HAL
open science

A Survey of Surface Reconstruction from Point Clouds

Matthew Berger, Andrea Tagliasacchi, Lee Seversky, Pierre Alliez, Gael Guennebaud, Joshua Levine, Andrei Sharf, Claudio Silva

To cite this version:

Matthew Berger, Andrea Tagliasacchi, Lee Seversky, Pierre Alliez, Gael Guennebaud, et al.. A Survey of Surface Reconstruction from Point Clouds. Computer Graphics Forum, 2016, pp.27. 10.1111/cgf.12802 . hal-01348404v2

HAL Id: hal-01348404

https://inria.hal.science/hal-01348404v2

Submitted on 21 Dec 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Volume 0 (1981), Number 0 pp. 1–27 COMPUTER GRAPHICS *forum*

A Survey of Surface Reconstruction from Point Clouds

Matthew Berger[1] Andrea Tagliasacchi[2,3] Lee M. Seversky[1]

Pierre Alliez[3] Gaël Guennebaud[4] Joshua A. Levine[5] Andrei Sharf[6] Claudio T. Silva[7]

[1]Air Force Research Laboratory, Information Directorate [2]École Polytechnique Fédérale de Lausanne (EPFL)
[3]University of Victoria [4]Inria Sophia Antipolis-Méditerranée [5]Inria Bordeaux Sud-Ouest [6]Clemson University
[7]Ben-Gurion University [8]New York University

Abstract

The area of surface reconstruction has seen substantial progress in the past two decades. The traditional problem addressed by surface reconstruction is to recover the digital representation of a physical shape that has been scanned, where the scanned data contains a wide variety of defects. While much of the earlier work has been focused on reconstructing a piece-wise smooth representation of the original shape, recent work has taken on more specialized priors to address significantly challenging data imperfections, where the reconstruction can take on different representations – not necessarily the explicit geometry. We survey the field of surface reconstruction, and provide a categorization with respect to priors, data imperfections, and reconstruction output. By considering a holistic view of surface reconstruction, we show a detailed characterization of the field, highlight similarities between diverse reconstruction techniques, and provide directions for future work in surface reconstruction.

1. Introduction

Advances made by the computer graphics community have revolutionized our ability to digitally represent the world around us. One subfield that has blossomed during this revolution is surface reconstruction. At its core, *surface reconstruction* is the process by which a 3D object is inferred, or "reconstructed", from a collection of discrete points that sample the shape. This survey compiles the major directions and progress of the community that addresses variants of the basic problem, and it reflects on how emerging hardware technology, algorithmic innovations, and driving applications are changing the state-of-the-art.

Surface reconstruction came to importance primarily as a result of new techniques to acquire 3D point clouds. Early on, these technologies ranged from active methods such as optical laser-based range scanners, structured light scanners, and LiDAR scanners to passive methods such as multi-view stereo. These devices fundamentally changed the way we accomplished engineering and rapid prototyping tasks, and they have improved hand-in-hand with technologies for computer-aided design.

Computer graphics took an immediate interest in such technology, following one of its longstanding goals: the modeling, recognition, and analysis of the real world. Moreover, current applications have made use of such scanners in all fields of data-driven science, spanning from the micro- to macro-scale. A more recent trend has seen the massive proliferation of point clouds from inexpensive commodity real-time scanners such as the Microsoft Kinect. This has impacted varied fields including automotive design, engineering, archaeology, telecommunications, and art.

In many ways, it is these new acquisition methods that pose the most significant challenge for surface reconstruction. Each distinct method tends to produce point clouds containing a variety of properties and imperfections. These properties, in conjunction with the nature of the scanned shape, effectively distinguish the class of reconstruction methods that exist today. This diverse set of techniques ranges from methods that assume a well-sampled point cloud, generalize to arbitrary shapes, and produce a watertight surface mesh, to methods that make very loose assumptions on the quality of the point cloud, operate on specific classes of shapes, and output a non-mesh based shape representation.

Our survey presents surface reconstruction algorithms from the perspective of *priors*: assumptions made by algorithms in order to combat imperfections in the point cloud and to

© 2016 Author version

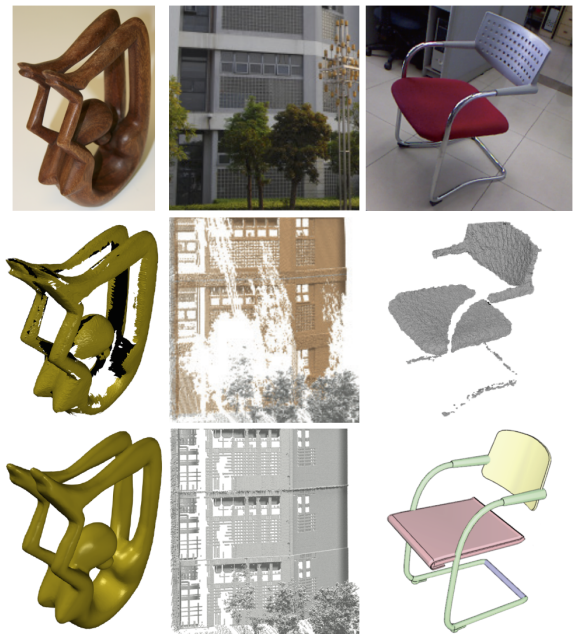

Figure 1: *Surface reconstruction has grown in diversity in recent years resulting in techniques taking on specialized priors.* ROSA *[TZCO09], shown on the left, uses volumetric smoothness to aid in reconstruction. Non-local consolidation [ZSW*10], shown in the middle, uses global regularity in the form of structural repetition. Part composition [SFCH12], shown on the right, uses data-driven techniques.*

eventually focus what information about the shape is reconstructed. Without prior assumptions, the reconstruction problem is ill-posed; an infinite number of surfaces can pass through (or near) a given set of data points. Assumptions are usually imposed on the point cloud itself, such as sampling density, level of noise, and misalignment. But just as importantly they are also imposed on the scanned shape, such as local surface smoothness, volumetric smoothness, absence of boundaries, symmetries, shape primitives, global regularity, and data-driven assumptions. In some instances, requirements are made on knowledge of the acquisition, such as scanner head position, as well as RGB images of the object. In other cases, the user is involved in prescribing high-level cues for reconstruction. All of these factors permit the regularization of the otherwise ill-posed problem of surface reconstruction, particularly when processing a point cloud containing severe imperfections. Figure 1 depicts several different priors used to reconstruct surfaces from challenging point clouds.

The utility of classifying surface reconstruction in terms of priors also helps to constrain expectations and prioritize desirables for reconstruction output in application dependent ways. For example, in the field of archaeology, dense, good-coverage scans might be available (e.g. the Digital Michaelangelo project [LPC*00]), allowing for standard smoothness priors that enable the high-detail reconstruction. Such great detail enables more than just digitization, but also the preservation of culture in ways that will transform how we study art [SCC*11]. In urban planning, reconstructing fine details such as individual bricks on a building might be both unnecessary and impossible given incomplete terrestrial LiDAR scans. However, in such contexts global regularity priors enable detail completion as well as interactive, realistic proxies for missing details such as plants and vegetation [BAMJ*11].

Historically, priors have evolved in conjunction with the types of point clouds being processed. For instance, surface smoothness priors were developed primarily to handle small objects acquired from desktop scanners. Mobile, real-time scanners have enabled the dynamic acquisition of more general scenes, rather than single objects, prompting more specialized structural and data-driven priors. Since priors tend to be coupled with the type of acquisition, we argue that our perspective of surface reconstruction is beneficial for understanding how to process future types of acquired point clouds.

Organization. Our survey is organized as follows. Sections 2–3 provide background on surface reconstruction. In Section 2 we examine the basic characteristics of a reconstruction algorithm, namely:

- **Point Cloud Artifacts:** the imperfections of the point cloud that the method is able to effectively handle.
- **Input Requirements:** the types of inputs associated with a point cloud required by the algorithm.

Section 3 provides an overview of the different priors from the perspective of the type of data produced through acquisition, the shape classes that tend to be acquired, and the type of output produced. We use all of these considerations as a way of examining surface reconstruction methods, starting with traditional surface smoothness priors in Section 4, and delving into specialized priors in Sections 5–10. In Table 1.1 we provide a summary of surface reconstruction methods by prior, characterizing their input and output, as well as their level of robustness to various artifacts. We discuss methods for evaluating surface reconstruction in Section 11, and conclude in Section 12 with a discussion on future trends.

1.1. Survey Scope and Related Works

There are many variants to surface reconstruction. This survey focuses on those relating to the reconstruction from point clouds of *static* objects and scenes acquired through *3D scanners*, wherein the point cloud contains a considerable level of *imperfection*. Furthermore, we concentrate on methods that *approximate* the input point cloud. For clarity, we contrast this organization relative to other important types of surface reconstruction:

Urban reconstruction. Our survey covers a wide variety of reconstruction methods, with *urban reconstruction* from

© 2016 Author version

Method	Point Cloud Artifacts					Input Requirements				Shape Class	Reconstruction Output
	nonuniform sampling	noise	outliers	misalignment	missing data	unoriented normals	oriented normals	scanner information	RGB image		
Surface Smoothness											
Tangent Planes [HDD*92]	○	○								general	implicit field
RBF [CBC*01]	○				○		✓			general	implicit field
MLS [ABCO*03]	○	○				✓				general	point set
MPU [OBA*03]	○	○			○		✓			general	implicit field
Poisson [KBH06]	○	●	○	○	○		✓			general	implicit field
Graph Cut [HK06]	○	○	○	○	○					general	volumetric segmentation
Unoriented Indicator [ACSTD07]	○	●	○	○	○	✓				general	implicit field
LOP [LCOLTE07]	●	●	○	○						general	point set
Dictionary Learning [XZZ*14]	●	●	●	○						general	mesh
Visibility											
VRIP [CL96]	○	●			○			✓		general	implicit field
TVL1-VRIP [ZPB07]	○	●	○	○	○			✓		general	implicit field
Signing the Unsigned [MDGD*10]	○	●	●		○	✓				general	implicit field
Cone Carving [SSZCO10]	○	○			●		✓	✓		general	implicit field
Multi-Scale Scan Merge [FG11]	●	●			○			✓		general	implicit field
Volumetric smoothness											
ROSA [TZCO09]	○	○			●		✓			organic	skeleton curve
Arterial Snakes [LLZM10]	○	○			●		✓			man-made	skeleton curve
VASE [TOZ*11]	○	○			●			✓		general	implicit field
l_1 Skeleton [HWCO*13]	○	○			●					organic	skeleton curve
Geometric Primitives											
Primitive Completion [SDK09]	○	○	○		●		✓			CAD	volumetric segmentation
Volume Primitives [XF12]	○	○	○		●			✓		indoor environment	interior volume
Point Restructuring [LA13]	○	○	○	○	○	✓		✓		general	volumetric segmentation
CCDT [vKvLV13]	○	○	○		○	✓		✓		urban environment	volumetric segmentation
Global Regularity											
Symmetry [PMW*08]	○	○			●	✓				architectural	point set
Nonlocal Consolidation [ZSW*10]	●	○	○		●	✓				architectural	point set
2D-3D Facades [LZS*11]	○	○			●	✓			✓	architectural	point set
Globfit [LWC*11]	●	●	○	●		✓				man-made	primitive relations
RAPter [MMBM15]	●	●	○	●	○					indoor environment	primitive relations
Data-driven											
Completion by Example [PMG*05]	○	○			●		✓			general	point set
Semantic Modeling [SXZ*12]	○	○			●	✓			✓	indoor scene objects	deformed model
Shape Variability [KMYG12]	○	○			●	✓				indoor scene objects	deformed model
Part Composition [SFCH12]	○	○			●	✓			✓	man-made	deformed model parts
Interactive											
Topological Scribble [SLS*07]	○	○			●		✓			general	implicit field
Smartboxes [NSZ*10]	●	○	○		●	✓				architectural	primitive shapes
O-Snap [ASF*13]	○	○	○		●	✓				architectural	primitive shapes
Morfit [YHZ*14]	○	○			●					general	skeleton + mesh

Table 1: *A categorization of surface reconstruction in terms of the type of priors used, the ability to handle point cloud artifacts, input requirements, shape class, and reconstruction output. Here ○ indicates that the method is moderately robust to a particular artifact and ● indicates that the method is very robust. ✓indicates an input requirement and ✓-indicates optional input.*

point clouds among them. We note that [MWA*13] surveys urban reconstruction more broadly: 3D reconstruction from images, image-based facade reconstruction, as well as reconstruction from 3D point clouds. Although there exists some overlap between the body of surveyed material, we cover these methods in a different context, namely the priors that underlay the reconstruction methods and how they address challenges in point cloud reconstruction.

Interpolatory reconstruction. An important field of surface reconstruction methods are those that *interpolate* a point cloud without any additional information, such as normals or scanner information. Delaunay-based methods are quite common in this area. The basic idea behind these methods is that the reconstructed triangulated surface is formed by a subcomplex of the Delaunay triangulation. A comprehensive survey of these methods is presented in [CG06], as well as the monograph of [Dey07]. A very attractive aspect of such methods is that they come with provable guarantees in the geometric and sometimes topological quality of the reconstruction provided the sampling of the input surface is sufficiently dense. Nonetheless, these requirements can be too severe for point clouds encountered in the wild, thus rendering the methods impractical for scanned, real-world scenes containing significant imperfections. We do not cover these methods, as our focus on reconstruction emphasizes how challenging artifacts are dealt with, though we note that there are some recent interpolatory approaches which are equipped to handle moderate levels of noise – see [DMSL11] for a scale-space approach to interpolatory reconstruction.

Dynamic reconstruction. Another recent advance in scanning techniques has enabled the acquisition of point clouds that vary dynamically. These devices promise the ability to capture more than just a single static object, but one that changes over time. Such techniques are still in their infancy, so they have not as yet explored their full range of potential applications. While we focus our survey on the broad spec-

© 2016 Author version

trum of techniques associated with static point clouds, it is interesting to note that already a prior-oriented viewpoint (e.g. incompressibility [SAL*08] and gradual change [PSDB*10]) has fueled a better understanding of shape in space-time.

Single-view reconstruction from images. Very recent work has considered the problem of surface reconstruction from single RGB images [SHM*14, HWK15]. These techniques are data-driven, in that they rely on large shape collections to estimate either depth or a full 3D model strictly from a single image. Due to the rapid growth in this new area, and the fact that our primary focus is on reconstruction from point clouds, we have decided not to cover this area in the survey.

2. Surface Reconstruction Fundamentals

Surface reconstruction methods must handle various types of imperfections and make certain requirements on input associated with the point cloud. Here we summarize these properties in order to cover the basic principles underlying surface reconstruction.

2.1. Point Cloud Artifacts

The properties of the input point cloud are an important factor in understanding the behavior of reconstruction methods. Here we provide a characterization of point clouds according to properties that have the most impact on reconstruction algorithms: *sampling density*, *noise*, *outliers*, *misalignment*, and *missing data*. See Figure 2 for a 2D illustration of these artifacts.

Sampling density. The distribution of the points sampling the surface is referred to as sampling density. Sampling density is important in surface reconstruction for defining a *neighborhood* – a set of points close to a given point which captures the local geometry of the surface, such as its tangent plane. A neighborhood should be large enough so that the points sufficiently describe the local geometry, yet small enough so that local features are preserved. Under uniform sampling density, a neighborhood may be constructed at every point in the same manner. For instance, one can define a neighborhood at a point $\mathbf{p} \in P$ via an ε–ball, defined as the set of points $N_\varepsilon(\mathbf{p}) \subset P$ such that each $\mathbf{y} \in N_\varepsilon(\mathbf{p})$ satisfies $\|\mathbf{p} - \mathbf{y}\| < \varepsilon$, under a single value ε used at all points.

3D scans typically produce a *nonuniform* sampling on the surface, where the sampling density spatially varies. This can be due to the distance from the shape to the scanner position, the scanner orientation, as well as the shape's geometric features. See Figure 2(b) for an illustration of nonuniform sampling on a curve. To capture the local variation in sampling density, a common approach is to use the k nearest neighbors (knn) at a given point for neighborhood construction. Another alternative is to use a spatially-varying ε–ball, commonly defined as a function of a point's knn neighborhood [GG07].

More sophisticated sampling density estimation techniques use reconstruction error bounds [LCOL06] and kernel methods [WSS09]. The method of [LCOL06] finds the neighborhood size at each point by bounding the error of a moving least squares surface approximation [ABCO*03], where the selected ε minimizes this error bound. The work of [WSS09] formulates the error of a moving least squares surface approximation in terms of kernel regression, where the optimal neighborhood size may be defined via a point-wise error or the error in the support region of the defined kernel.

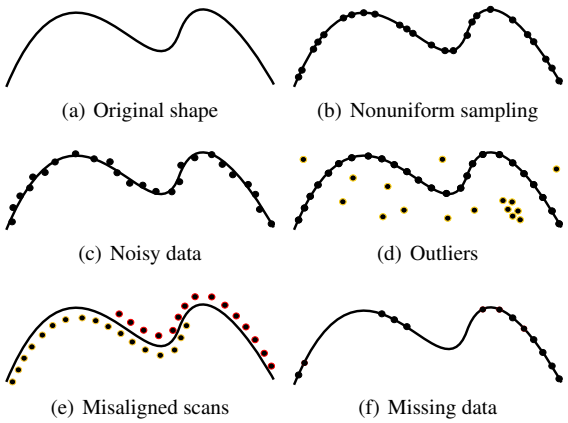

Figure 2: *Different forms of point cloud artifacts, shown here in the case of a curve in 2D.*

Noise. Points that are randomly distributed near the surface are traditionally considered to be noise – see Figure 2(c). The specific distribution is commonly a function of scanning artifacts such as sensor noise, depth quantization, and distance or orientation of the surface in relation to the scanner. For some popular scanners, noise is introduced along the line of sight, and can be impacted by surface properties, including scattering characteristics of materials. In the presence of such noise, the typical goal of surface reconstruction algorithms is to produce a surface that passes near the points without overfitting to the noise. Robust algorithms that impose smoothness on the output [KBH06], as well as methods that employ robust statistics [OGG09], are common ways of handling noise.

