Quality-driven Poisson-guided Autoscanning

Shihao Wu[1] Wei Sun[1] Pinxin Long[1] Hui Huang[1*]
Daniel Cohen-Or[2] Minglun Gong[3] Oliver Deussen[4] Baoquan Chen[5]
[1]Shenzhen VisuCA Key Lab / SIAT [2]Tel-Aviv University
[3]Memorial University of Newfoundland [4]University of Konstanz [5]Shandong University

Figure 1: *Our robot-based, Poisson-guided autoscanner can progressively, adaptively, and fully automatically generate complete, high quality, and high fidelity scan models.*

Abstract

We present a quality-driven, Poisson-guided autonomous scanning method. Unlike previous scan planning techniques, we do not aim to minimize the number of scans needed to cover the object's surface, but rather to ensure the high quality scanning of the model. This goal is achieved by placing the scanner at strategically selected Next-Best-Views (NBVs) to ensure progressively capturing the geometric details of the object, until both completeness and high fidelity are reached. The technique is based on the analysis of a Poisson field and its geometric relation with an input scan. We generate a confidence map that reflects the quality/fidelity of the estimated Poisson iso-surface. The confidence map guides the generation of a viewing vector field, which is then used for computing a set of NBVs. We applied the algorithm on two different robotic platforms, a PR2 mobile robot and a one-arm industry robot. We demonstrated the advantages of our method through a number of autonomous high quality scannings of complex physical objects, as well as performance comparisons against state-of-the-art methods.

Keywords: 3D acquisition, autonomous scanning, next-best-view, Poisson surface reconstruction

Links: DL PDF WEB VIDEO

*Corresponding author: Hui Huang (hhzhiyan@gmail.com)

1 Introduction

Scanning real physical objects is becoming more and more commonplace, which stimulates much research on accurate shape reconstruction from raw scanned data [Berger et al. 2014]. One of the main data capturing challenges via laser scanning is to fully cover the entire surface of an object with a defined accuracy. The coverage problem is difficult to solve due to physical occlusions, less than ideal material properties or potentially imprecise movement of the scanner [Tagliasacchi et al. 2009; Mullen et al. 2010; Huang et al. 2013]. Efficiently obtaining a complete surface requires a good planning based on a priori knowledge of the scanned object. When the shape is unknown, for instance in a robot-operated setting, an autonomous scanner can be strongly assisted by an on-the-fly geometric analysis of the acquired data to guide the completion of the full coverage [Scott et al. 2003].

Obtaining guidance on where to position the sensor next in order to increase the information gain is known as the Next-Best-View (NBV) planning problem, first proposed by Connolly [1985]. As we shall elaborate in the next section, the common approach to the NBV problem is to analyze the temporal boundaries of the patches reconstructed from the partial scan, aiming to complete the surface. Alternatively, the partial scan can be analyzed by considering a global watertight reconstruction and its ambient space. This approach requires a non-trivial analysis of a 3D point cloud of an unknown shape. In [Pauly et al. 2004], measures of uncertainty and variability of a given point cloud surface are presented. Their analysis mainly focus on density, regularity and curvature. Such measures may form an estimate about where more scans are needed to improve the quality of the *current* data. However, their analysis does not account for geometric completeness and topological aspects, and in particular holes or occluded regions that need to be covered by additional scans.

In this paper, we introduce a method for point cloud analysis that accounts for both completeness and quality of the reconstruction.

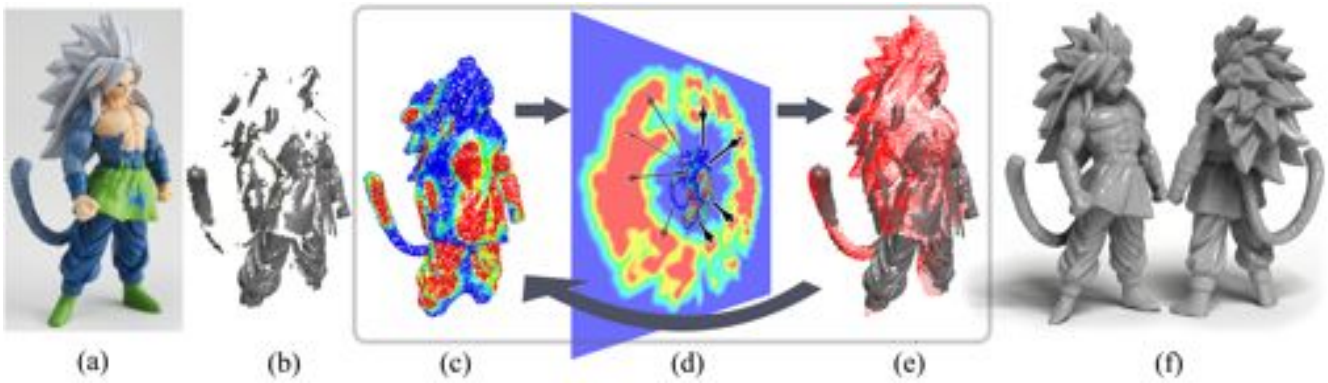

Figure 2: *Autoscanning overview: given an incomplete point cloud (b) obtained by a blind scanning of an unknown object (a), we first reconstruct a Poisson iso-surface and estimate its confidence map (c), where high confidence areas are shown in red and low confidence in blue. A 3D viewing vector field (VVF) is then generated to determine a set of next-best-views (NBVs). A slice of the VVF is visualized in (d), where black arrows show the NBVs. Scanning from these NBVs captures more points (red in (e)) in low confidence areas. The scanning process is iterated until convergence to a high quality reconstruction (f).*

Unlike previous scan planning methods, we do not aim at minimizing the number of scans needed to cover the whole surface. Instead our objective is to ensure high fidelity scanning of the model. The core idea is to *validate* and *assess* a watertight surface reconstructed from a Screened Poisson equation [Kazhdan and Hoppe 2013], which leads to a confidence map defined on an estimated iso-surface. We then build an ambient viewing field, which suggests us a set of NBVs with five degrees of freedom for covering low confidence areas on the iso-surface. By positioning the scanner at the suggested NBVs, the surface of the input object is *progressively* scanned at high fidelity.

We have implemented and experimented our autoscanning technique on two types of robots. One is a mobile robot PR2 from Willow Garage (Figure 1) and the other is a one-arm industry robot from Reis Robotics (model RV20-6, see Figure 10). Using these two robotic platforms we tested our algorithm on a number of real and complex objects with initially unknown shapes. To compare our method with state-of-the-art NBV techniques we perform a qualitative comparison on a real object and a quantitative evaluation on a virtual model; see Figures 11, 9 and 12.

2 Related work

Autoscanning of a scene or an object is a challenging view planning problem since in many cases a priori knowledge of the object or scene geometry is unavailable. Typically, planning is computed on-the-fly and based on the information collected so far from previous scans. Computing the potentially best view for the next scan is called the Next-Best-View (NBV) problem. It has been studied since 1980s [Connolly 1985], but still remains active. The problem is known to be NP-hard; however, it can be reduced to the well studied coverage problem that can be solved efficiently [Low and Lastra 2006; Chen et al. 2011].

Acquiring and reconstructing of indoor/outdoor scenes as well as navigating in such environments pose fundamental problems for autonomous robotics system [Low and Lastra 2006; Blaer and Allen 2007; Chen et al. 2011; Zhou and Koltun 2013]. KinectFusion based systems [Newcombe et al. 2011; Chen et al. 2013] recently gained popularity for real-time dense surface mapping, tracking, and reconstruction with mobile low-cost depth cameras. Kim et al. [2013] presents a guided real-time scene scanning interface in which the user receives interactive feedback based on 3D shape retrieval operations supported by a model repository. The work focuses on the coverage and exploration of scenes at a coarse level. Li et al. [2013] presents a system to allow users scanning themselves with a single Kinect by rotating in the same pose for a few different views, this allows the scanner to cover the full body. However, scanning and registration in this case highly depends on having the human body shape as a prior. Some NBV techniques [Wenhardt et al. 2007; Dunn and Frahm 2009] build 3D models from images, based on camera movements and on utilizing the model's covariance structure as well as texture appearance. In our work, we instead aim at high fidelity 3D scanning of *unknown* objects that allow automatic reconstruction of objects with rich geometric details and a complex topology.

For scanning of single objects, Scott et al. [2003] classify the related NBV algorithms into two main categories: model-based [Trucco et al. 1997] and non-model-based methods. The challenge of the latter case, i.e., automatically finding NBVs for unknown objects, has been widely studied in the robotics literature [Chen et al. 2011]. The pioneer works are occlusion-guided methods [Maver and Bajcsy 1993; Pito 1996] or based on a volumetric analysis [Banta et al. 2000; Vásquez-Gómez et al. 2009]. However, these techniques restrict the search space for the NBV onto a sphere or a generalized cylinder. Such simplification helps to reduce the computation cost but limits the precision, thus level of fidelity, when scanning scanning complex, self-occluded, or feature-rich objects.

