data   (d) Misaligned scans
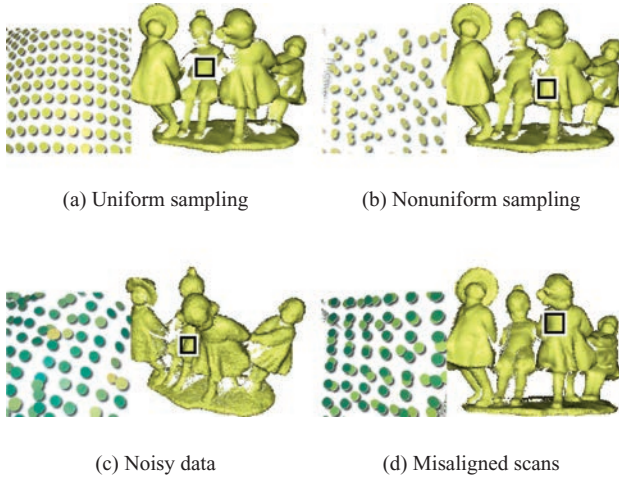
Fig. 6. Common characteristics of 3D scans. These point clouds were generated using our synthetic scanner, illustrating our capability to replicate common scan properties. In the noise and misalignment insets we have color mapped the points by their distance away from the implicit shape, with yellow being far and green being close. The Dancing Children model is courtesy AIM@SHAPE.

Systematic error is the result of imprecise range measurement due to the peak detection algorithm. Our range scans are generated by synthesizing random error, while reproducing systematic error by performing standard peak detection.

Random error synthesis. We synthesize random errors by generating a series of *radiance images*, where each image is the result of a single laser stripe projection onto the implicit surface. To this end, given a pinhole camera at position **c** and a baseline configuration, we first generate the noise-free range data by ray tracing the implicit surface. We reject all points that are not visible from the laser position, a function of the baseline distance. This provides us with a set of pixels containing geometry $P = \{\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_n\}$ and their corresponding points $X = \{\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n\}$.

We now project laser stripes onto the range geometry. We model each laser stripe projection according to a cylindrical projection, parameterized by laser position **l**, field of view of the laser stripe α, and triangulation angle θ. The triangulation angle is defined with respect to an initial laser stripe plane. We may then define the laser stripe frustum as the volume enclosed by the two planes $\{\mathbf{l}, \theta - \frac{\alpha}{2}\}$ and $\{\mathbf{l}, \theta + \frac{\alpha}{2}\}$. A point is considered to be contained within the frustum if it is within positive distance to both planes. The inset depicts a 2D illustration of this configuration, where the red points of the green curve are considered to be within the laser's frustum.

For a single laser stripe, we gather all range geometry that is contained within the stripe. This defines the set of "active" pixels to which the laser stripe contributes. We then determine the noise-free radiance at pixel $\mathbf{p}_i$ due to a laser stripe at triangulation angle θ by [Curless and Levoy 1995]

$$L_\theta(\mathbf{p}_i) = |\mathbf{n}_i \cdot \omega| e^{\frac{-2.0(d(\mathbf{x}_i))^2}{\beta^2}}. \tag{6}$$

Here, $\mathbf{n}_i$ is the normal of the implicit surface at $\mathbf{x}_i$, ω is the unit vector pointing towards the laser position from $\mathbf{x}_i$, $d : \mathbb{R}^3 \to \mathbb{R}$ is the closest distance to the center of the laser frustum, and β is

the width of the frustum at $\mathbf{x}_i$. Here we assume that the surface is purely diffuse, hence the BRDF is reduced to a constant factor that we omit.

In practice, diffuse surfaces suffer noise in the form of laser speckle, where surface roughness contributes to variations in the reflectance [Baribeau and Rioux 1991]. We observe that this form of noise is more significant further away from the center of the laser stripe frustum. We model this as normally distributed additive noise, where the variance is the distance away from the center of the laser stripe.

$$\tilde{L}_\theta(\mathbf{p}_i) = L_\theta(\mathbf{p}_i) + \eta \epsilon_\sigma(\mathbf{x}_i) \tag{7}$$

Here, η is a user-specified noise magnitude, and ϵ is a random variable normally distributed with variance σ, the distance from the center stripe. In addition, we also allow for smoothing of the noisy radiance image by convolving $\tilde{L}_\theta$ with a Gaussian kernel of a user-specified bandwidth.

Systematic error. For each corrupted radiance image $\tilde{L}_\theta$, we next perform peak detection in order to find each pixel's laser stripe plane. From the laser stripe plane, depth is obtained by triangulation. A common assumption in many peak detection algorithms is that the radiance profile, either over space or time (i.e., triangulation angle), is Gaussian [Curless and Levoy 1995]. However, in the presence of depth discontinuities, curved surfaces, and noise, this assumption is violated, producing systematic error.

To this end, we consider all radiance images $\tilde{L}_\theta$ defined for each triangulation angle $\theta \in \{\theta_1, \theta_2, \ldots \theta_m\}$, where m is the number of laser stripes. For each pixel, we consider its radiance profile as θ increases. We fit a Gaussian to this radiance profile via the Levenberg-Marquardt method. This Gaussian provides us with a mean, which determines the stripe plane, as well as a peak magnitude and variance, both of which are used for rejecting low-confidence range data.

Please see Appendix B for the full list of scanning parameters and common parameter settings.

5.2 Validation

It is important to verify that the range scans we are producing contain artifacts found in real scans. To this end, we validate our synthetic scans by comparing them to data acquired by commercial scanning systems. We illustrate our capability of replicating noise and missing data artifacts, which arguably have the greatest impact on surface reconstruction. We are not interested in exactly reproducing scans produced by commercial scanning systems. Most systems perform postprocessing which is far beyond the scope of our scanning simulation. Instead, we show that our scanning simulation is expressive enough to generate a range of scan artifacts, while still capable of generating artifacts of a commercial scanner with proper scan parameters. To perform validation, we first 3D-print a given implicit surface, then scan the printed model, and lastly register the real scan to the implicit surface in order to compare against our synthetic scan.

We have manufactured the Gargoyle model by 3D-printing, through the company Shapeways [Shapeways 2011]. The minimum detail at which models may be manufactured through Shapeways is 0.2mm. We then scan the model with an optical triangulation-based scanner, namely the NextEngine scanner [NextEngine 2011]. The scanner has a maximum accuracy of 0.127mm at its finest resolution. For surfaces at an optimal distance from the scanner, with normals roughly aligned to the scanner's optical axis, we found this to be true. However for a complex shape like the Gargoyle, as we
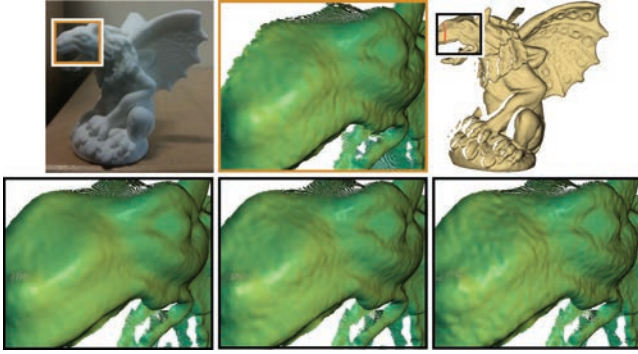
Fig. 7. Comparison of noise profiles between our scanning simulation in increasing noise magnitude (bottom), and a NextEngine scan (top-center). Note that real scanner noise takes the form of bumps aligned in the direction of the laser scan projection (top-right), and our synthetic noise is able to capture this anisotropic noise over varying noise magnitude. The Gargoyle model is courtesy VClab, ISTI - CNR.
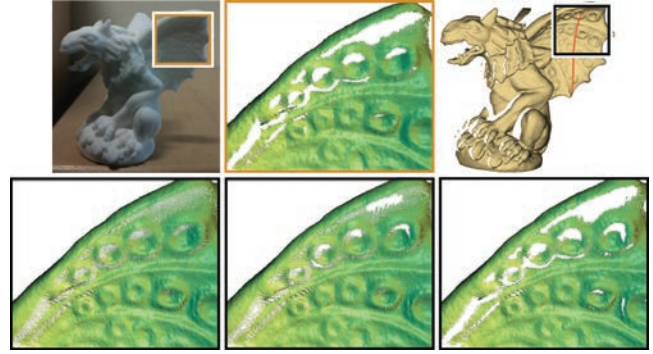

Fig. 8. A comparison of missing data between our scanning simulation in increasing peak threshold (bottom), and a NextEngine scan (top-center). Note the similarities in regions of missing data between our scan (bottom-right) and the NextEngine scan, primarily due to the grazing angle at which the laser strikes the surface, resulting in a low level of radiance. The Gargoyle model is courtesy VClab, ISTI - CNR.

will demonstrate, the accuracy can vary and the noise magnitude becomes greater than the shape's resolution.

To compare a real scan to a synthetic scan, we first register the real scan to the implicit surface. We perform ICP under a rigid-body deformation in order to best align the real scan to the implicit surface. As the NextEngine does not provide specifics on their CCD sensor, we take the depth image and use the camera calibration toolbox [Bouguet 2010] to obtain the intrinsic and extrinsic camera parameters. We feed these camera parameters into our synthetic scanning system to obtain a comparable range scan. We note that a small nonrigid deformation might be preferable to a rigid-body deformation for registration due to small nonlinear camera deformation artifacts [Brown and Rusinkiewicz 2007]. However, this adversely impacts camera calibration and hence is unsuitable for our purposes.

