Eurographics Symposium on Geometry Processing 2010
Bruno Levy and Olga Sorkine
(Guest Editors)

Volume 29 (2010), Number 5

High Fidelity Scan Merging

J. Digne[1], J.-M. Morel[1], N. Audfray[2] and C. Lartigue[2]

[1]CMLA, ENS Cachan, CNRS, UniverSud, 61 Avenue du Président Wilson, F-94230 Cachan
[2] LURPA, ENS Cachan, Univ. Paris Sud 11, 61 Avenue du Président Wilson, F-94230 Cachan

Abstract

For each scanned object 3D triangulation laser scanners deliver multiple sweeps corresponding to multiple laser motions and orientations. The problem of aligning these scans has been well solved by using rigid and, more recently, non-rigid transformations. Nevertheless, there are always residual local offsets between scans which forbid a direct merging of the scans, and force to some preliminary smoothing. Indeed, the tiling and aliasing effects due to the tiniest normal displacements of the scans can be dramatic. This paper proposes a general method to tackle this problem. The algorithm decomposes each scan into its high and low frequency components and fuses the low frequencies while keeping intact the high frequency content. It produces a mesh with the highest attainable resolution, having for vertices all raw data points of all scans. This exhaustive fusion of scans maintains the finest texture details. The method is illustrated on several high resolution scans of archeological objects.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation—Digitizing and scanning

1. Introduction

Recent high precision triangulation laser scanners can scan surfaces of medium size objects with a precision of less than 10μ. Yet, although each scan has a very high precision, this precision can be lost again when merging multiple scans and meshing them together. This loss of precision entails a loss of visible texture, which explains the smooth and glassy aspect of most rendered scanned objects. On the other hand the merging of the multiple scans (often called *super-resolution*) is absolutely necessary. A patch of the object may well be acquired tens and even hundreds of times on well exposed parts. Indeed, many sweeps with varying trajectories are necessary to acquire the less exposed parts of the object. The main goal of the merging considered here is not to gain more detail and texture or to denoise the data point cloud by super-resolution: recent triangulation scanners yield scan sweeps with excellent quality. Unfortunately this quality is at risk of being damaged by the merging procedure itself. Thus, more trivially, the goal is to secure that the texture of each scan is not lost again due to slight matching errors which force a smoothing before a joint meshing. Fig. 1 illustrates the problem. With two overlapping shifted scan grids, as seen in (a), the aliasing risk is high. Meshing each scan separately yields two almost identical surfaces and textures (b, c). Nevertheless, a joint meshing (d) provokes strong tiling and aliasing effects, due to very small local offsets between both scans, in spite of the fact that they have been globally well registered. The challenge is therefore to merge both scans in such a way that the rendering quality does not decrease! The numerical problem is made more acute by two facts. First, not just two, but up to hundred scans may overlap in some region. Second, scans boundaries appear everywhere, as illustrated in fig. 2 and make the fusion near these boundaries still more problematic.

Each point of each scan has three-dimensional coordinates given either in a global coordinate system if the acquisition device is calibrated, or in a local coordinate system if the device is not calibrated. In the case of non-calibrated devices, the scans must be registered in a common coordinate system, and the registration problem becomes a rigid transform estimation. This problem has been widely investigated and has found efficient solutions [BM92] [RL01]. Yet if the scans had some internal local warping (which is usually the case), the rigid transform framework is not sufficient. A whole theory of non-rigid scan registration has therefore been developed [BR04], [BR07]. If the acquisition device is well calibrated the delivered scans are well registered, up to a given precision. Yet, as we already mentioned, a tiny resid-

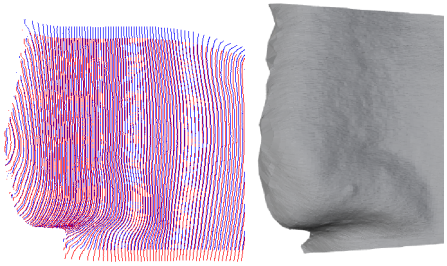

(a) Overlapping scans (b) Mesh of Scan 1

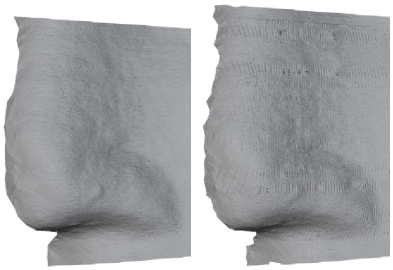

(c) Mesh of Scan 2 (d) Mesh of 1&2

Figure 1: *Example of two overlapping scans, points of each scanned are first meshed ((c)-(d)) separately. The result can be compared to the meshing of points of both scans together (d)*

