

A Hierarchical Extension to 3D Non-parametric Surface Relief Completion

Toby P. Breckon¹ and Robert B. Fisher²

¹*School of Engineering, Cranfield University, Bedford, MK43 0AL, UK.*

²*School of Informatics, University of Edinburgh, Edinburgh, EH8 9AB, UK*

Abstract

Typical stereo and laser scan based 3D acquisition approaches are essentially limited to 2.5D capture. The resulting 3D completion problem, to derive missing information in 2.5D scenes from limited contextual information, has received increasing attention in literature. Here we propose a hierarchical extension to our recent nonparametric approach for the 3D completion of surface relief detail to allow the resolution of inconsistencies arising in the global structure of an area completed with this technique. We test our approach over a range of surface types and contrast the presence of global discontinuities in the resulting completion with those of the earlier approach.

1 3D Surface Completion

Over recent years the problem of visual completion has become of increasing interest in both 2D and 3D imaging. A large proportion of work in this area can be set in the context of human abilities within the domain of perceptual psychology [6]. Specific to the completion problem faced in 3D computer vision is the 2.5D limitation of current 3D sensing approaches. With a single laser scan or stereo capture it is not possible to capture all faces of a 3D scene. Essentially, from this uni-directional viewpoint we have a 2.5D scene capture - 3D information but in one single direction. The result is 3D surface models that are inherently 2.5D in nature or 3D models which require construction from a combination of multiple uni-directional captures. In this latter case, additional capture and viewpoint combination incur significant additional work [3].

Following on from our work in [7] we propose a hierarchical approach that builds up the missing information within unknown scene portions as an iterative process from global structure to localised surface detail. This in turn follows the paradigm of *visual propagation* as the basis for the completion of 3D scenes via the propagation of knowledge from known to unknown scene

portions [6]. Specifically we propose a multi-level hierarchical approach for the nonparametric synthesis of 3D surface relief from an original sample surface region to a co-joined target region. This approach uses a 3D extension of non-parametric texture synthesis [13] in conjunction with knowledge derived about the underlying shape continuity of the surface [28,8,11]. Whilst the latter can be achieved via a variety of approaches [28,8,11,10,21,33,17,31] we rely on a simple geometric fitting approach to provide the required surface knowledge to illustrate our technique. As shown by our results, the resulting approach overcomes the limitations in global surface structure continuity (e.g. Figure 3) found in earlier work [7].

Prior work in this area can be divided into two areas - completion via *good continuation*, and completion via *structure propagation*. Work within the area of good continuation concentrates on the smooth completion of underlying surface. Several approaches have been proposed within this context including localised geometric and algebraic surface fitting [28,8,11,9,18], volumetric distance-field based techniques [10,32,23,15], cross-triangulation patching and refinement [21,33,31] and spatial occupancy based approaches [24,17]. Whilst these techniques concentrate on the completion of the global underlying surface shape more limited attention has been paid to the additional topic of propagating surface detail - thus limiting them in application with real-world surfaces [6]. By contrast work in the area of structure propagation concentrates on the concept of “*completion by example*”.

Structural propagation techniques, by the nature of the problem, frequently draw upon the success of 2D work in the analogue problem of 2D texture synthesis. The approach proposed here is no different. Notably the prior work of [27] pursues a patch-based “*copy and paste*” completion approach akin to [12] that, whilst well suited to smooth surface or irregular/anisotropic relief completion, relies on non-rigid alignment (warping) that limits its application to structured relief (as found in architecture, e.g. Fig. 2D). The brittleness of the patch-wise “*copy and paste*” approach can also lead to uncharacteristic ‘tiling’ artifacts over large areas and an inability to adaptively complete in scenarios where no suitable propagate patch exists [27]. By contrast the “*example-based*” completion work of [25] 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 $2\frac{1}{2}$ D surface. This is clearly limited by suitable *a priori* knowledge and, by its use of similarity matching, relates more to the concept of *completion via recognition* than the *completion via generalisation* we aim for here.

The related works of [2] and [19] are similar respectively to [25] and [27] but vary primarily in underlying surface representation, and the graphics work on geometric texturing [4,20] covers a related but not explicitly completion-orientated problem domain. Both [4] and [20] are so far limited to arbitrary

geometric texturing rather than explicit relief completion where the issues of plausibility and good continuation must additionally be considered. A body of graphics driven relief texturing work (e.g. [26]) facilitates the illusive texturing of such 3D surfaces through the use of artificial bump mapping techniques - the underlying 3D surface remains unchanged. However, recent work [36], using colour as an additional constraint and 2D texture synthesis as primary driver does explicitly consider completion but fails to offer an approach suitable for regular surface relief with additional global completion constraints. Related work, in a similar vein to our earlier non-parametric work of [7], has also recently been performed by [37] also primarily targets geometric texturing but using texture patches rather than point samples. Prior work in 3D texturing [27] has shown similar limitations to such patch based approaches [37] as found in [7]. In this work we propose an extension to our own non-parametric approach in this area [7], itself driven by the paradigm of the 2D synthesis work of [13], to aid to overcome similar limitations found with this approach - namely a shortcoming in the maintenance of such global structures over large completion areas.

A full overview of computer vision work in 3D completion and related theories from perceptual psychology is given in [6].

2 Non-parametric Surface Completion - a brief overview

The main contribution of this paper is a multi-stage, multi-pass extension to our previous work using non-parametric completion as an approach for realistic surface relief completion [7]. We initially present an overview of this base technique (with direction to further reading more detail), examples of the results it achieves (Section 2.2) and its limitations (Section 2.3) before moving forward to examine a viable hierarchical Level-of-Detail (LoD) extension to overcome these limitations (Section 3).

As originally presented in [7] our proposed non-parametric surface completion approach comprised of 2 stages:- 1) the completion of the underlying surface geometry (“*smooth*” completion), 2) the completion of surface relief over the underlying surface geometry completed in stage 1. An example of this two stage process is shown in Figure 2D where we see the completion of the architectural surface detail over the underlying geometric surface completion of the 2.5D Pisa Tower surface.

The first stage of our approach essentially forms a requirement for the underlying “*good continuation*” of the surface. This can be achieved via a number of approaches outlined in the literature (e.g. [10,21,33,17,31]). Specifically, we utilise a simplistic geometric completion approach based on the prior work of

[28,11,8]. This facilitates the automatic best-fit model selection to the original 2.5D surface data and provides an accurate surface parametrisation for completion via good geometric continuation [28,11,8] (e.g. Figure 2A). More complex surface structures can thus be processed via a route of prior segmentation into simple geometric sub-parts or an alternative approach [10,21,33,17,31] without loss of generality for the secondary application of the surface relief completion.

Non-parametric surface relief completion essentially mimics its 2D counterpart in texture synthesis [13] but with specific consideration for in-place synthesis directly on a 3D surface mesh itself. In this way it differs significantly from related work in geometric texturing [4,20] where alternative interim non-surface based representations are used for the key processing steps. The original approach in 2D is in essence very simple:- successively grow a 2D texture outward from an initial seed (the *target*), one pixel at a time, by comparing the neighbourhood of the pixel being synthesised to each and every such neighbourhood in a given *sample* texture. Select a neighbourhood that matches well from the sample and directly copy the corresponding pixel value from the chosen sample area to the target position. In 2D image textures these neighbourhoods are simply defined as $w \times w$ square image neighbourhoods 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 [13]. Matching is based upon the traditional correlation approach of the normalised sum of squared difference metric (SSD) between two pixel neighbourhoods (i.e. the textured pixels surrounding the target and those surrounding each sample pixel). From the set of matches found, a random selection is then made from the top $\eta\%$ of matches as to which value to copy to the target. The result is an effective technique, surpassing the abilities of other parametric approaches, using a statistical non-parametric model akin to the robotics paradigm “*the world is its own best model*”.

2.1 3D Nonparametric Completion

In essence basic aspects of non-parametric 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 neighbourhood becomes the set of nearest neighbours to a surface point and the actual pixel values being synthesised become displacement vectors mapping discrete points on a textured surface (i.e. the relief) to the geometric surface derived from prior fitting (Figure 1).

The main input to our non-parametric completion process is thus the geometrically completed version of the 3D surface. The original textured points (with relief), are the sample points, $s \in \text{samples}$, whilst those forming the geometrically completed “smooth” portion form the target points, $t \in \text{targets}$, (Fig.

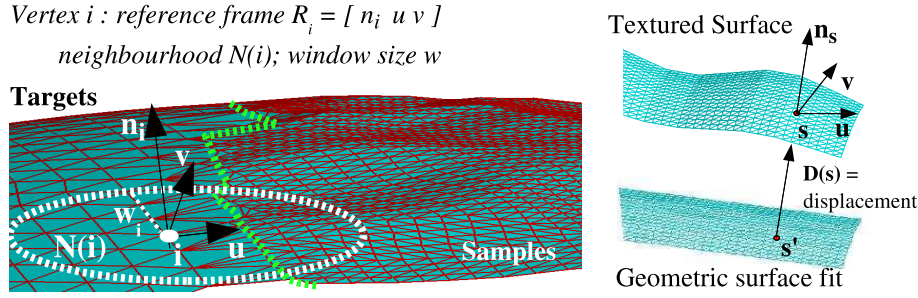


Figure 1. 3D vertex neighbourhoods, $N(i)$ and displacement vector, $D(s)$, within context

1). Each point has an associated surface normal, n , and an associated relief displacement vector, $\vec{D}(s)$, derived from earlier geometric surface fitting (Fig. 1).

2.1.0.1 Algorithm Outline Our non-parametric algorithm adapts to 3D by considering vertex neighbourhoods on the 3D surface in place of the pixel neighbourhoods of [13]. Each vertex neighbourhood, $N(i)$, is the set of vertices lying within a radius of w edge connections from the target vertex being textured (Fig. 1) where w forms the window size parameter synonymous to that of the earlier 2D approach. The algorithm now proceeds by finding the best sample region matching the textured portion of a target vertex’s neighbourhood.