Outliers. Points far from the true surface are classified as outliers. Outliers are commonly due to structural artifacts in the acquisition process. In some instances, outliers are randomly distributed in the volume, where their density is smaller than the density of the points that sample the surface. Outliers can also be more structured, however, where high density clusters of points exist far from the surface, see Figure 2(d). This can occur in *multi-view stereo* acquition, where view-dependent specularities can result in false correspondences. Unlike noise, outliers are points that should not be used to infer the surface, either explicitly through detection [LCOLTE07], or implicitly through robust methods [MDGD*10].

© 2016 Author version

Misalignment. The imperfect registration of range scans results in misalignment, see [vKZHCO11] for a survey on registration techniques. Misalignment tends to occur for a registration algorithm when the initial configuration of a set of range scans is far from the optimal alignment – see Figure 2(e) for a 2D illustration. In scanning scenarios where we are only concerned with the acquisition of a single object, it is common for the object to rotate in-place with respect to the sensor for each scan; hence, the amount of misalignment is bounded since the initial scan alignment may be estimated from the known rotations. In SLAM-based RGBD-mapping [HKH*12], however, *drift* can occur. Drift is the gradual accumulation of misregistration errors between sequential scans. This can manifest as substantial misalignment between scans of a single object taken at very different times, a phenomenon known as *loop closure*.

Imperfections due to misalignment require techniques which differ from handling standard noise. For instance, under a Manhattan world prior [VAB12] a scene is composed of planar primitives aligned along three orthogonal axes – hence planar primitives from erroneously rotated scans can be robustly "snapped" onto one of these axes. Under the prior of repetitive relationships between geometric primitives [LWC*11], a misaligned scan can be corrected if it fails to conform to the remaining discovered repetitions.

Missing data. A motivating factor behind many reconstruction methods is dealing with missing data. Missing data is due to such factors as limited sensor range, high light absorption, and occlusions in the scanning process where large portions of the shape are not sampled. Although the aforementioned artifacts are continually improved upon, missing data tends to persist due to the physical constraints of the device. We note that missing data differs from nonuniform sampling, as the sampling density is zero in such regions – see Figure 2(f).

Many methods deal with missing data by assuming that the scanned shape is *watertight* [CBC*01, Kaz05, KBH06, HK06, ACSTD07]. Here the goal can be to handle the aforementioned challenges where data exists, and infer geometry in parts of the surface that have not been sampled. Other methods are focused on handling missing data by trying to infer topological structures in the original surface at the possible expense of retaining geometric fidelity, for instance, finding a surface that is homeomorphic to the original shape [SLS*07]. For significant missing data, other approaches seek the reconstruction of higher-level information such as a skeleton [TZCO09], shape primitives [SDK09], symmetry relationships [PMW*08], and canonical regularities [LWC*11].

2.2. Point Cloud Input

Reconstruction methods have different types of input requirements associated with a point cloud. The bare minimum requirement of all algorithms is a set of 3D points that sample the surface. Working with the points alone, however, may fail to sufficiently regularize the problem of reconstruction for certain types of point clouds. Other types of input can be extremely beneficial in reconstruction from challenging point clouds. We consider the following basic forms of inputs commonly associated with point clouds: surface normals, scanner information, and RGB imagery.

2.2.1. Surface Normals

Surface normals are an extremely useful input for reconstruction methods. For smooth surfaces the normal, uniquely defined at every point, is the direction perpendicular to the point's tangent space. The tangent space intuitively represents a localized surface approximation at a given point. Surface normals may be *oriented*, where each normal is consistently pointing inside or outside of the surface, or may lack such a direction. Normals that are oriented provide extremely useful cues for reconstruction algorithms – see [CBC*01, KBH06].

Unoriented normals. Normals that do not possess direction – the input normal at every point can be expected to be pointing either in the inside or the outside of the surface – are considered to be unoriented normals. This information can be used in a number of ways: determining planar regions in a point cloud [SWK07], the projection of a point onto an approximation of the surface [ABCO*03], or the construction of an unsigned distance field [AK04].

Unoriented normals are typically computed directly from the point cloud alone. A popular and simple method for computing the normal at a given point $\mathbf{p} \in P$ is to perform principal component analysis (PCA) in a local neighborhood of $\mathbf{p}$. The method of [HDD*92] estimates the normal as the smallest eigenvector of the covariance matrix constructed over a local neighborhood of points, obtained via an ε-ball or its k nearest neighbors. PCA defines a total least-squares plane fitting estimation of the tangent plane, and as such can be sensitive to imperfections in the point cloud, such as the sampling density and noise. The work of [MNG04] analyzes the accuracy of this form of normal estimation, where they show that the angle between the true normal $\mathbf{n}$ and the PCA-estimated normal $\hat{\mathbf{n}}$, with probability $1-\varepsilon$, is:

$$\sphericalangle(\mathbf{n},\hat{\mathbf{n}}) \leq C_1 \kappa r + C_2 \frac{\sigma_n}{r^2\sqrt{\varepsilon\rho}} + C_3 \frac{\sigma_n^2}{r^2}, \qquad (1)$$

where κ is the curvature at $\mathbf{p}$, r is the radius for the r–ball used in constructing the neighborhood, σ_n is the noise magnitude under zero-mean i.i.d. noise, ρ is the sampling density, and C_1, C_2, and C_3 are constants independent of these quantities. Note the trade-off between noise and the neighborhood size: as σ_n vanishes, r should be small, but as noise is introduced, r should be large in order to combat the effects of noise.

Other methods for normal estimation consist of using a weighted covariance matrix [PMG04] and higher-order approximations via osculating jets [CP05]. Similar to standard PCA, these methods require a local neighborhood of points, hence in the presence of point cloud imperfections it can

© 2016 Author version

be quite challenging to find the optimal neighborhood. As a result, these estimation methods can produce rather noisy normals, and so reconstruction algorithms must be robust to this.

Oriented normals. Normals that have consistent directions, either pointing in the inside or the outside of the surface are referred to as being oriented. Knowledge of the exterior and interior of the surface has proven extremely useful in surface reconstruction. It can be used to construct a signed distance field over the ambient space, where up to a sign, the field takes on positive values in the exterior and negative values in the interior. The surface is then represented by the zero crossing of the signed distance field. Other methods generalize this to implicit fields and indicator functions, but the basic idea of trying to construct the exterior and interior remains the same, see [CBC*01, OBA*03, KBH06] to name a few.

There are numerous ways to compute oriented normals. If the original 2D range scans are known, then the 2D lattice structure provides a way of performing consistent orientation since one always knows how to turn clockwise around a given vertex. For instance, if we denote the point in a range scan at pixel (x,y) as $\mathbf{p}_{x,y}$, then one can take the normal at $\mathbf{p}_{x,y}$ simply as the cross product between $(\mathbf{p}_{x+1,y} - \mathbf{p}_{x,y})$ and $(\mathbf{p}_{x,y+1} - \mathbf{p}_{x,y})$. If the point cloud is noisy, then this method can produce rather noisy normals, since it does not aggregate points in overlapping scans.

If scanner information is absent altogether, then one must orient the points exclusively from the unoriented normals. The well-known method of [HDD*92] achieves this by constructing a graph over the point cloud (e.g. through each point's knn neighborhood) and weights each edge w_{ij} for points $\mathbf{p}_i$ and $\mathbf{p}_j$ based on the similarity between the respective points' unoriented normals $\mathbf{n}_i$ and $\mathbf{n}_j$ as $w_{ij} = 1 - |\mathbf{n}_i \cdot \mathbf{n}_j|$. A minimal spanning tree is then built, where upon fixing a normal orientation at a single point serving as the root, normal orientation is propagated over the tree. The method of [HLZ*09] adjusts the weights w_{ij} by prioritizing propagation along tangential directions – if the estimated tangent space between two points $\mathbf{p}_i - \mathbf{p}_j$ is perpendicular to the normal directions, then this indicates these two points are valid neighbors, rather than belonging to opposite sides of the surface.

Although these methods are able to deal with nonuniform sampling, noise, and misalignment to a certain degree, they still remain sensitive to imperfections in the point cloud, and as a result can leave some normals unoriented or pointing in the wrong direction – see Figure 3. The impact on surface reconstruction largely depends on the distribution of incorrect orientations: if randomly distributed, then methods may treat this as spurious noise, but if incorrect orientations are clustered together over large regions, then this form of structured noise can be difficult to handle.

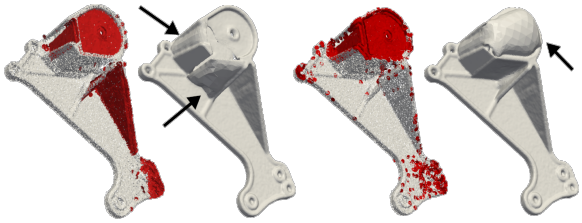

Figure 3: *The impact of incorrect normal orientation. On the left we show the result of normal orientation via [HDD*92], where red splats indicate incorrect orientation. The results of running Poisson surface reconstruction [KBH06] on this point cloud are shown in mid-left, where we indicate unwanted surface components due to the clustered normal flips. Similarly, on the right we show the orientation results of [LW10], and the corresponding results of [KBH06].*

2.2.2. Scanner Information

The scanner from which the point cloud was acquired can provide useful information for surface reconstruction. Its 2D lattice structure permits the estimation of sampling density which can be used to detect certain forms of outliers – points whose lattice neighbors are at a far greater distance than the sampling density are likely outliers. However, caution must be taken in distinguishing outliers from sharp features.

Scanner information may also be used to define the *confidence* of a point, which is useful in handling noise. Certain scanners (e.g. LiDAR) provide confidence measures in the form of the reflectivity measured at each point. One can also derive confidence through *line of sight* information. Line of sight is the collection of line segments between each point in the point cloud and the scanner head position from which that point was acquired. In active scanning systems, i.e. laser-based scanners, if the angle between the line of sight and the surface normal is large, this can result in noisy depth estimation, i.e. poor laser peak estimation [CL96], implying low confidence.

Note that line of sight also defines a region of space marked as lying outside of the shape. Combining line of sight from multiple scans refines the bounding volume in which the surface lies – this is known as the *visual hull*. This information is particularly useful when handling incomplete data – it can infer that there exists a large concavity in the shape [TOZ*11].

2.2.3. RGB Imagery

Different acquisition modalities that complement depth acquisition can be of great assistance. RGB image acquisition is a very common modality that accompanies numerous sensors, such as the Microsoft Kinect. In the case of the Kinect, the RGB camera is co-located with the IR camera, hence assuming the two are calibrated, it is straightforward to identify corresponding depth and RGB values at a pixel level.

© 2016 Author version

RGB images are most useful for reconstruction when they are able to complement depth information that is either not as descriptive as its visual appearance, or simply not measured by the data. For instance, if a color image and 3D scan are at a wide baseline, hence containing very different views, then segmented parts of the image can be used to infer 3D geometry in the original scan [NSC14].

3. The Role of Priors

The development of priors is largely driven by emerging data acquisition technologies. Acquisition methods set expectations for the class of shapes that can be acquired and the type of artifacts associated with the acquired data, consequently informing the type of output produced by reconstruction algorithms and the fidelity of the reconstruction. In this section we provide an overview of each prior, discussing the type of inputs each prior expects and subsequent output produced, as well as the typical shape classes and acquisition methods that characterize these scenarios.

3.1. Surface Smoothness

The *surface smoothness* prior constrains the reconstructed surface to satisfy a certain level of smoothness, while ensuring the reconstruction is a close fit to the data. Perhaps the most general form is *local smoothness*, which strives for smoothness only in close proximity to the data. The output of such approaches are typically surfaces that smooth out noise associated with the acquisition, while retaining boundary components where there exists insufficient sampling (or simply no sampling) [HDD*92, ABCO*03]. Due to their generality, these methods can be applied to many shapes and acquisition devices, yet absent of additional assumptions, handling severe artifacts poses a significant challenge.

In contrast, *global smoothness* seeks higher order smoothness, large-scale smoothness, or both. High order smoothness relates to the variation of differential properties of the surface: area, tangent plane, curvature, etc. Large-scale herein relates to the spatial scale where smoothness is enforced – not just near the input. It is common for these methods to focus on reconstructing individual objects, producing watertight surfaces [CBC*01, KBH06]. As a result, this limits the class of shapes to objects that can be aquired from multiple views, captured as completely as possible. Desktop scanners capable of scanning small (i.e. 1 inch) to mid-sized (i.e. several feet) objects are commonly used to produce such point clouds. Laser-based optical triangulation scanners, time-of-flight (TOF) scanners, and IR-based structured lighting scanners are all representative devices for such scenarios. Furthermore, due to the sensor's close proximity to an object and its high resolution capabilities, an emphasis is commonly placed on reconstructing very fine-grained detail.

Piecewise smooth priors seek to preserve sharp (i.e. nonsmooth) features in the shape [FCOS05, ASGCO10]. For this prior the acquisition device does not differ much from global or local smoothness priors, but rather the class of shapes are restricted to those which contain a set of sharp features – i.e. CAD models, man-made structures, etc.. Surface smoothness priors are covered in detail in Section 4.

3.2. Visibility

The *visibility* prior makes assumptions on the exterior space of the reconstructed scene, and how this can provide cues for combating noise, nonuniform sampling, and missing data. Scanner visibility is a powerful prior, as discussed in Section 2.2.2, as it can provide for an acquisition-dependent noise model and be used to infer empty regions of space [CL96]. This enables the filtering of strong, structured noise – a common characteristic of multi-view stereo outputs – for watertight reconstruction of individual objects [ZPB07]. Recent work has extended the visibility prior to *scene reconstruction* in an interactive setting by relaxing the watertight constraint and only maintaining a surface in close proximity to the data [NDI*11]. The Microsoft Kinect and Intel's RealSense are two representative scanners which enable the interactive acquisition of geometry. Such scanners permit the reconstruction of very large spaces, i.e. building interiors, but often at the expense of geometric fidelity compared to static, more constrained scanning setups. Methods which employ visibility priors are discussed in more detail in Section 5.

3.3. Volume Smoothness

The *volume smoothness* prior imposes smoothness with respect to variations in the shape's volume. This has shown to be quite effective when faced with large amounts of missing data [TZCO09, LYO*10]. Some volume priors assume the watertight reconstruction of an individual object with an emphasis on topological accuracy, where it is assumed that the well-sampled acquisition of an object is prohibited, primarily due to self-occlusions and limited mobility in sensor placement. This can be seen in man-made objects composed of such materials as coils or metal wires, where the shape can be described as a complex arrangement of generalized cylinders [LLZM10]. Other techniques such as [LYO*10] focus on extracting the skeleton structure of a shape from significant missing data. This can be seen in organic shapes, in particular trees, which tend to be scanned via LiDAR sensors in uncontrolled outdoor environments, and as a result many branches and leaves may only be partially scanned. Volume smoothness methods are covered in Section 6.

3.4. Primitives

The *geometric primitive* prior assumes that the scene geometry may be explained by a compact set of simple geometric shapes, i.e. planes, boxes, spheres, cylinders, etc.. In cases

© 2016 Author version

where we are concerned with watertight reconstruction of individual objects, the detection of primitives [SWK07] can subsequently be used for primitive extrapolation for reconstruction when faced with large amounts of missing data [SDK09]. CAD shapes naturally fit this type of assumption since they are usually modeled through simpler geometric shapes, however once scanned, the point clouds may be highly incomplete due to complex self occlusions. Furthermore, certain CAD models are mechanical parts whose physical materials may be unfavorable for scanning with many devices (be it IR, TOF, etc..), hence structured noise may result. The primitive prior can thus be useful for robustly finding the simpler shapes underlying the noise-contaminated point cloud [SWK07]. Note, however, that any fine-grained details are likely to be treated as noise with the primitive prior, if they are unable to be represented as a union of smaller primitives.

Indoor environments are another shape class ideal for primitives, as they may be summarized as a collection of planes or boxes. A typical application of this is in the reconstruction of a floor plan of a building. Such environments may be captured by LiDAR scanners [XF12], where obtaining full scans can be difficult due to large-scale scene coverage. Hence planes can be a useful prior for completion from incomplete data, in addition to denoising [JKS08, XF12]. Similar to CAD shapes, however, fine geometric details may be lost as a result, though hybrid methods have been developed to preserve detail while simultaneously extracting primitives [LA13, vKvLV13]. Primitive prior methods are discussed in detail in Section 7.

3.5. Global Regularity

The *global regularity* prior takes advantage of the fact that many shapes – CAD models, man-made shapes and architectural shapes – possess a certain level of regularity in their higher-level composition. Regularity in a shape can take many forms, such as a building composed of facade elements, building interiors composed of regular shape arrangements, or a mechanical part consisting of recurring orientation relationships between sub-parts.

For instance, facades can often be described in terms of repeating parts, for instance a collection of windows, where the parts possess some regularity in their arrangement, i.e. a uniformly-spaced grid of windows. Facade acquisition, however, is often faced with substantial noise and incomplete measurements, as typical acquisition devices – i.e. LiDAR or multi-view stereo – can only take measurements at far distances and suffer from occlusion with parts of the environment. Hence, if such regularity was detected in the input data, it can be used to model the rest of the facade for denoising and filling in missing parts [ZSW*10, LZS*11].

A major artifact addressed by global regularity in building interiors and mechanical parts is scan misalignment. For instance, building interiors may be captured by the registration of depth scans in real-time scanning devices, but imperfections in the registration can manifest in drifting, see Section 2.1 on misalignment. Enforcing regularity on angles between detected planes can be highly beneficial in correcting for drift [OLA15, MMBM15]. The acquisition of mechanical parts through standard desktop 3D scanners can similarly result in poorly registered depth scans, yet finding canonical relationships in extracted geometric primitives can be useful in correcting for these errors [LWC*11]. For mechanical parts, it is common for the reconstruction objective to produce a watertight reconstruction of an individual object, whereas for building interiors the objective is primarily plane detection, not necessarily a watertight reconstruction. However, due to the large scale of building interiors, there may exist more evidence for regularity compared to a mechanical part, whose small size inherently limits the potential number of relationships, so we see a trade-off in reconstruction fidelity and regularity. Global regularity methods are covered in detail in Section 8.

