THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Three-Dimensional Surface Relief Completion Via Nonparametric Techniques

Citation for published version:

Breckon, TP & Fisher, RB 2008, 'Three-Dimensional Surface Relief Completion Via Nonparametric Techniques', *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 30, no. 12, pp. 2249-2255. <https://doi.org/10.1109/TPAMI.2008.153>

Digital Object Identifier (DOI):

[10.1109/TPAMI.2008.153](https://doi.org/10.1109/TPAMI.2008.153)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Early version, also known as pre-print

Published In:

IEEE Transactions on Pattern Analysis and Machine Intelligence

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Three-Dimensional Surface Relief Completion via Nonparametric Techniques

Toby P. Breckon and Robert B. Fisher

Abstract—Common 3D acquisition techniques, such as laser scanning and stereo capture, are realistically only 2.5D in nature. Here, we consider the automated completion of hidden or missing portions in 3D scenes originally acquired from 2.5D (or 3D) capture. We propose an approach based on the nonparametric propagation of available scene knowledge from the known (visible) scene areas to these unknown (invisible) 3D regions in conjunction with an initial underlying geometric surface completion.

Index Terms—Image processing, occlusion, range data, surface fitting, texture.

1 THE PROBLEM OF COMPLETION

THE general problem of visual completion has been widely studied both in terms of human abilities in perceptual psychology and artificial computer vision approaches [1]. Here, we concentrate specifically upon the problem of 3D visual completion as it pertains to the 2.5D limitation of current 3D sensing technologies. With such capture technology, it is impossible to capture all faces of a 3D scene from a single unidirectional viewpoint (i.e., a single laser range scan or stereo image). This results in models which are inherently 2.5D in nature or 3D models which are constructed from a combination of multiple unidirectional scans or images. In the latter case, the additional capture and subsequent viewpoint combination incurs additional cost in terms of time and computation (e.g., [2]).

The same sensing limitation is also true of human stereopsis, where visual completion, as an aspect of our extending visual reasoning system, allows us to perceive how an object may appear from an alternative viewpoint (e.g., an opposing backward view) [1]. Ideally, a similar method of visual completion is required to facilitate the completion of a 3D model, from a single 2.5D or limited 3D capture, akin to that present in human visual reasoning. Based on our prior work in the examination of human capabilities in this area, we propose an approach based on the paradigm of completion via *visual propagation*—completion via the propagation of knowledge from known to unknown scene portions.

In our work, such propagation is performed on two levels: geometric surface continuation and nonparametric relief synthesis. Our approach uses a combination of global surface fitting to derive an initial underlying geometric surface completion, together with a 3D extension of nonparametric texture synthesis in order to provide the propagation of localized structural 3D surface detail. We define surface detail as the *3D surface relief*—the unique tactile nature of a 3D surface characterized by localized displacement over its area (e.g., tree bark, brickwork, or sandpaper). It is this fine geometric surface detail that we desire to complete in order to achieve plausible 3D completion.

- T.P. Breckon is with the Applied Mathematics and Computing Group, School of Engineering, Cranfield University, Whittle Building, Cranfield, Bedfordshire MK43 0AL, UK. E-mail: toby.breckon@cranfield.ac.uk.
- R.B. Fisher is with the School of Informatics, University of Edinburgh, Informatics Forum, 10 Crichton St., Edinburgh EH8 9AB, UK. E-mail: rbf@inf.ed.ac.uk.

Manuscript received 1 June 2007; revised 8 Jan. 2008; accepted 2 June 2008; published online 4 June 2008.

Recommended for acceptance by L. Van Gool.

For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number TPAMI-2007-06-0322.

Digital Object Identifier no. 10.1109/TPAMI.2008.153.

In contrast, the majority of prior work in this area has primarily concentrated on the concept of “*good continuation*” of existing smooth surfaces rather than paying attention to any relief detail present [2], [3], [4], [5], [6]. A variety of methodologies have been proposed for this problem, including localized geometric and algebraic surface fitting [7], [8], [9], [10], volumetric distance-field-based techniques [2], [11], [12], cross-triangulation patching and refinement [3], [5], [6], and spatial occupancy-based approaches [13], [4]. While these techniques concentrate on the completion of underlying surface geometry, limited attention has been paid to additionally propagating surface detail—thus limiting them to a subset of real-world surfaces [1]. We propose a novel method for the “*completion by example*” of 3D surface relief from a limited 2.5D surface.