Firstly, the set of target vertices currently lying on the textured/un-textured surface boundary are identified as the current target list, L . The first target vertex, $t \in L$, is then matched, using neighbourhood based matching, against every available vertex $s \in samples$. Following the traditional route of Efros/Leung [13], a match is then randomly selected from the best $\eta\%$ of this set, based upon 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 labelled 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. Once L is exhausted, the next set of target vertices on the textured/untextured boundary are identified, based on the updated vertex labelling, and the process is continued until all $t \in targets$ are labelled 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 neighbours prior to processing [13].

Throughout the process, progress is monitored over each target list, L , con-

structed. 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%), in this occurrence, to relax the acceptable error constraint (as per [13]) and thus allow relief synthesis to hopefully progress over the next iteration of L . Further discussion relating to parameters e, w and η is presented in [7].

2.1.0.2 Orientated Neighbourhood Matching From our outline discussion of the nonparametric surface relief completion one key element remains - the orientationally consistent matching of target neighbourhood, $Nt(t)$, and to sample vertex neighbourhoods, $Nt(s)$. Matching is performed using an adaptation of the pixel-wise Sum of Squared Difference metric (SSD, [13]) based on the projection of neighbourhood vertices onto the surface at each sample point as follows:

$$SSD_{shape} = \sum_{v_t}^{Nt(t)'} g_{v_t} \min_{\Delta_j \in triangles(v_2)} (dist(v_t, \Delta_j)^2) \quad (1)$$

where $Nt(t)'$ is the textured target neighbourhood, $Nt(t)$, rigidly transformed to the surface position s and thus vertex $v_t \in Nt(t)'$. To ensure consistent vertex matching, independent of relative vertex density, we match vertices, $v_1 \rightarrow v_2$, $v_1 \in Nt(t)$ $v_2 \in Nt(s)$, based on their relative projected positions on the common geometric surface model, embodied in the displacement vector associated with every vertex, $v'_i = v_i - \vec{D}(v_i)$ [7]. The SSD is calculated based on the distance of each target vertex, $v_t \in Nt(t)'$, directly to the triangulated surface at s (not just the closest vertex) - i.e. the minimum squared distance to any surface triangle, Δ_j , that has v_2 as a vertex, $\Delta_j \in triangles(v_2)$ [7]. A 2D Gaussian kernel weight, g_{v_i} is additionally used in Equation 1 to weight the SSD vertex matches, $v_1 \rightarrow v_2$, relative to the distance $t \rightarrow v_t$, $v_t \in N(t)$ (i.e. spatial proximity to t) [13,7]. A Gaussian kernel (size $\sigma = 1$) is normalised over the spatial dimension of $Nt(t)$ defined by parameter w .

The use of this form of projective matching (Equation 1) requires the ability to rigidly transform a given target neighbourhood $Nt(t)$, to an arbitrary surface position (i.e. sample position s) whilst maintaining a consistent orientation of the relief (i.e. $Nt(t)$ is matched to $Nt(s)$ such that both texture neighbourhoods are commonly orientated in a consistent frame of reference). In order to achieve this a local frame of reference, $R_{V_i} = [n, u, v]$ (Fig. 1), is derived for each and every surface vertex, v_i , by either deterministically finding mutually perpendicular vectors, $\vec{u} \vec{v}$, to the known surface normal, \vec{n} , [7] or an iterative localised neighbour-to-neighbour sweep approach based on [35]. Rigid transformation of any target vertex neighbourhood, $Nt(t)$, to any sample neighbourhood, $Nt(s)$ is thus facilitated by transformation of $R_s R_t^{-1}$. As t

is itself un-textured whilst s is textured, the natural misalignment (owing to the presence/absence of texture) is additionally avoided by further transforming $Nt(t)$ to the corresponding un-textured position of s on the underlying surface - s' , calculated using the known displacement vector at s , $\vec{D}(s)$, as $s' = s - \vec{D}(s)$ [7].

Overall, it is this use of exhaustive neighbourhood-base for each target, $t \in targets$, at each position on the sample surface region, $s \in samples$, results in a completion algorithm that is $\mathcal{O}(stw^2)$ for neighbourhood window size w [5].

2.1.0.3 Summary As we have outlined, we now have a viable 3D relief completion approach based on the adaptation of [13] to the consideration of vertex neighbourhoods upon surfaces in place of pixels within an image [7]. As we present in Section 2.2 (Figure 2), this approach can be successfully employed for relief completion over a wide variety of example surfaces. However, from our discussion here a key point to note is the localised view of the surface this approach follows using only localised surface knowledge (i.e. $w \times w$ vertex neighbourhoods). In Section 2.3 we will explicitly see the limitations this incurs upon this overall localised approach and, in presenting our main contribution of this work, outline a further hierarchical completion approach to overcome such limitations. *A full overview of the single-pass non-parametric discussed here is presented by the authors in [7]. It is re-presented here as a baseline to the contribution of Section 3.*

2.2 Exemplar Surface Completion Results

From the non-parametric surface completion technique outlined thus far we present a number of completion results in Figure 2A-E. In Figures 2A & B show the successful completion synthetic relief patterns over cylindrical surfaces whilst Fig. 2C shows 3D relief completed over a planar surface. Each case shows the completion of an initial surface patch (light) over a geometric completion (dark) of the underlying surface model. Surface completions based on real $2\frac{1}{2}$ D surface captures are presented in Figures 2D-E. In Figure 2D we see the completion of regular, anisotropic architectural features of Pisa Tower including the successful completion of regular windows, doors, struts (bottom) and building specific architectural features (top) from an initial $2\frac{1}{2}$ D scan of this scale model. Figure 2E shows the completion of anisotropic, irregular tree bark¹ texture over a geometrically completed cylinder. Here, we see the

¹ N.B. Linear bark structure running parallel to cylindrical surface axis.

successful completion of the bark structure despite the highly stochastic nature of the initial sample relief over the bark “*structure*” itself.

Based on this single-pass approach [7] successful completion results can be achieved over a wide range of surface relief types. These examples (Figure 2) are reproduced from [7] (where a full quantitative evaluation of this single-pass approach is performed) for comparison with the proposed enhancements presented here (Section 4) .

2.3 Limitations of a single-pass non-parametric approach

However, a key limitation of this single-pass approach is the inherent localness of a single-pass non-parametric technique. This reliance solely upon a measure of local consistency means that the global consistency of the surface relief completion can be compromised due to the accumulation of errors introduced over the iterative relief propagation process (Section 2). As a result, various artifacts are apparent in portions of surface relief completion. This limitation was highlighted in the original work [7] and is generally limited to relief completions over larger surface areas (as the error accumulates substantially) rather than use for isolated hole filling. Here we examine these limitations using a range of examples (Figure 2).

For example, whilst the Pisa Tower example (Figure 2D) shows no obvious visual difference between the completion and original in Figure 3 A/B, we see that on closer examination subtle anomalies exist with the completed surface portion - repeated structure and mis-matches 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. Both of these issues lead to accumulated error problems (Fig. 3 A/B) similar to the problem originally encountered in [13] with the loss of high-order information.

We additionally show the example of Figure 3C (completion of a plinth type surface) that has both cylindrical and right-angled components in its overall shape. In this simple case we can fit a cylindrical geometric completion to the available $2\frac{1}{2}$ D surface data despite its non-conformance to a general geometric cylinder. Generally we see the successful completion of this underlying shape geometry (Figure 3C, top down views) but a closer examination reveals some anomalies similar to those suffered earlier with the Pisa Tower (Figure 3A/B). Figure 3C shows a similar rear-side joining anomaly, due to accumulated error, as to that found in the Pisa Tower completion. This is similarly attributable to the accumulated noise issues but additionally suffers from the added potential for poor initial geometric fitting due to the non-conformance of the