A number of surface-based NBV algorithms have been recently proposed [Khalfaoui et al. 2013; Kriegel et al. 2013], which determine NBVs by analyzing the shape of the so-far scanned surface. In general, these techniques detect visibility limitations or open boundaries and guide the scanner to cover the surface accordingly. While these approaches produce complete scans for the input objects, they do not guarantee that the reconstruction captures all geometric details on the surfaces of the objects. For example, Kriegel et al. [2013] reconstruct a partial surface first and then determine successive NBVs solely based on holes or boundaries on the surface. This implicitly assumes that the areas already been scanned are of sufficient quality, which may not be true in many practical cases. In comparison, our approach assesses both the completeness and the local smoothness of the reconstructed surface, which results in high-fidelity scan models.

For point cloud assessment, our work is closely related to Pauly et al. [2004], which present uncertainty and variability analysis of

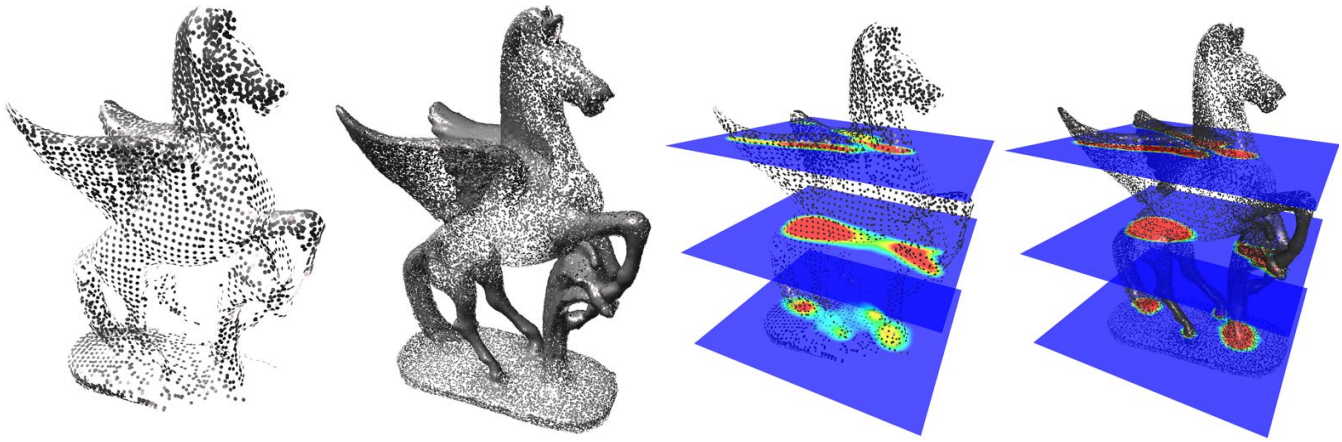

(a) Incomplete point cloud. (b) Complete point cloud. (c) Poisson field of (a). (d) Poisson field of (b).

Figure 3: *Visualization of the Poisson scalar fields for incomplete and complete point clouds. In areas with missing data the field is blurry, i.e., it has a low gradient.*

point-sampled geometry. They introduce a statistical representation for each point in space. It quantifies the likelihood that a surface obtained by least-square data fitting passes through that point. This likelihood map is combined with a corresponding confidence map that measures the quality of local normal estimations. In our work, we also investigate raw point clouds. However, we measure not only existing points, but also the predicted points from a Poisson field, and use such measurement to the acquisition of new points. This is conducted in an iterative process, through selecting next best scans and integrating new scans.

In a broader context, our work is also related to many recent publications on surface reconstruction from point cloud. The most notable works are methods for reconstructing watertight surface models [Wheeler et al. 1998; Carr et al. 2001; Kazhdan et al. 2006; Kazhdan and Hoppe 2013]. Other related works deal with the consolidation of raw point data [Lipman et al. 2007; Huang et al. 2009], surface completion [Davis et al. 2002; Sharf et al. 2004; Sagawa and Ikeuchi 2008; Tagliasacchi et al. 2009; Harary et al. 2014; Huang et al. 2013] or differentiating real holes from incomplete scanning data [Seversky and Yin 2012].

3 Overview

The core of our autonomous scanning system consists of an analysis of the data acquired by a scanner and the generation of a set of NBVs for the scanning robot. We evaluate the quality of a global, complete iso-surface model extracted from a Poisson field, rather than only analyzing existing points close to open boundaries. The evaluation is based on a confidence map, which is generated by assigning each surface sample a confidence score that reflects its measured quality; see Figure 2(c). The confidence measure is derived from two criteria: one measures the directional gradients of the Poisson field in the vicinity of the iso-surface and assesses the stability of the local topology. The other criterion analyzes to what extent the available scans support the extracted iso-surface.

Once the confidence map is computed on the iso-surface, we define a 3D viewing vector field (Figure 2(d)). The magnitude of the vectors reflect their potential for being the next scanning locations, and their directions suggest a rough scan orientation. A sparse set of positions having local maxima of vector magnitudes is then selected to form a set of NBVs for the next scanning iteration. For each selected position, its optimal scanning direction is computed. According to the current position of the robot arm, these NBVs are ordered into a sequence for the next scan path (Figure 2(e)).

Figure 2 illustrates the scanning pipeline, which starts with a *blind*, all-around scanning of the object and iteratively refines the model by additional scans from automatically selected NBVs. The iterative process continues until a high-quality model is obtained or the reconstruction cannot be further improved.

4 Iso-surface estimation and measurement

Given an object with unknown shape, blind scanning around the object give us an initial point cloud $Q = \{q_j\}_{j \in J} \subset \mathbb{R}^3$, e.g., in Figure 2(b). Our first goal is to reconstruct a tentative, watertight surface model from Q and estimate its reconstruction quality.

4.1 Point consolidation and Poisson reconstruction

Depending on the quality of the scanner, the initial raw point cloud may be corrupted due to noise and outliers, and may not be equipped with reliable normals. In this case, we apply a *point cloud consolidation* preprocessing step (includes noise filtering, outliers removal, normal estimation, and thinning), which facilitates the estimation of a robust iso-surface. Here a Weighted Locally Optimal Projection operator [Huang et al. 2009] is employed, which outputs a cleaned and oriented point set $P = \{(p_i, \mathbf{n}_i)\}_{i \in I} \subset \mathbb{R}^6$ that adheres faithfully to the underlying shape.

Having a consolidated point set P, we are able to generate a tentative model for the object using the so-called Screened Poisson surface reconstruction method [Kazhdan et al. 2006; Kazhdan and Hoppe 2013]. This approach interprets an inward pointing normal field of a surface as the gradient of its solid indicator function. Thus, a surface can be obtained by transforming the oriented points into a continuous vector field in 3D, and then extracting the appropriate iso-surface from an indicator function whose gradients best match the vector field (see Figure 3(c) and (d)). The resulting model is watertight and of high quality if the underlying points are dense and accurate enough. However, if there are areas with missing data or being insufficiently sampled, the reconstructed Poisson surface may contain topological errors and/or miss geometry details. Hence, our goal is to automatically identify these areas and apply additional s-

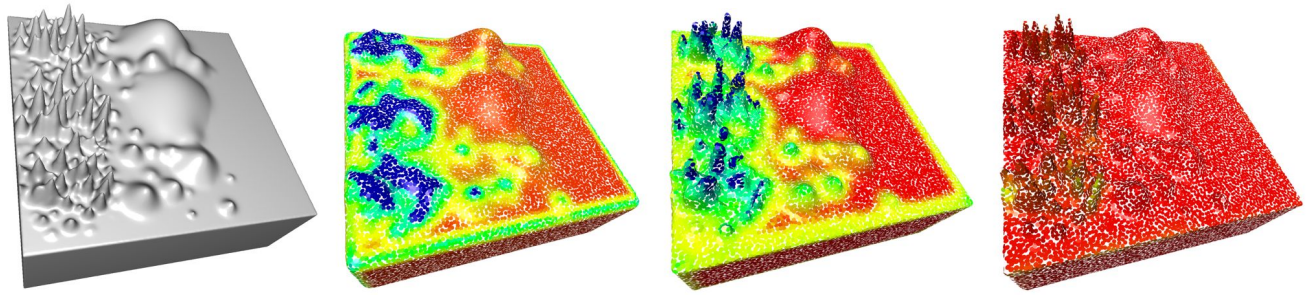

(a) A synthetic model. (b) One scan with f_s^0. (c) Five scans with f_s^5. (d) Ten scans with f_s^{10}.