ual mismatch can provoke strong artifacts similar to aliasing patterns (see fig. 1) and forbids a direct meshing of the union of all data point clouds. This problem is generally solved by applying a method which meshes an implicit zero level set of a distance function to the raw points. The distance function is approximated by its Fourier coefficients [Kaz05] or by radial basis functions [KBH06]. The problem is that these methods result in a serious loss of accuracy when the final result is compared with each scan separately.

This paper experiments on sets of scans of an object that have been either previously optimally registered by rigid or non-rigid methods, or registered through a high precision calibration of the acquisition tool. To demonstrate that no texture content will be lost, the goal is to mesh the complete point cloud. This means that all raw acquired points of all scans will be vertices of the mesh. This requirement guarantees a complete preservation of all the acquired information, including noise and fine textures. Of course such a mesh is not numerically economic, but it is necessary for two goals: first to get high quality rendering of complex shapes such as archeological objects, and second to precisely explore all remanent artifacts such as the holes, inherent in any scanning process. For scanning control purposes, it is anyway quite rewarding to be able to *see exactly* what has been scanned.

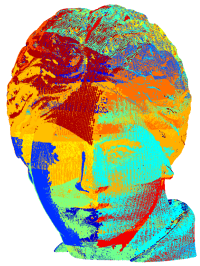

Figure 2: *Example of overlapping scans. This head is such a complex structure that not less than* 35 *scans were acquired to fill in most holes.*

2. Previous Work

2.1. Rigid Scan registration

When the scans contain no warp, the registration problem sums up to estimating a rigid transform between scan coordinate systems by minimizing the distance between the reference scan and the transformed scan. The estimation of the transform in the quaternion form was proposed in [Hor87]. The Iterative Closest Point (ICP) Registration procedure introduced in [BM92] and [CM92] matches a point from one scan to the closest point of the other scan and computes the transform based on these matchings. ICP was very successful and many variants were introduced ([RL01], [BS97]). A study of the optimal sampling for the ICP algorithm was proposed in [GIRL03].

Since ICP converges to a local minimum, the initial scan position must be close to optimal. Thus, some robust initial matches are needed to initialize the algorithm. This robust match search has also been investigated. The common idea is to find features easily identifiable on the scans (usually linked with highly curved points) and to match them. Spin Images, introduced in [JH99], represent potential feature points by images. The spin image associated with point p is the function that maps a 3D point to (α, β), where α is the distance to the tangent plane and β the distance to the line parallel to $\vec{n}(p)$ and passing through p. A linear correlation coefficient is used as a similarity coefficient between images to find the correspondences. In [SA01] another feature carrier is introduced. It is based on computing geodesic circles around a point and projecting them onto the tangent plane. This gives a $1D$ 2π-periodic function parameterized by the angle. Feature matching is done by sampling the contours and computing a similarity measure between the contours. Other popular descriptors were described in [YF99], [VSR01], [ZH99] or [KFR03].

Another approach was proposed in [AMCO08]: all coplanar 4-point sets that are approximately congruent are extracted on both shapes and matched using the fact that the distance ratios relatively to the intersection point are invariant to rigid motions.

Rigid scan registration methods assume that the scans must fit perfectly using only rigid transforms. Yet, if the scans have some warping, the method does not apply. Another problem is that the registration error might cumulate when registering multiple scans (see [BR07] for examples of bad registrations). These considerations led to modeling the registration transform by a non rigid transform as will be seen in the next subsection.

2.2. Non-rigid scan registration

The method for non rigid registration using thin plate splines [BR04] first applies a hierarchical ICP to find good features: the source mesh is iteratively divided through the middle of its longest axis and each half is realigned separately. This yields good feature correspondences at the cost of substantial discontinuities in the source mesh. These point matches are then used to compute the thin plate spline that best approximates all pairs of points. The spline being continuous, any discontinuity introduced in the scan splitting process is removed by the spline approximation. Once the scans are registered they must be merged. This is done using the VRIP method [CL96], which will be described below.