3.6. Data Driven

The *data driven* prior exploits the large-scale availability of acquired or modeled 3D data, primarily in scenarios where the input point cloud is highly incomplete. In these scenarios we are primarily focused on reconstruction of individual objects, or a collection of objects, since 3D databases tend to be populated with well-defined semantic object classes. For instance, we may be concerned with obtaining a watertight reconstruction of an individual object from only a single, or several, depth scans. A database of objects may be used to best match the incomplete point cloud to a complete model [PMG*05], or a composition of parts from different models [SFCH12]. In other cases we may be concerned with scene reconstruction, and using the database to complement the detected objects in the scene [SXZ*12, KMHG13].

Due to the generality of a data driven prior – the reconstruction quality is as good as the provided data – these methods can be applied to a wide variety of data acquisition scenarios and shape classes. Employing complete 3D models for reconstruction poses few limitations, so long as these models accurately reflect the scanned data. If not, then complex nonrigid registration methods are necessary to deform the models appropriately [NXS12, KMYG12]. Employing parts from different models can significantly expand the generalization of a model database to the input, however this assumes certain shape classes that can be described with respect to well-defined parts: most man-made objects fit this description, but organic shapes may be difficult to describe in terms of parts. To address this, an alternative data-driven prior is to use a *shape space*, or a compact means of representing shape variations, for which the input data should conform to. This prior can enable the reconstruction of more general scenes, not just individual objects but natural environments consisting of trees [BNB13] or flowers [ZYFY14]. Data driven methods are discussed in Section 9.

© 2016 Author version

3.7. User Driven

The *user driven* prior incorporates the user into the process of surface reconstruction, allowing for them to provide intuitive and useful cues for reconstruction. The specific form of interaction is largely driven by the type of shape being reconstructed, and how it was acquired. In performing watertight reconstruction of individual objects, often the focus is on topology recovery due to incomplete sampling from the sensor. Hence certain approaches focus on reconstruction of arbitrary shapes [SLS*07], while others admit interaction with a shape's skeletal model, relying on the volumetric smoothness prior to guide the reconstruction [YHZ*14]. In the reconstruction of architectural buildings, similar to the case of facades discussed in Section 3.5, scanning in outdoor environments can cause large gaps in the acquisition. Hence, user interaction can help in the discovery of global regularity, as well as how to apply the detected regularity to the rest of the point cloud [NSZ*10]. If finer-grained control of the reconstruction is desired, then the user can specify geometric primitives to model the building, with guidance provided by the relationships discovered in the input [ASF*13]. User driven methods are covered in detail in Section 10.

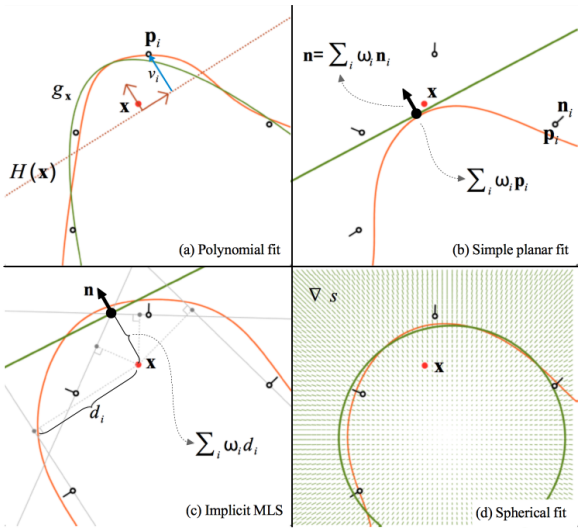

Figure 4: *Illustration in 2D of the principle of many MLS surface variants. The local approximations computed for the evaluation point* $\mathbf{x}$ *in red are show in green. The orange curves correspond to the reconstructed iso-contours.*

4. Surface Smoothness Priors

The surface smoothness prior can roughly be divided into *local smoothness*, *global smoothness*, and *piecewise smoothness*. These methods primarily vary based on the smoothness constraint and how it is prescribed in practice.

Notation. We first fix the notation for this section and all subsequent sections. We assume that we are given a point cloud P which is a sampling of a shape S. Individual points in P are indexed as $\mathbf{p}_i \in P$ for the i'th point. Many methods also require normals associated with the point cloud, where we define the *normal field* N as a set of normal vectors such that for each $\mathbf{p}_i \in P$ there is an accompanying normal $\mathbf{n}_i \in N$. The distinction between oriented and unoriented normals is made explicit for each method.

4.1. Local Surface Smoothness Priors

The pioneering method of [HDD*92] was hugely influential on the class of methods that impose local smoothness priors. This method approximates a signed distance field $\Phi : \mathbb{R}^3 \to \mathbb{R}$ by assigning, for each point in the ambient space $\mathbf{x} \in \mathbb{R}^3$, its signed projection onto the tangent plane of its closest point to P, denoted $\mathbf{p}_i$:

$$\Phi(\mathbf{x}) = (\mathbf{x} - \mathbf{p}_i) \cdot \mathbf{n}_i. \tag{2}$$

Note that the normal field N must be oriented in order to obtain an estimate of the signed distance field. The surface is then defined by the zero level set of Φ. Although straightforward to implement, this approach suffers from several issues. The method is very sensitive to the estimated normals – noisy normals, or worse inverted normal orientations, can give rise to very inaccurate signed distance estimates. Furthermore, in the presence of nonuniform sampling, choosing the closest tangent plane to define a signed projection can produce a rather noisy output. Subsequent methods based on local surface smoothness have focused on addressing such issues.

4.1.1. Moving least squares (MLS)

This class of methods approaches reconstruction by approximating the input points as a spatially-varying low-degree polynomial. Assuming that a scalar value v_i is associated to each input sample $\mathbf{p}_i$, the reconstructed signal at $\mathbf{x}$ is then defined as the value at $\mathbf{x}$ of the multivariate polynomial $g_\mathbf{x}$ which best approximates the neighborhood of $\mathbf{x}$ in a weighted least-square sense:

$$g_\mathbf{x} = \operatorname*{argmin}_g \sum_i \omega(\|\mathbf{x} - \mathbf{p}_i\|)(g(\mathbf{p}_i) - v_i)^2 , \tag{3}$$

where ω is a smooth decreasing weighting function giving larger influence to samples near the evaluating point. This weighting function plays the role of a low-pass filter. It is an essential ingredient of MLS to combat moderate levels of noise by allowing the weight function to have a larger spatial influence. For nonuniform sampling, it is necessary to define a weight function whose spatial support varies as a function of the sampling density – see Section 2.1 For surface reconstruction this approach needs to be adapted to address the lack of a natural parameterization for which v_i could be prescribed – see [CWL*08] for a survey on MLS.

Levin's method. In the first MLS formulation [Lev03, ABCO*03], the neighboring samples of the evaluation point

© 2016 Author version

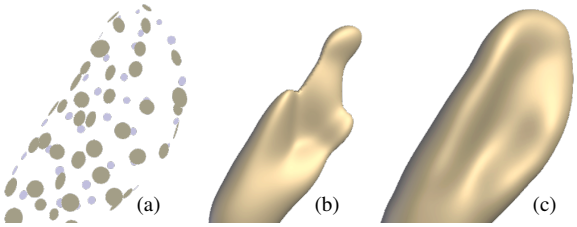

Figure 5: *When sampling density is insufficient to resolve local curvature (a), the plane fitting operation employed by moving least squares [AA04] becomes highly unstable (b). APSS [GG07] addresses this by locally fitting spheres instead of planes. Using spheres tackles the aforementioned problem while remaining computationally inexpensive.*

$\mathbf{x}$ are first parameterized with respect to a local tangent plane $H(\mathbf{x})$ obtained through a weighted PCA as in normal estimation methods (Section 2.2.1). In this parameterization, the neighborhood can be seen as a height-field with displacements v_i which is approximated by a low-degree bivariate polynomial $g_{\mathbf{x}}$ (Figure 4a). The *projection* of $\mathbf{x}$ is defined as the point closest to the polynomial approximation. The MLS surface is finally implicitly defined as the fixed points of the projection operator, suitably defined for points near the input.

Planar approximations. Authors rapidly observed that Levin's method can be significantly simplified by omitting the polynomial fitting step [AK04]. The MLS surface is then directly defined as the zero level set of the scalar field defined by the distance between the evaluation point and the best fitted plane. If oriented normals are available, then one might advantageously compute the normal $\mathbf{n}(\mathbf{x})$ of the best fitted plane as the weighted average of the neighboring normals:

$$\mathbf{n}(\mathbf{x}) = \frac{\sum_i \omega(\|\mathbf{x}-\mathbf{p}_i\|)\mathbf{n}_i}{\sum_i \omega(\|\mathbf{x}-\mathbf{p}_i\|)}\,, \tag{4}$$

For the sake of brevity, let us introduce the normalized weights $\omega_i(\mathbf{x}) = \omega(\|\mathbf{x}-\mathbf{p}_i\|)/\sum_j \omega(\|\mathbf{x}-\mathbf{p}_j\|)$, such that $\mathbf{n}(\mathbf{x}) = \sum_i \omega_i(\mathbf{x})\mathbf{n}_i$. This leads to the following scalar field definition [AK04, AA04] (see Figure 4b):

$$\Phi(\mathbf{x}) = \mathbf{n}(\mathbf{x})^T\mathbf{x} - \mathbf{n}(\mathbf{x})^T \sum_i \omega_i(\mathbf{x})\mathbf{p}_i\,. \tag{5}$$

Another variant called Implicit MLS (IMLS), constructs the implicit field as the weighted average of distances between $\mathbf{x}$ and the prescribed tangent planes [SOS04]:

$$\Phi(\mathbf{x}) = \mathbf{n}(\mathbf{x})^T\mathbf{x} - \sum_i \omega_i(\mathbf{x})\mathbf{n}_i^T\mathbf{p}_i = \sum_i \omega_i(\mathbf{x})\mathbf{n}_i^T(\mathbf{x}-\mathbf{p}_i)\,. \tag{6}$$

As seen in figure Figure 4c, the summation done with IMLS tends to expand/shrink the surface away from the input points, especially for large weighting support. Recent work has addressed this by replacing the neighboring points $\mathbf{p}_i$ in Equation 5 by the projection onto their respective tangent planes [AA09]. The Robust Implicit MLS (RIMLS) method [OGG09] discounts outliers by iteratively reweighting points based on their spatial and normal residual errors.

Spherical approximations. All the aforementioned MLS methods assume that it is possible to construct a well-defined tangent plane at each evaluation point, which may not exist for sparsely sampled data. In this case, a higher-order approximation such as algebraic point set surfaces [GG07], which uses an MLS definition with spheres for shape approximation, can be more robust (see Figure 5). APSS first fits a gradient field of the algebraic sphere s to the input (oriented) normals by solving a small linear least square problem(see Figure 4f):

$$\operatorname*{argmin}_s \sum_i \omega_i(\mathbf{x})\|\nabla s(\mathbf{p}_i) - \mathbf{n}_i\|^2\,. \tag{7}$$

This minimization fixes the linear and second order coefficients, and after integrating ∇s, a simple weighted average minimizing $\sum_i \omega_i(\mathbf{x})\|s(\mathbf{p}_i)\|^2$ permits to pick the best iso-sphere. More recently, it has been shown that algebraic spheres can also be robustly fitted to unoriented normals [CGBG13] by computing the spherical gradient field maximizing the sum of squared dot products with the prescribed unoriented normals:

$$\operatorname*{argmax}_s \sum_i \omega_i(\mathbf{x})\|\nabla s(\mathbf{p}_i)^T\mathbf{n}_i\|^2\,, \tag{8}$$

under the quadratic constraints $\sum_i \omega_i(\mathbf{x})\|\nabla s(\mathbf{p}_i)\|^2 = 1$.

4.1.2. Hierarchical methods

For this set of techniques, the reconstruction problem is approached as hierarchical partitioning. In the **multi-level partition of unity** (MPU) method [OBA*03], an octree-based partitioning is constructed top down: given a cell, the points inside and nearby the cell are approximated by either a bivariate quadratic polynomial or an algebraic trivariate quadric if the set of orientations spanned by the input normals is too large. If the residual error is too large then the cell is subdivided and the overall procedure is repeated. This results in a set of locally defined distance fields Φ which are smoothly blended to form a globally defined implicit surface:

$$\Phi(\mathbf{x}) = \frac{\sum_k \Phi_k(\mathbf{x})\omega(\|\mathbf{x}-\mathbf{c}_k\|/h)}{\sum_k \omega(\|\mathbf{x}-\mathbf{c}_k\|/h)}\,, \tag{9}$$

where $\mathbf{c}_k$ is the center of the cell, and h is the support radius which is proportional to the diagonal length of the cell. This blending needs signed implicit primitives to be consistent, and thus oriented normals are required. This approach resembles the weighted average of implicit planes achieved by the IMLS method with the major difference that primitives and weight functions are attached to cells instead of the individual input points. Another difference is that the level of smoothness and hence robustness to noise is indirectly adjusted by the error residual tolerance. Missing data can be partly addressed by allowing for the extrapolation and subsequent blending of spatially adjacent shape fits. However, such an extrapolation is prone to erroneous surface sheets. Those can be resolved by applying a diffusion operator on the collection of implicit

© 2016 Author version

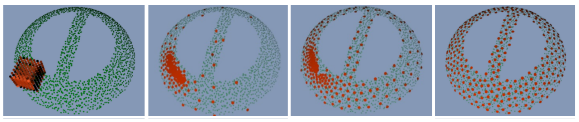

Figure 6: *Illustration of the LOP operator [LCOLTE07]: a set of samples shown in red is projected onto the multivariate median of the input points in green though attraction/repulsion forces.*

primitives, in order to perform smoothing directly on the MPU representation [NOS09].

4.1.3. Locally Optimal Projection (LOP)

Unlike the methods presented so far, this class of techniques do not strictly reconstruct a continuous surface, but rather fit an arbitrary set of points Q onto the multivariate median of the input point cloud P while ensuring that Q is evenly distributed [LCOLTE07] – see Figure 6. These methods do not require normal information, nor local parameterization, and they do not involve least square fits. This makes them particularly appealing to handle raw point clouds exhibiting noise, outliers and even misalignment. The median of a set of points is defined as the minimum $\mathbf{q}$ of $\sum_i \|\mathbf{q}-\mathbf{p}_i\|$. This definition is localized and extended to multiple target samples Q by attaching a fast decaying weight function ω to each sample and by performing a fixed-point iteration:

$$Q^{t+1} = \underset{Q}{\operatorname{argmin}}\, E_{data}(Q^t,P) + E_{spread}(Q^t,Q)\,, \qquad (10)$$

where

$$\begin{aligned} E_{data} &= \sum_j \sum_i \|\mathbf{p}_i - \mathbf{q}_j\| \omega(\|\mathbf{q}_j^t - \mathbf{p}_i\|) \\ E_{spread} &= \sum_j \lambda_j \sum_k \omega(\|\mathbf{q}_j^t - \mathbf{q}_k^t\|)\, \eta(\|\mathbf{q}_j - \mathbf{q}_k^t\|)\,) \end{aligned}$$

Here E_{data} attracts the samples to the local medians, while E_{spread} produces an even distribution by repulsing points too close to each other. Repulsion forces are defined by the function η, with typical choices $\eta(r) = 1/r^3$ and $\eta(r) = -r$. The coefficient λ_j permits to balance both the importance of these two energy and the relative influence of each particles. The method of [HLZ*09] tackles highly nonuniform sampling by incorporating weighted local densities into E_{data} and E_{spread}. More recently, this *discrete* formulation has been extended to a *continuous* one using a Gaussian mixture model to continuously represent the input point cloud density [PMA*14]. This formulation is almost an order of magnitude faster while leading to even higher quality results.

4.2. Global Surface Smoothness Priors

Radial basis functions (RBFs). RBFs are a well-known method for scattered data interpolation. Given a set of points with prescribed function values, RBFs reproduce functions

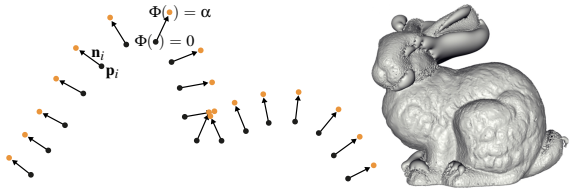

Figure 7: *(left) For RBFs, the scalar field to be optimized for should evaluate to zero at sample locations $\Phi(\mathbf{p}_i) = 0$, while at off surface constraints $\Phi(\mathbf{p}_i + \alpha\mathbf{n}_i) = \alpha$; this choice is appropriate as signed distance functions have unit gradient norm almost everywhere. The cluster of off-surface samples reveals how specifying constraints in regions of high curvature must be done carefully. (right) The surface reconstructed by RBF typically has severe geometric and topological artifacts when inconsistent off-surface constraints are provided.*

containing a high degree of smoothness through a linear combination of radially symmetric basis functions. For surface reconstruction, the method of [CBC*01] constructs the surface by finding a signed scalar field defined via RBFs whose zero level set represents the surface. More specifically they use globally-supported basis functions $\phi : \mathbb{R}^+ \to \mathbb{R}$. The implicit function Φ may then be expressed as:

$$\Phi(\mathbf{x}) = g(\mathbf{x}) + \sum_j \lambda_j \phi(\|\mathbf{x} - \mathbf{q}_j\|), \qquad (11)$$

where $g(\mathbf{x})$ denotes a (globally supported) low-degree polynomial, and the basis functions are centered at *nodes* $\mathbf{q}_j \in \mathbb{R}^3$. The unknown coefficients λ_j are found by prescribing interpolation constraints of function value 0 at $\mathbf{p}_i \in P$; see Figure 7. Off-surface constraints are necessary to avoid the trivial solution of $f(\mathbf{x}) = 0$ for $\mathbf{x} \in \mathbb{R}^3$. Positively (resp. negative) valued constraints are set for points displaced at $\mathbf{p}_i$ along $\mathbf{n}_i$ in the positive (resp. negative) direction. Interpolation is performed via the union of on-surface and off-surface constraint points as the set of node centers $\{\mathbf{q}_j\}$. The coefficients λ_i are found via a dense linear system in n unknowns, efficiently computed via fast multipole methods [CBC*01].