Figure 4: *Visualization of smoothness confidence scores f_s. The synthetic input model (a) is designed to contain a spatially-varying geometric details. The extracted Poisson iso-surface by a single scan (b) does not accurately represent the input model and the confidence scores properly identify the problematic areas. Adding more scans (c and d), the confidence scores increase (from blue to red) and converge after all details are captured and reconstructed well.*

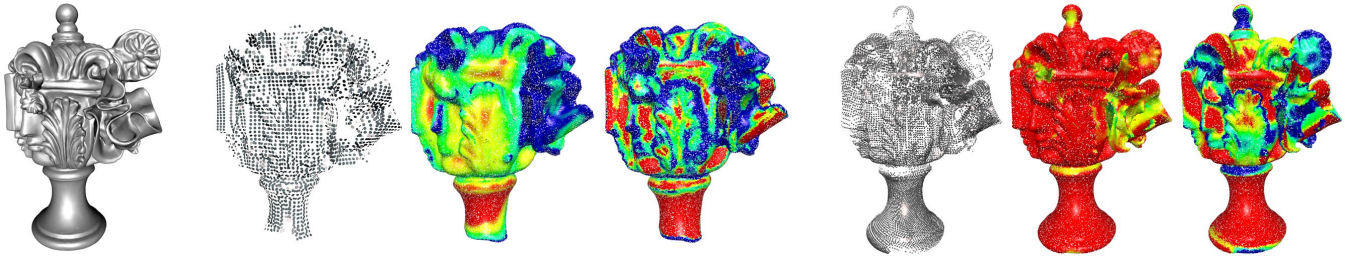

(a) Artwork. (b) Q^0. (c) f_g^0. (d) f_s^0. (e) Q^4. (f) f_g^4. (g) f_s^4.

Figure 5: *Given a synthetic input model (a), the point cloud obtained from the initial blind scanning is shown in (b) and the corresponding surface reconstructed is shown in (c-d), where the completeness (f_g) and smoothness (f_s) confidence scores are mapped, respectively. The surface reconstructed after 4 scanning iterations (e) is shown in (f-g) with confidence scores being mapped in the same order.*

cans so that the Poisson iso-surface extracted accordingly can be used to accurately reconstruct the surface of a real object.

4.2 Quality assessment on a Poisson iso-surface

To evaluate the quality of the obtained Poisson iso-surface, we first discretize it into a set of oriented points, $S = \{(s_k, \mathbf{n}_k)\}_{k\in K} \subset \mathbb{R}^6$, using Poisson-disk sampling [Corsini et al. 2012]. These points are regularly distributed on the iso-surface and hereafter we refer to them as *iso-points*. Two criteria are used in our approach for evaluating how well these iso-points approximate the true underlying surface of the input object. Firstly, we measure our confidence in the locations of the iso-points, which reside at the zero-crossings of the Poisson scalar field. If this scalar field varies too slowly, we cannot precisely locate the zero-crossings, which is an important indicator that more scans are needed. Moreover, reliable iso-points should be supported by a sufficient number of raw scan samples. Hence, for a given iso-point s_k, we locate the nearby points from Q and evaluate how consistent their locations are with respect to s_k.

The first criterion is based on the observation that Poisson surface reconstruction highly relies on the density and smoothness of the underlying points. The denser and smoother the input points are, the more rapidly the scalar field varies, and the closer the iso-points are to their exact locations on the surface. In areas with missing data the Poisson scalar field is blurry and the extracted iso-points are unreliable; see Figure 3. Accordingly, we define the *completeness* confidence score $f_g(s_k, \mathbf{n}_k)$ for iso-point s_k as

$$f_g(s_k, \mathbf{n}_k) = \Gamma(s_k) \cdot \mathbf{n}_k, \tag{1}$$

where $\Gamma(s_k)$ is the gradient of the scalar field at location s_k. It is worth noting that Γ is calculated during the Poisson reconstruction process and hence no additional computation is needed.

The second criterion directly evaluates a given iso-point s_k using the raw input points. Here we locate all raw points $\{q_j\}_{j\in\Omega_k}$ in the K-nearest neighborhood Ω_k of s_k. Theoretically, a large value of this parameter K leads to a smooth confidence map, while a small value may emphasize more the fine scale features. In practice, we find that this parameter does not have big effect on results. It is set to 100 (a relatively small value) throughout our experiments.

We then compute the *smoothness* confidence score $f_s(s_k, \mathbf{n}_k)$ for iso-point s_k using the following bilateral weighted sum:

$$\begin{aligned} f_s(s_k, \mathbf{n}_k) &= \sum_{j\in\Omega_k} \theta(\|s_k - q_j\|)\phi(\mathbf{n}_k, q_j - s_k), \quad (2)\\ \theta(\|s_k - q_j\|) &= e^{-\|s_k - q_j\|^2/(h_k/2)^2},\\ \phi(\mathbf{n}_k, q_j - s_k) &= e^{-\|\mathbf{n}_k^\top (q_j - s_k)\|^2/(h_k/4)^2}, \end{aligned}$$

where $\|\cdot\|$ is the ℓ_2-norm. The spatial weighting function $\theta(\|s_k - q_j\|)$ rapidly decays with the support radius $h_k = \max_{j\in\Omega_k}\{\|s_k - q_j\|\}$; $\phi(\mathbf{n}_k, q_j - s_k)$ is the orthogonal weighting function that measures how distant the Knn points are to the tangent plane at iso-point s_k along an orthogonal direction.

By definition, a high value of f_s indicates that the region is locally smooth and the scan quality is high, i.e., the raw scan points in Ω_k form a thin plate. A low f_s value, on the other hand, suggests that

either the scan is noisy or the surface contains rich geometric details and thus more scans are needed. Figure 4 visualizes the f_s values calculated for areas with different levels of geometric details and after different numbers of scans. It shows that areas that require additional scans are properly identified.

Figure 5 confirms that the two confidence scores f_g and f_s measure reconstruction quality from different perspectives. Right after the initial blind scans, there are areas that are considered smooth enough (i.e., with high f_s values in Figure 5(d)), but incomplete since they are close to holes or open boundaries. After four scan iterations, the model is mostly complete (i.e., with high f_g values in Figure 5(f)), but the low f_s values in areas with rich geometric details still demand additional scans. The final confidence scores are calculated using the product of these two indicators,

$$f(s_k) = f_g(s_k, n_k) f_s(s_k, n_k), \tag{3}$$

which has high values only at areas that are both complete and locally smooth. As a result, our approach not only suggests NBVs based on large areas of missing data, but also on the amount of surface details to be captured in order to meet given quality criteria.

5 Next-Best-Views (NBVs)

Having analyzed the Poisson surface and assigned confidence to each of its points, our next task is to determine good locations and orientations to place the scanners for future capturing, a.k.a. the NBVs. To efficiently plan the movement of the scanning robot, we select multiple NBVs for each scan iteration and use them to determine the next motion path of the robot, avoiding unnecessary back and forth motion. The selection of NBVs are guided by a viewing vector field (VVF), where the magnitude of each vector in the field indicates how suitable the corresponding location can serve as the NBV, and the vector direction suggests a scan orientation; see Figure 6 for a 2D illustration of the NBV algorithm.

5.1 Viewing vector field (VVF) generation

All scanners have their optimal working distance range $[d_n, d_f]$ from near to far. We cannot place a scanner too close to the surface, nor too far away. Based on the furthest working distance d_f, we select a bounding box B, which contains all points in the space whose distance to the closest point on the surface is smaller or equal to d_f. In particular, the Z-direction of B is constrained to be the vertical. To determine the X-Y directions, we project the already acquired 3D points onto the horizontal plane and compute the two principle axes. The space inside the bounding box is then our searching space for the NBVs. We quantize the space B into a 3D voxel grid ($100 \times 100 \times 100$ by default). The NBV scores are only evaluated at the centers of the voxels that are not occupied by the scanned object. For each empty voxel v_i, we cast a ray through each iso-point s_k to test their visibilities.

To evaluate the suitability of a given location v_i in B for serving as a viewpoint in the next scan, we compute a NBV score $g(v_i)$ using:

$$g(v_i) = \max_{s_k} \{w(v_i, s_k)(1 - f(s_k))\}, \tag{4}$$

where $f(s_k)$ is the aforementioned confidence score for a given iso-point s_k and normalized to the range $[0, 1]$. The weight $w(v_i, s_k)$ between s_k and the candidate location v_i is computed as a product of three terms: a distance-based term w_d, an orientation-based term w_o, and a visibility-based term w_v. That is:

$$\begin{aligned}
w(v_i, s_k) &= w_d(s_k, v_i) w_o(s_k, v_i) w_v(s_k, v_i), \\
w_d(s_k, v_i) &= e^{-(\|v_i - s_k\| - d_o)^2 / (d_n/4)^2}, \\
w_o(s_k, v_i) &= e^{-(1 + \mathbf{n}_k^\top \mathbf{d}(v_i, s_k))^2 / (1 - cos(\sigma))^2}, \\
w_v(s_k, v_i) &= \begin{cases} 1 & \text{if } s_k \text{ is visible from } v_i, \\ 0 & \text{otherwise,} \end{cases}
\end{aligned}$$

where the scanner optimal working distance $d_o = (d_n + d_f)/2$. The normalized vector $\mathbf{d}(v_i, s_k) = (v_i - s_k)/\|v_i - s_k\|$ denotes the direction that points from s_k to v_i. The parameter σ is set to 25^o by default. Intuitively, the score function $g(v_i)$ has high values when there is a low confident iso-point s_k that is facing v_i, visible from v_i, and whose distance to v_i is within the optimal distance range for scanning. Once the iso-point s_k that maximizes $g(v_i)$ is found, the score value $g(v_i)$ is stored together with the direction $\mathbf{d}(v_i, s_k)$. Computing the best s_k for all the voxels in B yields the viewing vector field (VVF); see Figure 6(b) for an illustration on the magnitudes of the vectors.