Noise validation. In our scanning simulation, noise is strongly dependent on laser stripe resolution, laser stripe field of view, noise magnitude, and image smoothing bandwidth. As NextEngine does not provide these parameters for their system, to compare noise against the NextEngine scanner we have best estimated the stripe resolution, field of view, and smoothing bandwidth, while varying the noise magnitude. See Figure 7 for the comparison. Note that real scanner noise is in fact anisotropic—a function of the baseline [Abbasinejad et al. 2009]. Hence we see "bumps" that are slightly aligned with the direction of the laser projection in the NextEngine scan. Our synthetic scans demonstrate this anisotropy as well. We show that by tuning the noise magnitude, we can produce a variety of noise profiles, including something similar to that of the NextEngine scanner.

Missing data validation. Missing data in a range scan is typically the result of the rejection of low-confidence range data. In our scanning simulation, this is related to the peak intensity threshold, where a small peak may indicate a poor Gaussian fit. To compare the missing data profile from our scanning simulation to that of the NextEngine scanner, we vary the peak detection threshold; see Figure 8. The NextEngine profile seems to correspond to a particular threshold. Our tunable parameter for peak detection is based on this threshold for the experiments in Section 7.

5.3 Scanning and Registration

Given that we have a means of acquiring range scans, next we must determine where to scan. It is extremely difficult to automate the process of positioning/orienting a scanner, as this is inherently a manual process. We assume an ideal environment in which we place the scanner at uniformly sampled positions over the bounding sphere of the object, such that the camera is oriented to look at the object's center of mass. Note that such acquisition systems are starting to gain popularity [Vlasic et al. 2009].

From these individual range scans, we next register them into a single coordinate system. First we overlap the scans by a prescribed amount. This is achieved by applying a small, random rotation to each scan, where initially each scan is in a common global coordinate frame. We then run locally weighted ICP [Brown and Rusinkiewicz 2007] to align the scans. Note that the amount of overlap effectively determines the quality of the alignment. Less overlap means a poorer initialization, and the optimization process may hit an undesirable local minimum causing misalignment.

5.4 Orientation

From the registered point cloud, we must assign an oriented normal to each point. To estimate a point's local tangent plane, and hence its unoriented normal, we gather its k-nearest-neighbor points and perform PCA. Note that this method may produce noisy tangent planes due to nonuniform sampling, noise, misalignment, and missing data.

We allow normal orientation to be performed using two different methods. The first method chooses the normal direction which has smallest angle with the vector formed from the scanner position to the sample point. The second method is the approach of Hoppe et al. [1992] which forms a minimum spanning tree over the point cloud to propagate normal directions. Both methods can produce normals oriented in the opposite direction, primarily due to noise in the estimated tangent planes. In particular, the method of Hoppe et al. [1992] may result in large regions of inverted normals due to sharp features and nonuniform sampling.

6. EVALUATION

In order to evaluate the quality of a surface mesh M output by a reconstruction algorithm against the input implicit surface Ω, we take the view of *discrete differential geometry* for defining error
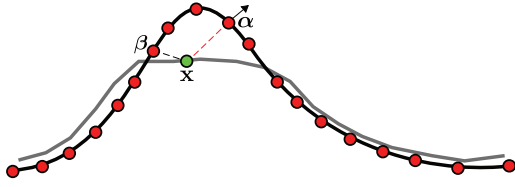
Fig. 9. A situation where the Φ mapping produces an incorrect shortest distance correspondence. The dashed red line indicates the normal line from $\boldsymbol{\alpha}$ to $\mathbf{x}$, giving us an inaccurate correspondence since $\boldsymbol{\beta}$ is closer to $\mathbf{x}$ than $\boldsymbol{\alpha}$. So we instead take $(\mathbf{x}, \boldsymbol{\beta})$ as a correspondence.

measures. As illustrated in Hildebrandt et al. [2006], pointwise plus normal convergence of a polyhedral surface to a smooth surface implies convergence in the metric, surface area, and Laplace-Beltrami operator. In their context, pointwise convergence is measured in terms of Hausdorff distance and normal convergence is measured as the supremum of the infinity norm over all normals. We take their basic framework and expand it to include other error measures, in order to provide a more informative evaluation.

6.1 Shortest Distance Maps

To measure error, we first construct shortest distance maps, which defines the correspondence for a given point on one surface as its closest point on the other surface. Let the closest point function $D \colon M \to \Omega$ map a point of the output polygonal surface M to its closest point on the implicit surface Ω, and define $\Phi = D^{-1}$ to be the inverse map of D. Notice that $\Phi(\boldsymbol{\alpha})$ is generally not the closest point on M to $\boldsymbol{\alpha} \in \Omega$, so that $\Phi(\boldsymbol{\alpha})$ may be undefined for some $\boldsymbol{\alpha}$ and Φ is not one-to-one everywhere.

We wish to sample Φ nearly uniformly on Ω, so as to be able to make accurate estimates of mean distance and normal error. This is a departure from the technique of METRO [Cignoni et al. 1998], which only requires dense sampling. We build a point sample on Ω using the particle system method of Meyer et al. [2007]. We empirically chose the uniform inter-particle distance for each shape so as to preserve features. Denote P_Ω as the resulting sample set.

For each sample point $\boldsymbol{\alpha} \in P_\Omega$, we shoot a ray in the normal direction towards M to get a candidate $\mathbf{x} = \Phi(\boldsymbol{\alpha})$ [Hildebrandt et al. 2006]. If the closest sample point to $\mathbf{x}$ in P_Ω, denoted $\boldsymbol{\beta}$, is indeed $\boldsymbol{\alpha}$, then we accept it as a closest point correspondence. Otherwise we use $(\mathbf{x}, \boldsymbol{\beta})$ as the correspondence; see Figure 9 for a 2D illustration. From this process we obtain a set of closest point correspondences.

$$C_\Omega = \{(\mathbf{x}, \boldsymbol{\alpha}) \mid \boldsymbol{\alpha} \in P_\Omega, \mathbf{x} = \Phi(\boldsymbol{\alpha})\} \qquad (8)$$

We also construct a dual map $\Psi \colon M \to \Omega$, using the same methodology. Instead of choosing an inter-particle distance during the sampling, we fix the number of sample points, since M may be arbitrarily complex. We denote P_M as the resulting uniformly spaced sample set on M. We thus obtain the following set of closest point correspondences.

$$C_M = \{(\boldsymbol{\alpha}, \mathbf{x}) \mid \mathbf{x} \in P_M, \boldsymbol{\alpha} = \Psi(\mathbf{x})\} \qquad (9)$$

## 6.2 Discrete Error Measures

Given these maps we define a variety of discrete error measures between Ω and M. Denoting $|S| = |C_\Omega| + |C_M|$, Hausdorff distance is approximated by

$$H(\Omega, M) = \max \left\{ \max_{(\mathbf{x}, \boldsymbol{\alpha}) \in C_\Omega} |\mathbf{x} - \boldsymbol{\alpha}|, \max_{(\boldsymbol{\alpha}, \mathbf{x}) \in C_M} |\boldsymbol{\alpha} - \mathbf{x}| \right\}. \qquad (10)$$

Further, mean distance is approximated by

$$\mu(\Omega, M) = \frac{1}{|S|} \left(\sum_{(\mathbf{x}, \boldsymbol{\alpha}) \in C_\Omega} |\mathbf{x} - \boldsymbol{\alpha}| + \sum_{(\boldsymbol{\alpha}, \mathbf{x}) \in C_M} |\boldsymbol{\alpha} - \mathbf{x}| \right). \qquad (11)$$

These measures depict error in very different ways, as the inset illustrates. Here the circle is the smooth shape, while the piecewise linear curve is the approximating mesh. Hausdorff distance will be large for the pair of shapes on the left, while mean distance will be rather small. For the pair of shapes on the right, the mean distance will be much larger than the pair of shapes on the left, while Hausdorff distance will be less.
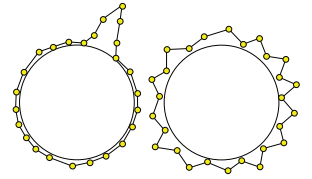
From these shortest distance correspondences, we have a method of measuring higher-order geometric properties, by comparing differential properties at the correspondences. This is analogous to defining pullbacks on Φ and Ψ. We measure normal angle deviations in a similar manner to distance measures. If we denote $\gamma(\boldsymbol{\alpha}, \mathbf{x}) = \angle(\mathbf{N}_\Omega(\boldsymbol{\alpha}), \mathbf{N}_M(\mathbf{x}))$, the maximum and mean angle deviation of point correspondences, respectively, are

$$H_N(\Omega, M) = \max \left\{ \max_{(\mathbf{x}, \boldsymbol{\alpha}) \in C_\Omega} \gamma(\boldsymbol{\alpha}, \mathbf{x}), \max_{(\boldsymbol{\alpha}, \mathbf{x}) \in C_M} \gamma(\boldsymbol{\alpha}, \mathbf{x}) \right\} \qquad (12)$$

$$\mu_N(\Omega, M) = \frac{1}{|S|} \left(\sum_{(\mathbf{x}, \boldsymbol{\alpha}) \in C_\Omega} \gamma(\boldsymbol{\alpha}, \mathbf{x}) + \sum_{(\boldsymbol{\alpha}, \mathbf{x}) \in C_M} \gamma(\boldsymbol{\alpha}, \mathbf{x}) \right). \qquad (13)$$

In practice we take $\mathbf{N}_M$ to be triangle normals, as opposed to more sophisticated normal estimation methods [Meyer et al. 2002]. Such methods are sensitive to the triangulation and typically assume smoothness in the normal field. As a result, the presence of sharp features can result in undesirable oversmoothing.