An extension of this method was introduced in [BR07]: the same thin plate spline non rigid deformation model is applied using an improved earlier point matching: features are found and matched in a process which rejects outliers. Other methods include [CR03] where the non rigid registration is done using as input soft assignment between point pairs. This yields a functional minimization comprising a term of soft assignment.

2.3. Super-resolution from several scans

Recently, the problem of achieving super-resolution from multiple scans has been raised, mostly for improving range image resolutions ([KMA06]) where various low resolution scans with the same depth direction are acquired and registered by ICP. Depth values at each point of a high resolution grid are then interpolated with depth values of points falling in the neighboring cells. [AKSA09] used rotating scans around the depth axis to build for each orientation high resolution range images. These high resolutions range images were registered and combined by considering the image gradient and the angle between the baseline and the image gradient to weight each point. Other methods building hybrid scanners to achieve high resolution include [NRDR05] where positions and normals are acquired and used to improve the resolution.

Once scans have been computed and registered, a mesh must be built to allow a fast visualization of the surface model. The next section reviews methods to build the mesh.

2.4. Meshing

Meshing methods can be divided into two categories: methods that approximate the point cloud and methods that mesh directly the point set. Approximating methods usually build a function defined on $\mathbb{R}^3$ whose 0-level set is the shape surface. A mesh is then built on the 0 level-set by the marching cubes algorithm [LC87]. These methods include [KBH06], [Kaz05], among others. A very interesting variant of these methods is the VRIP algorithm ([CL96]). VRIP considers an implicit function taking into account not only point positions but also their reliability. Nonetheless, two drawbacks common to the mentioned methods are the automatic filling in of holes, and the implicit low pass filtering performed by the level set method. These methods usually compute the distance to the surface as an average of the signed distance of the point to its k-nearest neighbors. Thus initial points are forgotten and *de facto* replaced by local averages. This removes noise in the cloud, but also loses fine details and textures.

Direct meshing methods include [ABK98] or [AB98]. We shall use the incremental ball-pivoting method [BMR*99], which is fast and does not fill holes. The method is based on pivoting a ball of fixed radius r around edges. Three points are triangulated if they lie on a ball with radius r and empty interior. The ball is then pivoted around all three edges of the triangle, until it meets a point and has still empty interior. If no such point is met then the edge is a hole border. The parameter r is a bound for the creation of triangles: no triangle edge can have length above $2r$. Thus low density areas are considered as holes.

In short, the dilemma is this: the approximating methods are not sensitive to a slight registration error, but can lose detail and texture. On the contrary, aliasing-like artifacts become visible when a direct meshing method attempts to keep all raw points of several scans. The best choice is to preserve the raw points and to apply a direct meshing method. But this requires eliminating all traces of inaccurate registration.

3. Scan Merging

3.1. Principle

The general idea behind the scan merging method experimented here is to preserve point positions in non overlapping areas, and to make a fusion of the scans on overlapping regions while keeping all raw points. The fusion involves a smooth-base/height-function decomposition for each scan. The decomposition of a surface as the sum of a smooth base and of a height function was proposed for a different purpose in [KST09], and [ZTS09], where the height function was used to segment the mesh and extract features as contours of the height function. The underlying idea is that a surface S can be decomposed into a smooth base B and a height function h, so that:

$$S = B + h$$

B can be seen as the low frequency surface and h can be seen as the high frequency term. Given several surfaces $S_1 = B_1 + h_1, S_2 = B_2 + h_2, \cdots, S_N = B_N + h_N$, the idea is to fuse the bases, but to keep exactly the h_i terms, thus preserving all fine details. In other terms, a common basis B for all surfaces must be found, the high frequencies of all scans adopting this common basis thereafter. This strategy is comparable to the one used for morphing applications in [PKG06].

The data merging using a high/low frequency decomposition has long been a classic method in image processing [BA83]. This article introduced the idea of separating each image into various frequency bands by a Gaussian pyramid. The low frequency bands were merged separately to obtain a smooth blending of different images. The method has been successfully used to create panoramas from multiple images [BL07] and texture 3D models [Bau02]. Two major differences are that in [BA83] all frequency bands are merged, whereas the method proposed here only merges the low frequencies while keeping the high frequencies intact. Another important difference is the usage of a nonlinear heat equation instead of a linear frequency decomposition.

The next section addresses the robust decomposition of a surface into a base and a height.