Prior work on the propagation of surface detail is limited and, in general, pursues an alternative methodology to our own [14], [15]. The work of Sharf et al. [14] pursues a patch-based “*copy and paste*” completion approach that, while well suited to smooth surface or irregular/anisotropic relief completion, relies on nonrigid alignment (warping) that limits its application to structured relief (as found in architecture, Fig. 9). Additionally, the brittleness of the patch-wise “*copy and paste*” approach prevents the deviation of suitable completions through the combination of multiple source regions on the original surface. This leads to uncharacteristic “*tiling*” artifacts over large areas and an inability to adaptively complete in scenarios where no suitable propagatable patch exists [14]. In [15], “*example-based*” completion is performed with reference to a database of similar a priori complete 3D objects from which similar surfaces are selected and blended to perform completion on the partial 2.5D surface. This is clearly limited to the availability of such suitable a priori knowledge and, by its use of similarity matching, relates more to the concept of *completion via recognition* than the *completion via generalization* we aim for here. Also of note is the related works in [16] and [14], similar, respectively, to [15] and [17] but varying primarily in surface representation, and graphics work on geometric texturing [18], [19]. Both [18] and [19] are so far limited to geometric texturing, using volumetric and geometry image representations, respectively, rather than explicit relief completion on surface meshes.

For a full overview of prior computer vision work in 3D completion, with reference to related theories from perceptual psychology, the reader is directed to [1].

2 COMPLETING 3D SURFACE RELIEF

Here, we pursue the completion of plausible 3D surface relief through a two-stage process: 1) the geometric completion of the smooth underlying surface model [7], [8], [9] and 2) the propagation of localized surface relief, from the existing surface, over this newly formed completion [20]. As an example, we complete both the geometric sphere and surface dimples of a 2.5D golf ball as shown in Fig. 1. Here, we see the successful plausible completion of the surface relief pattern (i.e., Stage 2, Figs. 1c and 1d) over a geometric completion (i.e., Stage 1, Fig. 1b) of the original 2.5D capture (Fig. 1a).

2.1 Geometric Completion

Our initial stage of geometric completion follows the prior work in [7], [9], [8] to produce a smooth surface completion that conforms to the geometric surface model that underlies the available 2.5D surface data P_{original} . The approach in [7], [9] is used to facilitate the automatic best-fit model selection to P_{original} from the set $\{\text{plane}|\text{circular cylinder}|\text{sphere}\}$ and to provide an accurate surface parameterization for completion via good geometric continuation [7], [9], [8] (e.g., Figs. 1b and 2a). These geometric surface extensions are then reproduced as complementary surface data points $P_{\text{completion}}$ at sample density d , where a suitable d is

Fig. 1. Completion of a 2.5D golf ball: (a) 2.5D original, (b) 3D geometric completion, (c) 3D relief completion, and (d) 3D original and completion.

derived from the surface triangulation properties as the minimum surface edge length present in the original surface [21].

For more complex surface structures, prior segmentation allows the separation of these surfaces into simple subparts prior to completion. As an alternative, the use of a more powerful initial smooth surface completion technique (e.g., [2], [3], [5], [4], [6]) would allow the application of our novel surface relief propagation technique (stage 2) to a wider class of surfaces beyond this geometrical assumption (e.g., spline surfaces).

2.2 Surface Relief Propagation

We discuss the second stage of our technique in two subparts. First is a background overview of nonparametric sampling, followed by details of adaptation to the completion of 3D surface relief.

2.2.1 Two-Dimensional Nonparametric Sampling

Nonparametric sampling was proposed as a method for texture synthesis in 2D images based on using a statistical nonparametric model and an assumption of spatial locality (Efros and Leung [20]). Unlike other parametric approaches in the texture synthesis arena which attempt to explicitly model the texture prior to synthesis, this approach samples directly from the texture sample itself—a kind of implicit modeling akin to the robotics paradigm “the world is its own best model.” As a result, it is “very powerful at capturing statistical processes for which a good model has not been found,” outperforming contemporary approaches [20], and, thus, highly suited to our work in 3D.

In 2D images, nonparametric sampling is very simple—it successively grows a texture outward from an initial seed area, one pixel at a time, based on finding the pixel neighborhood in the sample image that best matches that of the current target pixel (i.e., the one being synthesized). In operation, it exhaustively evaluates every possible sample position, i.e., pixel neighborhood, in the available sample texture to determine the best match to that surrounding the unknown target. It then copies the central pixel value from this best matching neighborhood as the new texture value at the target.