Comparison with METRO [Cignoni et al. 1998]. An alternative might have been to use METRO to compare the original meshes with the output meshes. However, there would have been two issues in using a mesh for both sampling and evaluation. First, there would have been cases in which many points of a synthetic scan all lie on one triangle, an undesirable artifact. Second, we would have needed to estimate normals on the input triangle mesh for comparison, introducing more error.

As an alternative, we could have used the implicit surface for sampling and an iso-surface of the implicit surface for evaluation. However, errors in iso-surfacing can lead to errors in the evaluation, since a reconstruction algorithm tries to recover the implicit surface, not its iso-surface.

6.3 Algorithms

We have chosen a wide variety of publicly available surface reconstruction algorithms to test our benchmark. For the sake of fair comparison, we have only used algorithms that take an oriented point cloud as input, and output an approximating surface. Here, we provide a categorization and brief description of each algorithm along with an abbreviation. This abbreviation is used to identify the algorithms in the experiments.

Indicator function. This class of algorithms reconstructs a three-dimensional solid O by finding the scalar function χ, known

as the indicator function, defined in $\mathbb{R}^3$ such that

$$\chi(\mathbf{x}) = \begin{cases} 1 & \mathbf{x} \in O \\ 0 & \mathbf{x} \notin O \end{cases}. \tag{14}$$

Here the surface Ω is then defined by ∂O. In practice, these approaches approximate χ by operating on a regular grid or an octree, and generate Ω by iso-surfacing the grid.

Poisson surface reconstruction (abbr. Poisson) [Kazhdan et al. 2006] solves for χ by noticing that $\nabla \chi$ should agree with the normal field $\mathbf{N}$ at ∂O. This amounts to inverting the gradient operator. Hence, χ is found by solving the Poisson equation.

$$\nabla \cdot \nabla \chi = \nabla \cdot \mathbf{V} \tag{15}$$

Here, $\mathbf{V}$ is the smoothed normal field defined throughout the volume. The Poisson equation is efficiently solved only near the surface by using an adaptive multigrid solver defined on the octree built on the point cloud. Note that use of an octree may result in low grid resolution in regions of missing data.

An alternative method of constructing the indicator function is to solve for it indirectly by projecting χ onto a basis, and then performing an inverse transform to obtain χ. By invoking the Stokes theorem, this projection need only be performed on ∂O.

$$\int_O \nabla \cdot \mathbf{F}(p) dp = \int_{\partial O} \langle \mathbf{F}(p), \mathbf{N}(p) \rangle \, dp \tag{16}$$

Here, $\mathbf{F}$ is a vector-valued function whose divergence $\nabla \cdot \mathbf{F}$ defines the basis. Note that these methods are equivalent to solving Eq. (15), where properties of the basis functions are used to simplify the original problem.

Fourier surface reconstruction (abbr. Fourier) [Kazhdan 2005] employs the Fourier basis as part of their solution. For efficiency they use the Fast Fourier Transform (FFT), hence requiring a regular grid and the grid resolution being a power of two.

Wavelet surface reconstruction (abbr. Wavelet) [Manson et al. 2008] employs a Wavelet basis for the solution of Eq. (16). In our experiments we use the 4-tap Daubechies basis. Due to the multiresolution structure of wavelets, they use an octree for the basis projection. Hence, similar to Poisson, this method may result in limited grid resolution over regions of missing data.

Point set surfaces. Point Set Surfaces (PSS) are defined based on Moving Least-Squares (MLS), where a projection operator is used to define a surface by its collection of stationary points. A point is considered stationary when its projection is the identity map. Originally defined for unoriented points, its definition is greatly simplified when considering points equipped with normals, and may be used for surface reconstruction by considering its implicit surface definition, rather than its projection operator.

Basic PSS methods use a weighted combination of linear functions to locally define the surface at every point. Borrowing terminology from Guennebaud and Gross [2007], we use two different definitions in our experiments: Simple Point Set Surfaces (abbr. SPSS) [Adamson and Alexa 2003] and Implicit Moving Least-Squares (abbr. IMLS) [Kolluri 2005]. The implicit surface definition of SPSS is

$$f(\mathbf{x}) = \mathbf{n}(\mathbf{x})^T (\mathbf{x} - \mathbf{c}(\mathbf{x})). \tag{17}$$

Here, $\mathbf{n}$ is a weighted average of normals in a neighborhood of $\mathbf{x}$, and $\mathbf{c}$ is the weighted centroid in a neighborhood of $\mathbf{x}$. The weights used in computing the normal and the centroid are derived from a smooth, positive function $w_{\mathbf{x}}$ defined with respect to $\mathbf{x}$, which gives

points closer to $\mathbf{x}$ larger influence. IMLS is defined as the implicit function

$$f(\mathbf{x}) = \frac{\sum_i w_{\mathbf{x}}(\mathbf{p}_i)(\mathbf{x} - \mathbf{p}_i)^T \mathbf{n}_i}{\sum_i w_{\mathbf{x}}(\mathbf{p}_i)}. \tag{18}$$

We note that IMLS is a weighted average of linear functions, whereas SPSS is a single linear function, whose centroid and normal is a weighted average of points and normals, respectively.

Algebraic Point Set Surfaces (abbr. APSS) [Guennebaud and Gross 2007] uses spheres defined algebraically as the shape function. Rather than directly obtaining the implicit function at a point, APSS fits a sphere to a neighborhood of points, requiring the solution of a linear least-squares system for every point. By using a higher-order function, the method can be more robust to sparse data than SPSS and IMLS.

For our experiments, the software package provided by Gaël Guennebaud contains implementations of SPSS, IMLS, and APSS. Each PSS is evaluated over a regular grid, and the reconstructed surface is obtained by iso-surfacing the zero level set. In the software, neighborhoods used to locally fit functions are estimated at each point based on the density of the input point cloud. In the presence of missing data this may produce holes in the output due to empty neighborhoods. This has an impact on evaluation, which we discuss in the experiments sections.

Multilevel partition of unity. In our own implicit surface definition we use a variant of Multilevel Partition of Unity (MPU) applied to polygon soup, and so we refer to Section 4.1 for details about the overall approach, noting that the construction of MPU with points is quite similar to that of polygons. In our experiments we use three variants. First we use the original approach of Ohtake et al. [2003] (abbr. MPU), where linear functions are used as low-order implicits. We opted not to use the fitting of sharp features, as we found its sharp feature detection to be rather sensitive and frequently produces erroneous fits. We also use the approach of Nagai et al. [2009] (abbr. MPUSm), which defines differential operators directly on the MPU function, though restricted to linear functions. In doing so, diffusion of the MPU function becomes possible, resulting in a more robust reconstruction method. Lastly, we also use the method by Ohtake et al. [2005b] (abbr. RBF), which uses compactly supported radial basis functions for locally defined implicit functions in the MPU construction. For all MPU methods a surface mesh is generated by first evaluating the MPU function over a regular grid, and iso-surfacing the zero level set to obtain the surface.

Scattered point meshing. The method of Ohtake et al. [2005a] (abbr. Scattered) is a departure from the aforesaid approaches. This method grows weighted spheres around points in order to determine the connectivity in the output triangle mesh. Quadric error functions [Garland and Heckbert 1997] are used to position points in the output mesh, which can result in a small amount of simplification in the output. Similar to the PSS methods, regions empty of data may produce holes in the output.

6.4 Algorithm Parameters

We provide a brief description of the most relevant parameters for each algorithm.

Resolution. As all the algorithms, except Scattered, contour a grid to obtain the surface they must contain sufficient grid resolution to adequately preserve all surface details. Our goal is to provide each algorithm with such a resolution, while maintaining fairness across algorithms which may use and define grids differently. To achieve this, for each implicit surface we first determine the resolution

which is necessary to extract the surface with minimal error. We find that across all shapes, a resolution of 350^3 provides for sufficient resolution to preserve surface details, hence for the PSS and MPU methods we set their resolution to 350.

For Fourier, Poisson, and Wavelet, grid resolution serves two purposes: iso-surface precision, and the accuracy and convergence of the Poisson solver (see Eqs. (15) and (16)). We experimentally found that sufficient resolution for iso-surfacing does not imply sufficient resolution for the solution to the Poisson equation. Hence for Fourier we set the grid resolution to 512, in order to reduce any smoothing resulting from the FFT. For Poisson and Wavelet, although an octree depth of 9 may appear most reasonable, we set it to 10. We find that this additional resolution in regions of high sampling density produces a more accurate iso-value for contouring, which makes a significant difference in contouring low-resolution regions of the octree.

Noise. Algorithms tend to handle noise according to their categorization. For indicator functions, noise may be combated by splatting the points into the grid using a large splat radius, as well as through lowering the grid resolution, effectively serving as a low-pass filter. PSS methods all contain a bandwidth which determines the extent of neighborhood influence. A large bandwidth results in more points for consideration in shape fitting and hence larger data smoothing. MPU methods and Scattered all contain error thresholds that determine the quality of a shape fit. In the presence of noise the tolerance may be increased to avoid overfitting. MPUSm also provide parameters specific to their diffusion method, for which we use author-suggested settings.

Discussion. In practice we set an algorithm's parameters based on the characteristics of the input point cloud, namely the noise level. As the point clouds of experiments 7.1–7.3 contain a constant level of noise, we have kept all algorithm parameters fixed throughout these experiments. The parameters were empirically determined by measuring each method's performance on a subset of the point clouds. Though one may fine-tune an algorithm's parameters to improve its performance with respect to a particular error metric, parameter insensitivity is an important indication of algorithmic robustness. Only in experiment 7.4, where noise varies, do we set algorithm parameters in accordance with the noise level.