3.2. Low/High frequency surface decomposition

Since the pioneering article [Tau95] it is known that mesh high frequencies are removed by the application of the intrinsic heat equation $\frac{\partial x}{\partial t} = \Delta x$. Yet, our scanned surfaces are given as point clouds and not as meshes. A numerical scheme of the heat equation for raw point clouds must be used. This question has been addressed in [BSW09] and [PKG06]. We shall use the simple implementation of the intrinsic heat equation proposed in [DMMSL09]: this paper proves that the intrinsic heat equation can be implemented by iterating a projection of each point onto the local cloud regression plane. Consider the projection operator T_r that projects each point p onto the regression plane of the neighbors of p enclosed in a ball of radius r. Then it can be proven that this motion is tangent to the intrinsic heat equation. The iteration of T_r yields a scale space (a representation of the shape at various smoothing scales). In all experiments r is set so that the ball $B(p,r)$ contains about 30 points at almost all points, and the number of scale space iterations n is set to 4. The first parameter (30) is fixed so that a reliable regression plane is always computed. The second parameter, namely the number of iterations 4, is chosen to guarantee a smooth enough basis in all cases. It can be increased without damage. When iterating the projection operator with an initial surface S_0, the surface S_t is iteratively smoothed. To each point p_t of S_t corresponds a point p_0 of S_0, and the height function can be taken to be the vector $h(p_t) = p_t - p_0$.

An alternative definition for the height would be the scalar function $h(p_t) = (p_t - p_0) \cdot n(p)$. Yet, the results with both height variants being fairly identical, the simplest definition was kept: it separates each data point into a smooth base point and a high frequency vector.

Figure 3: *The unmerged head with aliasing artifacts (left), its smooth base (middle) and the merged result (right)*

3.3. Finding a common smooth basis for all surfaces

Choosing a common basis for all scans is the next question. A natural constraint on the method is to keep fixed the points belonging to regions where only one scan is available. Finding the common basis then becomes straightforward: It is enough to apply the same number n of iterations of T_r with the same parameter r to all the sets after they have been put together. This global filtering assumes that the high frequency term of the set $S = \cup_i S_i$ contains the registration error: when filtering S the registration error is filtered away (see fig. 3).

3.4. Algorithm

The method is summarized in Algorithm 1. The algorithm is based on two applications of the intrinsic heat equation scheme (here the iterated projection on the regression plane) with the same parameters and the same number of iterations. All registered scans are given in the same global coordinate system. The first application (Line 2) is done on the separate scans yielding the intrinsic high frequencies of each scan. The second application (Line 6) is done on all scans together. When filtering all scans together (lines 5 and 6) the registration error is suppressed and we get a common low scale registration or basis, the set of points $b(p)$. Adding back to them the high frequency component $\overrightarrow{b_i(p)p}$ restores all details of all scans.

An important feature of the method is that each region $\mathcal{A}$ of the shape that has been acquired by one scan only is not altered. Indeed, inside such a region, applying the separate scale space or the common scale space is strictly equivalent, since there is only one scan in the neighborhood of the points of $\mathcal{A}$. Then the point is first filtered to $b_i(p) = b(p)$, and therefore moved back to its original position p at Line 8. So in areas with only one scan, point positions are not changed. The only effect of the algorithm is the merging of overlapping scans.

Algorithm 1: Scale Space Merging

Data: N point sets (scans) $(S_i)_{i=1\cdots N}$, a number of projection filter iterations n and a radius r
Result: Merged scans, Q

for $i = 1 \cdots N$ **do**
  Apply n steps of the projection filter T_r to the set S_i ;
  Store for each point $p \in S_i$ with corresponding filtered point $b_i(p)$ the high frequency vector $\vec{\delta}(p) = \overrightarrow{b_i(p)p}$;
end
$S \leftarrow \cup_{i=1}^{N} S_i$;
Apply for each $p \in S$ n steps of the projection filter T_r, yielding a point $b(p)$;
for $p \in S$ **do**
  $q = b(p) + \vec{\delta}(p)$;
  Add q to Q;
end
Return Q;

RMSE	Both lines	Line A	Line B
Before Merging	X	$9.95e-04$	$9.76e-04$
After Merging	$9.85e-04$	$9.94e-04$	$9.75e-04$

Figure 4: *Noise estimates on each separate scan A and B before and after their merging*

3.5. One-dimensional study

It is easy to illustrate the method in 1-D on simple $1D$ shapes. Our goal was to check that the proposed method superposes two simulated scans without any smoothing effect. To do so, two noisy straight lines A and B were synthesized from the same model and then merged by the algorithm. The noise of each set A, B, $A \cup B$ was estimated as the root mean square error to their regression lines before and after merging. The results in tab. 4 show that the merging did not cause any denoising. Indeed, the RMSE does not decrease by the merging procedure. Fig 5 shows another 1D example of the merging procedure where the bases are actually slightly different, in accordance with the real situation encountered on real scans.