An advantage to using globally-supported basis functions for surface reconstruction is that the resulting implicit function is globally smooth. Hence RBFs can be effective in producing a watertight surface in the presence of nonuniform sampling and missing data. However, when the input contains moderate noise, determining the proper placement of off-surface points can become challenging (see Figure 7-right).

The need for off-surface constraints is elegantly avoided by the Hermite RBF (HRBF) scheme [Wen05]. In addition to positional constraints $\Phi(\mathbf{p}_i) = 0$, HRBF imposes that the gradient of Φ interpolates the input normal $\mathbf{n}_i$ at $\mathbf{p}_i$. The application of HRBF to 3D datasets is discussed in [MGV11].

Indicator functions. These methods approach surface reconstruction by estimating a *soft labeling* that discriminates

© 2016 Author version

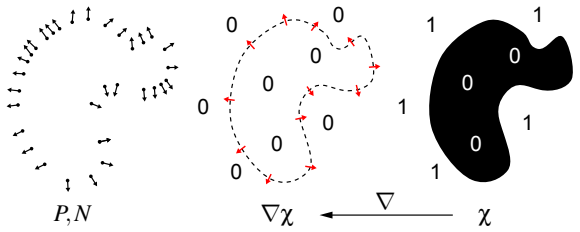

Figure 8: *Taking the gradient of the indicator function χ reveals its connection to the point cloud normals. Poisson reconstruction [KBH06] optimizes for an indicator function whose gradient at P is aligned to N.*

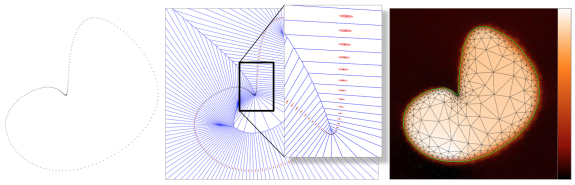

Figure 9: *The reconstruction of an unoriented point cloud (left) is achieved by seeking for an indicator function χ (right) whose gradients $\nabla\chi$ match the anisotropy of the given covariance matrices (middle, shown as pink ellipsoids). These tensor constraints can be estimated from the second order moments of the Voronoi diagram of P (middle) [ACSTD07].*

the interior from the exterior of a solid shape. Indicator function methods are an instance of *gradient-domain* techniques [PGB03]. For surface reconstruction, such a gradient-domain formulation results in robustness to nonuniform sampling, noise, and to a certain extent, outliers and missing data. This is accomplished by finding an implicit function χ that best represents the indicator function. The key observation in this class of methods is that, assuming a point cloud with oriented normals, χ can be found by ensuring that the gradient of the indicator function measured at the point cloud P is aligned with the normals N; see Figure 8. The indicator function thus minimizes the following quadratic energy:

$$\underset{\chi}{\operatorname{argmin}} \int \|\nabla\chi(\mathbf{x}) - \mathcal{N}(\mathbf{x})\|_2^2 \, d\mathbf{x} \tag{12}$$

The differential equation that describes the solution to this problem is a Poisson problem; it can be derived by applying variational calculus as $\Delta\chi = \nabla \cdot \mathcal{N}$. Once a solution χ of this equation is found, an appropriate iso-value corresponding to the surface is selected as the average (or median) of the indicator function evaluated at P. The implicit function's gradient being well-constrained at the data points enforces smoothness and a quality fit to the data and since the gradient is assigned zero away from the point cloud, χ is smooth and well-behaved in such regions. Furthermore, for small scan misalignment (see Figure 2(e)), normals tend to point in a consistent direction, which yields a well-defined gradient fit for the indicator function.

The approach of [Kaz05] solves the Poisson problem by transforming it into the *frequency* domain, where the Fourier transforms of $\Delta\chi$ and $\nabla \cdot N$ result in a simple algebraic form for obtaining the Fourier representation of χ. By operating in the frequency domain, however, it is necessary to use a regular grid in order to apply the FFT, hence limiting spatial resolution in the output. In order to scale to larger resolutions, the method of [KBH06] directly solves for χ in the spatial domain via a *multi-grid* approach, hierarchically solving for χ in a coarse-to-fine resolution manner. An extension of this method for *streaming* surface reconstruction, where the reconstruction is done on a subset of the data at a time, has also been proposed [MPS08].

A known issue with the approach of [KBH06] is that fitting directly to the gradient of χ can result in over-smoothing of the data [KH13, Fig. 4(a)]. To address this, the method of [KH13] directly uses the point cloud as positional constraints into the optimization, resulting in the *screened Poisson* problem:

$$\underset{\chi}{\operatorname{argmin}} \int \|\nabla\chi(\mathbf{x}) - \mathcal{N}(\mathbf{x})\|_2^2 \, d\mathbf{x} + \lambda \sum_{\mathbf{p}_i \in P} \chi^2(\mathbf{p}_i)$$

Setting a large value for λ ensures the zero-crossing of χ will be a tighter fit to the input samples P. While this can reduce over-smoothing, it can also result in over-fitting similar to interpolatory methods. The method of [CT11] also incorporates positional and gradient constraints, but also includes a third term, a constraint on the Hessian of the indicator function:

$$\begin{aligned} \underset{\chi}{\operatorname{argmin}} \sum_{\mathbf{p}_i \in P} \|\nabla\chi(\mathbf{p}_i) - \mathbf{n}_i\|_2^2 + \lambda_1 \sum_{\mathbf{p}_i \in P} \chi^2(\mathbf{p}_i) \\ + \lambda_2 \int \|H_\chi(\mathbf{x})\|_F^2 \, d\mathbf{x}, \end{aligned} \tag{13}$$

which can improve surface extrapolation in regions of missing data [KH13, Fig. 6(a)]. The main difference between the approaches is that [KH13] solves the problem via a finite-element multi-grid formulation, whereas [CT11] use finite-differences, due to the complexity in discretizing the Hessian term; In particular, the screened Poisson formulation [KH13] is up to two orders of magnitude faster than [CT11], see [KH13, Table 1].

As the above approaches rely on *oriented* normals they can only tolerate sparsely distributed normal orientation flips. Large continuous clusters of improper normal orientation significantly impacts these methods; see Figure 3 for an example. To address this limitation, the method of [ACSTD07] uses covariance matrices instead of normals to represent *unsigned* orientations leading to the following optimization problem:

$$\underset{\chi}{\operatorname{argmin}} \int \nabla\chi^T(\mathbf{x})\mathbf{C}(\mathbf{x})\nabla\chi(\mathbf{x}) + \lambda_1(\Delta\chi(\mathbf{x}))^2 + \lambda_2\chi(\mathbf{x})^2 \, d\mathbf{x}$$

In the energy above, the first term penalizes the misalign-

© 2016 Author version

ment of $\nabla\chi$ with the (unsigned) orientation represented in the principal component of tensor $\mathbf{C}$, computed via the second order moments of the union of neighboring Voronoi cells; see Figure 9. The anisotropy of the covariance acts as a notion of normal confidence, where isotropic tensors have little influence on the alignment. The second term is a *biharmonic* energy that measures the smoothness of $\nabla\chi$; it is this energy that favors coherent orientations of $\nabla\chi$, as incoherent orientations would result in a significant increase in energy. Finally, the third term fits to the data which improves conditioning.

Volumetric segmentation. These methods perform reconstruction via a *hard labeling* of a volumetric discretization, where the goal is to label cells as being either interior or exterior to the surface. The method of [KSO04] constructs a graph Laplacian from the Delaunay triangulation of P, where each node represents a tetrahedron of the triangulation and each edge measures the likelihood of the surface passing through the adjacent tetrahedra. The Laplacian eigenvector with smallest nonzero eigenvalue then smoothly segments tetrahedra into interior and exterior, as this eigenvector simultaneously seeks a smooth labeling and a partitioning with low edge weights. This approach has shown to be robust to noise and outliers without the use of normals, thanks to the robustness of spectral partitioning. Since it produces an explicit volume segmentation, it also ensures a watertight surface. However, in regions of missing data, the discretization from the Delaunay triangulation may be too coarse, giving a poor approximation to the surface [KSO04, Fig. 6].

The method of [HK06] first estimates an *unsigned* distance function $|\Phi(\mathbf{x})|$ from which the reconstructed surface $\mathcal{S}$ is extracted as the minimizer of the following energy:

$$\underset{\mathcal{S}}{\operatorname{argmin}} \int_{\mathcal{S}} |\chi(\mathbf{x})| \, d\mathbf{x} + \lambda \int_{\mathcal{S}} d\mathcal{S}$$

Provided viable sink/source nodes can be specified, energies of this kind can be minimized by graph cut optimization. The optimal cut results in a volumetric segmentation where the surface is localized in the proximity of samples, where $|\Phi(\mathbf{x})| \approx 0$. The solution is also smooth, as the second term of this energy enforces minimal surface area. A limitation of [HK06] is that source/sink nodes are specified by first defining a small *crust* on the exterior and interior through a dilation operation on point-occupied voxels, where the crust must reflect the topology of the shape for the graph cut to produce a topologically-correct surface. However, this method does not use normals, as it only needs to compute an unsigned distance. This results in robustness to nonuniform sampling, noise, and misalignment. Furthermore, the regularization allows for the method to handle missing data, where gaps are inpainted by minimal surfaces.

4.3. Piecewise Surface Smoothness Priors

Moving from the smooth, closed case to the *piecewise smooth case* (possibly with boundaries) is substantially harder as the ill-posed nature of the problem applies to each sub-feature of the inferred shape. The features of a piecewise smooth surface range from boundary components, sharp creases, corners, and more specific features such as tips, darts, and cusps. In addition, the inferred surface may be either a stratified manifold or a general surface with non-manifold features. Another difficulty stems from the fact that a feature is a notion that exists at specific scales, such that reconstruction and feature approximation cannot be decoupled.

4.3.1. Partitioning-based approaches

A first class of approaches consists in explicitly segmenting the input points with respect to sharp features. We distinguish methods depending on whether this segmentation is performed locally or globally.

The Robust MLS (RMLS) method [FCOS05] reproduces features by observing that points lying on two different smooth components can also be seen as outliers with respect to each other. They exploit this fact using a variant of the least median of squares (LMS) regression scheme where the summation in Equation 3 is replaced by a median. This strategy is used to obtain an initial robust approximation of a very small neighborhood which is then expanded as long as the residual is below a given threshold. Such a partitioning might not be consistent with respect to the evaluation point $\mathbf{x}$, thus leading to jagged edges. The Data-Dependent MLS (DDMLS) [LCOL07] method overcomes this by first computing a singularity indicator field (SIF) estimating the proximity of each point to a discontinuity as a lower bound of the derivatives of the unknown signal. Then, this SIF is used as weights to fit local univariate polynomials approximating the sharp creases. They are used to segment the neighborhoods into smooth components which are approximated by bivariate polynomials constrained to interpolate the crease curves if any. Corners are handled by applying this strategy recursively to the one-dimensional components. In [WYZC13], a feature preserving normal field is robustly computed using mean-shift clustering and a LMS regression scheme to find inlier clusters with scales from which a variant of RANSAC permits to both extract the best tangent planes and their associated connected points. This provides local partitions from which edge-preserving smoothing is performed by fitting multiple quadrics similarly to previous methods. However, the locality of the feature detection can still generate fragmented sharp edges.

To reduce crease fragmentation, some approaches favor the extraction of long sharp features. In [DHOS07], RMLS is used to detect sharp creases and grow an explicit set of polylines through feature points. In [JWS08], feature curves as extracted by robustly fitting local surface patches and computing the intersection of close patches with dissimilar normals. In both techniques, the extracted feature curves are then used to guide an advancing front meshing algorithm [SSFS06].

© 2016 Author version

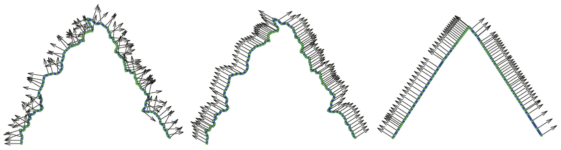

Figure 10: *Illustration of the ℓ_1 reconstruction method of [ASGCO10]: a sharp normal field is computed from which positions are optimized along the normal directions.*

4.3.2. Normal-field based approaches

Here we consider methods which do not make use of an explicit partitioning but rather tackle the sharp feature reconstruction problem by decoupling the computation of a sharp normal-field from the computation of the surface location.

For instance, the Robust Implicit MLS (RIMLS) method [OGG09] considers that points lying on two different smooth components can been seen as outliers in both the spatial and gradient domain. RIMLS uses $\varphi-$type M-estimators, replacing the standard least-square criterion of Equation 3 by a suitable norm robust to outliers – taken as the Welsch function. The sharpness of the reconstruction depends on the bandwidth used in the gradient domain, as well as the sharpness on the input normal field. For noisy input, the quality of the normal field is improved by applying a robust normal mollification. We note that this method produces an implicit surface which is still differentiable everywhere: it does not generate true discontinuities and since the analysis is performed locally, a reconstructed feature may exhibit a varying amount of sharpness (e.g., see [ASGCO10, Fig. 6]).

The approach of [ASGCO10] formalizes the sharp reconstruction problem as global ℓ_1 optimizations under the insight that the set of all neighboring normal differences should be sparse, with large differences reflecting sharp features. The method starts by reconstructing a normal field N preserving sharp features in the shape via ℓ_1 sparse reconstruction:

$$N = \underset{N}{\operatorname{argmin}} \sum_{i,j} \omega(\bar{\mathbf{n}}_i^T \bar{\mathbf{n}}_j) \|\mathbf{n}_i - \mathbf{n}_j\|, \tag{14}$$

subject to the constraint $\forall i \, \|\mathbf{n}_i - \bar{\mathbf{n}}_i\| \leq \gamma$ bounding the change in orientation with respect to the initial normals $\bar{\mathbf{n}}_i$. The surface location is then reconstructed assuming a local planar criteria leading to:

$$P = \underset{P}{\operatorname{argmin}} \sum_{i,j} \omega({\mathbf{n}_i}^T \mathbf{n}_j) \|{\mathbf{n}_{ij}}^T (\mathbf{p}_i - \mathbf{p}_j)\|, \tag{15}$$

where each reconstructed sample position $\mathbf{p}_i$ is constrained to move along its normal direction, that is: $\mathbf{p}_i = \bar{\mathbf{p}}_i + \alpha_i \mathbf{n}_i$. These two steps are depicted in Figure 10.

Similar to the RIMLS approach, the edge-aware resampling (EAR) method [HWG*13] employs a bilateral mechanism to smooth normals while separating them across sharp features. Then a variant of the LOP energy [LCOLTE07] is used to robustly smooth and resample the points away from sharp features, while a specific bilateral projection operator permits to upsample the point cloud in the vicinity of sharp edges. This operator is based on a planar fit similar to Equation 5 but using bilateral weights on normal differences to both estimate a robust normal direction and average neighbors. This procedure yields clean and dense point clouds which can then be safely and much more accurately reconstructed using, for instance, RIMLS.

4.3.3. Direct meshing

The recent approach of [DCSA*13] tackles the feature-capturing surface reconstruction problem by directly turning an input point cloud into a low triangle-count simplicial complex. It starts with a simplicial complex filtered from a 3D Delaunay triangulation of the input points. This initial approximation is iteratively simplified based on an error metric that measures, through optimal transport, the distance between the input points and the current simplicial complex, both seen as mass distributions. This approach exhibits both robustness to noise and outliers, as well as preservation of sharp features and boundaries.

4.3.4. Dictionary learning

The work of [XZZ*14] casts surface reconstruction as a problem of dictionary learning. Namely, the goal is to find a surface mesh whose vertices $\mathbf{V}$ are the dictionary atoms, and whose triangles are the sparse codes such that each point in the input cloud is approximated by a single point on the output triangle. This may be formulated as follows:

$$\min_{\mathbf{V},\mathbf{B}} \frac{1}{n} \sum_{i=1}^{n} \|\mathbf{p}_i - \mathbf{V}\mathbf{b}_i\|_2^q + E_{reg}(\mathbf{V},\mathbf{B}) \tag{16}$$

$$s.t. \; \forall_i \; \|\mathbf{b}_i\|_0 \leq 3, \; \|\mathbf{b}_i\| = 1, \; \mathbf{b}_i \geq 0, \tag{17}$$

subject to manifold constraints on the mesh, where E_{reg} favors a uniform distribution in $\mathbf{V}$. The constraints on each $\mathbf{b}_i$ enforce $\mathbf{p}_i$ to map to a point on some triangle. The data fitting term of $\|\mathbf{p}_i - \mathbf{V}\mathbf{b}_i\|_2^q$ uses a sparsity seeking l_q norm, with $q < 1$, which is shown useful in both dealing with noise and outliers, while preserving sharp features (see [XZZ*14, Fig. 3]). In practice, dictionary learning proceeds by alternating between fixing $\mathbf{V}$ and optimizing for $\mathbf{B}$ via edge swaps, and fixing $\mathbf{B}$ and finding $\mathbf{V}$ via sparse optimization.

5. Visibility Priors

Visibility has generally been used in three different ways – see Figure 11. The first class of methods considers how to use the visibility provided by the scanner that produced the point cloud – this is used primarily to obtain the *line of sight* associated with each sample; see Section 5.1. The second class of methods uses line of sight that is not provided from the scanner, but rather approximated from the exterior space;

© 2016 Author version

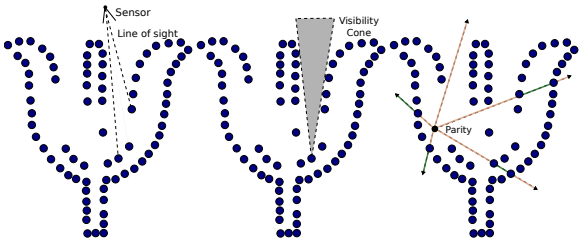

Figure 11: *Different sources of visibility information. Left: Scanner Visibility from line of sight information. Middle: Exterior Visibility using point visibility approximation. Right: Parity from volumetric surface parity estimation.*

see Section 5.2. The third class of methods uses visibility to approximate *parity* – the number of times a ray intersects a surface – in order to approximate the interior and exterior, as discussed in Section 5.3.