7. RESULTS

Our results are broken down into two main sets of experiments: one in which complex shapes are sampled with a variety of sampling settings, and another in which simple shapes are sampled with specific sampling settings. Please see Appendix B for reference to the types of units used throughout the results.

We have not used the maximum angle deviation as an error measure in our experiments. By using triangle normals as the normal field over a surface mesh, this measure can be quite high even when the mesh contains low error in all other measures. Since it did not distinguish between the algorithms, we omitted it.

Note that it is possible for these algorithms to produce surfaces containing multiple connected components. We extracted the largest connected component, in terms of surface area, as the surface for evaluation rather than all components. Unfortunately, this favors algorithms in which connected components are created far from the ground-truth surface over algorithms which create additional components near the surface. Hence, in addition to the error metrics, we have provided additional information on the algorithms including the number of connected components, the length of the boundary

Table I. Range of Scanning Parameters for Error Distribution Experiments

shape	res	scans	camera dist	peak	variance
Gargoyle	250–350	7–11	75–115	0.2–0.4	0.5–0.75
DC	250–350	7–11	75–115	0.2–0.4	0.5–0.75
Quasimoto	250–350	7–11	75–115	0.2–0.4	0.5–0.75
Anchor	175–225	8–12	60–100	0.2–0.4	0.5–0.75
Daratech	250–350	8–12	75–115	0.2–0.4	0.5–0.75

The range of scanning parameters used in the error distribution experiments. Here, *res* represents the image resolution of a single range scan, *scans* is the number of scans taken, *camera dist* is the camera distance away from the center of the object, *peak* is the radiance threshold at which to reject depth, and *variance* is the variance threshold at which to reject depth.

components, whether or not the surface is manifold, deviation from the true genus, and computation time.

7.1 Error Distributions

Our first set of experiments focuses on the performance of surface reconstruction algorithms restricted to a single shape. Given an input we sample it across a variety of scanner parameter settings and run all reconstruction algorithms across all point clouds. We then compute error metrics for each point cloud. For each algorithm, we aggregate the error metrics across all point clouds to obtain what we term *error distributions*.

We argue that error distributions are more effective for benchmarking reconstruction algorithms, rather than comparing algorithms with respect to a single point cloud. Each algorithm has its strengths and flaws for particular forms of data, and to sample a shape in such a way that it favors the strengths of certain algorithms provides an incomplete picture in the comparison of reconstruction algorithms.

To this end we generate samples by varying scanning parameters across typical use case settings. Namely, we vary sampling resolution, the number of range scans, the distance the camera resides from the object, peak threshold, and variance threshold. Please see Table I for the full range of parameters over all shapes. We have adapted certain parameter ranges to specific shapes in order to ensure adequate coverage in the point clouds, and to sufficiently capture shape details. To reproduce small imperfections commonly found in range data, we introduce a constant, modest amount of noise into the laser signal. We also slightly overlap the scans and register them, causing small misalignment errors. For each point cloud we randomly distribute camera positions uniformly on the bounding sphere of the object, rather than keeping their positions fixed.

See Figure 10 for the results of this experiment across all shapes, where the distributions take the form of box plots. The three error measures, mean distance, Hausdorff distance, and mean angle deviation, demonstrate the various strengths and weaknesses of the algorithms.

Smooth surfaces. The Gargoyle, Dancing Children, and Quasimoto shapes represent our class of shapes containing entirely smooth surface features. We find that the algorithms generally perform quite well on these shapes. However, the different error metrics point to subtle differences in performance. For instance, Wavelet tends to produce nonsmooth, rather bumpy surfaces, yet the reconstructed surface tends to stay close to ground truth, which is likely due to the use of wavelet bases in the presence of nonuniform or missing data. This nonsmoothness is depicted in the mean distance and angle deviation plots, yet its Hausdorff distance performance
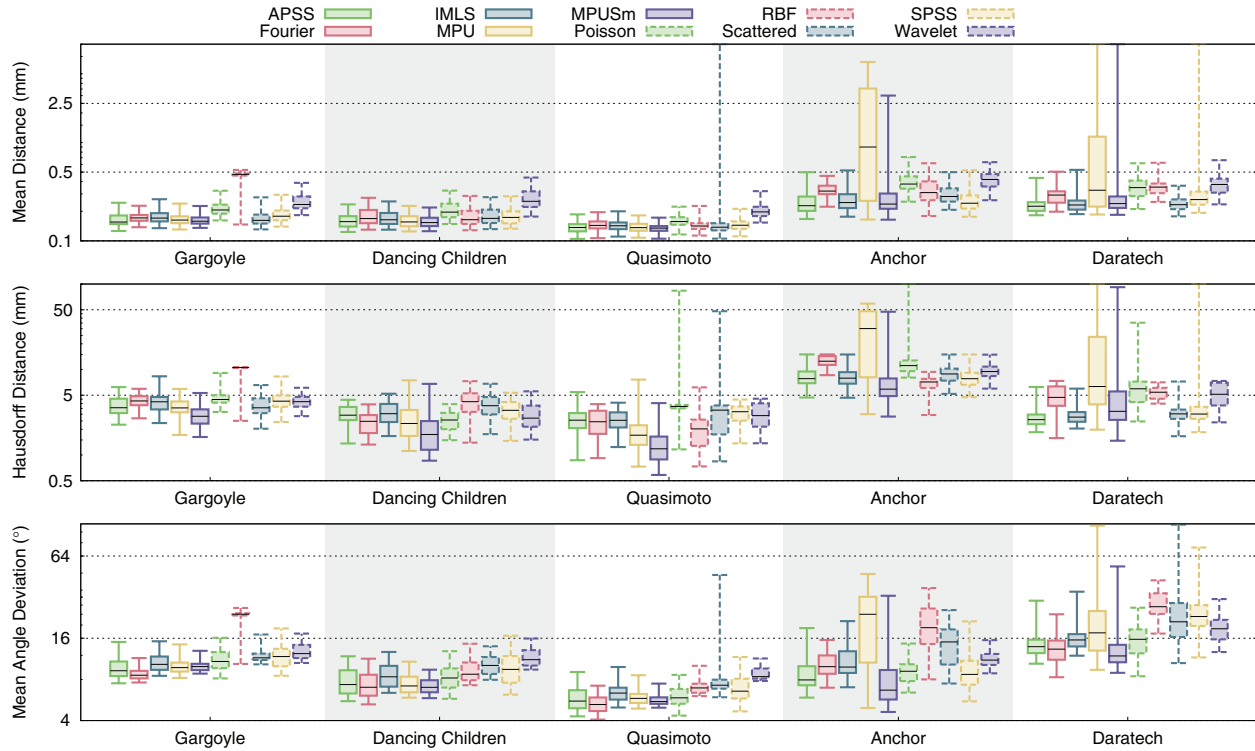
Fig. 10. Plots for all of the error distribution experiments. Each bar plot represents the distribution of a particular error measure for a given shape, sampled with a wide variety of scan parameters. The median provides a good indication of overall algorithmic performance for a given error measure, while the quartiles give an indication of algorithmic robustness.

is quite competitive, indicating it never strays too far from ground truth.

It is well-known that Poisson and Fourier tend to oversmooth the data, and in our experiments this is reflected in their rather large error in mean distance. However, in terms of Hausdorff distance and mean angle deviation they perform rather well, and are fairly consistent in their performance. This indicates that these algorithms are reliable in producing surfaces which remain close to the original, while also remaining close in differential quantities. We note that Fourier is more consistent than Poisson, as Poisson suffers a lack of resolution in regions of missing data.

While RBF performed well on the Dancing Children and Quasimoto models, on the Gargoyle model we see that it performed poorly across all metrics. The Gargoyle model is particularly difficult to sample as it has many concavities, where an undersampled concavity caused RBF to produce a thin crust throughout the interior of the shape rather than fill the hole.

We observe that for MPUSm its mean and Hausdorff distance is quite stable over all of the shapes. However, its normal error tends to not be consistent with the distance measures compared across different algorithms. We find that MPUSm produces slightly oversmoothed surfaces, and as such it can fail to capture fine-scale details. Hence, its normal error is comparatively larger. On the other hand, its smoothing of the MPU parameters can correct erroneous MPU shape fits, which likely explains its better performance over MPU in Hausdorff distance.

Sharp features. The Anchor and Daratech shapes are particularly difficult to reconstruct. As these are shapes with sharp features,

algorithms that only model smooth surfaces have difficulty with them. Additionally, these shapes have small topological features which are difficult to adequately scan due to occlusion. Hence we do not necessarily expect these algorithms to perform as well on these shapes as the others, and instead we use these shapes to measure robustness.

In observing MPU and MPUSm, we find instability in the presence of the Anchor and Daratech point clouds, where large spurious surface sheets are produced as a result of improperly fitting smooth shape functions to sharp features. However, note that the PSS methods perform much better, despite also using smooth shape functions. PSS methods fit shape functions at every point, hence the error will be contained locally if there exists a poor fit. MPU and MPUSm hierarchically fit a set of shape functions according to an error criterion, which can result in unbounded error if a poor fit exists. Interestingly, RBF performs quite well in distance, yet has rather large error in normals. The RBF interpolant tends to remain quite close to the surface, but produces spurious high-frequency details, hence the large normal deviations.