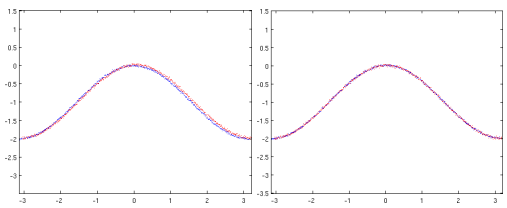

(a) Without Merging (b) With Merging

Figure 5: *Two noisy sine functions before and after merging*

4. Implementation and Results

The inputs of the merging algorithm were the outputs of our laser triangulation scanner. This device being accurately calibrated, the scans were in principle already registered so that no extra software registration was needed. Nevertheless, the ICP algorithm was applied to see if it could remove the aliasing and tiling artifacts. It didn't. The positions computed by ICP oscillated around the input scan positions, and the resulting meshes were no better. The registration process was implemented on a 1.5 Ghz processor with $48GB$ RAM. An octree structure was first built to allow for fast access to the neighbors of each given point. Table 6 gives the computation times for various shapes of various sizes with varying numbers of scans. Notice the high number of scans necessary to get a good covering of the object. It entails that several dozens of scans have to be merged on the more exposed parts.

Point set	points	scans	Time(s)	height
Dancer	5524627	94	321	$17cm$
Greek Mask	8961736	78	106	$12cm$
Nefertiti	15554528	115	819	$18cm$
Tanagra	17496999	160	1258	$22cm$

Figure 6: *Computation time for the proposed merging. It is significantly faster than the scanning time itself*

Figs 7, 8 and 9 present the results on these data. For all point sets, two different renderings will be displayed: the first one is a ball pivoting [BMR*99] meshing of all raw scan points without any merging. The scans were preregistered by the calibrated acquisition device and no software post-registration was needed. The second rendering is again a ball-pivoting meshing, but applied to the merged point set. The rendering was made using the POV-RAY ray-tracer. The conclusion is common to all experiments: even if the scans are actually very accurately registered, the tiny warps of the grids always create some aliasing visible as grid or tiling effects. After the merging procedure (which only slightly affects the low frequencies), these undesirable effects disappear almost completely. In the procedure more than 99.9% of the raw points were kept. Thus, the final result indeed is highly faithful to the raw scan. Yet a careful attention shows some remains of aliasing (Fig. 8, last column). The area of these is actually small, being inferior to the area of the holes. They could easily be removed by a selective local smoothing. Some of the bigger pieces, like Nefertiti, show no defect at all.

Comparison To better judge the texture preservation, the rendering of a scan alone (ground truth) was computed and compared to the rendering of all scans in the same region on Fig 10. This shows that the visual information loss after scale space merging is very low compared to the one due to a simple joint meshing.