These neighborhoods are defined as $w \times w$ square windows around each pixel, where w , the window size, is a parameter perceptually linked to the scale of the largest regular feature present in the texture [20]. Neighborhood matching is then based upon using the normalized sum of squared difference metric (SSD) between two pixel neighborhoods. From the set of all sample neighborhoods, the top η percent of matches are selected as those with the lowest SSD values (i.e., those with the minimal matching difference). From this reduced set, one is then randomly selected to provide the value at the target. As an additional constraint, the

Fig. 2. “Smooth” surface completion and displacement vectors. (a) Smooth surface completion. (b) Displacement vectors for the sample set. (c) Combined target and sample triangulation.

randomly selected match is only used to fill the target provided it has a normalized SSD value less than a specified error threshold ϵ related to the acceptable level of noise in the synthesized texture.

Overall, the technique [20] is highly successful and was inspirational to numerous subsequent nonparametric approaches to this problem. Here, despite a wealth of progression in the literature from the original work, we take the seminal approach of Efros and Leung [20] as the base case for the first extension of 2D texture synthesis to 3D surface relief completion.

2.2.2 Nonparametric 3D Completion

The basic aspects of nonparametric sampling map well from 2D to 3D: The 2D image becomes a 3D surface, the individual pixel becomes a point on that surface, a pixel neighborhood becomes the set of nearest neighbors to a surface point, and the actual pixel values being synthesized become displacement vectors mapping discrete points on a textured surface (i.e., the relief) to the geometric surface derived from prior fitting.

The main input to our nonparametric completion process is thus the geometrically completed version of the 3D surface $P_{\text{original}} \cup P_{\text{completion}}$. The originals P_{original} , labeled as textured,¹ are the sample points $s \in \text{samples}$, while those forming the completed “smooth” portion $P_{\text{completion}}$, labeled untextured, form the target points $t \in \text{targets}$ (Fig. 2a). In our subsequent discussion, P_{original} becomes our set of *samples*, with 3D surface relief, and $P_{\text{completion}}$ becomes our set of *targets*, requiring the completion of 3D surface relief that both belong to a common triangulated surface $\text{triangulation}(P_{\text{original}} \cup P_{\text{completion}})$ (Fig. 2c). Each point has an associated surface normal n and an associated relief displacement vector $\vec{D}(s)$ derived from earlier geometric surface fitting (Fig. 2b).

Algorithm outline. The nonparametric algorithm adapts to 3D by considering vertex neighborhoods on the 3D surface in place of the pixel neighborhoods of [20]. Each vertex neighborhood $N(i)$ is the set of vertices lying within a radius of w_i edge connections from the target vertex being textured (Fig. 3), where w_i forms the window size parameter synonymous with that of the earlier 2D approach. Here, however, this parameter is assigned on a per-vertex basis (i.e., w_i for vertex i) in order to facilitate the reduction of the vertex neighborhood on an individual per-vertex basis when no suitable match can be found, for a given target, at the current neighborhood level. Initially, all window size parameters are initialized to a globally specified window size w and only adjusted automatically as later required. The algorithm now proceeds by

1. Here, and in subsequent discussions, the terms textured/untextured are used to describe the presence/absence of 3D surface relief.

 Fig. 3. 3D vertex neighborhoods $N(i)$.

finding the best sample region matching the textured portion of a target vertex's neighborhood.

First, the set of target vertices currently lying on the textured/untextured surface boundary is identified as the current target list L . The first target vertex $t \in L$ is then matched, using neighborhood-based matching, against every available vertex $s \in \text{samples}$. Following the traditional route of Efros/Leung, a match is then randomly selected from the best η percent of this set, based upon the matching score (here, $\eta = 10$). Provided the matching score for the selection is below the specified acceptable error threshold parameter e , this choice is accepted and the current target vertex t is textured by mapping the displacement vector $\vec{D}(s)$ from the selected sample vertex s to t . The current target t is now labeled as textured and the algorithm proceeds to the next vertex in L . If the match is not accepted (or no match was possible), the vertex is simply skipped and returned to the pool of target vertices for future synthesis. In this specific case, the window size w_t associated individually with vertex t for future matching is reduced in size, $w_t = w_t - 1$, to facilitate matching on a scale of reduced constraint, global $\rightarrow$ local, where required (until a limit of $w_t = 1$).