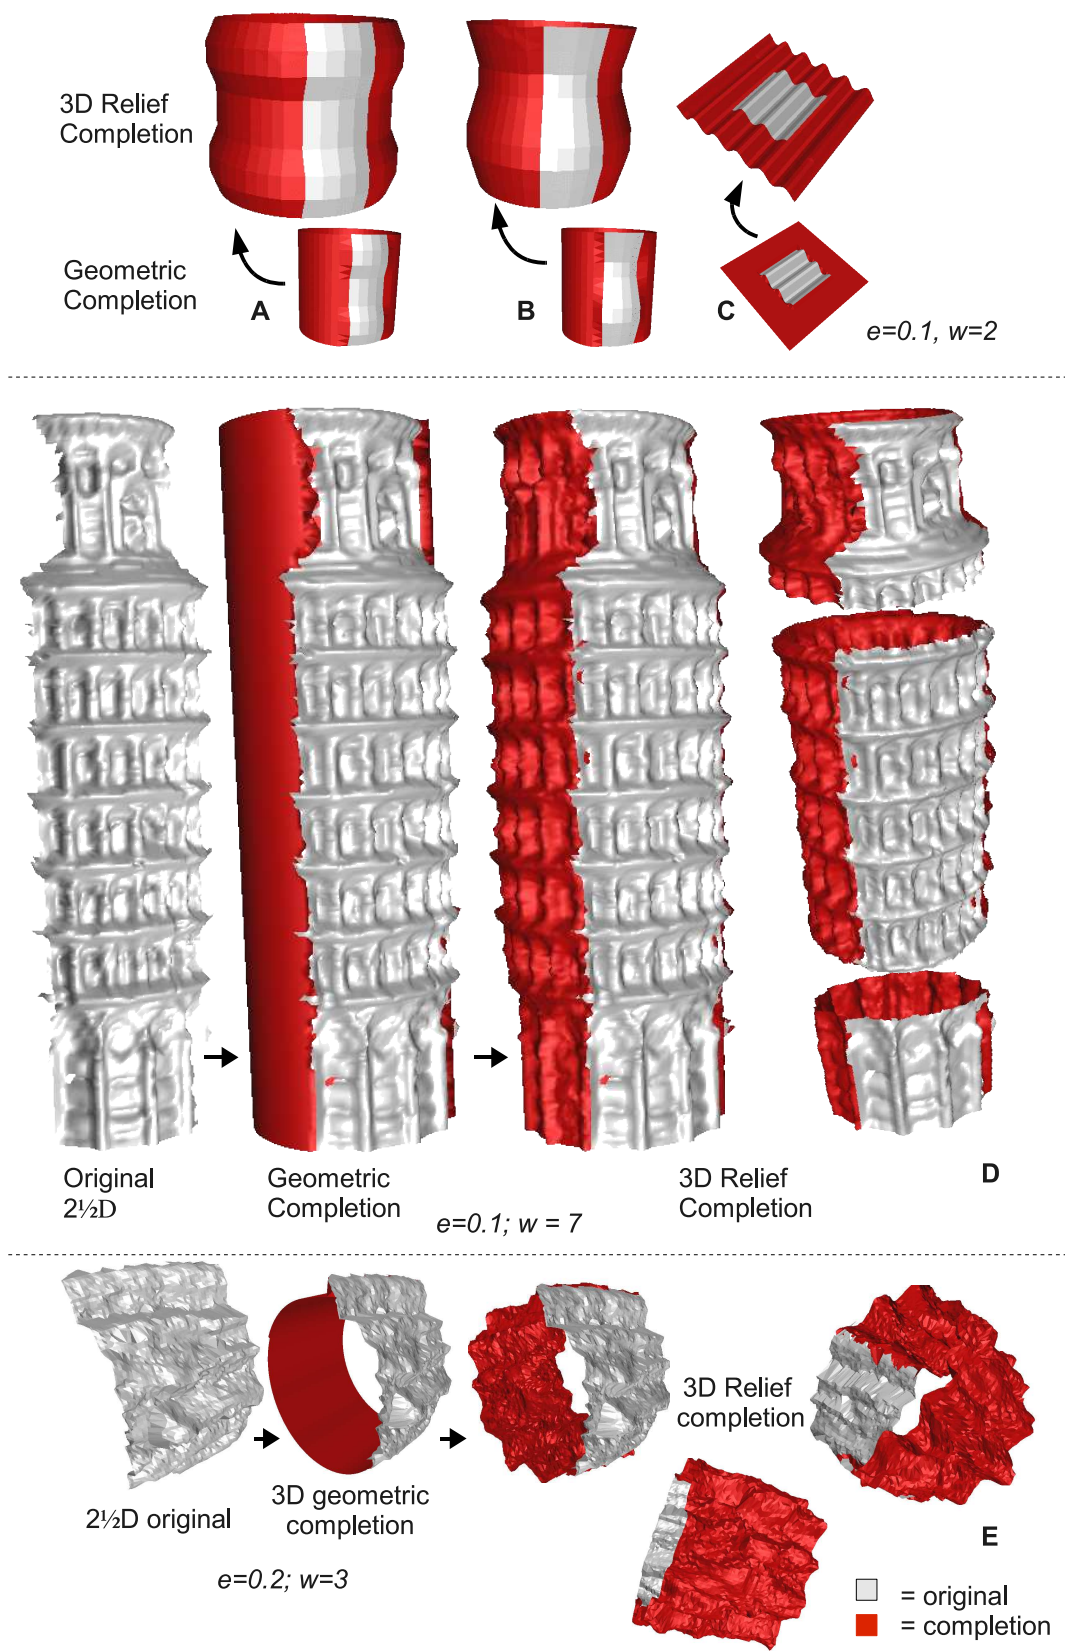


Figure 2. Examples of 3D relief completion using non-parametric surface completion

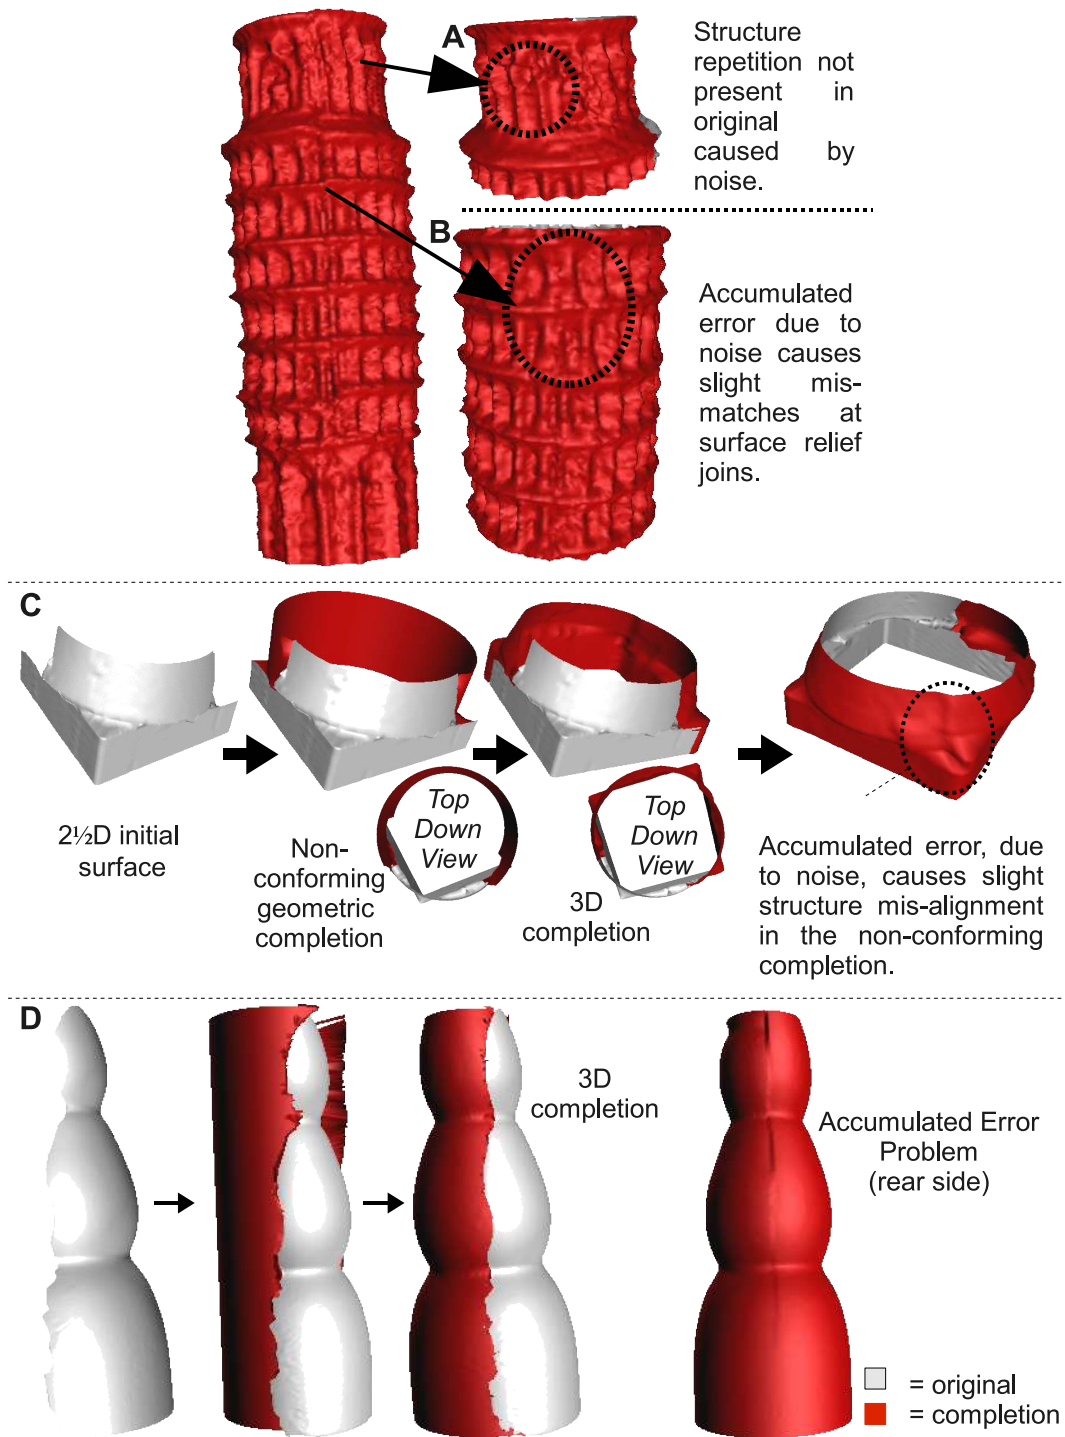


Figure 3. Limitations of 3D relief completion using non-parametric surface completion

original surface data. Finally, we also see that in the relatively simple (i.e. low frequency) surface geometry completion of example of Figure 3D similar global consistency artifacts are suffered.

These examples, and the successful correction of these artifacts in the presence of 3D surface completion, will form the basis for evaluation of the novel hierarchical extension to the core technique of Section 2 which is the main contribution of this paper.

3 A Hierarchical Extension to Non-parametric Surface Completion

We can additionally extend our mono-scale technique (presented thus far) to consider relief completion using a hierarchical level of detail (LoD) surface model. As in prior 2D texture synthesis extensions (e.g. [34]) we see the use of multi-scale techniques both as a mean of facilitating synthesis on a global to local scale and as a means of improving computational performance. We are primarily interested in two areas of improvement over the previous mono-scale approach:

- **Structural Accuracy:** Hierarchical completion allows the completion of successive levels of surface detail from coarse to fine (i.e. from global relief structures to local relief detail). Here we aim to use this separation of structure and detail to overcome the accumulated error problems associated with the mono-scale technique (e.g. Figure 3).
- **Efficiency:** A hierarchical completion technique will require vertex matching at multiple levels in the hierarchy. As a result the best match found at level $i - 1$ in the hierarchy can be used to “*guide*” the set of possible match candidates at successive level i . Reducing the possible set of match candidates in this way reduces the search space for non-parametric completion at a given level. Overall, a coarse match at the lowest hierarchical level will allow successive search space reduction and refinement at all subsequent levels leading to potential gains in efficiency.