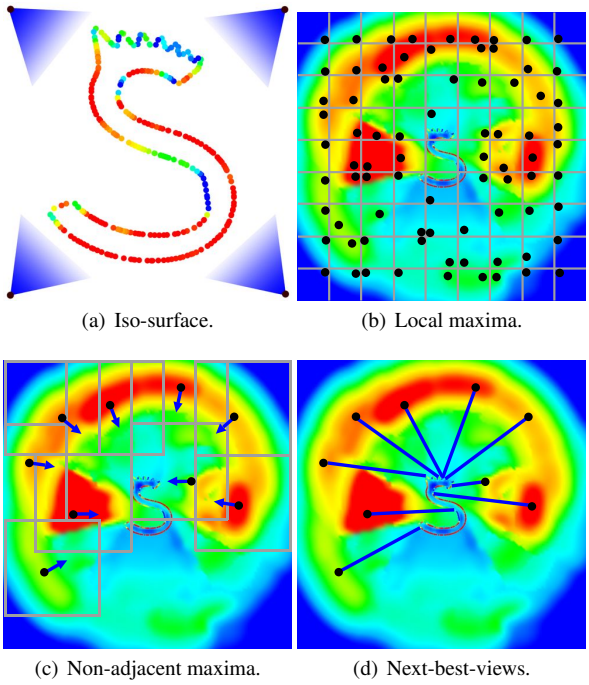

(a) Iso-surface. (b) Local maxima. (c) Non-adjacent maxima. (d) Next-best-views.

Figure 6: *A 2D illustration for our NBV selection algorithm. The input shape is blindly scanned from four locations, yielding the initial point cloud with its confidence evaluated (a). The magnitude of the corresponding viewing vector field is visualized in (b-d). The field is first partitioned into a grid of cells and the maximal magnitude location within each cell is computed (b). The non-adjacent locations with large enough magnitudes are then selected as the scanner's view positions (c). Finally the scanner's view directions are determined by optimizing the coverage of both low confident and high confident iso-points for accurate registration between the new and the existing scans (d).*

5.2 Field-guided view selection

Our next task is to select a set of view positions and directions for the next round of scanning. We start by picking locations that have

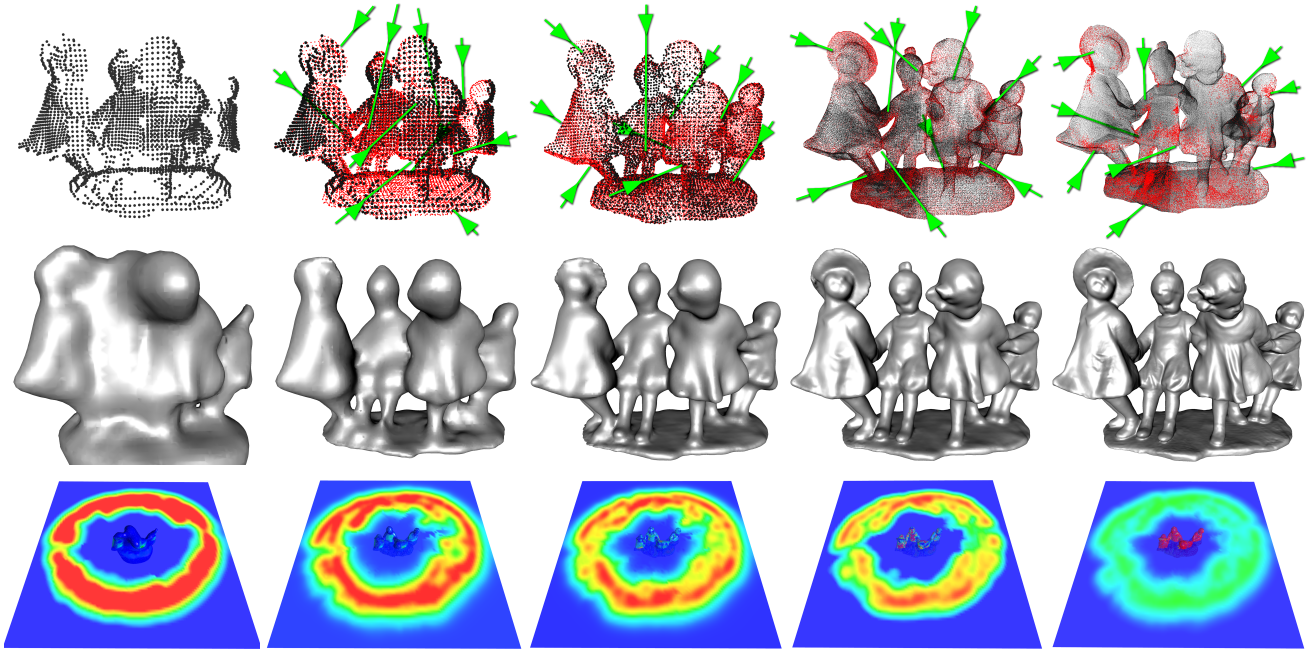

(a) Initial blind scan. (b) After 2 iterations. (c) After 3 iterations. (d) After 4 iterations. (e) After 5 iterations.

Figure 7: *For an unknown digital model, its reconstruction is progressively and rapidly enriched as more virtual scans are performed. At each scan iteration, multiple NBVs (indicated by green arrows) are computed and used for positioning the virtual scanner. The top row shows the obtained point clouds, where the red ones are points acquired during the current scan iteration. The middle row presents the Poisson reconstructed models; and the bottom row demonstrates the magnitude of the VVFs along a given slicing plane.*

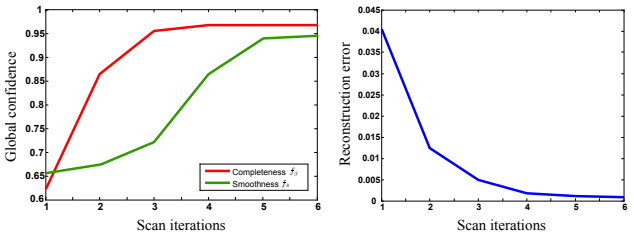

Figure 8: *Quantitative evaluation on reconstructed models shown in Figure 7 that are obtained after different scan iterations.*

local maximal magnitudes in the VVF and use them as the scanner's view positions. The field is partitioned into a grid of cells and we determine the locations with maximal vector magnitude within each cell. We then select a small subset of these local maximal vectors as the basis for optimizing the scanner's view directions. When selecting the subset, we avoid picking locations that are too close to each other and select locations only in non-adjacent cells.

To efficiently compute the subset of non-adjacent locally maximal magnitude locations, a greedy algorithm is used. With all the per-cell local maxima vectors forming a set $V = \{v_i\}$, the algorithm starts with filtering out all vectors whose magnitude $g(v_i)$ is lower than a given threshold. It then iteratively selects v_i with the largest $g(v_i)$ value from V and moves it to a second set V^*. Once a vector v_i is moved, all vectors from the cells adjacent to v_i are removed from V to prevent them being added to V^*. The process repeats until the set V is empty.

As shown in Figure 6(c), the above greedy approach ensures that the location selected each time is not in the neighborhoods of previous selections and has the highest vector magnitude among the rest. Thus, the set V^* gives us the positions for the NBVs. The vector direction $\mathbf{d}(v_i, s_k)$ at each selected location points toward a low confident iso-point s_k. Our next goal is to select an optimal view direction that allows the scanner to cover as many low confident iso-points as possible, and at the same time cover enough high confident iso-points to ensure an accurate point cloud registration between the existing and the new scans. Thus, instead of using $\mathbf{d}(v_i, s_k)$ to set the scanner's orientation directly, we select a locally optimal iso-point $s_{k'}$ in the vicinity of s_k and let the scanner point to $s_{k'}$, i.e., using direction $\mathbf{d}(v_i, s_{k'})$. The location $s_{k'}$ is computed by maximizing the following objective function:

$$\begin{aligned} &\underset{s_{k'} \in \Upsilon_k}{\operatorname{argmax}}\ \eta(s_{k'})\zeta(s_{k'}), \qquad (5)\\ &\eta(s_{k'}) = \max_{j \in \Upsilon_{k'}} \{w(v_i, s_j) f(s_j)\},\\ &\zeta(s_{k'}) = \sum_{j \in \Upsilon_{k'}} w(v_i, s_j)(1 - f(s_j)), \end{aligned}$$

where Υ_k is a set holding all the iso-points within the local neighborhood (10% of the scan range d_f by default) of s_k. Intuitively, function η returns high values when there are confident iso-points within the local neighborhood, whereas function ζ is high when there are many low confident iso-points. Maximizing the product of these two terms yields a neighborhood that contains high confident iso-points and, at the same hand, covers as many low-confident points as possible; see Figure 6(d).