Once L is exhausted, the next set of target vertices on the textured/untextured boundary is identified, based on the updated vertex labeling and the process is continued until all $t \in \text{targets}$ are labeled as textured. To ensure target vertices are processed in the order of most to least constrained, L is always sorted by decreasing number of textured neighbors prior to processing [20]. In practice, L is first randomly shuffled prior to sorting to ensure variation in the ordering of vertices with equally textured neighborhoods.

Throughout the process, progress is monitored over each target list L constructed. If no match selections are accepted over an entire iteration of L , the algorithm would reach an impasse due to the constraining value of e . To avoid this problem, the acceptable error threshold e is raised slightly (10 percent) in this occurrence to relax the acceptable error constraint (as per [20]) and thus allow relief synthesis to hopefully progress over the next iteration of L .

Overall, our use of exhaustive neighborhood-base for each target $t \in \text{targets}$ at each position on the sample surface region $s \in \text{samples}$ results in a completion algorithm that is $\mathcal{O}(stw^2)$ for neighborhood window size w [21].

Neighborhood orientation. The remaining key element in our algorithm outline is the matching of textured target neighborhoods (Fig. 3) to vertices in the sample region. In order to perform matching at different positions on the surface, with varying localized surface orientation, we require knowledge of local surface orientation frames $[n, u, v]$ (Figs. 4 and 5), with which to consistently align different surface neighborhoods.

Matching is performed using an adaptation of the pixel-wise SSD [20] based on the projection of neighborhood vertices onto the surface at each sample point. To compute this match between target vertex t , with textured neighborhood vertices $Nt(t)$ and a sample vertex s with textured neighborhood $Nt(s)$, $Nt(t)$ must be first transformed rigidly into the coordinate system of s . This is based on the local reference frames at s and t , denoted R_s and R_t , respectively, which, combined with the positional translations

 Fig. 4. Vertex displacement vector $\vec{D}(s)$.

Fig. 5. Consistent surface orientation ("L" pattern shows U/V axis directions). (a) Original reference frames (inconsistencies). (b) After iterative reorientation. (c) Iterative reorientation on other surfaces.

given by t and s , facilitate the transformation of $Nt(t)$ relative to s as $Nt(t)'$. However, as t is itself untextured while s is textured, the natural misalignment (due to the presence/absence of texture) has to be avoided by transforming to the corresponding untextured position of s on the underlying surface (s'), calculated using the known displacement vector at s , $\vec{D}(s)$, as $s' = s - \vec{D}(s)$ (Fig. 4). Overall, we have a resulting, $t \rightarrow s'$, homogeneous coordinate transformation as follows:

$$Nt(t)' = \begin{bmatrix} [R_s] & s' \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} [R_t] & t \\ 0 & 0 & 0 & 1 \end{bmatrix}^{-1} Nt(t). \quad (1)$$

Equation (1) first transforms the target neighborhood $N(t)$ into a neutral frame (from its orientation R_t at position t), which is then transformed into the orientation of the sample frame R_s at adjusted position s' . In order to estimate this spatial transformation, the reference frames R_s and R_t are required. These are derived deterministically based on finding mutually perpendicular vectors $\vec{u}$ $\vec{v}$ to the surface normal $\vec{n}$ by standard algebraic manipulation.

For the geometric surfaces considered here, this method ensures at least localized consistency, but some global irregularities still exist (e.g., Fig. 5a). These problems of global inconsistency can be solved by either initially reorientating the localized reference frames or augmenting the algorithm to match the target neighborhood at R different rotational orientations around the normal axis at each vertex.

In the former, each reference frame is reorientated to minimize the rotational transform to either 1) its neighbor in an iterative sweep process starting from an initial seed or 2) a global frame of reference. While 2) is applicable to surfaces with orientable descriptors (e.g., cylinders), it may suffer from singularities on

Fig. 6. Point matching via surface projection. (a) Nearest point matching space gives misleading distance-based match metrics. (b) Matching based on common surface projection gives more realistic distance-based match metrics.

closed surfaces such as spheres—in these instances, 1 is more applicable. The alternative, at the expense of computation, is to simply match the target neighborhood to every sample region at R different rotational orientations around the normal axis, where additional parameter R specifies the divisions of 2π giving a set of rotations (e.g., $R = 4$ gives four orientations at $0, \frac{\pi}{2}, \pi$, and $\frac{3\pi}{2}$).

Here, we use a combination of initial reorientation using either an iterative localized neighbor-to-neighbor sweep approach based on [22] or uniform orientation based on global reference vector on a per-surface basis as required. Fig. 5b shows the use of the iterative reorientation technique [22] to remove the global orientation inconsistencies previously highlighted in Fig. 5a. Fig. 5c shows its more general application to the consistent orientation of more complex surface examples.