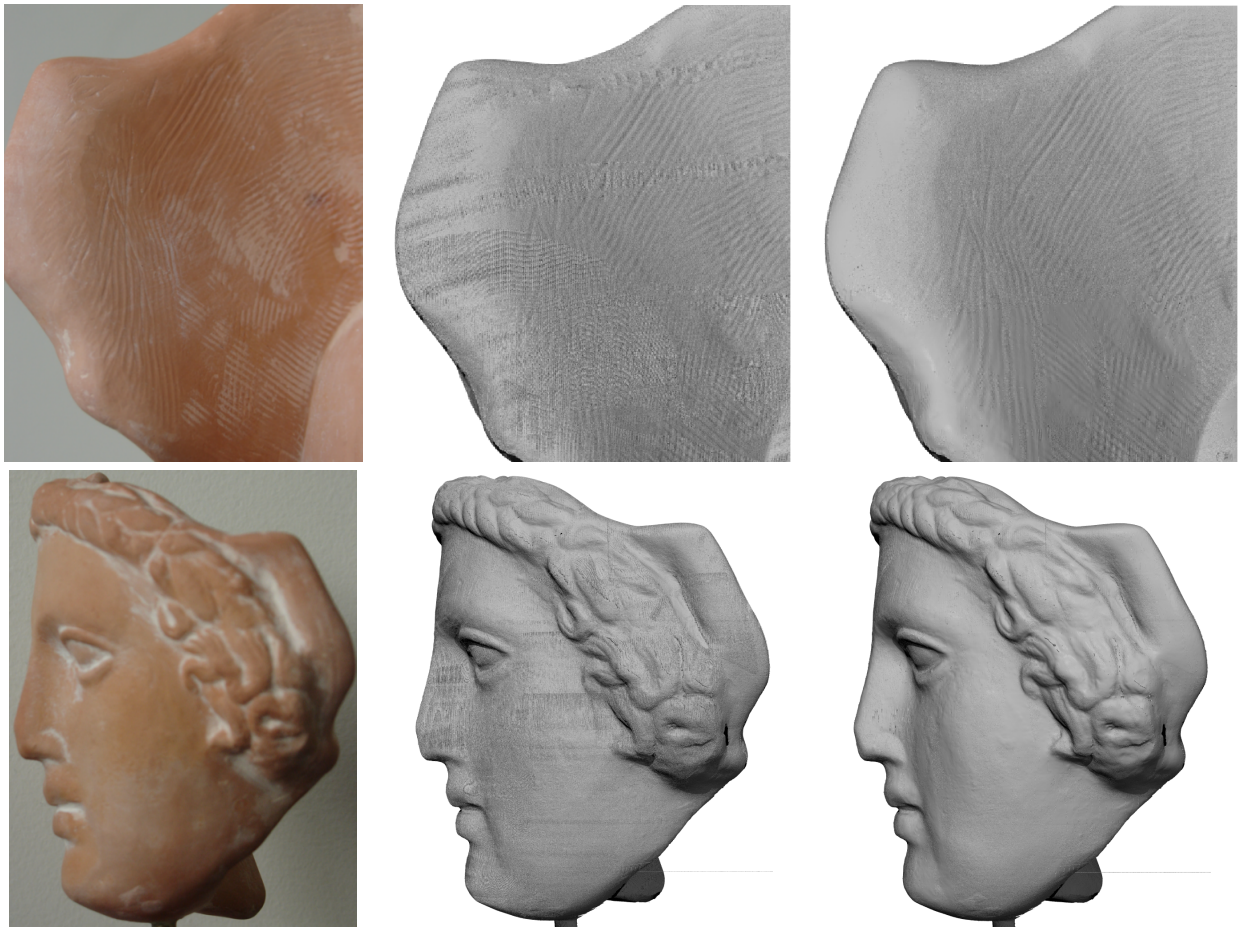

Figure 7: *Merging of the mask scans seen from the back side (top row) and left side (bottom row). Left: picture, middle: without merging, right: with merging*

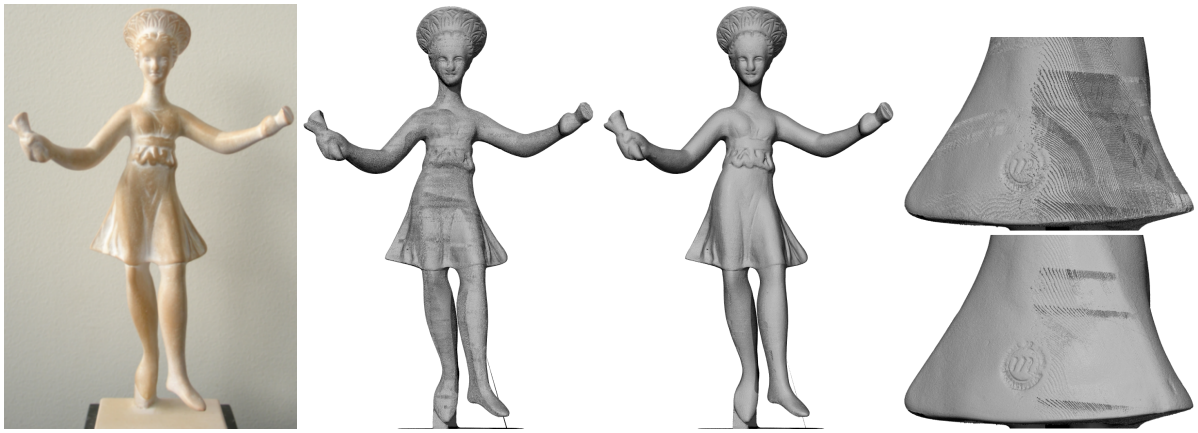

Figure 8: *Merging of the Dancer With Crotales. From left to right: picture, without merging, with merging, an example of merging failure taken from the back of the object (top: unmerged, bottom merged)*