5.3 On-the-fly registration and reconstruction

When the scanner can be precisely positioned at the specified NBV, the newly scanned raw points can be directly added to the existing

Operation	Time (sec)	Percentage
Robot motion/capture	463	65%
Scan consolidation/registration	178	25%
NBV computation	71	10%

Table 1: *The average time needed for a single scan iteration, which usually covers $5 - 15$ NBVs. The distribution of processing time among different operations are also listed.*

#Fig	#Pt	#Iter	#NBV	Time (min)
Figure 1	1828K	8	105	124
Figure 2	513K	3	35	41
Figure 10	406K	2	23	26
Figure 12	543K	2	25	29
Figure 13	2117K	10	127	154
Figure 14	376K	3	36	37
Figure 15	1777K	8	96	115
Figure 17	714K	4	41	49
Figure 18	4140K	15	213	235

Table 2: *The overall time took for scanning each of the physical models presented in the paper. #Pt: number of points in the final point set surface; #Iter: number of scan iterations applied; #NBV: number of NBVs selected in all iterations.*

point cloud, i.e., no registration is needed. However, in practice, misalignments do occur due to the precision of the robot movement. To compensate this misalignment, we employ a weighted ICP algorithm [Masuda 2002], where the weights are determined based on the confidence map of the iso-surface. That is, points in higher confident regions are given higher weights during the rigid registration optimization.

Once the new scan is registered with the existing ones, the newly obtained points need to be inserted into the point cloud. Note that these points not only come from areas requiring more scans, but also from areas already sufficiently scanned. If our final goal is to generate a 3D mesh using approaches such as Poisson surface reconstruction, we can simply add all these points to the cloud. However, if our goal is to generate a point set surface that faithfully depicts the object, the excess of points can be redundant. Thus, we selectively insert each new point p_i based on a probability:

$$\mathcal{P}(p_i) = (1 - f(s_k))^{\rho}, \tag{6}$$

where s_k is the closest iso-point to p_i. By definition, the lower the confidence value $f(s_k)$ is, the more likely that point p_i will be inserted into the point cloud. This allows points being added in low confident regions to increase the scan quality, while keeping highly confident regions unchanged. The parameter ρ (set to 3 by default) controls how likely the new points are concentrated at low confident regions. Figure 3(b) illustrates the effects of selective point insertion. It confirms that, in our final point cloud, the point density is much higher at areas with rich geometric details, such as wings and legs, than in smooth regions, e.g., the main body. This adaptive sampling rate is achieved as a result of the additional scans applied to obtain similar high level of confidence at feature areas.

In Figure 7, we show a progressive virtual scanning on a digital model with unknown shape based on our NBV computation. The coordinates of the model have been normalized to the range $[0, 1]$. Figure 8 further plots how the global confidence measure (the average of the confidence values on all points) and the reconstruction error (the closest distance to the original digital model) vary during the scanning process. It shows that as more scans are performed, the global confidence on the Poisson reconstructed surface is higher and the reconstruction error decreases monotonically, ensuring the convergence. In practice, we terminate the scanning once the overall confidence does not improve any more. That is, the acquisition completes when the global confidence difference between two consecutive iterations is below a threshold (0.005 by default). For example, in Figure 7, the scanning and reconstruction successfully completes after the sixth iteration.

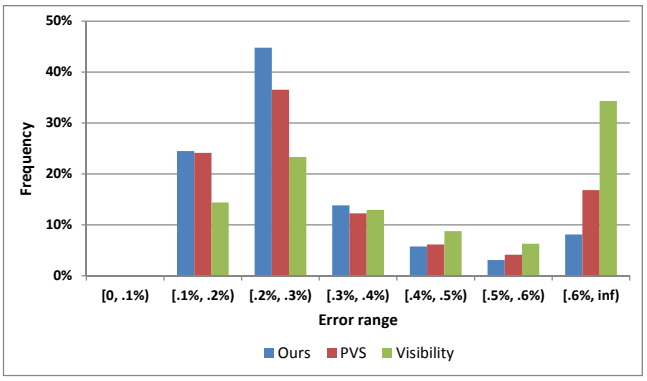

Figure 9: *Quantitative evaluation on the reconstructed surfaces shown in Figure 11. The error distributions measured using the ground truth model are plotted in histogram, which shows that a higher percentage of points has low errors using our approach than employing visibility-based and PVS approaches.*

In addition, the plot in Figure 8 also demonstrates that the gradient-based confidence measure has the largest increases after the second and the third scan iterations, whereas the smoothness-based confidence measure increases more dramatically after the fourth and the fifth iteration. This suggests that the initial scan iterations improve the completeness of the reconstructed model, whereas the later ones add geometry details. This finding is consistent with the results shown in Figure 5, where after the first 4 iterations, the gradient-based confidence measure is high across the rather complete surface, whereas some surface details are still missing and the smoothness-based confidence measure remains low.

6 Results

We have applied the presented autoscanning technique to two different types of robotic platforms. The first is a mobile robot PR2 from Willow Garage, which is equipped with two 7-DOF arm manipulators. For fast data acquisition, we mount an Artec Spider scanner to the PR2's right hand and attach a 3D printed resin table to its left hand (see Figure 1). This allows the robot to turn the object that is placed on the table, rather than having to move its arm around the object. As a result, the user only needs to specify the position and orientation of the scanner for the first scan. The acquired partial scan can be used to update the bounding box B of the object, which is then used to adjust the distance between the robot arm and the turn table to avoid collision. The second platform is a one-arm 6-DOF industry robot RV20-6 from Reis Robotics. We mount a grasping tool at the end of the arm to hold the scanner and use a motorized turntable to rotate the model; see Figure 10.

The two autoscanning platforms are tested using a number of highly detailed and self-occluding physical objects, such as the action figures in Figures 1 and 2 and the animal models in Figures 10, 12, and 13. During scanning, we utilize prior knowledge on the geometry of the robot arm and the supporting table to remove any captured points associated with them from the raw scan data. The default parameter settings specified throughout the paper are used

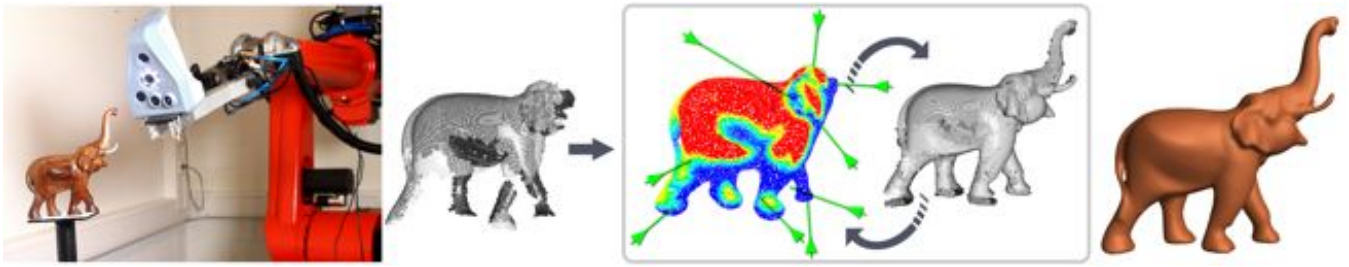

Figure 10: *An elephant model being scanned by a one-arm industry robot. The trunk is completely missing in the initial blind scan. Through iteratively scanning from the automatically selected NBVs, the final model captures all geometric details.*

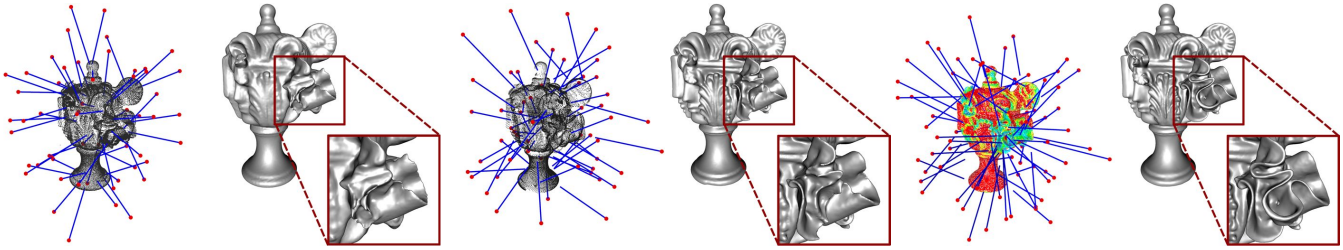

(a) Visibility (85K pts). (b) PVS (86K pts). (c) Ours (74K pts).