Neighborhood matching. Having now aligned our target neighborhood at s , the task now is to compute the SSD as a vertex matching problem between this transformed neighborhood $Nt(t)'$ and the textured surface vertices at s . Although this seems to be a simple 3D point matching problem, the presence of sampled surface texture means that simple euclidean space “nearest point” matching using the raw textured vertices can produce artificial matches in common scenarios, as shown in Fig. 6a. Although such problems could be overcome by enforcing a scheme of one-to-one minimal distance cross-matching between the sets, this relies on the assumption that the densities of both point sets are equal, which cannot be assured [21].

To ensure consistent vertex matching, independent of relative density, we match vertices $v_1 \rightarrow v_2$, $v_1 \in Nt(t)'$, and $v_2 \in Nt(s)$ based on their relative projected positions on the common geometric surface model, embodied in the displacement vector associated with every vertex $v'_i = v_i - \vec{D}(v_i)$ (Fig. 4). This effectively matches vertices based solely on their relative spatial position on the geometric surface rather than the relative textured-related depth, as shown in Fig. 6b. From these pairings in surface projected space $v'_1 \rightarrow v'_2$, the SSD is calculated based on the original vertex positions $v_1 \rightarrow v_2$ at s and s' .

It should also be noted that, here, we are *not* performing a neighborhood $Nt(t)'$ to closed neighborhood $Nt(s)$ match. $Nt(t)'$ is matched against the unrestricted set of textured vertices $N(s) = (i \in P_{original} \cup P_{completion} | label(i) = textured)$, with a viable match only being considered when all matching partners v_2 of $v_1 \in Nt(t)'$ are themselves also textured. The SSD is calculated based on the distance of each target vertex $v_1 \in Nt(t)'$ directly to the triangulated surface (not just the closest vertex)—i.e., the minimum squared distance to any surface triangle Δ_j that has v_2 as a vertex $\Delta_j \in triangles(v_2)$:

Fig. 7. Synthetic surface results.

$$SSD_{shape} = \sum_{v_1}^{Nt(t)'} g_{v_1} \min_{\Delta_j \in triangles(v_2)} \left(dist(v_1, \Delta_j)^2 \right). \quad (2)$$

Additionally, as in [20], a weight g_{v_1} , based on a 2D Gaussian kernel, is used in (2) to weight the SSD vertex matches $v_1 \rightarrow v_2$ relative to the distance $t \rightarrow v_1$, $v_1 \in N(t)$ (i.e., spatial proximity to t).

Parameter selection. Here, we formalize and briefly discuss the setting of the parameters detailed in the proposed approach:

- w : represents the neighborhood size of target vertex t forming the vertex neighborhood $N(t)$, as all vertices within radius w edge connections from t . As per [20], w is set in relation to the largest repetitive artifact of the sample region.
- e : represents an upper bound on the acceptable level of matching error in terms of SSD. The setting of e is purely related to the acceptable level of localized error.
- η : controls the size of the selection set (percentage, η percent) from which a sample vertex is then randomly selected. Although varying this parameter may be useful for the completion of natural surface relief, empirical results show that stochastic and manmade surfaces can equally be handled with a constant $\eta = 10$ value.
- R : represents the number of rotational positions around the normal axis, in the form $i \frac{2\pi}{R}$ for $i = \{0 \rightarrow R-1\}$, that the transformed target neighborhood $N(t)'$ is matched at. The setting is dependent on the requirement for isotropic relief completion and the presence of global inconsistencies in the surface orientation field, as previously discussed. By default, $R = 0$ for no rotations.

Although some automated parameter selection maybe possible, this is left as an area for future work.

3 RESULTS

Our evaluation is based upon visual comparison of the original and completed surface portions together with the statistical comparison of the Mean Surface Integral (MSI) between the relief surface and the underlying geometric surface fit on the original/completed surface portions (3):

$$MSI = \frac{1}{n} \sum_{i=0}^n \vec{D}(i) \cdot \vec{n}(i). \quad (3)$$

The MSI (3) is defined for n vertices, where $\vec{D}(i)$ is the surface relief displacement vector (original or propagated) and $\vec{n}(i)$ is the surface normal at vertex i (Fig. 4). The primary means of comparison is the percentage difference between this measure for the original and completed surface portions.