5.1. Scanner Visibility

The most common method for using scanner visibility is the merging of range scans. The approach of [CL96] incrementally constructs a truncated signed distance function (TSDF) localized to near the input points, built from individual TSDFs of each range scan. A scan's TSDF is determined as the distance between the scanner head and where a ray, originating at the scanner head, intersects the triangulated range scan. By summing the distance functions, and appropriately weighting each SDF by scanner confidence (c.f. Section 2.2.2), the aggregated TSDF is found at a point $\mathbf{x}$ as the TSDF $D(\mathbf{x})$ and weight function $W(\mathbf{x})$ representing certainty in the distance:

$$D(x) = \frac{\sum w_i(x) d_i(x)}{\sum w_i(x)} \quad W(x) = \sum w_i(x) \tag{18}$$

The confidence W is useful in combating noise inherent with laser scanning, see [CL96, Fig. 4].

The sign of the distance function, indicating the orientation of the surface can be obtained by performing *space carving* through line of sight information, via marking regions of space observed by the scanner as empty. The approach of [CL96] uses this information to extract geometry between regions marked empty and regions that are unseen, where the assumption is that unseen regions are the interior of the shape. The interface between seen and unseen regions are reflected through sign changes in the distance function.

For other forms of missing data, the approach of [CL96] will typically preserve the hole as it does not enforce any type of smoothness prior. It is possible to incorporate a minimal surface area regularization to encourage smoothness in regions of missing data, while using line-of-sight as a data-fitting term. Existing approaches solve such a formulation via level-set models [Whi98] and graph cuts [LB07]. The method of [LPK09] seeks an interior and exterior labeling of tetrahedra from a Delaunay triangulation of the point cloud, formulated as a graph cut problem using line of sight information. At each tetrahedron, the method accumulates evidence for belonging to the exterior through line of sight of all range scans, hence assuming outliers are randomly distributed, this method is robust to such defects; see [LPK09, Fig. 13].

For scans that contain a high level of misalignment and structured outliers, the method of [ZPB07] approaches range scan merging by using the l_1 norm for the data term, and the minimization of the signed distance gradient magnitude as the regularization term. This type of regularization, commonly known as *total variation* denoising, allows the algorithm to be robust to structured outliers and scan misalignment; see [ZPB07, Fig. 4].

The method of [FG11] considers the case when range scans have widely varying *scales* – range scans have very different sampling densities. In such cases, merging multiple scans of a coarse scale with a single scan at a fine scale can overly smooth out the fine-grained detail. [FG11] extends [CL96] by constructing a *hierarchical* signed distance field. This permits retaining the high resolution detail of fine-scale scans, while capturing the more general scene present in coarse-scale scans. This was recently extended in [FG14] by using a continous level of scale, where the surface is represented as a sum of Gaussian basis functions centered at the points, and each Gaussian's bandwidth is determined by the point's scale.

5.1.1. Dynamic Reconstruction

With recent advancements in interactive depth acquisition, such as the Microsoft Kinect or Intel's RealSense, the TSDF has become the predominant representation for performing dynamic reconstruction with such devices. This can be seen in the pioneering KinectFusion work [NDI*11], which extends [CL96] by incrementally fusing acquired depth scans into a TSDF, using the implicit representation to improve the registration of depth scans into the global scene. Follow up work has focused on improving the scalability, for instance through sparse voxel hashing [NZIS13], as well as improving the registration through extraction and matching of stable and robust contour-based features [ZK15], helping to minimize drift in long capture periods. Recent work [ZDI*15] has demonstrated how to use the high-resolution images captured in these devices to improve the geometric fidelity of the fused low-resolution depth images via shading-based refinement. Using a Lambertian reflectance model, they simultaneously estimate a spatially-varying albedo, scene luminance, and signed distance directly on the TSDF, which is demonstrated to provide robustness compared to other image or mesh-based representations (see [ZDI*15, Fig. 8]).

5.2. Exterior Visibility

It is possible to exploit visibility even in the absence of explicit information from the scanner. Given a chosen camera

© 2016 Author version

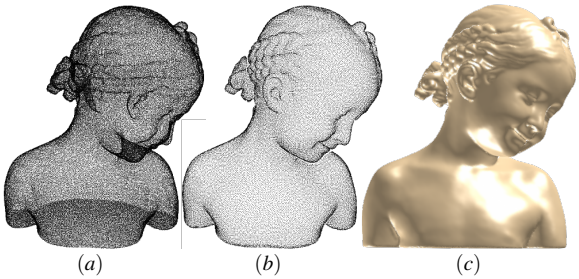

Figure 12: *The point cloud "hidden point removal" operator from [KTB07] applied to an input (a) determines the subset of visible points as viewed from a given viewpoint (b). Given this labeling, a view-dependent on-the-fly reconstruction (c) can be obtained by retaining the topology of well shaped triangles from the convex hull of the spherical inversion.*

position, *point set visibility* [KTB07] determines the portion of the point cloud that is not self-occluded. First, a spherical inversion of the point cloud with respect to the given query point is computed. Then, visible points are simply identified as those that lie on the convex hull of this set – see Figure 12. While [MTSM10] extended this method to handle moderate levels of noise, the input point cloud must respect strict sampling assumptions to produce satisfactory results.

Occlusion culling. The method of [CCLN10] builds upon these ideas and reconstructs a watertight surface by carving the space *occluded* by the point cloud when observed by a sufficiently large and randomly sampled set of directions. Similarly to [KTB07], the input cloud has to satisfy certain stringent sampling conditions. Conditions on sampling are relaxed in [CLCL11] where inconsistencies are detected by observing that if one point's *Voronoi pole* [AB99] lies in the exterior, the other Voronoi pole should be in the interior. If both are occluded or visible via [KTB07], this indicates an inconsistency. Unfortunately, since the method uses Voronoi poles, which cannot always be robustly estimated in the presence of missing data, its applicability remains limited.

Cone carving. The method of [SSZCO10] hypothesizes that each point in the cloud must have been *observed* from the scanner head. It computes high-likelihood visibility cones originating at each sample and takes the boundary of the union of all cones as an approximation to the surface. This method can be used to infer the geometry in large regions of missing data for challenging scenarios, i.e. two thin, spatially close, and undersampled surface sheets – producing topologically clean surfaces.

5.3. Parity

An alternative way of using visibility is to define a measure of *parity*. Assuming a closed surface, the parity for a given ray (point and direction) is defined as the number of times

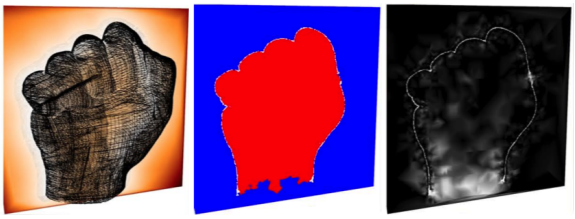

Figure 13: *The approach of [MDGD*10] first computes a robust unsigned distance function (left), and constructs an interior/exterior labeling (middle), and associated confidence (right) of the labeling. Note that low confidence is associated with regions of missing data, such as the bottom of the scan.*

the ray intersects the surface – if an odd number of times, this indicates the point lies in the interior, otherwise the point is in the exterior. This general idea can be extended to a point cloud, giving rise to a notion of *uncertainty* in whether or not a point belongs in the exterior or interior.

The approach of [MDGD*10] uses parity for estimating the sign of a robust unsigned distance function. Namely, the authors first construct a scalar field defined by the following unsigned distance function:

$$d_U(x) = \sqrt{\frac{1}{K} \sum_{p \in N_K(x)} ||x-p||^2} \tag{19}$$

A crust is then built around the surface from d_U, followed by stochastic ray shooting of the crust throughout the volume in order to estimate the sign of d_U. The uncertainty in this estimate is given by $2max(e,o)/r - 1$ relating the number of even e or odd 0 intersections with the crust and the number of rays r. This uncertainty estimate is used in constructing an implicit function, consisting of a data-fitting term and a smoothness term, such that high smoothness constraints will be assigned to regions that have high uncertainty (i.e. high disagreement in parity). Figure 13 shows the unsigned distance function for a challenging point cloud, along with its sign estimate and confidence in sign.

This approach is highly robust to noise, outliers, and missing data, and since its regularization is spatially-varying according to the uncertainty in parity, it will not over smooth the data where it exists. However, since smoothness is enforced via a Laplacian regularization, this could still result in poor behavior in regions of missing data, giving the incorrect topology. The method of [SY12] addresses this by performing space carving, guided by a parity estimate, to only carve out space where there does not exist highly confident interior regions. This can better retain topological features such as tunnels, where smoothness priors may erroneously over smooth and fill these regions in.

A drawback of [MDGD*10] is that the noise/outlier level must be known apriori via specification of K. The method of [GCSA13] extends [MDGD*10] by adapting to vary-

© 2016 Author version

ing noise levels within the point cloud. The key idea behind [GCSA13] is to find the smallest neighborhood K such that the apparent dimension defined via d_U matches that of a 2-manifold. From this, the method produces a sign estimate over a random collection of line segments in the volume. To determine the parity for each line segment, rather than using a crust as in [MDGD*10], they look at all local minima in the unsigned distance along the segment, flip the function according to the local minima, and of all possible flipped minima choose the one that is smoothest.

Tomographic reconstruction. The recent work of [NOS15] simulates the process of computed tomography to perform surface reconstruction. Namely, for a given set of random exterior views, a sinogram image is generated which measures, at each pixel, the thickness of the volume for a ray cast through the pixel. Parity is used as a means to determine valid pixels, as missing data may result in erroneous thickness measures. From this set of sinograms, tomogram reconstruction is then performed to recover the surface.

6. Volumetric smoothness

In order to handle challenging forms of missing data, a common way to regularize surface reconstruction is to enforce that the *local shape thickness* of a surface (i.e. a measurement of its local volume) varies smoothly. For watertight shapes, local thickness is measured by the radii of maximally inscribed spheres of its *medial axis transform*. However, as the medial axis is an alternative full representation of the shape, determining the medial axis over which to perform these measurements is an inherently difficult problem – as difficult as the reconstruction problem itself.

Skeletal regularizers. The *ROSA* method from [TZCO09] assumes that the medial axis of a shape can be approximated by curves instead of surfaces, that is, by a *curve-skeleton*. For organic geometry, a reconstruction of the skeleton can be obtained even in the presence of missing data by exploiting the redundancies of *local rotational symmetry*. ROSA proceeds in three-steps (see Figure 14): first, the distance from the cloud to the curve-skeleton is cylindrically parameterized using the curve-skeleton itself as a parameterization domain; then, an inpainting of the distance function is performed in image space, the inpainted distance can be converted to a 3D sample by using the inverse parameterization; finally, the inpainted point cloud can be processed by one of the algorithms in Section 4. It is important to note that a cylindrical parameterization prior constrains the class of shapes for which a reconstruction is possible to one having a star-shaped cross section. In particular, joint regions do not satisfy this requirement for which consequently a traditional surface smoothness prior must be employed. While the skeleton extraction method in [TZCO09] suffered the limitation of requiring oriented normals, subsequent research showed how it is possible to extract skeletons directly from unstructured point clouds [CTO*10] and even in the presence of outliers [HWCO*13].

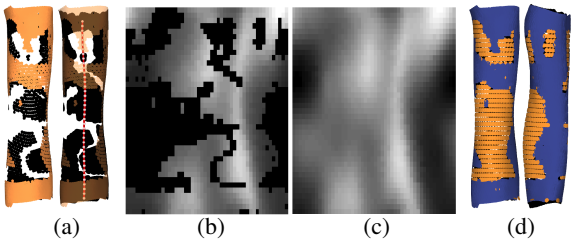

Figure 14: *(a) Curve-skeleton prior for reconstruction from an incomplete point cloud [TZCO09]. (b) The cylindrically parametrized point-to-skeleton distance where missing data is marked by black pixels. (c) In parameterization space, the distance is inpainted into regions of missing data by minimizing a thin-plate energy. (d) The inpainting and the inverse parameterization allows the insertion of additional samples to complete the geometry.*

Man-made skeletal geometry. To handle higher amounts of missing data, the approach of [LLZM10] presents a deformable model for reconstructing skeletal curves of *man-made* shapes composed of a collection of tubular components such as metal rods and canes. The *arterial snakes* model is obtained by sweeping a fixed-topology cross section through the input point cloud. The sweep operation is guided by a smooth version of the ROSA vector field [TZCO09] and starts simultaneously in all regions of high confidence (i.e. tubular areas away from joints). At first, simple circular cross sections are created by performing simple least squares circular fits. Sweeping is alternatively executed with a topology correction step to deal with colliding sweeping fronts and to model joint regions. The simple circular cross-sections are then optimized to fit the input point cloud while simultaneously satisfying a number of priors including curve-skeleton smoothness, planarity of cross sections and centricity of skeletal curves.

Organic skeletal geometry. Tubular components exhibiting piecewise smooth radii variations are also suitable to model organic and *natural* geometry, in particular trees [RFL*05, NFD07, LYO*10]. In tree reconstruction, biological constraints are exploited to simplify the problem and increase resiliency to data imperfections. The focus is on the reconstruction of tree branches and recovering the full topological skeletal structure of the tree. The skeleton is assumed to be a directed acyclic graph rooted at the ground plane; limbs are typically piecewise smooth and their thickness almost everywhere smoothly varying, where a *pipe-model* [RFL*05] controlling thickness variations can be used at branching locations. While a pair of orthogonal images have been shown sufficient to hallucinate the 3D information with the help of some user interaction [NFD07], recent research has attempted to model the tree structure directly from 3D point clouds [LYO*10].

© 2016 Author version

Medial priors. Several methods exist for imposing volumetric smoothness on the medial axis of the shape, supporting a much broader class of shapes. The method of [BS12] approaches reconstruction by segmenting the point cloud into volumetric regions and in each region taking the union of balls to obtain a coarse surface representation. The union of these regions then serves as an initial surface for the deformable model method of [SLS*06], to recover fine details. Key to their segmentation is a distance measure defined directly on the point cloud that robustly measures the likelihood of a medial ball being generated by any pair of points. This method is robust to noise and missing data, particularly when there exists nearby surface sheets, but can fail in regions where parts of the surface corresponding to medial sheets are missing. The method of [TOZ*11] extends the ideas in [TZCO09] to modeling local volume thickness directly from the (internal) *medial axis*. As incomplete point clouds do not possess a well defined medial axis, [TOZ*11] proceeds by using a level-set optimization where the intermediate solution (i.e. an iso-surface) is watertight, thus allowing the computation of medial measurements. Volumetric smoothness is then achieved by inserting an energy, in addition to the previously discussed fitting and surface smoothness terms, that attempts to perturb the surface to achieve a provided local thickness measurement. The target thickness value is obtained by diffusing the computed medial radii values on the object surface. This technique can prevent the formation of unnecessary holes in thin surfaces due to under-sampling, as the formation of a topological feature would correspond to a quickly vanishing medial radii. Furthermore, since the medial axis encodes local reflectional symmetry, this allows for information to be effectively propagated throughout the surface permitting the reconstruction of challenging data like the geometry of the highly concave areas in a vase – see [TOZ*11, Fig.5]. While highly general, the instability of the medial axis to surface-perturbations and the complexity of its computation limit the applicability of the method.

7. Geometric Primitives

Geometric primitives assume that a surface can be described as a composition of simple, canonical geometric shapes. This takes the form of surface primitives, volume primitives, and hybrid method which employ primitives with other priors.

Detecting primitives. The method of [SWK07] is an effective method for detecting geometric primitives in point sets. It uses RANSAC to robustly find planes, spheres, cylinders, cones, and torii, through an efficient means of sampling points for fitting and evaluating scores, both based on locality sensitive methods. Importantly, this method produces primitives that partially match the point cloud – the collection of these shapes can then be used for reconstruction. We note that although this method can detect a small set of easily parameterizable shapes, efficient pose detection methods for arbitrary shapes can also be used [DUNI10].

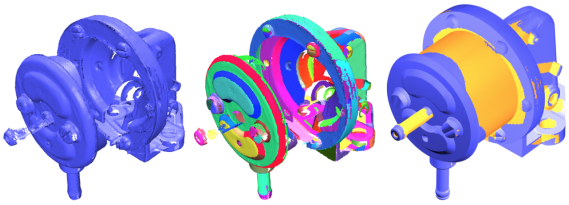

Figure 15: *(left) CAD models are often obtained by constructive solid geometry as a composition of simple primitives: planes, spheres, cones, etc. (middle) Randomized search [SWK07] can be used to detect such primitives in the point cloud data even in the presence of noise, outliers and missing data. (right) The primitives can then be extrapolated to obtain a watertight surface from incomplete data [SDK09].*

Primitive consolidation. The work of [JKS08] takes a set of detected plane primitives and performs reconstruction by aligning and merging the boundaries of adjacent primitives. More specifically, the boundaries of the plane primitives are extracted and an optimal configuration of boundaries is found by imposing a data-fitting term to the original boundary as well as a term that favors the snapping of boundary points and corner points of neighboring planes. This method can reconstruct CAD and architectural models alike, producing a surface mesh that retains the detected primitive structures. However, the method requires that adjacent primitive boundaries should be geometrically close to each other, which may not be satisfied if primitive detection is noisy, or if data are missing. The method of [SDK09] resolves this by explicitly extrapolating shape primitives (of all kinds) and forming the resulting output as the intersection of the extrapolated primitives. This extrapolation of primitives is formulated as a graph cut problem, where in addition to a standard minimal surface area term, a data fitting term is used that ensures the surface normal at a given point (the edge in the graph) is aligned with all intersecting primitives at that point. This does not constrain the primitives in a local manner: primitives whose boundaries are far away can eventually meet up and intersect with this method, as illustrated in Figure 15. The approach of [NSF14] is similar to [SDK09], but employs higher-order potentials in a conditional random field model to complete extracted contours which enforce scene relationships between wall, ceiling, floor, and internal plane primitives.