Figure 11: *Comparison with the visibility-based algorithm [Khalfaoui et al. 2013] (a) and the PVS approach [Kriegel et al. 2013] (b) on the virtual model shown in Figure 5(a). Both approaches converge after 55 scans, since no more barely visible or boundary points can be detected. Our approach (c) takes 67 scans under the default setting. However, through utilizing the selective point insertion strategy, the point cloud of our result contains fewer points than the previous two approaches and yet captures more details.*

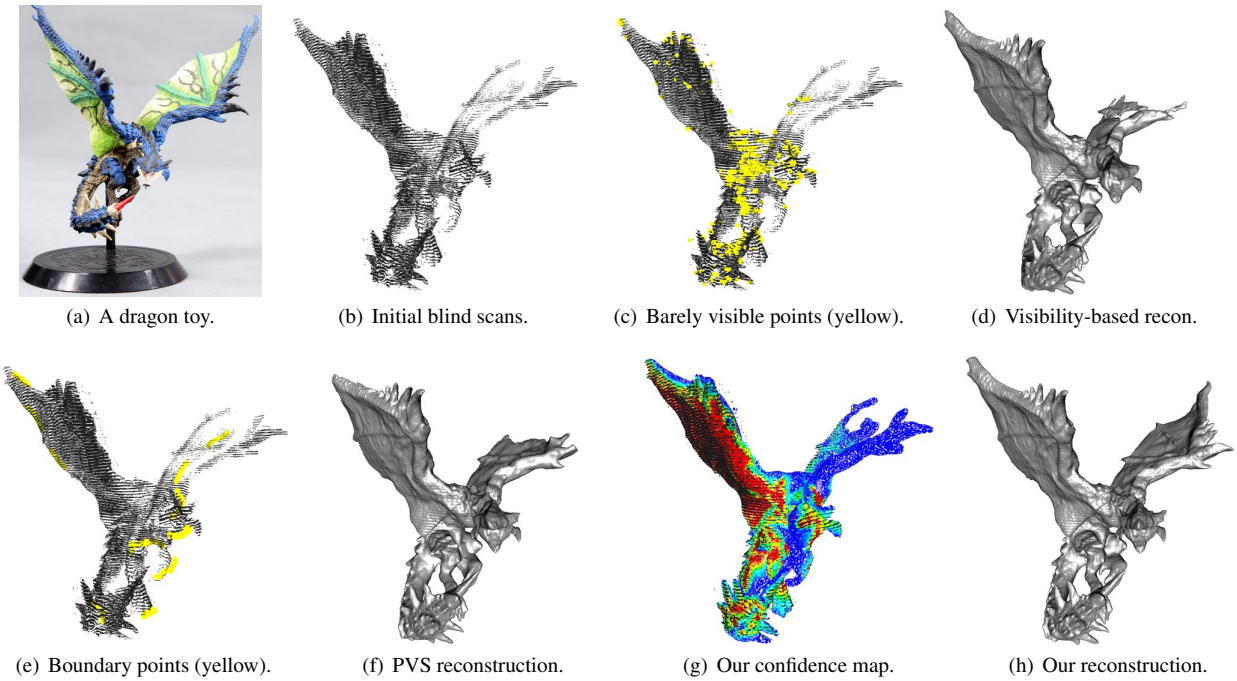

(a) A dragon toy. (b) Initial blind scans. (c) Barely visible points (yellow). (d) Visibility-based recon.

(e) Boundary points (yellow). (f) PVS reconstruction. (g) Our confidence map. (h) Our reconstruction.

Figure 12: *Comparison on a real object under the same number of scans. The initial blind scans (b) do not fully capture the input model (a). The visibility-based approach [Khalfaoui et al. 2013] selects NBVs based on barely visible points (c). The surface reconstructed using the first 25 NBVs is incomplete with the left wing missing. The PVS approach [Kriegel et al. 2013] performs better through detecting open boundaries (e), but also yields an incomplete model after using 25 NBVs (f). Our approach produces a complete model using the same number of NBVs (h).*

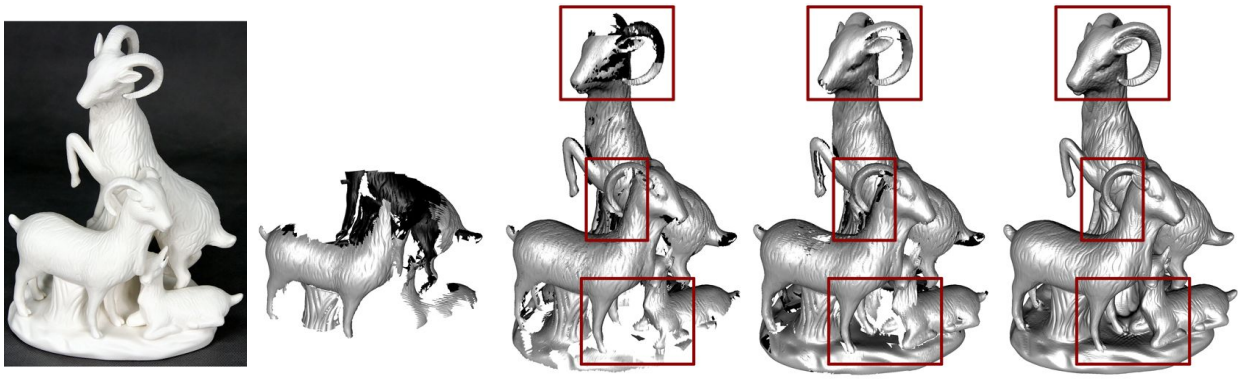

(a) A plaster figure. (b) Initial blind scans. (c) Curvature-based result. (d) Density-based result. (e) Our result.

Figure 13: *Reconstruction results under different confidence measures. For a given real object (a), the initial blind scans (b) do not cover the full object. Evaluating local surface curvature and using them to guide the NBV selection cannot produce a complete model after 20 scan iterations (c). Using local point density to guide NBV selection generates a more complete model (d) using the same number of scan iterations, but holes still exist at areas highlighted. Selecting NBVs using the presented confidence measures yields a complete model with rich surface details after only 10 scan iterations (e).*

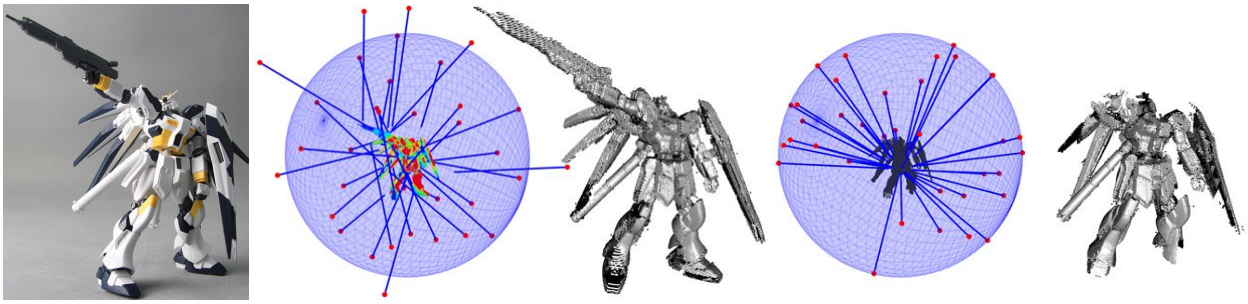

(a) A complex object. (b) Our selected NBVs with the scan result. (c) Sphere-based NBVs with the scan result.

Figure 14: *Comparison between two models obtained using the same number of scans. Scanning from the first 36 NBVs adaptively selected for the input object (a) results in a complete model (b). The raised weapon and the feet are missing from the model (c) obtained by scanning from 36 directions that are selected by a sphere-based NBV approach [Vásquez-Gómez et al. 2009].*

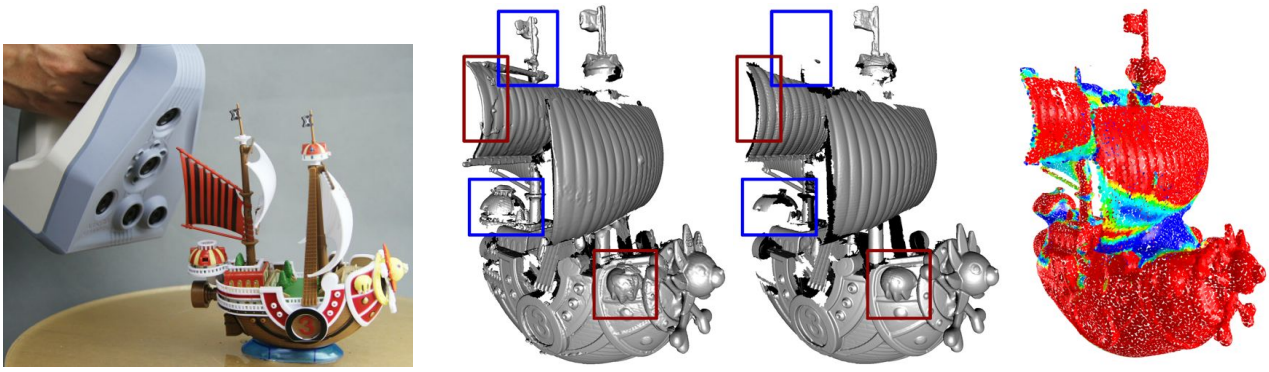

(a) A pirate boat. (b) Manual scans. (c) Autonomous scans. (d) Confidence map of (c).