Topology. Overall, we find that the PSS methods and Scattered tend to perform quite well in the error metrics. However, these are also methods that produce holes in the presence of insufficient data. To demonstrate the performance of these algorithms in terms of topology, we also show how these algorithms behave in their number of connected components, total length of boundary components, whether or not the reconstructed mesh is manifold, and the deviation from the true genus, averaged over all point clouds, and shapes; see Table II. As shown, Fourier and Poisson tend to outperform

Table II. Summary of Error Distribution Experiments

algorithm	comps	bndry	manifold	genus	time
apss	47.37	140.86	0.50	1.82	36.02
fourier	1.54	0.00	1.00	0.49	28.70
imls	38.48	194.65	0.74	1.66	34.11
mpu	100.69	9.71	0.49	0.79	12.83
mpusmooth	2.88	2.93	0.91	0.67	17.83
poisson	1.54	0.44	1.00	0.63	36.83
rbf	51.73	6.30	0.82	13.55	34.78
scattered	1.90	214.21	1.00	7.47	4.48
spss	174.53	143.14	0.26	3.98	33.53
wavelet	1.35	0.04	1.00	0.71	2.13

Additional information for experiment 1, averaged across all point clouds and shapes. Here, *comps* refers to number of connected components, *bndry* is the length of boundary components, *manifold* is whether or not a mesh is manifold, 1 being manifold and 0 otherwise, *genus* refers to the amount which deviates from the actual genus, and *time* is in seconds.

these methods in all categories. With respect to the PSS methods, this demonstrates that they tend not to produce topologically clean implicit functions, likely due to their local nature. However, note that their genus remains relatively close to ground truth, whereas the local method Scattered performs poorly in genus.

7.2 Sparse Sampling

It is common in range scan data for certain areas of the surface to be sampled less densely than others. Here we investigate how reconstruction algorithms behave as data sparsity varies, where we treat sparsity as a controllable parameter. We are interested in observing how these algorithms infer the surface *between* the given input points.

In this experiment we only vary the sampling resolution. We fix the number of scans and camera positions such that the shape is sufficiently covered, that is, no missing data. We use the analytical normals of the surface, and no noise or misalignment. We use such clean input in order to restrict the problem to only data inference. We use the bumpy sphere as the test shape, as the coarse-scale features of the surface make data inference plausible.

See Figure 11 for plots of the experiment. MPUSm was unable to smooth its spherical covering on half of the point clouds due to the extreme sparsity, so we have omitted it from this experiment. From the distance measures we immediately see a partitioning of the algorithms: IMLS, Poisson, SPSS, and Wavelet all tend to behave rather poorly, while the other algorithms perform well. This is expected for Poisson and Wavelet, as the resolution of the output is proportional to the input size. However, it is interesting to observe the significant improvement of APSS over IMLS and SPSS, indicating that fitting spheres to sparse data is more advantageous than trying to fit planes to the data.

We also see that Fourier demonstrates remarkable robustness to sparse data. Fourier performs best among all algorithms when the data is very sparse, whereas APSS, MPU, RBF, and Scattered perform rather poorly on such data, though they perform better as resolution increases. However, observe that as the sampling resolution becomes somewhat dense, the distance error in APSS, MPU, and RBF steadily decreases while Fourier remains stagnant. This is a consequence of Fourier's inherent data smoothing. The algorithms that fit shape functions to the data comparatively improve their fits when the resolution increases.

7.3 Missing Data

Missing data will almost always be present in scanned data, simply due to concavities in the shape which can not be reached by the scanner or insufficient scanning due to physical restraints of the scanner. Here we generate incomplete point clouds by treating missing data as a controllable parameter, where we vary the peak threshold at which range is rejected. We note that this is quite common for scanners, since the accuracy of the scanner suffers when the angle at which the optical axis and the normal becomes large, and the preferred option may be to reject unacceptably noisy points.

Similar to the previous experiment, here we fix the number of scans and camera positions, and use no additive noise, in order to isolate missing data as the primary challenge in the input. We then vary the peak threshold at which to reject samples from 0.8 to 0.4, where 1 is the expected peak. We have used the bumpy sphere and Mailbox shapes, in order to observe the behavior of these algorithms in the presence of missing data on both smooth and sharp features.

See Figure 12 for plots of the experiment. We find that all of the indicator function methods perform quite well across both shapes with the notable exception of Wavelet, which fails to converge to the limit surface as missing data decreases. We credit the robustness of indicator function methods to the fact that they are global methods that do not attempt to fit shape functions.

Indeed, methods that fit shape functions have rather erratic behavior, particularly in the Mailbox shape. MPU, MPUSm, and RBF are quite unstable, producing spurious surface sheets as missing data is introduced. When the neighborhood of an edge is sampled on one side but not the other, extraneous surfaces may appear.

Scattered and the PSS methods tend to produce holes in surface regions where there are no samples. Similar to MPU, MPUSm, and RBF, the PSS methods can produce poor local shape fits in the presence of sharp features and missing data. However, the error is contained locally, for similar reasons discussed in Section 7.1.

7.4 Noise

Finally we consider how robust reconstruction algorithms are to noise in the range data. We consider two scan parameters that have a significant impact on noise, noise magnitude and laser frustum field of view. The effect of noise magnitude is fairly clear, however, we note that the thickness of the laser has a significant impact. The thicker the laser, the more difficult peak detection becomes at depth discontinuities, resulting in samples being detected where there is no surface.

To this end, we have taken the Spiral shape and sampled it with varying noise magnitudes and varying laser thickness. We sufficiently sample it so that missing data or sparsity is not an issue, and compute normals directly from the points, allowing for improper orientation if direction propagation is incorrect. For each algorithm and each point cloud we also manually set the parameters to perform best, considering the scale of the noise. For the PSS and indicator function methods, such parameter settings are quite intuitive as they are based on sampling density bandwidths. However, for all other methods a maximum error tolerance effectively determines the amount of smoothing performed, which can be quite sensitive.

See Figure 13 for plots of the noise experiments. Note that Fourier and Poisson, in terms of all error metrics, are quite robust in the presence of noise. This is likely due to the global nature of these methods, where smoothing the data is a natural consequence. As observed by its large variance, RBF performs rather poorly in the presence of noise. Indeed, the necessity to produce dipoles for
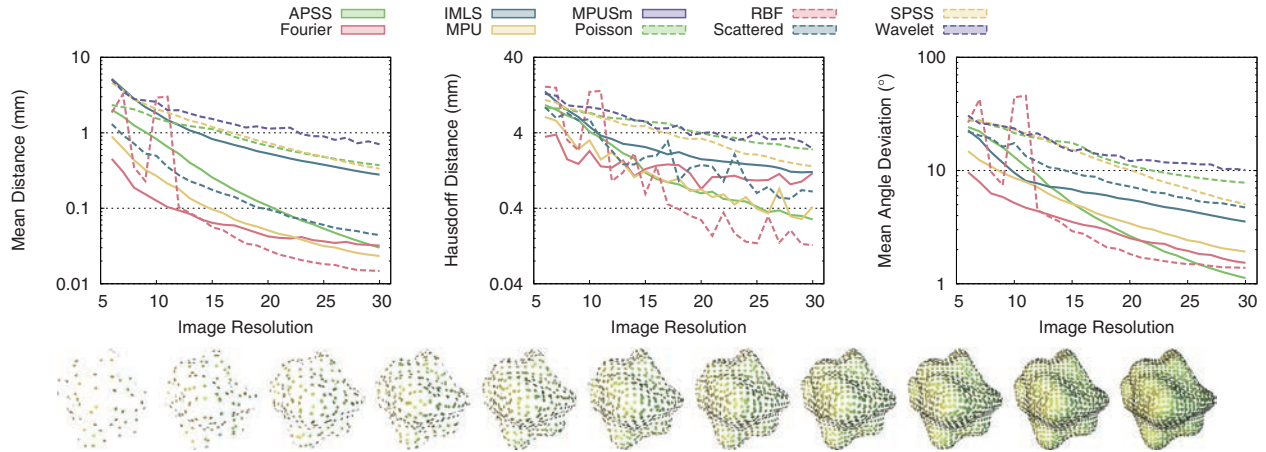
Fig. 11. Plots for the sparsity experiment, where we have sampled the bumpy sphere in increasing image resolution. The bottom row depicts a subset of these point clouds in decreasing sparsity. This experiment demonstrates how well these algorithms infer the surface from a sparse sampling.
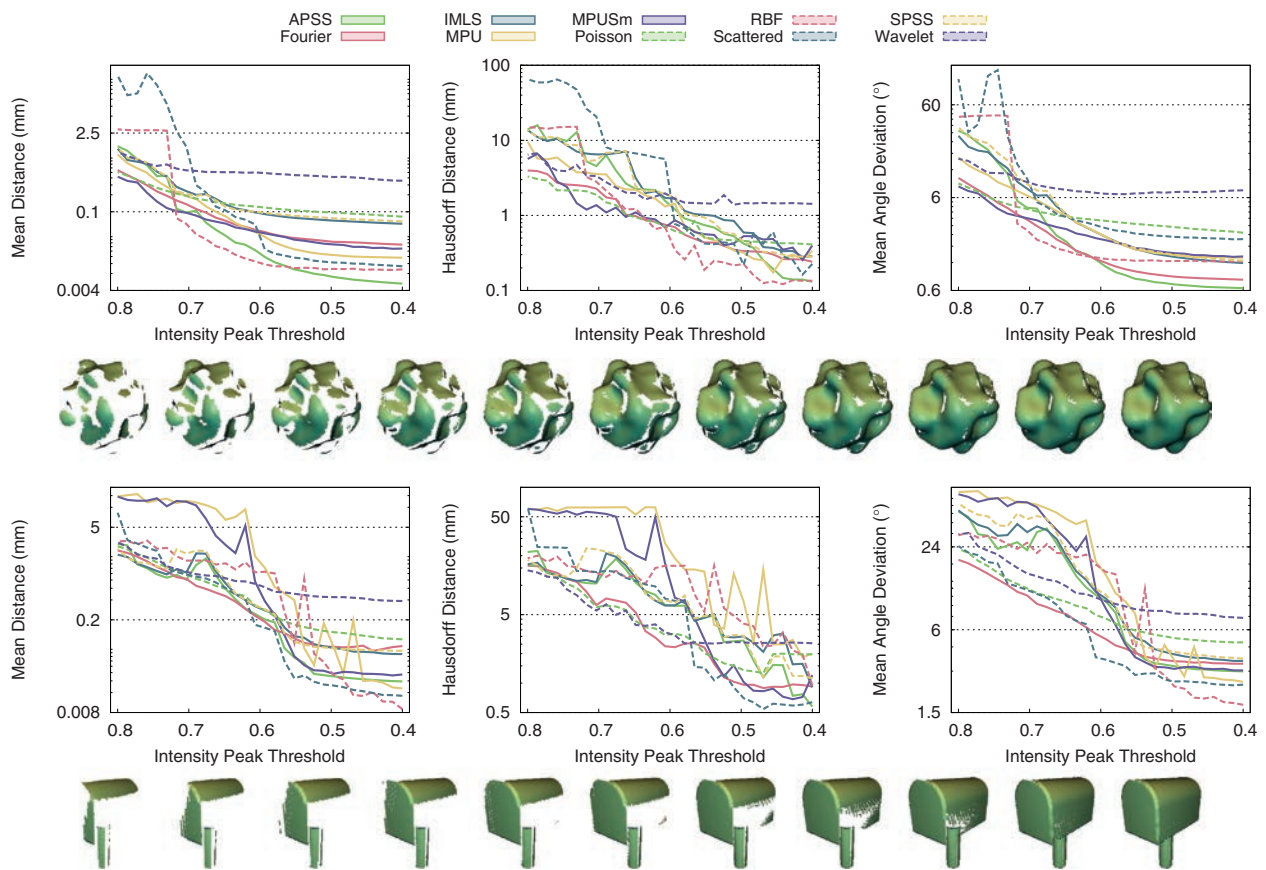