TABLE 1
Statistical Comparison of 2.5D Surface Completions

Surface Description	Figure	Parameters	Original (MSI)	Completion (MSI)	% Difference
Cylinder (2 bumps)	7A	$w=2; e=0.1$	0.186269	0.187597	0.71%
Cylinder (wave)	7B	$w=2; e=0.1$	0.247124	0.251569	1.80%
Plane (wave)	7C	$w=2; e=0.1$	1.191870	1.196030	0.35%
Golf ball	1	$w=7; e=0.01$	0.065876	0.064562	-1.99%
Pisa Tower (full)	9	$w=7; e=0.1$	0.923983	1.012520	9.58%
Tree bark (cylinder)	11	$w=3; e=0.2; R=2$	1.188590	1.365380	14.87%
Occluded Bumps	8	$w=7; e=0.1$	3.444780	4.152700	20.55%
Pisa Tower hole	12	$w=9; e=0.1$	1.30763	1.11035	-15.09%
Golf ball hole	13	$w=5; e=0.2$	0.06675	0.06137	-8.07%

3.1 General Results

Figs. 7A and 7B show the successful completion synthetic relief patterns over cylindrical surfaces, while Fig. 7C shows 3D relief completed over a larger planar area. Each case shows the completion of an initial surface patch (light) over a geometric completion (dark) of the underlying surface model. A statistical comparison of these results (Table 1) shows small statistical differences that correlate with the visual comparison of Fig. 7. These can be attributed to uniform undersampling of the original surface and any resulting aliasing effect.

Surface completions based on real 2.5D surface captures are now presented in Figs. 1, 8, 9, 10, and 11. First, we consider our prior example of a 2.5D golf ball (Fig. 1). Here, local surface orientation was derived using the iterative reorientation process (Fig. 5). Fig. 1 clearly shows successful regular tiling of equi-sized dimples over the geometric sphere completion (Fig. 1) and is further supported by the statistical results in Table 1.

In contrast to this isotropic texture (golf ball, Fig. 1), we see the completion of regular, anisotropic architectural features of Pisa Tower (Fig. 9). Here, we see the successful completion of regular windows, doors, struts (bottom), and building-specific architectural features (top) from an initial 2.5D scan of this scale model. Fig. 9 (top right) shows no obvious visual difference between the completion and original. However, as shown in Fig. 10, we see that, on closer examination, subtle anomalies exist with the

completed surface portion—repeated structure and mismatches on highly constrained surface joins. These are attributable to the effects of noise on the process that originate from the original surface capture and the quality of the underlying geometric surface fit. Noise in the original surface capture is difficult to completely eradicate without affecting fine surface detail and the quality of a geometric surface fit difficult to quantify against only partial (i.e., 2.5D) surface data that, due to its relief, will vary considerably from any underlying surface model. Both of these issues lead to accumulated error problems (Fig. 10) similar to the problem originally encountered in [20] with the loss of high-order information. It is hoped that these limitations can be addressed through the introduction of hierarchical surface techniques in future work. From the statistics of Table 1, we see a ~ 10 percent difference over the completion of the tower. This variation is to be expected on such a surface where clear variations exist within the surface relief itself (e.g., contrast of top, middle, and bottom sections), but, overall, the statistics again show differences not apparent to visual inspection.

Fig. 11 shows the completion of anisotropic, irregular tree bark² texture over a geometrically completed cylinder. Here, we see the successful completion of the bark structure despite the highly stochastic nature of the initial sample relief over the bark “structure” itself. In this case, the best results, both visually and statistically, were achieved by increasing the acceptable error bound e to 0.2 and additionally specifying a rotational matching parameter of 2 (default = 0) so that each sample match is attempted at rotational positions 0 and $\frac{\pi}{2}$ around the surface normal. This $\frac{\pi}{2}$ rotational matching selection reflects the linear nature of the underlying structure in the bark surface relief—each match is attempted in its original orientation and the corresponding reflection around this linear axis. The statistical comparison of Fig. 11 shows a difference of ~ 15 percent between the original and completed surface portions (Table 1). Such variation is to be expected in stochastic surface relief of this nature and is, as with the previous examples, not visually detectable.

Finally, we show a classical example of 2.5D occlusion resolution in Fig. 8, where we see the successful completion of surface relief (Fig. 8 (lower)) that was occluded during the original capture (Fig. 8 (upper)). Here again, despite a large statistical difference (~ 20 percent, Table 1), we see the successful completion of this regular, isotropic surface relief that contains both high-level (bumps) and lower-level (surface noise) detail. As Fig. 8 shows, the occlusion resolution abilities of this technique contrast sharply with the earlier work of [7], [9].