Here we adopt the progressive mesh representation of Hoppe [14], a continuous LoD mesh representation, to our purposes and adapt it to allow us to construct an increasing LoD surface hierarchy of m discrete levels of detail. Specifically we design a surface hierarchy that maintains approximately uniform vertex density over the surface, respecting any oversampling performed on the original portion, and contains $k \frac{n}{m}$ vertices for $k = \{1 \dots m\}$ levels, M_k . In summary this provides us with M_m discrete LoD for a surface and known vertex mappings, in terms of vertex addition/subtraction relationships, between adjacent levels, $M_{k-1} \rightarrow M_k$. A general overview of LoD surfaces is presented in [22] with further details of its adaptation to the discrete LoD model used here.




Figure 4. Calculation of displacement vectors from surface hierarchy

3.1 Hierarchical Surface Relief Propagation

Based on our hierarchical surface model we can outline an extension of the earlier non-parametric completion to hierarchical surfaces. In essence this adaptation of our non-parametric completion approach can be described very simply :- *at each level in a given surface hierarchy perform non-parametric completion using the completion of the previous level to guide the process.*

At each level in this discrete hierarchical model, M_k , we set the displacement vectors, $\vec{D}(v_j)$, for the each sample (i.e. original) vertex $s_j \in M_k$ based upon the projected displacement onto the M_{k-1} surface in the hierarchy (see Figure 4). This represents the current successive level of detail added to the original surface portion in the hierarchical transition from level M_{k-1} to M_k and forms the set of relative displacement vectors from LoD M_{k-1} to M_k . It is this additional difference in detail, relating to the current level in the hierarchy, that we wish to propagate from the original sample vertices to the set of target vertices (i.e. 3D geometric completion) with non-parametric completion at the current level. In order to achieve this we must maintain distinct sets of samples and targets at each LoD by vertices in lower LoD inheriting their label from parent vertices present at higher LoD [14]. This gives a distinct sample and target set for each LoD M_k upon which to propagate the displacement vectors for that specific LoD in the hierarchy.

For the first level in the hierarchy, M_1 , where no prior surface exists the displacement vectors are calculated based on projection to the underlying geometric surface model for the surface as per the mono-scale approach discussed earlier. In this formulation, the geometric surface fit can be considered as level M_0 within the surface hierarchy forming the base case for the propagation of successive layers of surface detail in the form of displacement vectors.

Overall, based on this successive propagation of relative displacement vectors between levels in the surface hierarchy we are able to outline the hierarchical algorithm as follows:

- *surface* = original triangulated surface with vertices labelled as target/sample.
- *targets* = list of target vertices of *surface*.
- *samples* = list of sample vertices of *surface*.
- *m* = number of discrete levels in surface hierarchy
- *G* = geometric surface fit

hierarchicalGrowSurface(surface, targets, samples, m)

construct *m* level hierarchy H

for each level M_k in H

if M_k is lowest level of H then

set displacement vectors of M_k from *G*

else

set displacement vectors of M_k from M_{k-1}

end

targets' = target vertices of M_k

samples' = sample vertices of M_k

non-parametric completion(surface M_k targets', samples')

end for

Here “*non-parametric completion()*” explicitly refers to the mono-scale completion algorithm discussed earlier (i.e. non-hierarchical). Here this is performed on subsequent levels, M_k , in the hierarchy with the corresponding target/sample sets at that LoD.

3.2 Improving Efficiency

As identified earlier, the computation associated with mono-scale non-parametric completion is bound by $\mathcal{O}(stw^2)$ for *s* sample vertices and *t* targets with window radius *w*. A similar bound exists for the generic hierarchical approach proposed previously where we find an overall bound of $\mathcal{O}(stw^2m)$ for *m* levels in the surface hierarchy² (further analysis in [5]). Overall, as previously

² Assuming, as is plausible from [14], that the transition between levels in the surface hierarchy require only minimal computation and can be done in linear time., i.e. $\mathcal{O}(m(s+t))$.

discussed the exhaustive search in the space of all sample vertices for every target match can be seen to be prominent in this computational cost.

Here we attempt to reduce the required computation for a given hierarchical non-parametric completion based upon the following principle: *By using matching knowledge obtained at the M_{k-1} level as a prior, the search space for potential matches at the M_k level can be reduced.*

Although there are several architectures for the implementation of this principle, here we implement a scheme based on restricting the set of samples considered for matching, against a given target at level M_k , to those within the spatial vicinity of the match chosen at the M_{k-1} . The justification behind this method is that matching in higher levels in the hierarchy, M_k , will be based upon the correlation of relief features previously propagated at the lower levels, i.e. M_{k-1} and below. As we already know the general area of the potential best correlation from matching, and subsequent relief propagation, performed at the previous level M_{k-1} in the hierarchy this seems a natural location to find an appropriate match for the successive level of relief detail at the current level, M_k . With this concept in mind we restrict matching for a given target vertex at level M_k to the spatial vicinity of the sample surface from which the chosen match for that target was found at the previous level, M_{k-1} (or its expanded parent vertices in the hierarchical surface model [14]). It should be noted that this is not a hard restriction as it is a) only limited to a region, not a precise vertex location and b) can be ignored if no suitable match is available in that constrained sample area (i.e. within match tolerance e). Additionally, as the surface LoD changes, as we progress from level to level, the spatial restriction of matching similarly adjusts based both on the surface locale of the match found in the previous M_{k-1} LoD and in relation to vertex density as we effectively reduce the physical size of the restricted window in step with the increased LoD (i.e. vertex density) - i.e. for greater LoD and thus higher relief detail the physical area covered for a constant matching window w is reduced. This concept is termed *sample pyramids* (discussed in Section 3.3).

It should be noted that these restrictions are maintained, level-to-level, on a per-vertex basis and no region based restrictions based on adjacent target vertices matching to the same sample area are enforced. This latter concept, restricting the neighbours of a particular target vertex, $N(t)$, to match to neighbours $N(s)$ of the sample vertex s that matches target vertex t is not desirable as it amounts to verbatim “*copying*” of 3D surface relief (akin to [27,19,25]) rather than completion by a more generalised form of relief derivation.

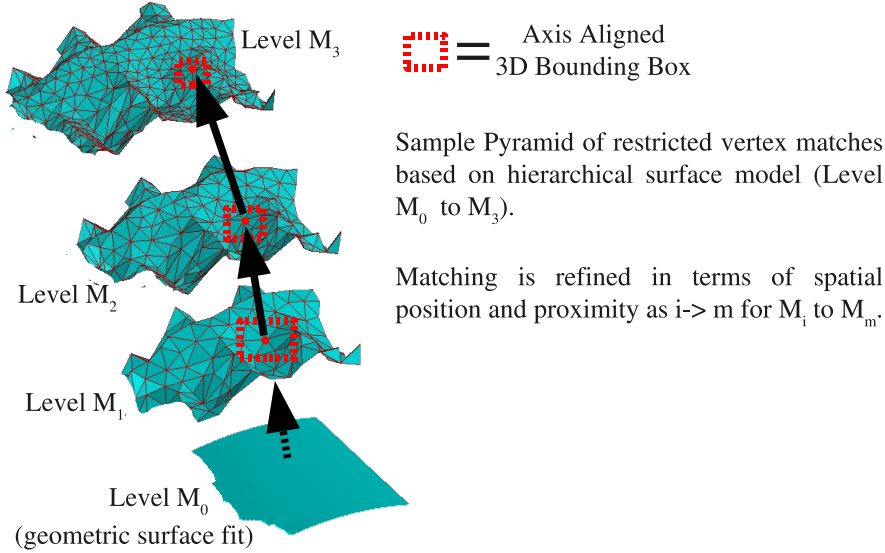


Figure 5. Sample pyramids for restricted matching

3.3 Sample Pyramid Approach

Our approach for implementing this restrictive sample matching is based on the concept of sample pyramids (Figure 5) whereby we successively restrict matching for targets at a given surface position, relative in each hierarchical level, to a specific region of the sample surface.

At the lowest level in the hierarchy, with a total of $\frac{s+t}{m}$ vertices present, unrestricted matching is performed between this reduced set of target and sample vertices. Based on the matches identified, a spatial bounding box identifying a portion of the sample surface is assigned to every target vertex completed at this level, M_1 . These bounding boxes are subsequently mapped, in the same manner as surface relief, to the target vertices at next level in the hierarchy, M_2 , by simple augmentation of the hierarchical surface model [14]. The result is that at level M_2 every target vertex has an associated spatial bounding box representing a region of the sample surface to which matching for this target can be suitably restricted. As, in general, the bounding box region can be assumed to be a subset of the whole sample surface the matching space for each target vertex at level M_2 is subsequently reduced.

This can now be repeated for successive levels in the hierarchy such that in all cases a target vertex at level M_k has an associated bounding box, derived from the match found at level M_{k-1} , that reduces the potential search space for a suitable matching sample vertex. For higher levels in the hierarchy, the spatial bounding box derived at level M_{k-1} will itself be based upon a match found using the restricted set of samples within a bounding box derived from matching at the M_{k-2} level in the surface hierarchy and so forth down to the base level, M_1 . This introduces the concept of sample pyramids - from an

Matching for t , in terms of rigid transformation, is limited to the set of sample vertex positions defined by the associated 3D bounding box of sample vertices defined by earlier matching at the previous level in the hierarchy. However, the neighbours of t , i.e. members of $N(t)$, are free to match to vertices outside of this restricted area in the overall target neighbourhood to sample surface matching construction at a given level.