Augmenting primitive information. Although the method of [SDK09] can robustly handle missing data, it can be sensitive to noisy primitives which may fail to define a coherent model when extrapolated. The work of [CLP10] instead uses line of sight information to help penalize poorly extrapolated primitives. Namely, this work takes the set of primitives as well as an additional set of primitives formed near the boundaries of the input primitives and constructs a cell complex reflecting the extrapolation of the primitives. An energy, similarly solved via graph cuts is then formed, where the data-

© 2016 Author version

fitting term uses line of sight information to penalize facets in the complex lying in regions marked empty via space carving. The method of [RKMP13] uses the point cloud, line of sight, and edge features in corresponding RGB images to infer constraints for finding the boundary of each input shape primitive, consequently producing the reconstructed surface. In addition to using planar primitives and line of sight, a recent approach [BdLGM14] augments the set of detected primitives by ghost plane hypotheses supported by the occluding edges of primitives. Surface reconstruction is then formulated as a discrete optimization problem based on both detected and hypothesized planes. Departing from the usual area-based regularization, the reconstructed surface is regularized with respect to the length of sharp creases and number of corners, which is shown to better capture low surface complexity of man-made environments such as buildings.

Volumetric primitives. In the case of indoor scene reconstruction, an alternative to surface primitives is to employ *volumetric primitives* to model the interior space. In [XF12] the volume is modeled by fitting cuboids to the empty space defined by the boundaries of the scan data. A set of 2D constructive solid geometry (CSG) vertical slices are built by incrementally adding and removing candidate rectangles that best model the interior – a function of line of sight information. A similar process is used to stack up these slices to build a set of volumetric primitives, producing a 3D CSG model that composes the interior.

Hybrid methods. A limitation of primitive-based methods is that they do not degrade gracefully if certain portions of the shape are poorly explained by a primitive, or if the shape detection process is imperfect. The method of [LA13] resolves this by introducing a *hybrid* approach to reconstruction: shape primitives are used to resample the point cloud and enforce structural constraints in the output, such as sharp features between adjacent primitives and corners, while a visibility-driven prior is employed in regions where a primitive fit is not found. A similar approach was proposed in [vKvLV13], where planar polygons of sufficient fitting quality are extracted and a conforming, constrained Delaunay triangulation is constructed on the polygons and the remaining points so that the polygons are preserved in the triangulation. A visibility-driven graph cut problem is then solved, where the extracted polygon primitives are retained, while the rest of the points use line of sight information for reconstruction.

8. Global Regularities

The global regularity prior typically takes the form of three basic properties properties: *symmetry*, *repetition*, and *canonical relationships*. Commonly associated with *high-level shape analysis* [MWZ*13], these priors have also shown to be of great use in handling severe defects in a point cloud.

8.1. Symmetry

Symmetry is a well-studied problem in shape analysis [MPWC13]. Symmetry detection is focused on finding transformations on the shape that maps the entire shape, or a subset of the shape, onto itself. Finding such transformations can be extremely useful for surface reconstruction in handling noise and missing data.

The method of [PMW*08] finds repeating elements (small subsets of the point cloud) that can be mapped onto one another by *local* similarity transformations. They show that the repetition of elements in a point cloud manifests as a lattice structure in a suitable transformation space. In particular, partial matches become prominent in this transformation space, hence repeating elements of varying levels of missing data can be robustly detected and used to reconstruct incomplete regions. The method of [LCDF10] finds symmetries in incomplete point clouds by constructing an affinity matrix that measures *how symmetric* all pairs of points are. The key insight made by [LCDF10] is that this matrix should be block-diagonal for many types of symmetries – i.e. rotational, bilateral, intrinsic. By considering powers of this matrix, the authors demonstrate how incomplete matches become more pronounced, allowing for a wide range of detected symmetries in challenging point clouds containing noise, outliers, and missing data. These simpler forms of symmetry can be generalized to a notion of *subspace symmetries* [BWM*11], where a symmetry group is defined by a set of local transformations as well as a low-dimensional shape space, in order to handle more general types of shapes.

8.2. Structural Repetition

In certain cases it is difficult to find repeating elements in a point cloud through symmetry transformations. Instead, *directly* seeking repeating elements in a transformation-free manner can provide us with more flexibility in the reconstruction process.

The method of [ZSW*10] utilizes this observation for reconstruction of building facades in noisy terrestrial LiDAR scans, where occlusions from vegetation or other objects result in significant missing data. For a given type of facade element, each elements' planes are detected via [SWK07] and the individual elements are registered at a per-plane level. Once registered, denoising is performed across all elements via the individually registered planes and the consolidated facade element is projected back onto each instance for reconstruction. The mutual use of information across all elements allows one to robustly remove noise and fill in missing data.

A drawback to the approach of [ZSW*10] is the strict requirement of user interaction. This limitation was addressed in the work of [SHFH11] by adaptively partitioning facades. The approach of [WS12] takes the consolidated point cloud of [ZSW*10] and segments it into depth layers and uses a grammar definition to individually segment each depth layer

© 2016 Author version

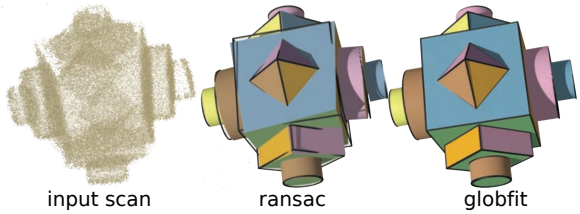

Figure 16: *From a set of misaligned scans shown on the left, the primitives extracted via [SWK07] (middle) retain the misalignment. Globfit [LWC*11] (right) is able to correct misalignment by enforcing consistent canonical relationships across primitives.*

into facades via the optimal sequence of grammar derivations. Facades may be appropriately extruded at each depth layer to obtain a polygonal representation of the building, at the possible expense of detail in the geometry due to the lack of expression in the shape grammar.

Another means of detecting regularity in incomplete scans is to find regularity in associated RGB imagery, and propagate this information back to the 3D scan to perform reconstruction. The approach of [LZS*11] achieves this by decomposing the RGB image into depth layers via the 3D scan, and upon detecting symmetries with respect to each layer, consolidates all element instances to robustly denoise and fill in missing data across the instances. A recent approach [NSC14] uses structural relationships in an RGB image to "lift" parts in the image which were not scanned to recover 3D geometry.

8.3. Canonical Relationships

Another useful prior on global regularities is the canonical intra-relationship between parts of a scene, or parts of a shape. Such relationships can be parallel or coplanar parts, recurring orthogonality between planes, concentric or co-axial parts, and regularity in orientation. This often arises in CAD models due to fabrication restrictions and budget considerations, as well as urban environments due to functional constraints.

Manhattan constraints. Perhaps the simplest form of a canonical relationship is the *Manhattan-world* (MW) assumption: all planar primitives in a scene belong to one of three mutually orthogonal planes. This can simplify facade reconstruction, as in the aforementioned methods of [ZSW*10, SHFH11, WS12, LZS*11]. In [VAB12], MW is used for building reconstruction by first classifying points by shape type – wall, edge, convex corner, or concave corner – and clustering points of a similar type. After constructing MW-aligned bounding boxes on all clusters, volume regions are found via parity, where interior regions of consistent parity are considered to belong to the building's volume. As edges and corners are detected via relationships between walls, this method is robust to missing data, but may be sensitive to noise for adjacent wall configurations.

Consolidating relationships. The method of [LWC*11] reconstructs CAD shapes consisting of a much richer variety of canonical relationships compared to MW. Namely, starting from an initial set of detected primitives [SWK07], parallel, orthogonal, angle-equality, and distance-equality relationships are individually detected and carefully selected so as to not cause any relationship conflicts. By enforcing these relationships, structured noise such as scan misalignment can be effectively handled – see Figure 16.

To handle large-scale scenes, the method of [OLA15] focuses on enforcing parallel and orthogonality constraints in the detection of planes. They progressively extract and regularize planes according to these constraints in order to confidently extract relationships early on, helping to resolve uncertain relationships later in the optimization. Regularization is performed by finding a single plane jointly over all parallel segments and their corresponding orthogonal segments. The RAPter method [MMBM15] similarly seeks a regular arrangement of planes, however as input, they additionally take a user-prescribed set of angles which subsets of pairs of planes should conform to. Plane relationships are found by finding ordered pairs of planes such that one plane is able to "explain" another (i.e. through a rotation). Their problem formulation is specifically designed not to penalize plane arrangements explained by a small minority of points, which would otherwise be dominated by more prominent relationships.

Canonical building relationships. The work of [LWC*11] was extended to the case of reconstruction of buildings from 2.5D scans in [ZN12]. The basic observation in this approach is that there exists three fundamental type of relationships in buildings: roof-roof relationships that consist of orientation and placement equalities, roof-roof boundary relationships that consist of parallelism and orthogonality relationships, and boundary-boundary relationships that consist of height and position equality. Upon finding the relationships via clustering (i.e., clustering similar angles, equality, etc..), they are used to inform the primitive fitting method so that the primitives simultaneously fit to the data and to the relationships.

9. Data-driven priors

The previously discussed priors may not always be appropriate and in practice, certain shapes may simply fail to adhere to these priors. A more flexible method of specifying a prior is through a *data-driven* means: using a collection of known shapes to help perform reconstruction. Data-driven priors often take the form of rigid and non-rigid deformation of retrieved objects or object parts, as well as class-specific shape spaces.

Scene reconstruction by rigid retrieval. The method of [SXZ*12] approaches this problem by first semantically segmenting the point cloud and then finding a complete model to replace each segment. More specifically, given a set of

© 2016 Author version

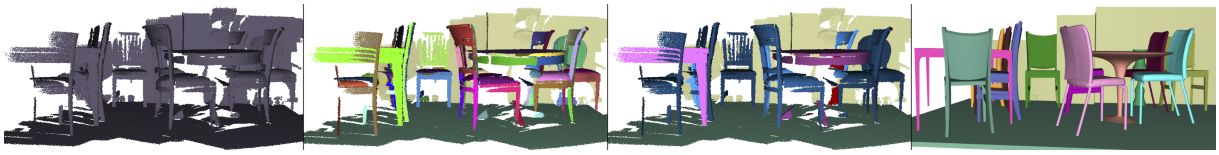

Figure 17: *From the incomplete and cluttered scan on the left, the approach of [NXS12] first oversegments the point cloud (mid-left), then iteratively merges segments which agree on class labels (mid-right), and finally deforms a set of detected models to best resemble their segmented objects (right).*

scans and RGB images, a conditional random field class labeling problem is formulated, balancing two objectives: a data fitting term based on a database of class-labeled objects and a smoothness term favoring local consistency in appearance and geometry. A training set of point clouds is built from virtual range scans of objects in various poses. This allows the model to be robust with regard to missing data, as the input should map to one of these poses. A random forest classification is built over this training set allowing for the closest complete object to an incomplete scan to be retrieved.

Recent work has extended this idea to the real-time scanning regime. The method of [KMHG13] interactively matches a merged point cloud against a large collection of models in a database, all of which have been virtually-scanned to best reflect the missing components of the input data. Geometric descriptors based on the distribution of surface normals are then quickly computed and efficiently matched, to permit interactive retrieval. The method of [LDGN15] has extended this to the more general setting of matching arbitrary models in scene-based scanning via matching directly against the TSDF of the captured scene, rather than the merged point cloud, eliminating the segmentation requirement of [KMHG13].

Scene reconstruction by non-rigid retrieval. A natural extension of reconstruction by retrieval is to consider non-rigid transformations of the template geometry to the input data. This is addressed in [NXS12] where upon finding a certain semantic class for a segmented object in the point cloud, every model is non-rigidly deformed via localized scale deformations. The best match is identified as the model with the smallest registration residual. This method approaches classification differently by building a semantically-labeled segmentation through incremental selection of oversegmented surface patches. A patch is chosen if the resulting merged object has high confidence in its label. This is particularly effective in noisy, outlier-ridden highly cluttered environments – see Figure 17. The authors of [KMYG12] extend these ideas by noting that in *indoor environments* it is common to have the same object in multiple poses. Their technique incorporates a deformation model directly into the segmentation and classification problem, rather than as a post-processing step. A deformation model is learned over multiple incomplete scans for a given object, allowing the object to be identified by incomplete observations of its parts. Given an input scan, it is first over-segmented and then iteratively merged into parts, where parts are matched against learned local deformation modes of a model. Part relationships are then used to verify the global quality of a match. Compared to [NXS12,SXZ*12], this permits the reconstruction of a broad set of objects, lessening the need for a large object database.

Object reconstruction by part composition. A disadvantage of the above approaches is that the granularity of the reconstruction is at the level of a whole model, that is, combining parts from different models is not possible. The approach of [SFCH12] overcomes this by combining individual parts from different objects to best compose the input scan. Namely, starting from a database of segmented models, 3D data is combined with RGB imagery to find candidate parts from the database matching the input. In particular, the use of RGB data can help find parts that are completely missing in the 3D scan. The best combination of candidates that closely match the geometry, while consisting of a small intersection with each other, composes the final model.

Reconstruction in shape spaces. A different line of data-driven methods constrains the reconstruction to a *shape space*, greatly regularizing the problem. For instance, [BNB13] performs foliage reconstruction by interleaving the registration of a set of segmented leaves to a single exemplar leaf with the construction of a statistical shape model from the reconstructed leaves in order to refine their geometry. The method of [ZYFY14] reconstructs flower petals by instead starting from a predefined shape space of exemplar petals and maps each segmented petal into the space by enforcing botany-specific priors (consistent roots, similarity in adjacent petals) with respect to the model parameters. Facial expressions comprise a compact shape space, where the method of [WBLP11] utilizes this by taking raw depth scans of facial performance capture in real time and performs reconstruction by finding a linear combination of a set of apriori defined facial shape meshes that best match the input.

10. User-Driven Methods

Incorporating the user in the process of surface reconstruction has shown to be extremely beneficial in dealing with challenging point clouds. Many user-driven methods seek a balance between level of intuition and user feedback which is useful for the reconstruction technique. Successful methods tend to tightly integrate the reconstruction algorithm with the

© 2016 Author version

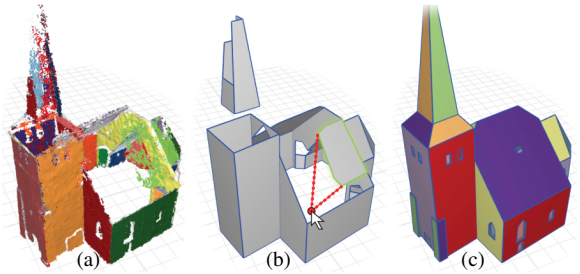

Figure 18: *Given a point cloud, planar primitives identified by RANSAC (a) may result in coarse and incomplete geometry (b). By exploiting the user's high-level knowledge while remaining faithful to the input data (b) a constrained optimization allows to recover a high-quality model (c).*

form of user interaction. Recent methods couple the scanning process with the reconstruction method, where the user is interactively involved in the acquisition.

Topology cues. The method of [SLS*07] presents a topology-driven interaction scheme. The approach seeks watertight and topologically correct reconstructions through the automatic detection of topologically weak regions. These low-confidence regions are then presented to the user to be resolved, via *scribbles* on a 2D tablet, which translate to interior and exterior constraints, or potentially no constraints if the user deems the region valid. The reconstruction is then updated, and the process repeats through further user edits.

Morfit [YHZ*14] allows the user to interactively edit a skeletal representation of the surface to better reflect the overall topology of the surface. Starting from an initial representation [HWCO*13], the user can break and mend branches in the skeleton to accurately model the shape's topology. Critical to the technique is an efficient scheme to reconstruct the surface from the skeleton. For each branch in the skeleton, this is achieved by interpolating well-sampled curve profiles across the remainder of the branch, while simultaneously fitting the interpolated profiles to the input data.

Structural repetition cues. For point clouds which exhibit a high amount of structural repetition, yet whose structure is challenging to extract, it is advantageous for the user to provide structural cues to guide the algorithm. In [NSZ*10] the authors present a technique to rapidly reconstruct architectural models, such as those acquired from the scanning of large urban environments. The key idea of the approach is to enable the user to define and manipulate simple geometric building blocks in the form of axis-aligned rectangular cuboids named *smartboxes*. The user sequentially places the smartboxes into the scene, where contextual regularities and symmetries between boxes are automatically captured and used to expedite the fitting process. The final placement of the user manipulated primitives is determined through an interactive optimization procedure that automatically adjusts the location, orientation, and sizes of the box primitives by considering how the cuboid fits the data and its relationship in context with previously placed boxes.

Primitive relationship cues. Another important structural cue is with respect to *primitive relationships*: if a shape is composed of a set of adjoining primitives – planes composing a building – then user feedback on primitive arrangements can be highly beneficial in challenging point clouds. The approach of [ASF*13] follows this by reconstructing a closed polygonal model through *snapping* each polygon into alignment with neighboring primitives. This is achieved by solving a combined optimization problem involving both *local* and *global* spatial constraints – see Figure 18. The user can provide two forms of feedback. A *polygon edit mode* allows the user to refine existing polygons by editing their boundaries and merging multiple disconnected polygons. A *polygon sketching mode* allows the user to provide new polygons for regions where automatic plane detection failed due to defects in the data. For both modes, the user has to only provide coarse edits, as the automatic snapping optimization is used to align polygon boundaries based on both local and global relations between primitives.

Interleaved scanning and reconstruction. Recent work seeks to integrate the user with *both* scanning and reconstruction, where these processes are interleaved. The method of [YSL*14] allows for the user to interact with the objects being scanned, where assuming that the user actions result in smooth non-rigid motions of the objects, the method can interactively update the reconstruction as the user scans and modifies the scene. This allows for the acquisition and reconstruction of "hidden" parts of objects, as the user may now expose them. Recent work [NFS15] has focused on real-time reconstruction of non-rigid scenes by updating a TSDF through graph-based deformations of the captured geometry. The approach of [WSL*14] prompts the user for next-best views as the user is scanning the object. This is accomplished through a Poisson-based uncertainty field, estimating regions in which the reconstruction algorithm is least confident.