Fig. 8. Completion of occluded bump surface.

2. N.B. Linear bark structure running parallel to cylindrical surface axis.

Fig. 9. Completion of Tower of Pisa.

Figs. 11 and 8 both use surface orientation based on a global reference vector (surface normal/axis) with a general error bound of 0.1 based on noise present in data acquisition. This bound was tightened in specific cases, such as the highly constrained golf ball surface (Fig. 1, $e = 0.01$).

Despite extensive precomputation and memorization, this technique remains computationally very expensive. In the simple synthetic completion cases of Fig. 7, average runtime for each surface completion example was 23.15 seconds for the cylindrical cases (720 samples on average, 561 targets) and 217.9 seconds (3.6 minutes) for the planar cases (1,681 samples, 1,408 targets). It is recognized that the computational demand does present some limitations in its practical usage, as it was in the original 2D work

of Efros and Leung [20], but a number of options exist. Notably, the original sample set s can be subsampled as required or, alternatively, the technique lends itself well to explicit parallelism. Most importantly, *completion is a task that is performed once for a given surface or object rather than repetitively.*

3.2 Comparison to Ground Truth

We additionally perform a comparison of our technique against hypothetical scenarios where we have original (i.e., ground truth knowledge) for the unknown scene portion.

First, we consider the completion of a large region of the upper part of the Pisa Tower surface (Fig. 12). The visual comparison of Fig. 12 shows a very high level of visual similarity between the resulting surface completion and the ground truth surface data with the correct completion of vertical/horizontal structures and window indentations. Statistical comparison, however, reveals a ~ 15 percent difference (Table 1) between completion and ground

Fig. 10. Pisa Tower errors. (a) Structure repetition not present in original caused by noise. (b) Accumulated error due to noise causes slight mismatches at surface joins.

Fig. 11. Tree bark completion.

Fig. 12. Pisa Tower ground truth comparison.

truth—a difference not immediately apparent to the viewer. The completions produced are similar enough to the original surface to make them appear plausible, but differ enough for us to tell they are different surfaces. We find, as we would expect, that we can produce a plausible “approximation” of the missing surface data, derived from similarly available structure, but we cannot make an accurate “copy” of it.

In our other example, the golf ball (Fig. 13), we similarly see a completion result with a similar statistical difference against the ground truth (~ 8 percent, Table 1) to that achieved for overall 2.5D to 3D completion (~ 2 percent, Table 1). However, when examined in detail (Fig. 13, left bottom) and compared in terms of the color-coded geometric displacement map (Fig. 13, center/left bottom), the effects of noise on the completion are apparent. Overall, despite the presence of noise in the fine surface detail, the overall completion matches well, visually and statistically, with the ground truth data.

Here, in a slightly more constrained scenario, the best results were achieved by raising the error threshold to $e = 0.2$ and using a smaller window size parameter $w = 5$ than in the earlier full 2.5D to 3D completion.

4 SUMMARY

Overall, we propose a two-stage surface completion technique—geometric surface completion [7], [9] followed by an adaptation of nonparametric sampling [20] to the 3D completion of surface relief. From the results detailed, we see the successful completion of both smooth surfaces (e.g., Fig. 7), regular/irregular/isotropic, and anisotropic surface relief (e.g., Figs. 1, 11, and 8). Some limitations in the approach exist (e.g., Fig. 10) because accumulated error over larger areas are apparent and work in this area is left for future investigation. In summary, the generality of this “completion by example” approach, together with the avoidance of “tiling” or related implausible completion artifacts, surpasses current abilities in smooth surface completion [7], [8], [9], [3], [13], [4], [6], [12], isolated surface hole filling [10], [11], [5], and surface-context-based relief “copying” techniques [14], [17], [15], [16]. Future work will investigate both the identified limitations and the generalization of the technique to more general underlying surface shapes.

REFERENCES

- [1] T. Breckon and R. Fisher, “Amodal Volume Completion: 3D Visual Completion,” *Int’l J. Computer Vision and Image Understanding*, vol. 99, no. 3, pp. 499–526, Sept. 2005.
- [2] J. Davis, S. Marschner, M. Garr, and M. Levoy, “Filling Holes in Complex Surfaces Using Volumetric Diffusion,” *Proc. First Int’l Symp. 3D Data Processing, Visualization, and Transmission*, June 2002.

Fig. 13. Golf ball ground truth comparison.