Fig. 12. Plots for the missing data experiments on the bumpy sphere (top row) and Mailbox (bottom row). We generate missing data by varying the peak intensity threshold at which the range is rejected. Note the differences in performance between the shape with smooth features and the shape with sharp features, as missing data is varied. The Mailbox model is courtesy of Dey et al. [2003].

RBF becomes especially problematic in the presence of noise and outliers.

We observe that MPU and MPUsm are somewhat robust in the presence of noise given their small variance in Hausdorff distance, though interestingly we see significant differences between them in the two different distance measures. The smoothing performed via MPUSm tends to expand the surface outward, resulting in poor mean distance, yet it never strays too far from
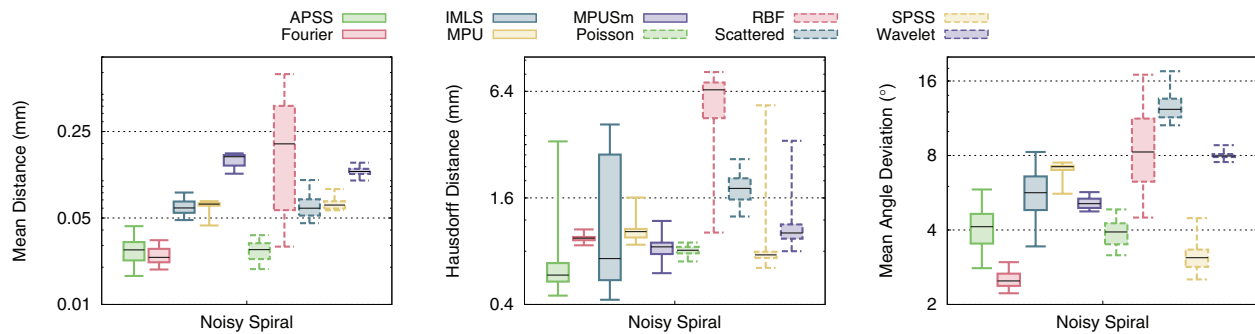
Fig. 13. Noise experiments for the Spiral shape. Here we vary noise level and laser thickness, and aggregate this into distributions. A small variance in a distribution is a good indication of robustness to noise.

Table III. Noisy Spiral Statistics

algorithm	comps	bndry	manifold	genus	time
apss	221.60	0.71	1.00	0.00	50.59
fourier	1.00	0.00	1.00	0.00	27.24
imls	193.16	4.76	1.00	0.00	48.62
mpu	1.20	0.00	1.00	0.00	7.13
mpusmooth	1.08	0.06	1.00	0.00	23.08
poisson	1.00	0.00	1.00	0.00	30.90
rbf	12.48	4.69	0.92	0.30	18.90
scattered	1.08	0.00	1.00	0.44	3.11
spss	257.20	1.13	1.00	0.00	48.18
wavelet	1.00	0.00	1.00	0.00	2.26

Additional information for the noisy Spiral experiments, averaged across all point clouds.

ground truth, hence its good behavior in terms of Hausdorff distance.

The PSS methods all tend to smooth out noise and remain robust to outliers. However, far away from the surface their behavior tends to be quite poor; see Table III. They tend to produce many extraneous connected components, as well as boundary components.

7.5 Discussion

Our small-scale experiments tend to correlate well with the results of the error distribution experiments. For instance, the unstable behavior of RBF in the presence of sparse and missing data manifests in its unstable behavior across the Gargoyle model, which is difficult to adequately sample due its numerous concavities. Likewise, the behavior of MPU and to a lesser extent MPUSm in the presence of missing data on the Mailbox correlates with their large variance in the Anchor and Daratech, indicative of the fact that they have trouble reconstructing sharp features. Observe that the stable behavior of Fourier in the small-scale experiments correlates well with its relatively small variance in the distribution plots.

Our experiments point toward a number of deficiencies in the state of surface reconstruction. Our results demonstrate the remarkable robustness of methods based on computation of the indicator functions, yet these methods tend to oversmooth the data, reflected in their poor performance in mean distance across complex shapes. Developing an algorithm based on the indicator function, which does not oversmooth the data, would be very useful. Conversely, although MLS methods perform rather well in terms of mean and

Hausdorff distance across the complex shapes, they demonstrate poor far-field behavior. We think that combining MLS methods with global constraints of some nature may rectify these issues.

Our benchmark should also prove useful for recent methods which resample point clouds with large missing data [Tagliasacchi et al. 2009; Cao et al. 2010; Shalom et al. 2010]. Although we have produced such point clouds in order to test robustness, it would be interesting to see how well these more recent resampling methods perform quantitatively.

All told, our benchmark consists of 351 point clouds across eight shapes, providing rich data for surface reconstruction developers. For our first set of experiments, we have 48 point clouds for each shape. Over 10 algorithms, this amounts to a total of 2400 different reconstruction outputs, and over both distance and normal correspondences we have a total of 4800 correspondence mappings. We think that this construction of a distribution of point clouds for a given shape could be used in other areas, for instance, potentially *learning* surface reconstruction, by using the point clouds and ground-truth data as training data.

Limitations. While the surfaces in our benchmark cover a broad range of shapes, they are by no means exhaustive. As surface reconstruction becomes more specialized, such as the reconstruction of large-scale architectural buildings [Nan et al. 2010], we envision our benchmark to expand to these specific forms of surfaces. Our implicit shape representation should easily be able to accommodate other types of shapes.

Although we have generated a large variety of point cloud data with our sampling scheme, we are keeping fixed certain settings which may be worth further exploration. For instance, we assume a diffuse BRDF in the scanning simulation, where it may be interesting to consider different forms of surface reflectance, and even spatially varying BRDFs. Though laser-based optical triangulation scanners are quite popular, other forms of scanning may be worth simulating in order to replicate different acquisition artifacts, such as time-of-flight scanners.

8. CONCLUSIONS

We have presented a benchmark for the evaluation and comparison of surface reconstruction algorithms, restricted to the class of algorithms that take an oriented point cloud as input and produce an approximating surface as output. Central to our benchmark is a mechanism for simulating point cloud data acquired from laser-based scanners. We use a broad class of implicit surfaces as reference shapes to sample, which allows us to obtain accurate quantitative measurements.

Our extensive experiments enable us to observe a wide range of behaviors across existing algorithms. For instance, global methods such as those which reconstruct the indicator function are very robust in the presence of noise, while more local methods such as MPU and MLS methods produce highly accurate reconstructions in the presence of clean data. The experiments point towards potential future work in surface reconstruction by illustrating the specific advantages and disadvantages in existing approaches.

By publicly releasing our data and code, researchers will now be able to benchmark their algorithms against existing algorithms and see where they stand. Additionally, our modeling and sampling methods will allow researchers to generate surfaces and point clouds tailored towards their interests. Hence we envision our benchmark to grow over time, continually incorporating data provided by the surface reconstruction community.

APPENDIX

A. CLOSED-FORM SOLUTION OF POLYGONAL WEIGHT FUNCTIONS

In this section, we detail the closed-form solution for Eq. (2) that is used in the formation of our implicit functions. The basic idea is to cast the integral into the local coordinate system of the triangle, and perform integration in terms of polar coordinates, analogous to the construction of Green coordinates [Lipman and Levin 2010].