Figure 15: *Comparison between models generated by manual scans and our autonomous scanner. Due to fine geometry details, complex occlusions, and challenging surface properties, an experienced user took over 200 scans and spent more than three hours to register them into a model (b). Our PR2 autoscanner achieves a comparable result (c) fully automatically with about 100 scans in less than two hours. The autoscanner missed some parts that are captured by the user (highlighted in blue boxes). However, it registers the scans more precisely, resulting an overall smoother model (e.g., areas in red boxes). It is worth noting that our approach, using the confidence map, properly identifies areas that require additional scans (d), but the model cannot be further improved due to limitations of the accessibility.*

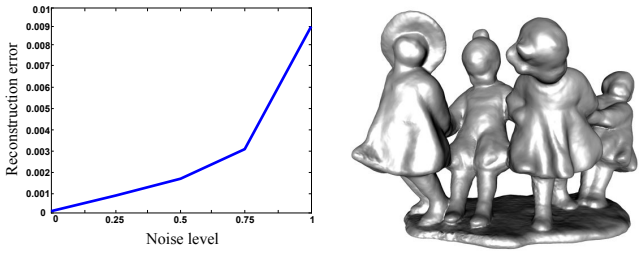

Figure 16: *Left: reconstruction errors resulted by different levels of simulated noise when scanning an unknown digital model; see also Figure 7. Right: the final model reconstructed under the scanning with strongest noise corruption.*

for acquiring all models presented. The processing time and the time distribution among different tasks are given in Tables 1 and 2.

The experiments show that blindly scanning complex objects often results in missing parts, e.g., the feet of the action figures and the trunk, wing, or horn of the animals. Through evaluating the confidence of the reconstructed iso-surface, we detect these missing areas and apply additional scans. As demonstrated, the final scanned models accurately capture the original objects.

Comparison to state-of-the-art techniques To evaluate the performance of our approach, we compare our algorithm with two recent NBV techniques [Khalfaoui et al. 2013; Kriegel et al. 2013] that represent the state-of-the-art. Kriegel et al. consider the boundaries of the surface reconstructed from the points scanned so far and select the NBVs that can best scan the boundary areas. Khalfaoui et al. select NBVs based on visibility information, which guides the scanner to cover the surface progressively. Implementing these two approaches on our robotic platform indeed allows complete models being automatically captured. However, neither approach considers the quality of the reconstructed surface in NBV calculation. As demonstrated in Figure 11, this leads to reconstructed models being much less detailed than ours. In addition, our approach selects NBVs by locating topological uncertain areas from the results of a Screened Poisson surface reconstruction. This makes it possible to obtain a complete model with fewer scans than detecting surface visibility or boundaries locally (see Figure 12), even though minimizing the number of scans is not our design goal.

To measure the fidelity of the reconstructed models, we conduct a quantitative evaluation using the input digital model as our ground truth. Here the error is measured using the closest distance between the reconstructed surface and the ground truth. The error histogram plotted in Figure 9 confirms that our reconstruction is more accurate than the competing approaches.

Comparison to more straightforward approaches To justify that our confidence measure is not unnecessarily complex, we compare it with two more straightforward approaches. The first one is density-based, where the NBVs are selected to cover areas with low point density. The second one uses importance-driven sampling, where areas with high-frequency content, detected using local curvature, attract more scans. Figure 13 shows that neither approach is as effective as the presented confidence measure. Moreover, the corresponding results do not give noticeable improvement with more scans because they keep scanning the low density or high curvature regions from less optimal viewing points. Our method can achieve better results with fewer scans instead.

Figure 14 further compares the model acquired using our selected

(a) A ceramic figure. (b) PR2 scan. (c) Texture mapping.

Figure 17: *Capturing both geometry and color information for the textured modeling and photo realistic rendering.*

NBVs with the one using a simpler alternative [Vásquez-Gómez et al. 2009]. The latter approach restricts the NBV searching space to a spheric surface (see Figure 14(c)). Such simplification helps to reduce computational costs, but at the same time, introduces limitations in handling complex, self-occluded, or feature-rich objects.

In addition, Figure 15 compares the models obtained by our autoscanner and by manual scanning. The latter was done using the Artec Studio software and with interactive feedback until an experienced user felt that further scans can no longer improve the model. It shows that our automatic approach can achieve a comparable model in less time and with much less user effort. In addition, since the autoscanner performs all scans from known NBV locations, the final registered result contains much fewer artifacts. Nevertheless, the autoscanner does not capture as much geometry as an experienced user is able to do. Our investigation shows that this is mainly due to the non-negligible physical size of the scanner when compared to the object size. As a result, placing the scanner at the selected NBVs cannot ensure that low confident areas are visible from all cameras and the structured lights on scanner, leading to imperfect depth acquisition of such areas.

Comparison to current commercial solutions There are already a few autoscanner products available in the market, such as Rexcan CS+ ($ 60,000, http://www.aniwaa.com/product/equality-tech-rexcan-cs-2-0mp), Matter and Form Scanner ($579, https://matterandform.net/scanner), MakerBotDigitizer ($800, http://www.aniwaa.com/product/makerbot-digitizer), Rubicon Rubicon 3D ($550, http://www.aniwaa.com/product/rubicon-3d), etc. These emerging techniques are inspiring, but in general their mounted scanners have limited degrees of freedom, e.g., only using sphere-based NBVs as shown in Figure 14, which make them hard to scan big objects with complex topology.

Inexpensive hardware options and noise handling Our PR2 ($400,000) can easily be replaced by a low-cost robot arm such as UR5 ($30,000, http://www.universal-robots.com), and the Artec Spider ($22,000) can exchange a cheaper scanner with the same resolution such as DAVID Structured Light Scanner SLS ($2,730, http://www.aniwaa.com/product/david-structured-light-3d-scanner). Using inexpensive hardware may cause less accurate scanner positioning and noisier scans. To see how robust our autoscanner system is, we add both random and systematic errors

(a) A church model. (b) Initial blind scans. (c) Next-best-scan of (b). (d) Next-best-scan of (c). (e) Final reconstruction.

Figure 18: *Simulating an outdoor scene scanning scenario using a scaled church model (a). The areas with missing data in the initial blind scans (b) are gradually covered in following scan iterations (e.g., c-d). The final reconstruction (e) is both complete and detail-preserving.*

simulated as in [Berger et al. 2013] into the virtual scanner, where the noise magnitude can be controlled by a user-specified parameter. We then run our NBV computation with point cloud consolidation preprocessing (Section 4.1) on the same digital model used in Figure 7, with noise magnitudes ranging from 0, i.e., no noise, to 1 that can heavily corrupt the accuracy of the virtual scanning. The resulted reconstruction errors are plotted at the left of Figure 16 and validate the robustness of our algorithm. The final model reconstructed under the strongest noise corruption is presented at the right, which still shows satisfactory quality.

More applications Figure 17 shows that both geometry and color of an object can be simultaneously captured using our autonomous scanners, making it easy to perform texture mapping onto the reconstructed surface. Nonetheless, we choose not to use color information for computing NBVs, because color information may not be accessible for all scanners and high color variation may not correlate with rich geometric details.

While the two robotic platforms used limit us to experiment with small objects, the presented algorithm is capable of suggesting NBVs for large-scale scanning as well. Figure 18 simulates an outdoor scene scanning scenario, similar to the setting of Allen's work [Blaer and Allen 2007]. It shows that a complex church model is progressively scanned and successfully reconstructed.

7 Discussion and Future work

We present an autonomous scanning technique. The key is to analyze the quality of a tentative watertight iso-surface model extracted from a Poisson field, and use the analysis to guide the selection of NBVs. Iteratively scanning and analyzing the object lead to a high quality surface reconstruction. In contrast to previous work, our method is *quality-driven* and selects NBVs through locating areas with geometrical and topological uncertainty in *a global manner*.

Please note that, due to the limitation of the scanner's accessibility (see Figure 15), we cannot guarantee a given quality requirement is satisfied everywhere on the model. This limitation in fact applies to virtual scanners as well (see Figures 7 and 11), since deep or complex cavities are hard or impossible to cover. Generally speaking, our approach is greedy, where at each iteration the scanner seeks to improve the quality of the regions with the currently lowest confidence. There could be more efficient sequences of scanning that can reach the same quality faster, with fewer scans, or shorter paths. Also our current approach assumes that the scan sensor used is isotropic. That is, we search for NBVs in 5D pose space, ignoring the effects of sensor rotation about its normal. This is one of the reasons that the scan acquired by our PR2 is not as complete as the manual scan in Figure 15. Incorporating more advanced sensor models [Gschwandtner et al. 2011] and searching in 6D pose space may further enhance the results using the proposed framework.

Humans are relatively good at high-level view planning for coverage of simple objects. However even experienced operators encounter considerable difficulties with topologically and geometrically complex shapes. Besides guiding fully automated scanning, our confidence map can also serve as visual feedback for user-operated capturing systems to facilitate and simplify the acquisition task. In the future, we would like to explore such semi-automatic interactive scanning mechanisms.

Acknowledgments

The authors would like to thank all the reviewers for their valuable comments and feedback. This work is supported in part by grants from NSFC (61232011, 61103166, 61379091), National 863 Program (2012AA011801), Shenzhen Innovation Program (CXB201104220029A, KQCX20120807104901791, ZD201111080115A, JCYJ20130401170306810, JSG-G20130624154940238), NSERC and Israel Science Foundation.