11. Evaluation of Surface Reconstruction

Given the wide diversity in reconstruction methods, the manner in which one reconstruction is evaluated compared to another may differ. In this section we look at different evaluation criteria used in the surface reconstruction literature.

Geometric Accuracy. A common method of evaluation is to directly compare the geometry of the reconstruction output to the *ground truth* surface from which the scan was obtained. Hausdorff distance, mean distance, as well as measuring error in normals are common geometric error measures in this scenario. However, it is often challenging to obtain the notion of a ground truth surface from a physical shape. Hence, computational representations of shapes are typically used as the ground truth [Kaz05, MPS08], where *synthetic scanning* of the digital representation can be used in place of

© 2016 Author version

an actual scanner [BLN*13]. In some cases, a direct comparison to ground truth data is insufficient when targeting reconstruction under an error tolerance or comprising several levels of details (LODs). This suggests evaluating instead the complexity-distortion tradeoff, or the capability to generate LODs that are both controllable via intuitive parameters and meaningful for the targeted applications. The evaluation criteria comprise the coherence of LODs across the scene, the ability to incrementally refine the geometry, and the level of abstraction provided by the LODs. See [VLA15] for recent work in this vein for city-scale aerial reconstruction. The process of abstraction creates recognizable visual depictions of known objects through compact descriptions involving a handful of characteristic primitives such as curves, icons or solids. Abstraction thus goes beyond simplification as it involves filtering, smoothing, and reinforcing regular structures.

Topological Accuracy. Another important evaluation criteria is the recovery of higher-level information of the shape and in particular, its topology. Certain methods are concerned with reconstructing a shape with the correct genus [SLS*07], while other methods that focus on recovering a skeletal representation of the shape are more concerned with the topology of the underlying skeletal structure – recovering important branches and junctions in the skeleton. Such a topological structure is of particular importance for structural shape editing applications [LLZM10]. However, we note that most skeleton-based methods are often concerned with qualitative evaluation, hence it can be difficult to compare different skeleton extraction methods.

Structure Recovery. Beyond geometry and topology it is also sometimes desirable to recover the structure during reconstruction. Beyond the simple notion of scene decomposition, the term structure has a broad meaning, ranging from the dimension of geometric entities (manifolds, stratified manifolds, non-manifold configurations) to adjacency relationships through canonical geometric relationships (parallelism, co-planarity, orthogonality, concentricity, co-axiality) and regularities (repetitions, symmetries). In addition, controlling the structure encompasses recovery, preservation, and reinforcement. Structure is especially relevant when dealing with large-scale scenes, not just individual objects, where scenes are composed of a collection of objects which may have structural interrelationships. Structure as well as global regularities are also a means to improve robustness and resilience to missing data and go beyond reconstruction to consolidation and abstraction.

Reproducibility An important consideration in evaluating the quality of a reconstruction method is its level of reproducibility. Perhaps the simplest means of determining reproducibility is whether or not certain methods are made publicly available or have been implemented by a third party, as this can be an important indicator of implementation complexity and algorithm robustness. For instance, Poisson surface reconstruction [KBH06] is a widely used surface reconstruction method as the code is highly stable and reliable. The issue of reproducibility and provenance is well studied in other areas, including visualization and scientific workflows [SFC07, FKSS08], and it has been shown to be useful for studying 3D model construction [DKP11]. Given the increasing complexity of reconstruction algorithms, the issue of reproducibility is likely to be of increasing importance.

12. Conclusions

The area of surface reconstruction has grown from methods that handle limited defects in point clouds while producing detailed reconstructions, to methods that handle substantial artifacts and produce high-level surface representations. Our survey provides insight into this wide array of methods, highlighting strengths and limitations that currently exist in the field. In doing so, our survey should also point the way towards future work across all of the different priors – making potential connections across input assumptions, point cloud properties, and shape classes that have not been previously considered.

Hints and solvers. In our survey of recent work we observed how the surface reconstruction problem is often tackled through either increasingly sophisticated solvers or richer reconstruction hints that make the problem easier to solve. For example, the availability of oriented normals requires only a linear solve through the Poisson reconstruction approach [KH13], while unoriented normals require solving for a generalized eigenvalue problem [ACSTD07]. Other hints such as generalized parity requires two linear solves: a first solve to consolidate the local hints as sign guesses, and a second solve to recover a signed implicit function [GCSA13].

Innovations in acquisition. As 3D acquisition methods continue to increase in variety and popularity, surface reconstruction will continue to be an important component in acquiring real-world shapes. To provide some historical context, consider the rise in accessibility of the optical laser-based triangulation scanner: since such a device provides line of sight information, this resulted in a whole category of visibility priors (Section 5) such as VRIP [CL96]. A more recent example is the Microsoft Kinect: the real-time capture of depth and RGB imagery has resulted in a new slate of techniques for reconstructing indoor scenes through data-driven means (Section 9) such as [SXZ*12, NXS12, KMYG12], since objects in indoor environments tend to satisfy such priors. As novel acquisition sensors and modalities are developed, it is likely that surface reconstruction will become even more specialized in order to handle the nuances of the acquisition type. In this setting our survey should prove useful in the development of novel priors that need to handle such new acquisition methods.

Acquisition ubiquity. Beyond the increasing variety of sensors, we are also witnessing a rapid evolution of the acquisition paradigms. The acquisition of our physical world can

© 2016 Author version

now be complemented by exploiting the massive data sets shared online, referred to as *community data*. We also predict a future where geometric data are acquired through disseminated sensors, yielding *disseminated data*. This evolution translates into a paradox: despite expectations that technological advances should improve quality, these data are hampered with high variability and unprecedented amount and variety of defects. In addition, we are observing a trend brought on by the speed of technological progress: while many practitioners use high-end acquisition systems, an increasing number of them turn to *consumer-level* acquisition devices, willing to replace an accurate albeit expensive acquisition by a series of low-cost acquisitions – see recent work on 3D acquisition from mobile phones [TKM*13]. These new acquisition paradigms translate into a lower control over the acquisition process, which must be compensated by an increased robustness of the algorithms and structural or physical *a priori* knowledge.

Big data and online algorithms. The scale of acquired data is also quickly growing: we no longer deal exclusively with individual shapes, but with entire *scenes*, possibly at the scale of entire cities with many objects defined as structured shapes. Recovering not just the geometry but also the structure of such large scale scenes is a stimulating scientific challenge. Last, we envision a future where the common on-disk paradigm must be replaced by *online* algorithms that perform reconstruction during acquisition. Recent works such as Kinect Fusion [NDI*11] and extensions [NZIS13] demonstrate the practicality of building such online systems. There are applications such as aero-reconstruction for disaster management where tight timing restrictions make an online reconstruction approach indispensable. In particular, we foresee a need to extend the surveyed priors into the online setting, in order to support such challenging problems in surface reconstruction.

Acknowledgments. Pierre Alliez is supported by an ERC Starting Grant "Robust Geometry Processing" (257474). Claudio Silva was partially supported by the National Science Foundation grant MRI-1229185, an IBM Faculty Award, and a grant from the Gordon and Betty Moore and Alfred P. Sloan Foundations. Joshua Levine is supported by the National Science Foundation award IIS-1314757.

References

[AA04] ALEXA M., ADAMSON A.: On normals and projection operators for surfaces defined by point sets. In *Proc. of the EG conf. on Point-Based Graphics* (2004). 10

[AA09] ALEXA M., ADAMSON A.: Interpolatory point set surfaces-convexity and hermite data. *ACM Trans. on Graphics* (2009). 10

[AB99] AMENTA N., BERN M.: Surface reconstruction by voronoi filtering. *Discrete & Computational Geometry* (1999). 16

[ABCO*03] ALEXA M., BEHR J., COHEN-OR D., FLEISHMAN S., LEVIN D., SILVA C.: Computing and rendering point set surfaces. *Trans. on Visualization and Computer Graphics* (2003). 3, 4, 5, 7, 9

[ACSTD07] ALLIEZ P., COHEN-STEINER D., TONG Y., DESBRUN M.: Voronoi-based variational reconstruction of unoriented point sets. In *Computer Graphics Forum (Proc. of the Symposium on Geometry Processing)* (2007). 3, 5, 12, 23

[AK04] AMENTA N., KIL Y. J.: Defining point-set surfaces. *ACM Trans. on Graphics* (2004). 5, 10

[ASF*13] ARIKAN M., SCHWÄRZLER M., FLÖRY S., WIMMER M., MAIERHOFER S.: O-snap: Optimization-based snapping for modeling architecture. *ACM Trans. Graph. (Proc. SIGGRAPH)* (2013). 3, 9, 22

[ASGCO10] AVRON H., SHARF A., GREIF C., COHEN-OR D.: ℓ_1-sparse reconstruction of sharp point set surfaces. *ACM Trans. on Graphics* (2010). 7, 14

[BAMJ*11] BENES B., ABDUL-MASSIH M., JARVIS P., ALIAGA D. G., VANEGAS C. A.: Urban ecosystem design. In *Symposium on Interactive 3D Graphics and Games* (2011), pp. 167–174. 2

[BdLGM14] BOULCH A., DE LA GORCE M., MARLET R.: Piecewise-planar 3d reconstruction with edge and corner regularization. In *Proc. of the EG/SIGGRAPH Symposium on Geometry processing* (2014). 19

[BLN*13] BERGER M., LEVINE J. A., NONATO L. G., TAUBIN G., SILVA C. T.: A benchmark for surface reconstruction. *ACM Trans. on Graphics* (2013). 23

[BNB13] BRADLEY D., NOWROUZEZAHRAI D., BEARDSLEY P.: Image-based reconstruction and synthesis of dense foliage. *ACM Trans. Graph. (Proc. SIGGRAPH)* (2013). 8, 21

[BS12] BERGER M., SILVA C. T.: Medial kernels. In *Computer Graphics Forum (Proc. of Eurographics)* (2012). 18

[BWM*11] BERNER A., WAND M., MITRA N. J., MEWES D., SEIDEL H.-P.: Shape analysis with subspace symmetries. *Computer Graphics Forum* (2011). 19

[CBC*01] CARR J. C., BEATSON R. K., CHERRIE J. B., MITCHELL T. J., FRIGHT W. R., MCCALLUM B. C., EVANS T. R.: Reconstruction and representation of 3d objects with radial basis functions. In *Proc. of ACM SIGGRAPH* (2001). 3, 5, 6, 7, 11

[CCLN10] CHEN Y.-L., CHEN B.-Y., LAI S.-H., NISHITA T.: Binary orientation trees for volume and surface reconstruction from unoriented point clouds. *Computer Graphics Forum* (2010). 16

[CG06] CAZALS F., GIESEN J.: Delaunay triangulation based surface reconstruction. In *Effective Computational Geometry for Curves and Surfaces*. Springer, 2006. 3

[CGBG13] CHEN J., GUENNEBAUD G., BARLA P., GRANIER X.: Non-oriented MLS Gradient Fields. *Computer Graphics Forum* (2013). 10

[CL96] CURLESS B., LEVOY M.: A volumetric method for building complex models from range images. In *Proc. of ACM SIGGRAPH* (1996). 3, 6, 7, 15, 23

[CLCL11] CHEN Y.-L., LEE T.-Y., CHEN B.-Y., LAI S.-H.: Bipartite polar classification for surface reconstruction. *Computer Graphics Forum* (2011). 16

[CLP10] CHAUVE A.-L., LABATUT P., PONS J.-P.: Robust piecewise-planar 3d reconstruction and completion from large-scale unstructured point data. In *CVPR* (2010). 18

[CP05] CAZALS F., POUGET M.: Estimating differential quantities using polynomial fitting of osculating jets. *Computer Aided Geometric Design* (2005). 5

© 2016 Author version

[CT11] CALAKLI F., TAUBIN G.: Ssd: Smooth signed distance surface reconstruction. *Computer Graphics Forum* (2011). 12

[CTO*10] CAO J., TAGLIASACCHI A., OLSON M., ZHANG H., SU Z.: Point cloud skeletons via laplacian based contraction. In *Proc. of IEEE Shape Modeling International* (2010). 17

[CWL*08] CHENG Z.-Q., WANG Y.-Z., LI B., XU K., DANG G., JIN S.-Y.: A survey of methods for moving least squares surfaces. In *Proceedings of the Fifth Eurographics/IEEE VGTC conference on Point-Based Graphics* (2008), Eurographics Association. 9

[DCSA*13] DIGNE J., COHEN-STEINER D., ALLIEZ P., DE GOES F., DESBRUN M.: Feature-preserving surface reconstruction and simplification from defect-laden point sets. *Journal of Mathematical Imaging and Vision* (2013), 1–14. 14

[Dey07] DEY T. K.: *Curve and surface reconstruction: algorithms with mathematical analysis*. Cambridge University Press, 2007. 3

[DHOS07] DANIELS J. I., HA L. K., OCHOTTA T., SILVA C. T.: Robust smooth feature extraction from point clouds. In *Proc. of Shape Modeling and Applications* (2007). 13

[DKP11] DENNING J. D., KERR W. B., PELLACINI F.: Meshflow: interactive visualization of mesh construction sequences. In *ACM Trans. on Graphics* (2011). 23

[DMSL11] DIGNE J., MOREL J.-M., SOUZANI C.-M., LARTIGUE C.: Scale space meshing of raw data point sets. In *Computer Graphics Forum* (2011). 3

[DUNI10] DROST B., ULRICH M., NAVAB N., ILIC S.: Model globally, match locally: Efficient and robust 3d object recognition. In *CVPR* (2010). 18

[FCOS05] FLEISHMAN S., COHEN-OR D., SILVA C. T.: Robust moving least-squares fitting with sharp features. *ACM Trans. Graph. (Proc. SIGGRAPH)* (2005). 7, 13

[FG11] FUHRMANN S., GOESELE M.: Fusion of depth maps with multiple scales. In *Proc. of ACM SIGGRAPH Asia* (2011). 3, 15

[FG14] FUHRMANN S., GOESELE M.: Floating scale surface reconstruction. *ACM Trans. Graph. (Proc. SIGGRAPH)* (2014). 15

[FKSS08] FREIRE J., KOOP D., SANTOS E., SILVA C. T.: Provenance for computational tasks: A survey. *Comp. in Science & Engineering* (2008). 23

[GCSA13] GIRAUDOT S., COHEN-STEINER D., ALLIEZ P.: Noise-adaptive shape reconstruction from raw point sets. *Computer Graphics Forum (Proc. of the Symposium on Geometry Processing)* (2013). 16, 17, 23

[GG07] GUENNEBAUD G., GROSS M.: Algebraic point set surfaces. *ACM Trans. Graph. (Proc. SIGGRAPH)* (2007). 4, 10

[HDD*92] HOPPE H., DEROSE T., DUCHAMP T., MCDONALD J., STUETZLE W.: Surface reconstruction from unorganized points. In *Computer Graphics (Proc. SIGGRAPH)* (1992). 3, 5, 6, 7, 9

[HK06] HORNUNG A., KOBBELT L.: Robust reconstruction of watertight 3d models from non-uniformly sampled point clouds without normal information. In *Computer Graphics Forum (Proc. of the Symposium on Geometry Processing)* (2006). 3, 5, 13

[HKH*12] HENRY P., KRAININ M., HERBST E., REN X., FOX D.: Rgb-d mapping: Using kinect-style depth cameras for dense 3d modeling of indoor environments. *The International Journal of Robotics Research* (2012). 5

[HLZ*09] HUANG H., LI D., ZHANG H., ASCHER U., COHEN-OR D.: Consolidation of unorganized point clouds for surface reconstruction. *ACM Trans. on Graphics* (2009). 6, 11

[HWCO*13] HUANG H., WU S., COHEN-OR D., GONG M., ZHANG H., LI G., CHEN B.: l_1-medial skeleton of point cloud. *ACM Trans. Graph. (Proc. SIGGRAPH)* (2013). 3, 17, 22

[HWG*13] HUANG H., WU S., GONG M., COHEN-OR D., ASCHER U., ZHANG H. R.: Edge-aware point set resampling. *ACM Trans. on Graphics* (2013). 14

[HWK15] HUANG Q., WANG H., KOLTUN V.: Single-view reconstruction via joint analysis of image and shape collections. *ACM Trans. Graph. (Proc. SIGGRAPH)* (2015). 4

[JKS08] JENKE P., KRÜCKEBERG B., STRASSER W.: Surface reconstruction from fitted shape primitives. In *Proc. of Vision modeling and Visualization* (2008). 8, 18

[JWS08] JENKE P., WAND M., STRASSER W.: Patch-graph reconstruction for piecewise smooth surfaces. *Proc. of Vision modeling and Visualization* (2008). 13

[Kaz05] KAZHDAN M.: Reconstruction of solid models from oriented point sets. In *Proc. of the EG/SIGGRAPH Symposium on Geometry processing* (2005). 5, 12, 22

[KBH06] KAZHDAN M., BOLITHO M., HOPPE H.: Poisson surface reconstruction. In *Proc. of the EG/SIGGRAPH Symposium on Geometry processing* (2006). 3, 4, 5, 6, 7, 12, 23

[KH13] KAZHDAN M., HOPPE H.: Screened poisson surface reconstruction. *ACM Trans. on Graphics* (2013). 12, 23

[KMHG13] KIM Y. M., MITRA N. J., HUANG Q., GUIBAS L.: Guided real-time scanning of indoor objects. In *Computer Graphics Forum* (2013). 8, 21

[KMYG12] KIM Y. M., MITRA N. J., YAN D.-M., GUIBAS L.: Acquiring 3d indoor environments with variability and repetition. *Proc. of ACM SIGGRAPH Asia* (2012). 3, 8, 21, 23

[KSO04] KOLLURI R., SHEWCHUK J. R., O'BRIEN J. F.: Spectral surface reconstruction from noisy point clouds. In *Proc. of the EG/SIGGRAPH Symposium on Geometry processing* (2004). 13