For a given evaluation point $\mathbf{x}$ and triangle t composed of the vertices $\mathbf{p}_1$, $\mathbf{p}_2$, and $\mathbf{p}_3$, and normal $\mathbf{n}$, we project $\mathbf{x}$ onto the plane of t.

$$\tilde{\mathbf{x}} = \mathbf{x} + \langle \mathbf{p}_1 - \mathbf{x}, \mathbf{n} \rangle \mathbf{n} \tag{19}$$

Now, for a given $\mathbf{p} \in t$, $|\mathbf{x} - \mathbf{p}|^2 + \epsilon^2 = |\tilde{\mathbf{x}} - \mathbf{p}|^2 + |\mathbf{x} - \tilde{\mathbf{x}}|^2 + \epsilon^2 = |\tilde{\mathbf{x}} - \mathbf{p}|^2 + \lambda_1$, where $\lambda_1 = |\mathbf{x} - \tilde{\mathbf{x}}|^2 + \epsilon^2$ and is constant throughout the integration. We can now rewrite the integral as

$$\int_{\mathbf{p} \in t} w(\mathbf{x}, \mathbf{p}) d\mathbf{p} = \sum_{t_i} sgn(t_i) \int_{\mathbf{p} \in t_i} \frac{d\mathbf{p}}{(|\tilde{\mathbf{x}} - \mathbf{p}|^2 + \lambda_1)^2}, \tag{20}$$

where t is broken up into t_1, t_2, t_3, formed from the triangles composed of $\tilde{\mathbf{x}}$ and $\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3$, and sgn represents the orientation of the triangle: positive if oriented properly, and negative otherwise. See the left image of Figure 14 for an illustration of this decomposition.

Without loss of generality we consider a single triangle t_1. We now convert this integral into polar coordinates.

$$\int_{\mathbf{p} \in t_1} \frac{d\mathbf{p}}{(|\tilde{\mathbf{x}} - \mathbf{p}|^2 + \lambda_1)^2} = \int_{\theta=0}^{\theta=\beta} \int_{r=0}^{R(\theta)} \frac{r \, dr \, d\theta}{(r^2 + \lambda_1)^2}$$
$$= -\frac{1}{2} \int_0^\beta \frac{d\theta}{R(\theta)^2 + \lambda_1} + \frac{\beta}{2\lambda_1}$$

The integration is centered with respect to $\tilde{\mathbf{x}}$, where β is the angle in t_1 opposite $\tilde{\mathbf{x}}$, and $R(\theta)$ is the length parameterized by θ. See the middle image of Figure 14.

In order to have a clean parameterization of the length $R(\theta)$, we break up the integral into two parts by considering the orthogonal projection of the point $\tilde{\mathbf{x}}$ onto its opposing edge, $\hat{\mathbf{x}}$, and breaking t_1 into: $t_1^1 = \langle \tilde{\mathbf{x}}, \mathbf{p}_2, \hat{\mathbf{x}} \rangle$ and $t_1^2 = \langle \tilde{\mathbf{x}}, \hat{\mathbf{x}}, \mathbf{p}_3 \rangle$. Without loss of generality we consider t_1^1, and we obtain: $R(\theta) = \frac{|\tilde{\mathbf{x}} - \hat{\mathbf{x}}|}{\cos(\theta)}$; see the
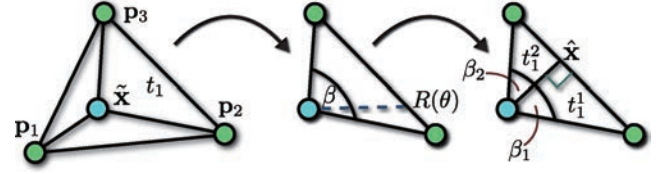

Fig. 14. We illustrate the decomposition of the integration of polygonal weight functions. We first decompose integration into three separate triangles (left), for such a single triangle perform integration in polar coordinates (middle), followed by breaking up the integration into simpler components through orthogonal projection onto the opposing edge (right).

right image of Figure 14. Hence the integral becomes

$$\int_0^{\beta_1} \frac{d\theta}{R^2(\theta) + \lambda_1} = sgn(t_1^1) \int_0^{\beta_1} \frac{\cos^2(\theta)}{|\tilde{\mathbf{x}} - \hat{\mathbf{x}}|^2 + \lambda_1 \cos^2(\theta)}$$
$$= sgn(t_1^1) \left(\frac{\beta_1}{\lambda_1} - \frac{|\tilde{\mathbf{x}} - \hat{\mathbf{x}}|^2}{\lambda_1} \int_0^{\beta_1} \frac{d\theta}{|\tilde{\mathbf{x}} - \hat{\mathbf{x}}|^2 + \lambda_1 \cos^2(\theta)} \right),$$

where $sgn(t_1^1)$ is the sign of the orientation of the triangle, which may be negative if $\hat{\mathbf{x}}$ projects outside of t_1. Applying the double angle formula to the preceding integral we obtain

$$= \int_0^{\beta_1} \frac{d\theta}{\left(|\tilde{\mathbf{x}} - \hat{\mathbf{x}}|^2 + \frac{\lambda_1}{2} \right) + \frac{\lambda_1}{2} \cos(2\theta)},$$

Setting $b = |\tilde{\mathbf{x}} - \hat{\mathbf{x}}|^2 + \frac{\lambda_1}{2}$ and $c = \frac{\lambda_1}{2}$, we may apply the relevant antiderivative [Abramowitz and Stegun 1964] to obtain

$$\int \frac{d\theta}{b + c \cos(2\theta)} = \frac{1}{\sqrt{b^2 - c^2}} \tan^{-1} \left\{ \sqrt{\frac{b-c}{b+c}} \tan \theta \right\} + C.$$

B. DESCRIPTION OF SYNTHETIC SCANNER

Here we provide additional details on our synthetic scanner, as described in Section 5.1. To clarify the following discussion, we note that for each shape in our benchmark we have set its maximum dimension to be 70mm. Hence any scanning parameter based on distance is defined with respect to the bound of 70mm. Additionally, we place an upper bound on the radiance to be 1.

Our synthetic scanner is controlled by the following parameters.

—*Image resolution*. The image resolution, in conjunction with the number of scans used, effectively defines the resolution of the point cloud.

—*Baseline distance*. A small baseline distance magnifies depth errors in triangulation, while a large baseline results in greater occlusion. We have fixed our baseline to be with respect to the x-axis of the camera, though this may easily be adjusted to the y-axis by changing the laser sweep direction. We found that baseline distances ranging from 10mm to 150mm provide good variety in triangulation accuracy and occlusions.

—*Stripe frustum field of view*. The thickness of the laser stripe has an impact on peak detection, in appropriately fitting a Gaussian. By default, we set the field of view such that the number of pixels visible within a distance of 50mm from the camera is roughly 10, which is a function of the image resolution.

—*Stripe resolution*. The number of laser stripes to project impacts the resolution of the depth. By default, we set this to be the x resolution of the camera, in order to obtain sufficient coverage. Setting the stripe resolution to be lower than the x resolution

may result in some points not being affected by the laser stripes. By assigning a sufficiently large stripe frustum field of view, one may be able to obtain sufficient coverage.

—*Noise magnitude.* The magnitude of the noise corrupts the laser projection, making peak detection imprecise. Typical noise magnitudes we have used range from 0, or no noise, to 0.6, which can greatly corrupt the radiance signal.

—*Radiance smoothing bandwidth.* Smoothing the radiance image reduces noise, though at the potential cost of sacrificing the expected Gaussian laser profile. The bandwidth to use is largely dependent on the stripe frustum field of view and noise level. For instance, a thick laser with large noise magnitude will require a fairly large bandwidth to sufficiently smooth out the noise. We note that smoothing, in conjunction with additive noise, may result in a radiance signal with smaller peak magnitudes, which can impact the peak magnitude threshold.

—*Peak magnitude threshold.* For large thresholds this will reject parts of the surface whose radiance signal is determined weak by a pixel's corresponding Gaussian fit. This is a major cause of missing data. For a laser containing little or no noise, typical thresholds range from 0.8, which will result in only highly confident range data, to 0.1, which will result in the rejection of few points. Under noise and radiance smoothing, the peak threshold must be adjusted to account for an expected reduction in peak magnitude.

—*Variance threshold.* Range at depth discontinuities are likely to be rejected with this threshold. We set the variance with respect to the width of the laser, where by default we only reject range whose variance in the Gaussian fit is larger than twice that of the laser width. Similar to the peak magnitude threshold, the variance threshold is sensitive to the noise magnitude and smoothing bandwidth.

We note that in our experiments, although we have generated quite a large number of point clouds, we have hardly explored the full parameter space of our scanner. By publicly releasing our synthetic scanner software, surface reconstruction researchers and practitioners will be able to replicate specific scanning conditions of interest.

ACKNOWLEDGMENTS

REFERENCES

ABBASINEJAD, F., KIL, Y., SHARF, A., AND AMENTA, N. 2009. Rotating scans for systematic error removal. *Comput. Graph. Forum 28*, 1319–1326.

ABRAMOWITZ, M. AND STEGUN, I. 1964. *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables.* Vol. 55, Dover Publications.

ADAMSON, A. AND ALEXA, M. 2003. Approximating and intersecting surfaces from points. In *Proceedings of the Eurographics/ACM SIGGRAPH Symposium on Geometry Processing.* Eurographics Association, 230–239.

ALEXA, M., BEHR, J., COHEN-OR, D., FLEISHMAN, S., LEVIN, D., AND SILVA, C. 2003. Computing and rendering point set surfaces. *IEEE Trans. Vis. Comput. Graph. 9*, 3–15.

ALLIEZ, P., COHEN-STEINER, D., TONG, Y., AD DESBRUN, M. 2007. Voronoi-Based variational reconstruction of unoriented point sets. In *Proceedings of the Symposium on Geometry Processing.* Eurographics Association, 39–48.