References

BANTA, J. E., WONG, L. R., DUMONT, C., AND ABIDI, M. A. 2000. A next-best-view system for autonomous 3D object reconstruction. *IEEE Trans. Systems, Man & Cybernetics 30*, 5, 589–598.

BERGER, M., LEVINE, J. A., NONATO, L. G., TAUBIN, G., AND SILVA, C. T. 2013. A benchmark for surface reconstruction. *ACM Trans. on Graphics (Proc. of SIGGRAPH) 32*, 2, 20:1–20:17.

BERGER, M., TAGLIASACCHI, A., SEVERSKY, L. M., ALLIEZ, P., LEVINE, J. A., SHARF, A., AND SILVA, C. 2014. State of the art in surface reconstruction from point clouds. *Eurographics STAR*, 165–185.

BLAER, P. S., AND ALLEN, P. K. 2007. Data acquisition and view planning for 3D modeling tasks. In *Proc. IEEE Int. Conf. on Intelligent Robots & Systems*, 417–422.

CARR, J. C., BEATSON, R. K., CHERRIE, J. B., MITCHELL, T. J., FRIGHT, W. R., MCCALLUM, B. C., AND EVANS, T. R. 2001. Reconstruction and representation of 3D objects with radial basis functions. *SIGGRAPH*, 67–76.

CHEN, S., LI, Y., AND KWOK, N. M. 2011. Active vision in robotic systems: A survey of recent developments. *Int. J. Robotic Research 30*, 1343–1377.

CHEN, J., BAUTEMBACH, D., AND IZADI, S. 2013. Scalable real-time volumetric surface reconstruction. *ACM Trans. on Graphics (Proc. of SIGGRAPH) 32*, 4, 113:1–113:16.

CONNOLLY, C. 1985. The determination of next best views. In *Proc. IEEE Int. Conf. on Robotics & Automation*, vol. 2, 432–435.

CORSINI, M., CIGNONI, P., AND SCOPIGNO, R. 2012. Efficient and flexible sampling with blue noise properties of triangular meshes. *IEEE Trans. Visualization & Computer Graphics 18*, 6, 914–924.

DAVIS, J., MARSCHNER, S. R., GARR, M., AND LEVOY, M. 2002. Filling holes in complex surfaces using volumetric diffusion. In *3DPVT*, 428–441.

DUNN, E., AND FRAHM, J.-M. 2009. Next best view planning for active model improvement. In *Proc. Conf. on British Machine Vision*, 1–11.

GSCHWANDTNER, M., KWITT, R., UHL, A., AND PREE, W. 2011. Blensor: blender sensor simulation toolbox. In *Advances in Visual Computing*. 199–208.

HARARY, G., TAL, A., AND GRINSPUN, E. 2014. Context-based coherent surface completion. *ACM Trans. on Graphics 33*, 1, 5:1–5:12.

HUANG, H., LI, D., ZHANG, H., ASCHER, U., AND COHEN-OR, D. 2009. Consolidation of unorganized point clouds for surface reconstruction. *ACM Trans. on Graphics (Proc. of SIGGRAPH Asia) 28*, 5, 176:1–176:7.

HUANG, H., WU, S., COHEN-OR, D., GONG, M., ZHANG, H., LI, G., AND CHEN, B. 2013. L_1-medial skeleton of point cloud. *ACM Trans. on Graphics (Proc. of SIGGRAPH) 32*, 4, 65:1–65:8.

KAZHDAN, M., AND HOPPE, H. 2013. Screened poisson surface reconstruction. *ACM Trans. on Graphics 32*, 1, 29:1–29:13.

KAZHDAN, M., BOLITHO, M., AND HOPPE, H. 2006. Poisson surface reconstruction. *Proc. Eurographics Symp. on Geometry Processing*, 61–70.

KHALFAOUI, S., SEULIN, R., FOUGEROLLE, Y., AND FOFI, D. 2013. An efficient method for fully automatic 3D digitization of unknown objects. *Computers in Industry 64*, 9, 1152–1160.

KIM, Y., MITRA, N., HUANG, Q., AND GUIBAS, L. J. 2013. Guided real-time scanning of indoor environments. *Computer Graphics Forum (Proc. Pacific Conf. on Computer Graphics & Applications) 32*, 7.

KRIEGEL, S., RINK, C., TIM, B., NARR, A., SUPPA, M., AND HIRZINGER, G. 2013. Efficient next-best-scan planning for autonomous 3D surface reconstruction of unknown objects. *J. Real-Time Image Processing*, 1–21.

LI, H., VOUGA, E., GUDYM, A., LUO, L., BARRON, J. T., AND GUSEV, G. 2013. 3D self-portraits. *ACM Trans. on Graphics (Proc. of SIGGRAPH Asia) 32*, 6, 187:1–187:9.

LIPMAN, Y., COHEN-OR, D., LEVIN, D., AND TAL-EZER, H. 2007. Parameterization-free projection for geometry reconstruction. *ACM Trans. on Graphics (Proc. of SIGGRAPH) 26*, 3, 22:1–22:6.

LOW, K.-L., AND LASTRA, A. 2006. An adaptive hierarchical next-best-view algorithm for 3D reconstruction of indoor scenes. In *Proc. Pacific Conf. on Computer Graphics & Applications*.

MASUDA, T. 2002. Registration and integration of multiple range images by matching signed distance fields for object shape modeling. *Computer Vision and Image Understanding 87*, 1, 51–65.

MAVER, J., AND BAJCSY, R. 1993. Occlusions as a guide for planning the next view. *IEEE Trans. Pattern Analysis & Machine Intelligence 15*, 5, 417–433.

MULLEN, P., DE GOES, F., DESBRUN, M., COHEN-STEINER, D., AND ALLIEZ, P. 2010. Signing the unsigned: Robust surface reconstruction from raw pointsets. *Computer Graphics Forum (Proc. Eurographics Symp. on Geometry Processing) 29*, 5, 1733–1741.

NEWCOMBE, R. A., DAVISON, A. J., IZADI, S., KOHLI, P., HILLIGES, O., SHOTTON, J., MOLYNEAUX, D., HODGES, S., KIM, D., AND FITZGIBBON, A. 2011. KinectFusion: Real-time dense surface mapping and tracking. In *Proc. IEEE Int. Symp. on Mixed and augmented reality*, 127–136.

PAULY, M., MITRA, N. J., AND GUIBAS, L. J. 2004. Uncertainty and variability in point cloud surface data. In *Proc. Eurographics Conf. on Point-Based Graphics*, 77–84.

PITO, R. 1996. A sensor-based solution to the next best view problem. In *Proc. IEEE Int. Conf. on Pattern Recognition*, vol. 1, 941–945.

SAGAWA, R., AND IKEUCHI, K. 2008. Hole filling of a 3D model by flipping signs of a signed distance field in adaptive resolution. *IEEE Trans. Pattern Analysis & Machine Intelligence 30*, 4, 686–699.

SCOTT, W. R., ROTH, G., AND RIVEST, J.-F. 2003. View planning for automated three-dimensional object reconstruction and inspection. *ACM Computing Surveys 35*, 1, 64–96.

SEVERSKY, L. M., AND YIN, L. 2012. A global parity measure for incomplete point cloud data. In *Computer Graphics Forum (Proc. Pacific Conf. on Computer Graphics & Applications)*, vol. 31, 2097–2106.

SHARF, A., ALEXA, M., AND COHEN-OR, D. 2004. Context-based surface completion. *ACM Trans. on Graphics (Proc. of SIGGRAPH) 23*, 3, 878–887.

TAGLIASACCHI, A., ZHANG, H., AND COHEN-OR, D. 2009. Curve skeleton extraction from incomplete point cloud. *ACM Trans. on Graphics (Proc. of SIGGRAPH) 28*, 3, 71:1–71:9.

TRUCCO, E., UMASUTHAN, M., WALLACE, A., AND ANDROBERTO, V. 1997. Model-based planning of optimal sensor placements for inspection. In *Proc. IEEE Int. Conf. on Robotics & Automation*, vol. 13, 182–194.

VÁSQUEZ-GÓMEZ, J. I., LÓPEZ-DAMIAN, E., AND SUCAR, L. E. 2009. View planning for 3D object reconstruction. In *Proc. IEEE Int. Conf. on Intelligent Robots & Systems*, 4015–4020.

WENHARDT, S., DEUTSCH, B., ANGELOPOULOU, E., AND NIEMANN, H. 2007. Active visual object reconstruction using D-, E-, and T-optimal next best views. In *Proc. IEEE Conf. on Computer Vision & Pattern Recognition*, 1–7.

WHEELER, M. D., SATO, Y., AND IKEUCHI, K. 1998. Consensus surfaces for modeling 3D objects from multiple range images. In *Proc. Int. Conf. on Computer Vision*, 917–924.

ZHOU, Q.-Y., AND KOLTUN, V. 2013. Dense scene reconstruction with points of interest. *ACM Trans. on Graphics (Proc. of SIGGRAPH) 32*, 4, 112:1–112:8.