Figure 6. Restricted neighbourhood based matching

initial exhaustive matching search at M_1 the matching region for vertices in a particular spatial area are refined further and further to a smaller spatial surface area - a bounding box at level M_m . This pyramid analogy comes from the successive reduction from a wide base (i.e. exhaustive matching at M_1) to a narrow spatial set of matching candidates for a given target (i.e. bounding box at M_m). However, unlike traditional pyramid models in image processing, at each level within the hierarchy the medial axis of the pyramid is “adjusted” slightly depending on the new spatial position of the chosen match at every successive level (see Figure 5).

Additionally, by specifying the size of the bounding box in terms of our earlier neighbourhood size parameter, w , we allow for successive refinement of matching in terms of spatial proximity as well as spatial position to the match at M_{k-1} . This is facilitated by a constant neighbourhood size, w , that measures proximity in terms of edge connections meaning that the bounding box volume will reduce in real-terms (i.e. spatially) as the level in the hierarchy, and hence the vertex density, increases. Specifying the bound box in this way re-enforces our concept of sample pyramid as the spatial proximity of a match at level M_k to that of M_{k-1} reduces (i.e. narrows) as $k \rightarrow m$. Based on the fixed vertex neighbourhood size w , the number of sample vertices considered for matching to each target vertex remains roughly constant, albeit over a smaller spatial area as density increases, for level M_k $k \geq 2$ (see Figure 5).

An important point to clarify here is that although we are now matching a given target vertex, t , to a subset of the set of sample vertices we are still not performing purely neighbourhood (closed set) to neighbourhood matching. Whilst t itself will only be matched against the restricted set of samples (from the bounding box) the vertices of transformed neighbourhood of t , $N'(t)$, will still be free to match against vertices outside of this set (that naturally lie within the same region of the sample set to this position). This is illustrated in Figure 6.




Figure 7. Hierarchical completion with and without post-completion smoothing

3.4 Implementation

In terms of implementing this scheme, the bounding boxes are simply derived as the 3D axis-aligned bounds of the set of neighbouring sample vertices (i.e. those originally with relief) that surround a given sample chosen for matching within a distance of w edge connections at level M_{i-1} . At level M_k , a k -d tree data structure [1] representing the sample surface portion is simply queried with the bounding box, associated to the current target vertex t , in order to efficiently return the restricted set of sample vertices that exist within this specified volume at this level. These are then used for matching against the current target vertex where a successful match produces both a refinement to the surface relief at t and a subsequent bounding box restriction for propagation and use at level M_{i+1} .

However, this scheme produces two additional cases that have to be handled : a) when no suitable match can be found in the restricted sample set and b) when a specified bounding box volume contains no sample vertices at level M_i . The former arises from the outline of our non-parametric completion algorithm where we use our restrictive upper threshold on matching error, e , to reject poor quality vertex matches. Based on this formulation it is possible that no suitable match will be found in the restricted set of samples. Similarly, it is possible, although unlikely that the hierarchical surface expansion from M_{i-1} to M_i may result in a scenario where no sample vertices lie within the volume, specified from M_{i-1} , at the expanded version of the surface M_i . In both cases we simply remove the spatial restriction on the set of samples to be used for matching to this particular target vertex and perform exhaustive matching against the full set of sample vertices for this level. When a match is found a new, updated bounding box restriction is derived based on the match identified for use in subsequent levels.

3.5 Smoothing

A problem stemming from this algorithm as outlined is the propagation of surface noise in conjunction with the surface relief features. As was seen in the earlier examples, in relation to mono-scale completion, noise present in the original surface portion can also be expected to be propagated in addition to completed relief.

Whilst the effects of these noise elements have been shown to be generally minimal in the mono-scale case, apart from the effects of accumulated error, the effect over multi-scale is somewhat different. In the mono-scale case the noise is propagated at a single level and is generally only marginally greater than that present in the original surface portion. However, the multi-scale hierarchical approach performs relief propagation, and hence noise propagation as a by-product, at m levels within the algorithm. Although, the natural noise will be reduced in the lower levels of the surface hierarchy, in a similar manner to surface relief detail, its presence is still apparent due to the increased presence of surface aliasing at these levels which is in itself inherent in the use of a reduced resolution (i.e. multi-scale) surface model. Additionally the propagation of even the slightest aspects of noise in lower hierarchical levels can affect matching at that level and at all subsequent levels that depend on it. In both cases we get accumulated error - either over the surface at a given level (as before) or through the hierarchy at multiple levels. The combination of both, unchecked can lead to high levels of noise that lead to poor completion results.

In order to overcome the potential effects of this noise propagation we implement post-completion smoothing at levels M_1 to M_{m-1} in our m level surface model. Surface smoothing provides a fairly robust, and generally accepted, manner of removing noise from a surface within a signal processing type framework [29,30]. By employing a smoothing technique after the application of non-parametric completion at each level the aspects of noise in the completed relief can be reduced, prior to relief propagation, and their overall effect isolated. In addition to smoothing the target vertices, over which non-parametric completion has been performed, we must also equally smooth the sample vertices. This is carried out so that displacement vectors for the samples at level M_i , computed from smoothed level M_{i-1} , are representative of the actual relative difference in surface relief as it will be present on the completed targets at level M_{i-1} (i.e. with smoothing applied).

Overall, all the vertices at a given level M_1 to M_{m-1} are smoothed post non-parametric completion using a Gaussian weighted smoothing operator over the field of vertex displacement vectors. Each displacement vector, $\vec{D}(v_i)$, is smoothed based on a weighted mean of its neighbourhood, $N(v_i)$. This is

effectively traditional Gaussian smoothing applied to the fixed spatial position vector field to achieve similar local surface smoothing results to [29]. At the final level of hierarchical completion, M_m , smoothing is not performed as any noise propagated at this level is synonymous both to that present inherently in the original surface portion and that present in a completion performed using the mono-scale technique. We illustrate the benefits of this post-completion smoothing over the potential ill effects of noise propagation in the comparison of Figure 7A/B where we see the clarity of surface structure clearly improved with the use of post-completion smoothing.

3.6 Complexity Analysis

In the ideal case³ this hierarchical approach reduces the computational complexity bound of completion to $\mathcal{O}(\frac{stw^2}{m^2})$ based on the unrestricted, exhaustive matching performed at level M_1 and matching at all subsequent levels, $m - 1$, being suitably reduced by the matching restrictions introduced. However, for cases in subsequent levels where no suitable match can be found within the restricted set (cases *a*) and *b*) discussed earlier) the target vertex in question must be matched to all possible samples resulting in $\mathcal{O}((\frac{s}{m})kw^2)$ where this occurs for k targets at level M_i . However, in the worse case scenario where this occurs for all targets at the top level, M_m (i.e. $k = t, i = m$), we have no asymptotic performance gain compared to our generic technique, i.e. $\mathcal{O}(stw^2m)$ (in fact lesser performance in practice due to the overheads of hierarchy). The hope is, however, that the average case run-time will be closer to the ideal case, with limited exhaustive cases, and will thus be a significant improvement over the mono-scale technique.

In general the spatial separation of given any n target vertex locations on the surface could lead to cases where explicit parallel completion could be performed by first spatially limiting any concurrent completion of target locations in the same surface locale (that hence “interact” in terms of our surface completion approach). This is left as an area for future work.

4 Evaluation

For consistency we adopt the evaluation framework of [7] using both visual comparison of the original and completed surface portions together with the statistical comparison of the Mean Surface Integral (MSI) between the relief

³ i.e. with no cases of restricted matching failure that cause the use of unrestricted matching for a given target vertex.




Figure 8. Hierarchical completion of synthetic surface

surface and the underlying geometric surface fit on the original/completed surface portions (Eqn. 2).

$$MSI = \frac{1}{n} \sum_{i=0}^n \vec{D}(i) \cdot \vec{n}(i) \quad (2)$$

The MSI (Eqn. 2) is defined for n vertices where $\vec{D}(i)$ is the surface relief displacement vector (original or propagated/completed) and $\vec{n}(i)$ is the surface normal at vertex i (Fig. 1). The primary means of comparison is the percentage difference between this measure for the original and completed surface portions for the use of both the single-pass and hierarchical extension of 3D nonparametric completion.

Our evaluation of the hierarchical technique is based on the visual and statistical comparators outlined previously together with an additional CPU run-time comparison. The uniform test environment for this run-time comparison is a Linux (2.6) workstation with an Intel Pentium4 1.8Ghz CPU and 1Gb of RAM. We evaluate both the plausibility of the hierarchical completion results and the relative performance of the proposed techniques.

A statistical comparison of the hierarchical algorithm and the original mono-scale approach is given in Table 1. Additionally we present a number of visual comparisons in Figures 8 - 16.

Firstly we return to our set of synthetic surfaces originally considered in Figure 2A/B/C to examine the hierarchical completion of three examples from this set (Figure 8 A-C) in comparison to the earlier mono-scale results. In Figure 8 we see three exemplar surfaces completed using two levels of surface hierarchy ($m = 2$). Here the number of levels is limited by the size of these surfaces, specifically the level to which the (limited) number of vertices can be effectively reduced.