[KTB07] KATZ S., TAL A., BASRI R.: Direct visibility of point sets. *ACM Trans. Graph. (Proc. SIGGRAPH)* (2007). 16

[LA13] LAFARGE F., ALLIEZ P.: Surface reconstruction through point set structuring. *Computer Graphics Forum (Proc. of Eurographics)* (2013). 3, 8, 19

[LB07] LEMPITSKY V., BOYKOV Y.: Global optimization for shape fitting. In *CVPR* (2007). 15

[LCDF10] LIPMAN Y., CHEN X., DAUBECHIES I., FUNKHOUSER T.: Symmetry factored embedding and distance. *ACM Trans. on Graphics* (2010). 19

[LCOL06] LIPMAN Y., COHEN-OR D., LEVIN D.: Error bounds and optimal neighborhoods for mls approximation. In *Computer Graphics Forum (Proc. of the Symposium on Geometry Processing)* (2006). 4

[LCOL07] LIPMAN Y., COHEN-OR D., LEVIN D.: Data-dependent MLS for faithful surface approximation. In *Computer Graphics Forum (Proc. of the Symposium on Geometry Processing)* (2007). 13

[LCOLTE07] LIPMAN Y., COHEN-OR D., LEVIN D., TAL-EZER H.: Parameterization-free projection for geometry reconstruction. *ACM Trans. Graph. (Proc. SIGGRAPH)* (2007). 3, 4, 11, 14

[LDGN15] LI Y., DAI A., GUIBAS L., NIESSNER M.: Database-assisted object retrieval for real-time 3d reconstruction. *Computer Graphics Forum (Proc. of Eurographics)* (2015). 21

[Lev03] LEVIN D.: Mesh-independent surface interpolation. *Geometric Modeling for Scientific Visualization* (2003), 37–49. 9

[LLZM10] LI G., LIU L., ZHENG H., MITRA N. J.: Analysis, reconstruction and manipulation using arterial snakes. *Proc. of ACM SIGGRAPH Asia* (2010). 3, 7, 17, 23

© 2016 Author version

[LPC*00] LEVOY M., PULLI K., CURLESS B., RUSINKIEWICZ S., KOLLER D., PEREIRA L., GINZTON M., ANDERSON S. E., DAVIS J., GINSBERG J., SHADE J., FULK D.: The digital michelangelo project: 3d scanning of large statues. In *Proc. of ACM SIGGRAPH* (2000), pp. 131–144. 2

[LPK09] LABATUT P., PONS J.-P., KERIVEN R.: Robust and efficient surface reconstruction from range data. *Computer Graphics Forum* (2009). 15

[LW10] LIU S., WANG C. C.: Orienting unorganized points for surface reconstruction. *Computers & Graphics* (2010). 6

[LWC*11] LI Y., WU X., CHRYSATHOU Y., SHARF A., COHEN-OR D., MITRA N. J.: Globfit: consistently fitting primitives by discovering global relations. *ACM Trans. Graph. (Proc. SIGGRAPH)* (2011). 3, 5, 8, 20

[LYO*10] LIVNY Y., YAN F., OLSON M., CHEN B., ZHANG H., EL-SANA J.: Automatic reconstruction of tree skeletal structures from point clouds. *Proc. of ACM SIGGRAPH Asia* (2010). 7, 17

[LZS*11] LI Y., ZHENG Q., SHARF A., COHEN-OR D., CHEN B., MITRA N. J.: 2d-3d fusion for layer decomposition of urban facades. In *ICCV* (2011). 3, 8, 20

[MDGD*10] MULLEN P., DE GOES F., DESBRUN M., COHEN-STEINER D., ALLIEZ P.: Signing the unsigned: Robust surface reconstruction from raw pointsets. *Computer Graphics Forum (Proc. of the Symposium on Geometry Processing)* (2010). 3, 4, 16, 17

[MGV11] MACÊDO I., GOIS J. P., VELHO L.: Hermite radial basis functions implicits. *Computer Graphics Forum* (2011). 11

[MMBM15] MONSZPART A., MELLADO N., BROSTOW G. J., MITRA N. J.: Rapter: Rebuilding man-made scenes with regular arrangements of planes. *ACM Trans. Graph. (Proc. SIGGRAPH)* (2015). 3, 8, 20

[MNG04] MITRA N. J., NGUYEN A., GUIBAS L.: Estimating surface normals in noisy point cloud data. *International Journal of Computational Geometry & Applications* (2004). 5

[MPS08] MANSON J., PETROVA G., SCHAEFER S.: Streaming surface reconstruction using wavelets. *Computer Graphics Forum* (2008). 12, 22

[MPWC13] MITRA N. J., PAULY M., WAND M., CEYLAN D.: Symmetry in 3d geometry: Extraction and applications. In *Computer Graphics Forum (STAR Proceedings of Eurographics)* (2013). 19

[MTSM10] MEHRA R., TRIPATHI P., SHEFFER A., MITRA N. J.: Visibility of noisy point cloud data. *Computers & Graphics* (2010). 16

[MWA*13] MUSIALSKI P., WONKA P., ALIAGA D. G., WIMMER M., GOOL L., PURGATHOFER W.: A survey of urban reconstruction. In *Computer Graphics Forum (STAR Proceedings of Eurographics)* (2013). 3

[MWZ*13] MITRA N. J., WAND M., ZHANG H., COHEN-OR D., BOKELOH M.: Structure-aware shape processing. In *Computer Graphics Forum (STAR Proceedings of Eurographics)* (2013). 19

[NDI*11] NEWCOMBE R. A., DAVISON A. J., IZADI S., KOHLI P., HILLIGES O., SHOTTON J., MOLYNEAUX D., HODGES S., KIM D., FITZGIBBON A.: Kinectfusion: Real-time dense surface mapping and tracking. In *IEEE international symposium on Mixed and augmented reality (ISMAR)* (2011), IEEE. 7, 15, 24

[NFD07] NEUBERT B., FRANKEN T., DEUSSEN O.: Approximate image-based tree-modeling using particle flows. *ACM Trans. Graph. (Proc. SIGGRAPH)* (2007). 17

[NFS15] NEWCOMBE R. A., FOX D., SEITZ S. M.: Dynamicfusion: Reconstruction and tracking of non-rigid scenes in real-time. In *CVPR* (2015). 22

[NOS09] NAGAI Y., OHTAKE Y., SUZUKI H.: Smoothing of partition of unity implicit surfaces for noise robust surface reconstruction. *Computer Graphics Forum (Proc. of the Symposium on Geometry Processing)* (2009). 11

[NOS15] NAGAI Y., OHTAKE Y., SUZUKI H.: Tomographic surface reconstruction from point cloud. *Computers & Graphics* (2015). 17

[NSC14] NAN L., SHARF A., CHEN B.: 2d-d lifting for shape reconstruction. In *Computer Graphics Forum* (2014). 7, 20

[NSF14] NATHAN SILBERMAN D. S., FERGUS R.: Instance segmentation of indoor scenes using a coverage loss. In *ECCV* (2014). 18

[NSZ*10] NAN L., SHARF A., ZHANG H., COHEN-OR D., CHEN B.: Smartboxes for interactive urban reconstruction. *ACM Trans. Graph. (Proc. SIGGRAPH)* (2010). 3, 9, 22

[NXS12] NAN L., XIE K., SHARF A.: A search-classify approach for cluttered indoor scene understanding. *Proc. of ACM SIGGRAPH Asia* (2012). 8, 21, 23

[NZIS13] NIESSNER M., ZOLLHÖFER M., IZADI S., STAMMINGER M.: Real-time 3d reconstruction at scale using voxel hashing. *ACM Transactions on Graphics (TOG)* (2013). 15, 24

[OBA*03] OHTAKE Y., BELYAEV A., ALEXA M., TURK G., SEIDEL H.: Multi-level partition of unity implicits. *ACM Trans. Graph. (Proc. SIGGRAPH)* (2003). 3, 6, 10

[OGG09] OZTIRELI C., GUENNEBAUD G., GROSS M.: Feature preserving point set surfaces based on non-linear kernel regression. In *Computer Graphics Forum* (2009). 4, 10, 14

[OLA15] OESAU S., LAFARGE F., ALLIEZ P.: Planar shape detection and regularization in tandem. 8, 20

[PGB03] PÉREZ P., GANGNET M., BLAKE A.: Poisson image editing. *ACM Trans. Graph. (Proc. SIGGRAPH)* (2003). 12

[PMA*14] PREINER R., MATTAUSCH O., ARIKAN M., PAJAROLA R., WIMMER M.: Continuous projection for fast l1 reconstruction. *tog* (2014). 11

[PMG04] PAULY M., MITRA N. J., GUIBAS L.: Uncertainty and variability in point cloud surface data. In *Proc. of the EG conf. on Point-Based Graphics* (2004). 5

[PMG*05] PAULY M., MITRA N. J., GIESEN J., GROSS M. H., GUIBAS L. J.: Example-based 3d scan completion. In *Proc. of the EG/SIGGRAPH Symposium on Geometry processing* (2005). 3, 8

[PMW*08] PAULY M., MITRA N. J., WALLNER J., POTTMANN H., GUIBAS L. J.: Discovering structural regularity in 3d geometry. *ACM Trans. Graph. (Proc. SIGGRAPH)* (2008). 3, 5, 19

[PSDB*10] POPA T., SOUTH-DICKINSON I., BRADLEY D., SHEFFER A., HEIDRICH W.: Globally consistent space-time reconstruction. *Computer Graphics Forum 29*, 5 (2010), 1633–1642. 4

[RFL*05] RUNIONS A., FUHRER M., LANE B., FEDERL P., ROLLAND-LAGAN A.-G., PRUSINKIEWICZ P.: Modeling and visualization of leaf venation patterns. *ACM Trans. Graph. (Proc. SIGGRAPH) 24*, 3 (2005), 702–711. 17

[RKMP13] REISNER-KOLLMANN I., MAIERHOFER S., PURGATHOFER W.: Reconstructing shape boundaries with multimodal constraints. *Computers & Graphics* (2013). 19

[SAL*08] SHARF A., ALCANTARA D. A., LEWINER T., GREIF C., SHEFFER A., AMENTA N., COHEN-OR D.: Space-time surface reconstruction using incompressible flow. *ACM Trans. on Graphics 27*, 5 (2008), 110. 4

© 2016 Author version

[SCC*11] SCOPIGNO R., CALLIERI M., CIGNONI P., CORSINI M., DELLEPIANE M., PONCHIO F., RANZUGLIA G.: 3d models for cultural heritage: Beyond plain visualization. *IEEE Computer 44*, 7 (2011), 48–55. 2

[SDK09] SCHNABEL R., DEGENER P., KLEIN R.: Completion and reconstruction with primitive shapes. *Computer Graphics Forum (Proc. of Eurographics)* (2009). 3, 5, 8, 18

[SFC07] SILVA C. T., FREIRE J., CALLAHAN S. P.: Provenance for visualizations: Reproducibility and beyond. *Comp. in Science & Engineering* (2007). 23

[SFCH12] SHEN C.-H., FU H., CHEN K., HU S.-M.: Structure recovery by part assembly. *Proc. of ACM SIGGRAPH Asia* (2012). 2, 3, 8, 21

[SHFH11] SHEN C.-H., HUANG S.-S., FU H., HU S.-M.: Adaptive partitioning of urban facades. *Proc. of ACM SIGGRAPH Asia* (2011). 19, 20

[SHM*14] SU H., HUANG Q., MITRA N. J., LI Y., GUIBAS L.: Estimating image depth using shape collections. *ACM Trans. Graph. (Proc. SIGGRAPH)* (2014). 4

[SLS*06] SHARF A., LEWINER T., SHAMIR A., KOBBELT L., COHEN-OR D.: Competing fronts for coarse–to–fine surface reconstruction. *Computer Graphics Forum (Proc. of Eurographics)* (2006). 18

[SLS*07] SHARF A., LEWINER T., SHKLARSKI G., TOLEDO S., COHEN-OR D.: Interactive topology-aware surface reconstruction. *ACM Trans. Graph. (Proc. SIGGRAPH)* (2007). 3, 5, 9, 22, 23

[SOS04] SHEN C., O'BRIEN J. F., SHEWCHUK J. R.: Interpolating and approximating implicit surfaces from polygon soup. *ACM Trans. on Graphics* (2004). 10

[SSFS06] SCHREINER J., SCHEIDEGGER C. E., FLEISHMAN S., SILVA C. T.: Direct (re)meshing for efficient surface processing. *Computer Graphics Forum (Proc. of Eurographics)* (2006). 13

[SSZCO10] SHALOM S., SHAMIR A., ZHANG H., COHEN-OR D.: Cone carving for surface reconstruction. *Proc. of ACM SIGGRAPH Asia* (2010). 3, 16

[SWK07] SCHNABEL R., WAHL R., KLEIN R.: Efficient ransac for point-cloud shape detection. *Computer Graphics Forum* (2007). 5, 8, 18, 19, 20

[SXZ*12] SHAO T., XU W., ZHOU K., WANG J., LI D., GUO B.: An interactive approach to semantic modeling of indoor scenes with an rgbd camera. *Proc. of ACM SIGGRAPH Asia* (2012). 3, 8, 20, 21, 23

[SY12] SEVERSKY L. M., YIN L.: A global parity measure for incomplete point cloud data. *Computer Graphics Forum* (2012). 16

[TKM*13] TANSKANEN P., KOLEV K., MEIER L., CAMPOSECO F., SAURER O., POLLEFEYS M.: Live metric 3d reconstruction on mobile phones. In *ICCV* (2013). 24

[TOZ*11] TAGLIASACCHI A., OLSON M., ZHANG H., HAMARNEH G., COHEN-OR D.: Vase: Volume-aware surface evolution for surface reconstruction from incomplete point clouds. *Computer Graphics Forum (Proc. of the Symposium on Geometry Processing)* (2011). 3, 6, 18

[TZCO09] TAGLIASACCHI A., ZHANG H., COHEN-OR D.: Curve skeleton extraction from incomplete point cloud. *ACM Trans. Graph. (Proc. SIGGRAPH)* (2009). 2, 3, 5, 7, 17, 18

[VAB12] VANEGAS C. A., ALIAGA D. G., BENES B.: Automatic extraction of manhattan-world building masses from 3d laser range scans. *Trans. on Visualization and Computer Graphics* (2012). 5, 20

[vKvLV13] VAN KREVELD M., VAN LANKVELD T., VELTKAMP R. C.: Watertight scenes from urban lidar and planar surfaces. *Computer Graphics Forum* (2013). 3, 8, 19

[vKZHCO11] VAN KAICK O., ZHANG H., HAMARNEH G., COHEN-OR D.: A survey on shape correspondence. In *Computer Graphics Forum* (2011). 5

[VLA15] VERDIE Y., LAFARGE F., ALLIEZ P.: *LOD Generation for Urban Scenes*. Tech. rep., 2015. 23

[WBLP11] WEISE T., BOUAZIZ S., LI H., PAULY M.: Realtime performance-based facial animation. 21

[Wen05] WENDLAND H.: *Scattered Data Approximation*. Cambridge University Press, 2005, ch. 16.2 - Hermite–Birkhoff interpolation. 11

[Whi98] WHITAKER R. T.: A level-set approach to 3D reconstruction from range data. *International Journal of Computer Vision 29*, 3 (1998), 203–231. 15

[WS12] WAN G., SHARF A.: Grammar-based 3d facade segmentation and reconstruction. *Computers & Graphics* (2012). 19, 20

[WSL*14] WU S., SUN W., LONG P., HUANG H., COHEN-OR D., GONG M., DEUSSEN O., CHEN B.: Quality-driven poisson-guided autoscanning. *Proc. of ACM SIGGRAPH Asia* (2014). 22

[WSS09] WANG H., SCHEIDEGGER C. E., SILVA C. T.: Bandwidth selection and reconstruction quality in point-based surfaces. *Trans. on Visualization and Computer Graphics* (2009). 4

[WYZC13] WANG J., YU Z., ZHU W., CAO J.: Feature-preserving surface reconstruction from unoriented, noisy point data. *Computer Graphics Forum* (2013). 13

[XF12] XIAO J., FURUKAWA Y.: Reconstructing the worlds museums. In *ECCV* (2012). 3, 8, 19

[XZZ*14] XIONG S., ZHANGY J., ZHENG J., CAI J., LIU L.: Robust surface reconstruction via dictionary learning. *Proc. of ACM SIGGRAPH Asia* (2014). 3, 14

[YHZ*14] YIN K., HUANG H., ZHANG H., GONG M., COHEN-OR D., CHEN B.: Morfit: Interactive surface reconstruction from incomplete point clouds with curve-driven topology and geometry control. *Proc. of ACM SIGGRAPH Asia* (2014). 3, 9, 22

[YSL*14] YAN F., SHARF A., LIN W., HUANG H., CHEN B.: Proactive 3d scanning of inaccessible parts. *ACM Trans. Graph. (Proc. SIGGRAPH)* (2014). 22

[ZDI*15] ZOLLHÖFER M., DAI A., INNMANN M., WU C., STAMMINGER M., THEOBALT C., NIESSNER M.: Shading-based refinement on volumetric signed distance functions. *ACM Trans. Graph. (Proc. SIGGRAPH)* (2015). 15

[ZK15] ZHOU Q.-Y., KOLTUN V.: Depth camera tracking with contour cues. In *CVPR* (2015). 15

[ZN12] ZHOU Q.-Y., NEUMANN U.: 2.5 d building modeling by discovering global regularities. In *CVPR* (2012). 20

[ZPB07] ZACH C., POCK T., BISCHOF H.: A globally optimal algorithm for robust tv-l1 range image integration. In *ICCV* (2007). 3, 7, 15

[ZSW*10] ZHENG Q., SHARF A., WAN G., LI Y., MITRA N. J., COHEN-OR D., CHEN B.: Non-local scan consolidation for 3d urban scenes. *ACM Trans. Graph. (Proc. SIGGRAPH)* (2010). 2, 3, 8, 19, 20

[ZYFY14] ZHANG C., YE M., FU B., YANG R.: Data-driven flower petal modeling with botany priors. In *CVPR* (2014). 8, 21

© 2016 Author version