Figure 9: *Merging of the Nefertiti (1st: picture, 2nd,4th: without merging, 3rd,5th: with merging)*

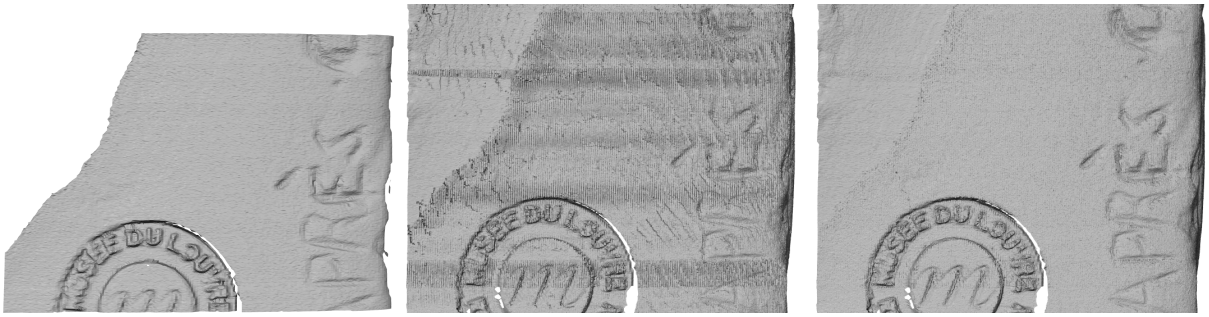

Figure 10: *Comparison of a rendering of a single scan (ground truth) and the merging of all scans that overlap in the same region. (Left: ground truth, middle: joint mesh of all scans without merging, right: joint mesh with merging)*

It is crucial to compare the raw merging method results with results obtained by the level set reconstruction method of the unmerged scans point set. The result of the level set method applied to the Tanagra head (fig. 11 b), obviously introduces an important smoothing and loses texture in comparison to the merging result (fig. 11, a). But even with that smoothing the result still keeps several artifact lines due to the scan offsets: these offsets become visible at the scans boundaries. See the nearly straight long lines on the surface, mostly vertical and horizontal. It can also be asked if an efficient denoising method could actually restore the raw set. Fig. 11-c, shows the result of the application of the bilateral filter [FDCO03] to the union of the scans. This iterated filtering was applied up to the point where aliasing artifacts were no more visible. Clearly, this entails a much too strong smoothing of detail and texture.

The scan merging is a very local method which is therefore computationally efficient (see Tab. 6). Yet, if the input data are not already well registered the merging could obviously fail. The method corrects the slight misalignments only in the normal direction. A tangential drift in the original registration could therefore cause a loss of sharpness or a loss of small details. Nevertheless, this degradation seems to pass unnoticed. Indeed, for last generation triangulation scanners like the one used for the experiments in this paper, the registration error is very small. For a point cloud with side-length $99mm$ the observed average point offset after merging was $0.081mm$, with standard deviation 0.012. The tangential offset could not be measured. The explanation of the relative visual success of the method is that even a tiny normal offset causes a dramatic change in triangles orientation, and therefore completely jeopardizes the visual quality of the triangulation. An equally small tangential offset seems to be visually undetectable. Thus, the merging method corrects the normal error, and makes the tangential error unnoticeable.