In the case of cylindrical completion (Figure 8 A & C) we find the resulting hierarchical completion gives similar results to the mono-scale approach in terms of statistical difference ($\sim 0-2\%$ difference, Table 1). In terms of run-time performance, the hierarchical technique provided a significant computational

Surface	Original			Hierarchical Completion					Mono-scale Completion				
	MSI	Fig.	Parameters	MSI	% diff.	CPU (s)	Fig.	Parameters	MSI	% diff.	CPU (s)		
Synth. Cylinder	0.24	8A	w=2; e=0.1; m=2	0.25	1.22%	14.9	2A	w=2; e=0.1	0.25	1.80%	39.5		
Synth. Plane	1.19	8B	w=2; e=0.2; m=2	1.13	-4.44%	145.1	2B	w=2; e=0.1	1.19	0.35%	290.4		
Synth. Cylinder	0.18	8C	w=2; e=0.1; m=2	0.18	0.49%	15.1	2C	w=2; e=0.1	0.18	0.71%	39.5		
Cyl. Band	0.21	9	w=6; e=0.1; m=4	0.21	-1.56%	14644.5	9	w=6; e=0.1	0.21	-0.08%	86520.8		
Pisa (small)	0.67	10B/C	w=4; e=0.1; m=7	0.67	-1.10%	4909.9	10A	w=4; e=0.1	0.68	1.03%	16249.6		
Pisa (top)	1.27	13	w=2; e=0.1; m=5	1.44	12.8%	1105.12	2D	w=7; e=0.1	1.34	5.51%	-		
Pisa (middle)	0.78	13	w=2; e=0.1; m=6	0.71	-8.8%	2166.9	2D	w=7; e=0.1	0.91	16.81%	-		
Pisa (bottom)	0.81	13	w=2; e=0.1; m=7	1.01	25.01%	777.0	2D	w=7; e=0.1	0.81	0.21%	-		
Tree Bark (cyl.)	1.18	16	w=3; e=1.0; r=2; m=3	1.23	4.25%	6490.8	2E	w=3; e=0.2; r=2	1.36	14.87%	59054.5		
Candlestick	5.13	14	w=2; e=0.1; m=3	5.42	5.49%	44739.0	3D	w=2; e=0.1	5.86	14.02%	-		
Plinth	1.88	15	w=2; e=0.2/1.0; m=4	1.86	-0.76%	140.6	3E	w=6; e=0.2	1.89	0.55%	3101.8		

Table 1. Statistical results of hierarchical completion




Figure 9. Accumulated error correction: simple example

saving (~ 25 seconds less $\approx 60\%$, Table 1). Over the entire set of example cylinders (720 sample vertices on average, 561 targets) the hierarchical approach achieved an average run-time of 14.9 seconds.

In contrast the hierarchical completion of our example synthetic planar surface (Figure 8B) fared less well. In the completion of Figure 8B we can see clear aliasing artifacts in the completion of an otherwise smooth surface. The presence of these artifacts is echoed in the difference in the statistical relief difference ($\sim 5\%$, Table 1) when compared to the mono-scale approach ($< 1\%$, Table 1). This is attributable to combination of limited original sampling for this surface example, itself with higher frequency relief than the cylinder examples (i.e. Figure 8 A & C), and the higher error threshold required to facilitate effective matching at level M_1 in the surface hierarchy. In the original mono-scale completion result (Figure 2A/B/C) the effects of aliasing, due to limited original surface sampling, were limited to a mild statistical difference in relief ($< 1\%$, Table 1). By contrast the introduction of reduced resolution, in terms of the surface hierarchy, increases the aliasing effects resulting in a highly coarse representation of this relief at level m_1 . This coarseness leads to a requirement for a higher error threshold, $e = 0.2$, to facilitate vertex matching at this level. The combination of this higher threshold and the amplification of the mild aliasing by an initial stage of highly coarse matching results in the artifacts present in Figure 8B.

However, the hierarchical technique facilitated a similar run-time improvement to that achieved in the cylindrical examples ($\sim 50\%$, Table 1). This is consistent with the results achieved over the wider set of synthetic planar surfaces where it took an average of 198.7 seconds per example [7].

In Figure 9 we specifically examine a very simple completion example where

we can clearly show the advantage hierarchical completion has in terms of producing plausibly acceptable completion results in terms of consistent global surface structure.

Figure 9 (top, left) depicts an original portion of a simple cylindrical surface with a single raised band around the centre. Unlike the earlier low-density, simple synthetic examples (e.g. Figure 8) this example is a simplistic real surface capture specifically captured at high resolution (0.5mm^2). Being a real surface capture this band contains subtle surface imperfections (i.e. fine detail) not present on a synthetic example.

The high point density of the surface coupled together with the physical imperfections in the surface band itself and additional capture noise create an ideal scenario for accumulated noise to propagate. As a given window size, w , will only cover a very small physical area of the surface (due to high sample resolution) completion decisions/actions made with this purely localised view of surface structure are highly likely to accumulate error. This can be caused either from capture noise or due to the imperfect the original surface and can in turn cause imperfections in the overall structure of the surface due to the global effects of this locally accumulated error. The hypothesis that is that a more global view of surface structure (e.g. using hierarchical completion in place of the earlier mono-scale approach) will overcome this problem.

From Figure 9 (top, right) we see that mono-scale completion of this example with a window size parameter $w = 6$ (reasonably large given level of surface relief detail present) does indeed cause an obvious structural anomaly in the relief completion at the rear of the surface. In this case it is the accumulation of noise due to fine surface sampling and capture noise that specifically causes this problem (Figure 9 top, right). This problem could possibly be solved with a larger window size parameter in order to reduce the localised view of the mono-scale completion but this incurs considerable computational cost. In addition, with a featureless surface such as this, suitable window size settings can be difficult to determine successfully. Instead we maintain the same set of parameters and introduce a more global view of surface structure by the use of hierarchical completion.

As shown in Figure 9 (bottom, right) the use of hierarchical completion overcomes the structural anomaly encountered in the earlier mono-scale completion (9 top, right). Here we introduce four levels of surface hierarchy ($m = 4$, chosen empirically) and see a structurally correct completion with only mild additional surface noise. This is reflected in the mild statistical difference between the mono-scale and hierarchical completions ($\sim 1\text{-}2\%$, Table 1). The hierarchical algorithm has a $\sim 83\%$ computational saving over the mono-scale technique (Table 1). Additional surface noise could be easily removed by additional post-processing (e.g. [16]) and overall we show the potential of hier-

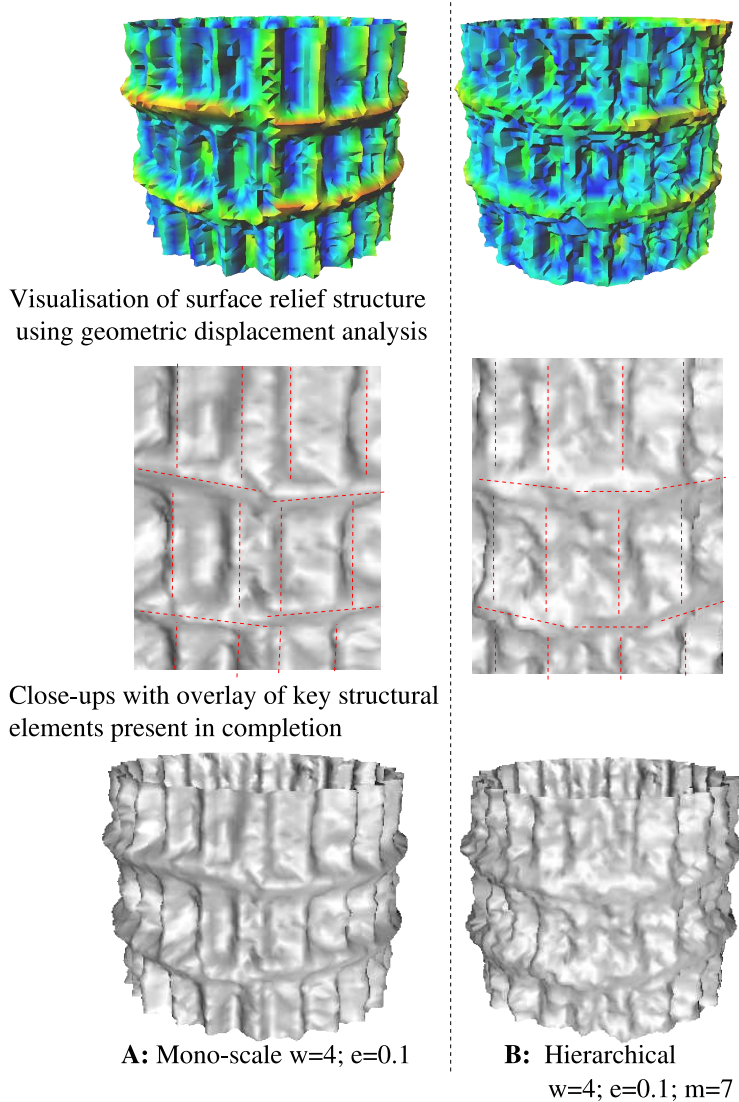


Figure 10. Hierarchical correction of Pisa Tower artifacts

archical approach to overcome the structural plausibility and computational limitations identified with the earlier mono-scale approach.

From these exemplar surfaces we now investigate the benefits of hierarchical completion to more general surface examples. Firstly let us consider a sub-part of the Pisa Tower example that we originally considered in Figure 2D. Here we consider part of the mid section with which the mono-scale completion approach originally had difficulty (Figure 3A/B). We isolate a sub-part of the original mid-section and complete it in turn with the mono-scale and hierarchical completion approaches (Figure 10). Using a smaller window size than in the earlier Pisa examples of (i.e. Figure 3A/B) we recreate the structural anomaly of Figure 3B in this example surface (Figure 10A, i.e. non-uniform window features, mis-aligned architectural rings).