- [3] P. Liepa, “Filling Holes in Meshes,” *Proc. First Eurographics/ACM SIGGRAPH Symp. Geometry Processing*, pp. 200–205, 2003.
- [4] T. Ju, “Robust Repair of Polygonal Models,” *ACM Trans. Graphics*, vol. 23, no. 3, pp. 888–895, 2004.
- [5] J. Wang and M. Oliveira, “A Hole-Filling Strategy for Reconstruction in Smooth Surfaces in Range Images,” *Proc. 16th Brazilian Symp. Computer Graphics and Image Processing*, pp. 11–18, 2003.
- [6] L. Tekumalla and E. Cohen, “A Hole-Filling Algorithm for Triangular Meshes,” Technical Report UUCS-04-019, School of Computing, Univ. of Utah, Dec. 2004.
- [7] F. Stulp, F. Dell’Acqua, and R. Fisher, “Reconstruction of Surfaces behind Occlusions in Range Images,” *Proc. Third Int’l Conf. 3D Digital Imaging and Modeling*, pp. 232–239, 2001.
- [8] U. Castellani, S. Livatino, and R. Fisher, “Improving Environment Modelling by Edge Occlusion Surface Completion,” *Proc. First Int’l Symp. 3D Data Processing, Visualization, and Transmission*, pp. 672–675, June 2002.
- [9] F. Dell’Acqua and R. Fisher, “Reconstruction of Planar Surfaces behind Occlusions in Range Images,” *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 24, no. 4, pp. 569–575, Apr. 2002.
- [10] P. Chalmoviansky and B. Juttler, “Filling Holes in Point Clouds,” *Math. Surfaces X*, M. Wilson and R. Martin, eds., pp. 196–212, Springer-Verlag, 2003.
- [11] J. Verdera, V. Caselles, M. Bertalmio, and G. Sapiro, “Inpainting Surface Holes,” *Proc. IEEE Int’l Conf. Image Processing*, vol. 3, pp. 903–906, 2003.
- [12] T. Masuda, “Filling the Signed Distance Field by Fitting Local Quadrics,” *Proc. Second Int’l Symp. 3D Data Processing, Visualization, and Transmission*, pp. 1003–1010, 2004.
- [13] F. Nooruddin and G. Turk, “Simplification and Repair of Polygonal Models Using Volumetric Techniques,” *IEEE Trans. Visualization and Computer Graphics*, vol. 9, no. 2, pp. 191–205, Apr.–June 2003.
- [14] A. Sharf, M. Alexa, and D. Cohen-Or, “Context-Based Surface Completion,” *ACM Trans. Graphics*, vol. 23, no. 3, pp. 878–887, 2004.
- [15] M. Pauly, N.J. Mitra, J. Giesen, M. Gross, and L. Guibas, “Example-Based 3D Scan Completion,” *Proc. Third Eurographics/ACM SIGGRAPH Symp. Geometry Processing*, pp. 23–32, 2005.
- [16] G.H. Bendels, R. Schnabel, and R. Klein, “Detail-Preserving Surface Inpainting,” *Proc. Sixth Int’l Symp. Virtual Reality, Archaeology and Cultural Heritage*, pp. 41–48, 2005.
- [17] V. Kraevoy and A. Sheffer, “Template Based Mesh Completion,” *Proc. Third Eurographics/ACM SIGGRAPH Symp. Geometry Processing*, pp. 13–22, July 2005.
- [18] P. Bhat, S. Ingram, and G. Turk, “Geometric Texture Synthesis by Example,” *Proc. Second Eurographics/ACM SIGGRAPH Symp. Geometry Processing*, pp. 41–44, 2004.
- [19] Y.-K. Lai, S.-M. Hu, D. Gu, and R. Martin, “Geometric Texture Synthesis and Transfer via Geometry Images,” *ACM Solid and Physical Modeling*, pp. 15–26, 2005.
- [20] A. Efros and T. Leung, “Texture Synthesis by Non-Parametric Sampling,” *Proc. Seventh IEEE Int’l Conf. Computer Vision*, pp. 1033–1038, Sept. 1999.
- [21] T.P. Breckon, “Completing Unknown Portions of 3D Scenes Using 3D Visual Propagation,” PhD dissertation, School of Informatics, Univ. of Edinburgh, 2006.
- [22] L. Wei and M. Levoy, “Texture Synthesis over Arbitrary Manifolds,” *Proc. 28th Ann. Conf. Computer Graphics and Interactive Techniques*, pp. 355–360, 2001.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.