AMENTA, N. AND BERN, M. 1999. Surface reconstruction by Voronoi filtering. *Discr. Comput. Geom. 22*, 4, 481–504.

AMENTA, N., CHOI, S., DEY, T. K., AND LEEKHA, N. 2002. A simple algorithm for homeomorphic surface reconstruction. *Int. J. Comput. Geom. Appl. 12*, 1–2, 125–141.

AMENTA, N., CHOI, S., AND KOLLURI, R. 2001. The power crust. In *Proceedings of the ACM Symposium on Solid Modeling and Applications.* ACM Press, New York, 249–266.

BARIBEAU, R. AND RIOUX, M. 1991. Influence of speckle on laser range finders. *Appl. Opt. 30*, 20, 2873–2878.

BOISSONNAT, J.-D. AND CAZALS, F. 2002. Smooth surface reconstruction via natural neighbour interpolation of distance functions. *Comput. Geom. 22*, 1–3, 185–203.

BOUGUET, J. 2010. Camera calibration toolbox for Matlab. http://www.vision.caltech.edu/bouguetj/calib_doc

BROWN, B. AND RUSINKIEWICZ, S. 2007. Global non-rigid alignment of 3-D scans. *ACM Trans. Graph. 26*, 3, 21:1–21:10.

CAO, J., TAGLIASACCHI, A., OLSON, M., ZHANG, H., AND SU, Z. 2010. Point cloud skeletons via Laplacian based contraction. In *Proceedings of the International Conference on Shape Modeling (SMI'10).* IEEE, 187–197.

CAZALS, F. AND GIESEN, J. 2006. Delaunay triangulation based surface reconstruction. In *Effective Computational Geometry for Curves and Surfaces,* J.-D. Boissonnat and M. Teillaud, Eds., Springer, 231–276.

CIGNONI, P., MONTANI, C., AND SCOPIGNO, R. 1998. A comparison of mesh simplification algorithms. *Comput. Graph. 22*, 1, 37–54.

CURLESS, B. AND LEVOY, M. 1995. Better optical triangulation through spacetime analysis. In *Proceedings of the 5th International Conference on Computer Vision.* IEEE, 987–994.

CURLESS, B. AND LEVOY, M. 1996. A volumetric method for building complex models from range images. In *Proceedings of the Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH'96).* 303–312.

DEY, T., WOO, H., AND ZHAO, W. 2003. Approximate medial axis for cad models. In *Proceedings of the 8th ACM Symposium on Solid Modeling and Applications.* ACM Press, New York, 280–285.

DEY, T. K. 2007. *Curve and Surface Reconstruction: Algorithms with Mathematical Analysis.* Cambridge University Press.

DEY, T. K., LI, G., AND SUN, J. 2005. Normal estimation for point clouds: A comparison study for a Voronoi based method. In *Proceedings of the Symposium on Point-Based Graphics.* Eurographics Association, 39–46.

FLEISHMAN, S., COHEN-OR, D., AND SILVA, C. T. 2005. Robust moving least-squares fitting with sharp features. *ACM Trans. Graph. 24*, 3, 544–552.

FRUEH, C., JAIN, S., AND ZAKHOR, A. 2005. Data processing algorithms for generating textured 3D building façade meshes from laser scans and camera images. *Int. J. Comput. Vis. 61*, 2, 159–184.

FUNKHOUSER, T., SHIN, H., TOLER-FRANKLIN, C., CASTANEDA, A., BROWN, B., DOBKIN, D., RUSIKIEWICZ, S., AND WEYRICH, T. 2011. Learning how to match fresco fragments. *J. Comput. Cult. Herit. 4*, 2, 7.

GARLAND, M. AND HECKBERT, P. S. 1997. Surface simplification using quadric error metrics. In *Proceedings of the 24th Annual Conference on*

Computer Graphics and Interactive Techniques (SIGGRAPH'97). ACM Press/Addison-Wesley Publishing, New York, 209–216.

GUENNEBAUD, G. AND GROSS, M. 2007. Algebraic point set surfaces. *ACM Trans. Graph. 26*, 3, 23:1–23:10.

HILDEBRANDT, K., POLTHIER, K., AND WARDETZKY, M. 2006. On the convergence of metric and geometric properties of polyhedral surfaces. *Geometriae Dedicata I23*, 1, 89–112.

HOPPE, H., DEROSE, T., DUCHAMP, T., MCDONALD, J., AND STUETZLE, W. 1992. Surface reconstruction from unorganized points. In *Proceedings of the 19th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH'92)*. ACM Press, New York, 71–78.

KAZHDAN, M. 2005. Reconstruction of solid models from oriented point sets. In *Proceedings of the Symposium on Geometry Processing*. Eurographics Association, 73–82.

KAZHDAN, M., BOLITHO, M., AND HOPPE, H. 2006. Poisson surface reconstruction. In *Proceedings of the Symposium on Geometry Processing*. Eurographics Association, 61–70.

KOLLURI, R. 2005. Provably good moving least squares. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'05)*. 1008–1017.

LEVOY, M., PULLI, K., CURLESS, B., RUSINKIEWICZ, S., KOLLER, D., PEREIRA, L., GINZTON, M., ANDERSON, S., DAVIS, J., GINSBERG, J., ET AL. 2000. The digital Michelangelo project: 3D scanning of large statues. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH'00)*. ACM Press/Addison-Wesley Publishing, 131–144.

LIPMAN, Y. AND LEVIN, D. 2010. Derivation and analysis of Green coordinates. *Comput. Methods Funct. Theory 10*, 1, 167–188.

MANSON, J., PETROVA, G., AND SCHAEFER, S. 2008. Streaming surface reconstruction using wavelets. *Comput. Graph. Forum 27*, 1411–1420.

MEYER, M., DESBRUN, M., SCHRODER, P., AND BARR, A. 2002. Discrete differential-geometry operators for triangulated 2-manifolds. *Vis. Math. 3*, 7, 34–57.

MEYER, M., KIRBY, R., AND WHITAKER, R. 2007. Topology, accuracy, and quality of isosurface meshes using dynamic particles. *IEEE Trans. Vis. Comput. Graph. 13*, 6, 1704–1711.

MITRA, N. J. AND NGUYEN, A. 2003. Estimating surface normal in noisy point cloud data. In *Proceedings of the 19th Annual Symposium on Computational Geometry (SCG'03)*. ACM Press, New York, 322–328.

NAGAI, Y., OHTAKE, Y., AND SUZUKI, H. 2009. Smoothing of partition of unity implicit surfaces for noise robust surface reconstruction. *Comput. Graph. Forum 28*, 1339–1348.

NAN, L., SHARF, A., ZHANG, H., COHEN-OR, D., AND CHEN, B. 2010. SmartBoxes for interactive urban reconstruction. In *ACM SIGGRAPH Papers*. ACM Press, New York, 1–10.

NEXTENGINE. 2011. NextEngine 3d laser scanner. http://www.nextengine.com

OHTAKE, Y., BELYAEV, A., ALEXA, M., TURK, G., AND SEIDEL, H. 2003. Multi-Level partition of unity implicits. *ACM Trans. Graph. 22*, 3, 463–470.

OHTAKE, Y., BELYAEV, A., AND SEIDEL, H. 2005a. An integrating approach to meshing scattered point data. In *Proceedings of the ACM Symposium on Solid and Physical Modeling*. ACM Press, New York, 61–69.

OHTAKE, Y., BELYAEV, A. G., AND SEIDEL, H.-P. 2005b. 3D scattered data interpolation and approximation with multilevel compactly supported rbfs. *Graph. Models 67*, 3, 150–165.

REGLI, W. AND GAINES, D. 1997. A repository for design, process planning and assembly. *Comput.-Aid. Des. 29*, 12, 895–905.

SCHARSTEIN, D. AND SZELISKI, R. 2002. A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *Int. J. Comput. Vis. 47*, 1, 7–42.

SEITZ, S., CURLESS, B., DIEBEL, J., SCHARSTEIN, D., AND SZELISKI, R. 2006. A comparison and evaluation of multi-view stereo reconstruction algorithms. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. Vol. 1, IEEE Computer Society, 519–528.

SHALOM, S., SHAMIR, A., ZHANG, H., AND COHEN-OR, D. 2010. Cone carving for surface reconstruction. *ACM Trans. Graph. 29*, 6, 150:1–150:10.

SHAPEWAYS. 2011. Shapeways. http://www.shapeways.com

SHEN, C., O'BRIEN, J., AND SHEWCHUK, J. 2004. Interpolating and approximating implicit surfaces from polygon soup. *ACM Trans. Graph. 23*, 3, 896–904.

SÜßMUTH, J., MEYER, Q., AND GREINER, G. 2010. Surface reconstruction based on hierarchical floating radial basis functions. *Comput. Graph. Forum 29*, 1854–1864.

TAGLIASACCHI, A., ZHANG, H., AND COHEN-OR, D. 2009. Curve skeleton extraction from incomplete point cloud. *ACM Trans. Graph. 28*, 3, 71:1–71:9.

TER HAAR, F., CIGNONI, P., MIN, P., AND VELTKAMP, R. 2005. A comparison of systems and tools for 3D scanning. In *3D Digital Imaging and Modeling: Applications of Heritage, Industry, Medicine and Land Workshop* (Session P.12).

VLASIC, D., PEERS, P., BARAN, I., DEBEVEC, P., POPOVIC, J., RUSINKIEWICZ, S., AND MATUSIK, W. 2009. Dynamic shape capture using multi-view photometric stereo. *ACM Trans. Graph. 28*, 5, 174:1–174:11.