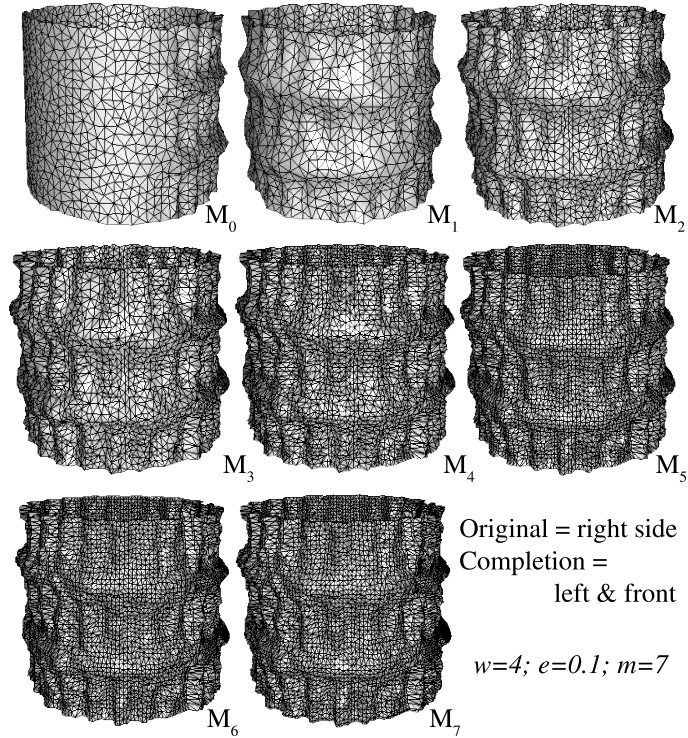


Figure 11. Stages of hierarchical completion for Figure 10

By maintaining the same window parameters we now introduce seven levels of hierarchy ($m = 7$) using our hierarchical completion approaches (Figure 10B). We see the correction of the structural anomaly with only a minimal statistical change in the surface MSI (~ 0.1 Table 1). This is further highlighted in Figure 10 with the use of overlays of the key structural elements (middle) and geometric displacement visualisation. It is clear to see that the anomaly is mostly corrected in the hierarchical cases albeit at the expense of additional surface noise in the resulting completion (despite the use of post-completion smoothing). This noise is less apparent in the less detailed hierarchical completion examples (e.g. Figures 9, 14 & 15) and may be a feature of applying hierarchical completion on a highly detailed surface such as the Pisa model.

This example (Figure 10) shows both how the introduction of hierarchy into the surface completion problem can : a) help overcome the problems of structural imperfections in real multi-level of detail surfaces and b) facilitate the use of a smaller window size parameter (here $w = 4$, in Figure 2D $w = 7$).

Figure 11 shows the level by level completion of this example (Figure 10). The realisation of the more global surface relief structure occurs almost initially from the first level of hierarchical completion, M_1 . After this successive levels of refining surface detail are added and subsequently smoothed ($M_2 \rightarrow M_6$) on top of this established surface relief structure on a global to local scale. Finally, the finest, most localised detail is added at level M_7 . By contrast we also show some final results for this example using less levels of hierarchy in

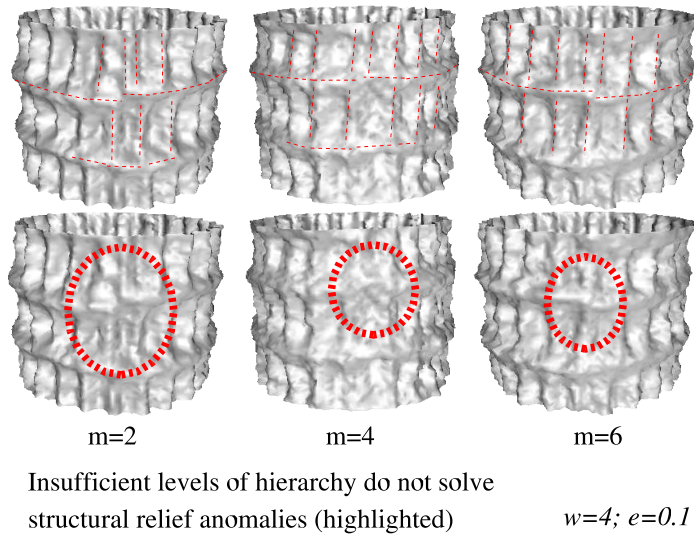


Figure 12. Figure 10 with less hierarchical levels

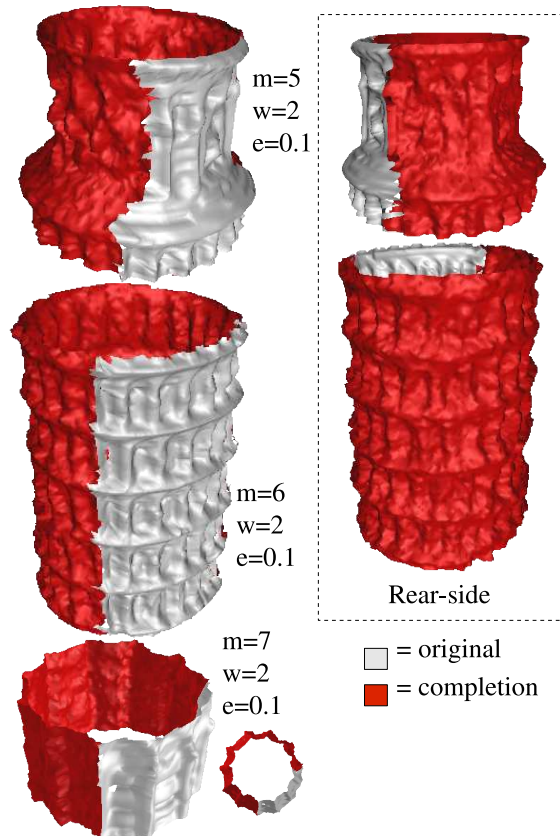


Figure 13. Hierarchical completion of Pisa Tower

Figure 12. This figure shows how inadequate levels of surface hierarchy fails to solve the problem of structural relief anomalies. In general determining a suitable choice of hierarchical LoD (parameter m) for a given surface is left as an area for future work.

We similarly re-consider the original Pisa Tower examples of Figure 2D in Figure 13. Here we see the use of hierarchical completion overcomes the structural anomalies previously highlighted in the mono-scale completion results (Figure 3 A/B) and additionally see the use of hierarchical completion offers a significant computational saving (Table 1). However, for the bottom part of the tower in isolation (Figure 13, bottom left) we see a significant increase in the statistical difference with the use of hierarchical completion ($\sim 25\%$, Table 1) despite a visual similarity from a side-on or top-down view (Figure 13, bottom left). Similar results are achieved for the isolated top section (Figure 13, top left) where a $\sim 7\%$ statistical increase is present when compared to mono-scale completion (Table 1). Similarly the mid-section of the tower (Figure 13, middle left) shows an improvement of $\sim 8\%$ in statistical difference. The same prevalence of noise in the resulting completion is suffered here in Figure 13 as was suffered in the original hierarchical Pisa example (Figure 10). We see significant loss of surface detail but overall structural consistency. The changes in MSI statistics comparisons in this example may be equally attributable to the presence of this noise in addition to differences in the more global structure of the surface itself. By varying the levels of hierarchy introduced for each isolated part (Figure 13) we are able to see the differing effects this may be having (statistically, not visually) - with more variation present, despite visual similarity, where more levels of hierarchy have been introduced (i.e. Figure 13, middle right). However, as suitability of a given number of hierarchical levels to a given surface relief is difficult to determine this may vary on a per-surface basis and no universal conclusion can be drawn. Further investigation of this issue is left as an area for future work.

In addition we also re-introduce the candlestick completion example previously considered in Figure 3D. Here we maintain the same parameter settings as the original mono-scale completion (Figure 3D) but introduce three levels of hierarchy ($m = 3$) in a bid to allow hierarchical completion to overcome the rear-side structural anomaly shown in Figure 3D. As shown in Figure 14, a significant reduction in the statistical difference against the original ($\sim 9\%$, Table 1) is also achieved. *Due to local computational resource management the original mono-scale completions, for the Pisa Tower and candlestick examples, were computed on a dedicated multi-job, multi-user compute server. As a result their CPU timing results are not comparable to those presented here for the hierarchical completion approaches and as such they are omitted.*

The example of Figure 15 shows the application of hierarchical completion to the non-conforming plinth example originally from Figure 3C. Figure 15 shows that hierarchical completion offer a visual improvement over the mono-scale completion with the removal of the rear-side structural anomalies encountered previously (shown in Figure 3C). Statistically we see a minor variation in the surface ($< 1\%$, Table 1) with a substantial computational gain ($\sim 95\%$ over mono-scale completion, Table 1). Notably these results were achieved using a




Figure 14. Hierarchical completion of candlestick

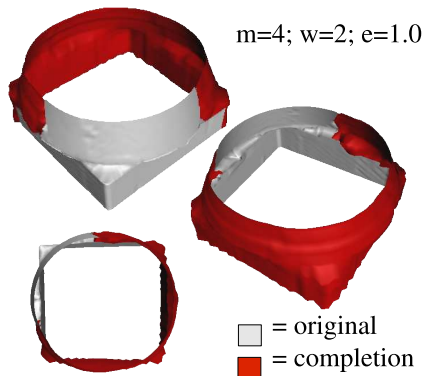


Figure 15. Hierarchical completion of non-conforming plinth

smaller hierarchical window size ($w = 2$, Figure 15) than required previously with mono-scale completion ($w = 6$, Figure 3C) and additionally a slight increase in error bound $e = 1.0$ was required to facilitate successful matching in the restricted case.

In our final example we return to quite a different type of surface than those previously examined in our evaluation - tree bark (Figure 2E). Here we are specifically interested in the completion of such a stochastic, natural surface relief texture using a methodology inherently suited towards highly structured relief types. Also we are interested in the additional effect that the use of restricted hierarchical completion may have on the plausibility of the resulting tree bark texture.

Maintaining the same window size and rotation parameters ($w = 3$, $r = 2$) as used in the original mono-scale completion of this surface example (shown Figure 3E) we apply hierarchical completion using a slightly increased error bound $e = 1.0$ to facilitate the coarseness of the lower levels of surface hier-

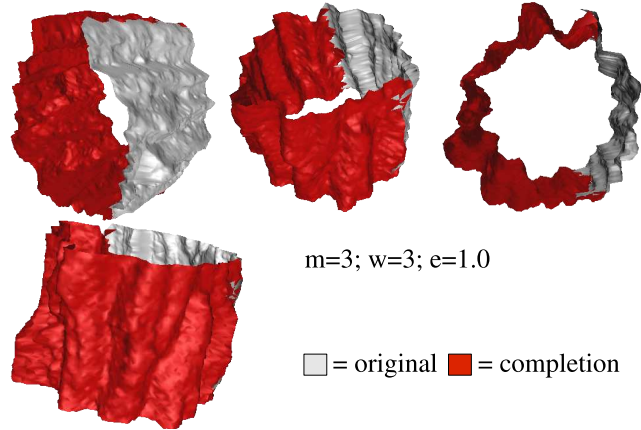


Figure 16. Hierarchical completion of tree bark

archy. From Figure 16 we see no immediate effects on visual plausibility from the use of either technique and the surface appears as visually plausible as the original mono-scale completion of Figure 3E. The careful observer may additionally argue that the hierarchical result appears slightly more plausible in terms of the apparent coherence of relief features (Figure 16) that are not immediately present in the earlier example (Figure 3E). This subjective visual opinion, which is difficult to verify for a stochastic surface such as this, is additionally supported by an improvement in the statistical difference with the use of hierarchical completion ($\sim 10\%$ improvement, Table 1). This may point to an underlying hierarchical nature for relief of this type.

Further investigation into the hierarchical completion of natural surface relief and possible underlying hierarchical traits in natural surface relief (and colour) is left as an area for future work. Overall, this example shows that the application of our hierarchical completion approaches is valid for stochastic “*non-structured*” such as tree bark and that restrictions on the diversity of the completion approach do not implausibly restrict the nature of the resulting completed surface. Indeed in this case hierarchical completion offers a $\sim 89\%$ improvement in computational cost (i.e. CPU run-time) over the mono-scale approach.

5 Summary

Overall we successfully propose a hierarchical extension to non-parametric 3D surface relief completion approach of [7] that facilitates the successful completion of 3D surface relief whilst maintaining the integrity of the global structure of the completed surface portion. This advances the prior state of the art in the automated completion of 3D surface relief detail and is successfully illustrated on a range of a range of real 3D surface scans of both artificial, structured and

natural surface relief types.

Notably, the introduction of such a hierarchical multi-scale approach allows issues of global surface structure to be resolved albeit at the potential expense of noise propagation for fine relief detail. In addition a considerable computational saving over prior work following this approach is illustrated [7]. Future work will investigate the effective attenuation of noise propagation with the hierarchical multi-scale surface relief synthesis and the use of surface completion via combined local and global learning approaches.

References

- [1] S. Arya and D. M. Mount. Approximate nearest neighbor queries in fixed dimensions. In *Proc. of the fourth Annual ACM-SIAM Symp. on Discrete algorithms*, pages 271–280, Philadelphia, PA, USA, 1993. Society for Industrial and Applied Mathematics.
- [2] G. H. Bendels, R. Schnabel, and R. Klein. Detail-preserving surface inpainting. In *Proc. 6th Int. Symp. on Virtual Reality, Archaeology and Cultural Heritage*, pages 41–48, 2005.
- [3] P.J. Besl and N.D. McKay. A method for registration of 3D shapes. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 14(2):239–256, 1992.
- [4] P. Bhat, S. Ingram, and G. Turk. Geometric texture synthesis by example. In *Proc. Eurographics/ACM SIGGRAPH symp. on Geometry processing*, pages 41–44, 2004.
- [5] T. P. Breckon. *Completing unknown portions of 3D scenes using 3D visual propagation*. PhD thesis, School of Informatics, University of Edinburgh, 2006.
- [6] T.P. Breckon and R.B. Fisher. Amodal volume completion: 3D visual completion. *International Journal of Computer Vision and Image Understanding*, 99(3):499–526, September 2005.
- [7] T.P. Breckon and R.B. Fisher. Three-dimensional surface relief completion via non-parametric techniques. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(12):2249 – 2255, December 2008.
- [8] U. Castellani, S. Livatino, and R.B. Fisher. Improving environment modelling by edge occlusion surface completion. In *Int. Symp. on 3D Data Proc. Vis. and Trans.*, pages 672–675, June 2002.
- [9] P. Chalmoviansky and B. Juttler. Filling holes in point clouds. In M.J. Wilson and R.R. Martin, editors, *Mathematics of Surfaces X*, number 2768 in LNCS, pages 196–212. Springer-verlag, 2003.

- [10] J. Davis, S. Marschner, M. Garr, and M. Levoy. Filling holes in complex surfaces using volumetric diffusion. In *Proc. First Int. Symp. on 3D Data Processing, Visualization, Transmission*, June 2002.
- [11] F. Dell’Acqua and R.B. Fisher. Reconstruction of planar surfaces behind occlusions in range images. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 24(4):569–575, 2002.
- [12] A.A. Efros and W.T. Freeman. Image quilting for texture synthesis and transfer. *Proc. 28th Conf. on Computer Graphics and Interactive Techniques*, pages 341–346, August 2001.
- [13] A.A. Efros and T.K. Leung. Texture synthesis by non-parametric sampling. In *Proc. IEEE Int. Conf. on Computer Vision*, pages 1033–1038, Corfu, Greece, September 1999.
- [14] H. Hoppe. Progressive meshes. In *Proc. of the 23rd Annual Conf. on Computer Graphics and Interactive Techniques*, pages 99–108, 1996.
- [15] M. Janaszewski, M. Couprie, and Babout L. Hole filling in 3d volumetric objects. *Pattern Recognition*, 43:3548–3559, 2010.
- [16] T.R. Jones, F. Durand, and M. Desbrun. Non-iterative, feature-preserving mesh smoothing. *ACM Trans. Graph.*, 22(3):943–949, 2003.
- [17] T. Ju. Robust repair of polygonal models. *ACM Trans. Graphics*, 23(3):888–895, 2004.
- [18] N. Kawai, T. Sato, and N. Yokoya. Efficient surface completion using principal curvature and its evaluation. In *Proc. International Conference on Image Processing*, pages 521–524, 2009.
- [19] V. Kraevoy and A. Sheffer. Template based mesh completion. In *Proc. 3rd Eurographics Symp. on Geometry Processing*, pages 13–22, July 2005.
- [20] Y.-K. Lai, S.-M. Hu, D.X. Gu, and R.R. Martin. Geometric texture synthesis and transfer via geometry images. *ACM Solid and Physical Modelling*, pages 15–26, 2005.
- [21] P. Liepa. Filling holes in meshes. In *Proc. of the Eurographics Symp. on Geometry Proc.*, pages 200–205, 2003.
- [22] D. Luebke, B. Watson, J.D. Cohen, M. Reddy, and A. Varshney. *Level of Detail for 3D Graphics*. Elsevier Science Inc., New York, NY, USA, 2002.
- [23] T. Masuda. Filling the signed distance field by fitting local quadrics. In *Proc. Int. Symp. on 3D Data Processing Visualization and Transmission*, pages 1003–1010, 2004.
- [24] F.S. Nooruddin and G. Turk. Simplification and repair of polygonal models using volumetric techniques. *IEEE Trans. Visualization and Computer Graphics*, 9(2):191–205, 2003.

- [25] M. Pauly, N. J. Mitra, J. Giesen, M. Gross, and L. Guibas. Example-based 3D scan completion. In *Symp. on Geometry Processing*, pages 23–32, 2005.
- [26] F. Policarpo, M. M. Oliveira, and J. L. D. Comba. Real-time relief mapping on arbitrary polygonal surfaces. *ACM Trans. Graph.*, 24(3):935–35, 2005.
- [27] A. Sharf, M. Alexa, and D. Cohen-Or. Context-based surface completion. *ACM Trans. Graphics*, 23(3):878–887, 2004.
- [28] F. Stulp, F. Dell’Acqua, and R.B. Fisher. Reconstruction of surfaces behind occlusions in range images. In *Proc. 3rd Int. Conf. on 3D Digital Imaging and Modeling*, pages 232–239, 2001.
- [29] G. Taubin. Curve and surface smoothing without shrinkage. In *Proc. of the Fifth Int. Conf. on Computer Vision*, page 852, 1995.
- [30] G. Taubin. Geometric signal processing on polygonal meshes. In *Proc. Eurographics*, 2000. State of the Art Report.
- [31] L.S. Tekumalla and E. Cohen. A hole-filling algorithm for triangular meshes. Technical Report UUCS-04-019, School of Computing, University of Utah, December 2004.
- [32] J. Verdera, V. Caselles, M. Bertalmio, and G. Sapiro. Inpainting surface holes. In *Proc. Int. Conf. on Image Processing*, volume 3, pages 903–906, 2003.
- [33] J. Wang and M.M. Oliveira. A hole-filling strategy for reconstruction in smooth surfaces in range images. In *16th Brazilian Symp. on Comp. Graphics and Image Proc.*, pages 11–18. IEEE Computer Society, 2003.
- [34] L. Wei and M. Levoy. Fast texture synthesis using tree-structured vector quantization. In *Proc. of the 27th Annual Conf. on Computer Graphics and Interactive Techniques*, pages 479–488, 2000.
- [35] L. Wei and M. Levoy. Texture synthesis over arbitrary manifold surfaces. In *Proc. of the 28th Annual Conf. on Computer Graphics and Interactive Techniques*, pages 355–360. ACM Press / ACM SIGGRAPH, 2001.
- [36] C. Xiao, W. Zheng, Y. Miao, Y. Zhao, and Q. Peng. A unified method for appearance and geometry completion of point set surfaces. *International Journal of Computer Graphics*, 23(6):433–443, 2007.
- [37] Kun Zhou, Xin Huang, Xi Wang, Yiyang Tong, Mathieu Desbrun, Baining Guo, and Heung-Yeung Shum. Mesh quilting for geometric texture synthesis. *ACM Trans. Graph.*, 25(3):690–697, 2006.