The proposed merging can be seen as a local non rigid registration. Therefore it can be compared to the result given by state of the art non rigid registration methods [BR07]. To perform the comparison, the problem arose that the scans did not systematically contain strongly identified features. Most scans of the mask point set were simply rejected by the non

(a) Merged Result

(b) Poisson Reconstruction (c) Bilateral filter

Figure 11: *Comparisons of the merging (a) with a level set reconstruction method ([KBH06]) of the unmerged scans point set (b) and a filtering of the unmerged scans point set (c). The level set method obviously introduces a serious smoothing, yet does not eliminate the scanning boundary lines. The bilateral filter, applied until all aliasing artifacts have been eliminated, over-smoothes some parts of the shape.*

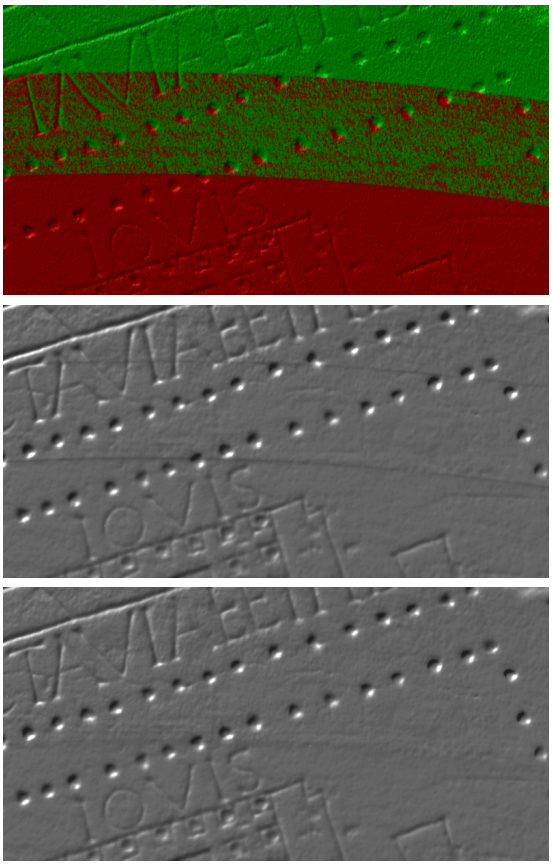

Figure 12: *Comparison of registration of two scans (colored in different colors on the top figure) using Global Non Rigid Alignment [BR07] (middle) and scale space merging (bottom). Meshes were reconstructed using [KBH06].*

rigid registration method described in [BR07]. In order to perform a serious comparison anyway, two sweeps of the fragment 31u of the Stanford FUR project were used. The computation times were, however, considerably different: it took more than 2*h*30 to register non rigidly these meshes. On the same computer, using only the raw points and not the meshes, the merging took only 84*s*. The final meshes were built using Poisson Reconstruction [KBH06] in both cases. The registration artifacts (two horizontal lines limiting the overlap area, fig 12) are much less visible with the scan merging than with the non rigid registration.

5. Conclusion

The main conclusion of the study is that it is possible to fuse multiple raw scans with minimal accuracy loss, provided an accurate previous registration has been performed. Future work will focus on the detection and handling of remaining holes, and on the automatic assessment of surface quality to replace a visual inspection.

Acknowledgments The authors would like to thank Prof. Marc Levoy for giving permission to use the fragment of the Stanford Forma Urbis Romae Project and Dr. Benedict Brown for his non rigid alignment implementation.

References

[AB98] AMENTA N., BERN M.: Surface reconstruction by voronoi filtering. In *SCG '98: Proceedings of the fourteenth annual symposium on Computational geometry* (New York, NY, USA, 1998), ACM, pp. 39–48.

[ABK98] AMENTA N., BERN M., KAMVYSSELIS M.: A new voronoi-based surface reconstruction algorithm. In *SIGGRAPH '98* (New York, NY, USA, 1998), ACM, pp. 415–421.

[AKSA09] ABBASINEJAD F., KIL Y. J., SHARF A., AMENTA N.: Rotating scans for systematic error removal. *Computer Graphics Forum 28* (2009), 1319–1326.

[AMCO08] AIGER D., MITRA N. J., COHEN-OR D.: 4-points congruent sets for robust surface registration. *ACM Transactions on Graphics 27*, 3 (2008), #85, 1–10.

[BA83] BURT P. J., ADELSON E. H.: A multiresolution spline with application to image mosaics. *ACM Trans. Graph. 2*, 4 (1983), 217–236.

[Bau02] BAUMBERG A.: Blending images for texturing 3d models. In *Proc. Conf. on British Machine Vision Association* (2002), pp. 404–413.

[BL07] BROWN M., LOWE D. G.: Automatic panoramic image stitching using invariant features. *Int. J. Comput. Vision 74*, 1 (2007), 59–73.

[BM92] BESL P. J., MCKAY N. D.: A method for registration of 3-d shapes. *IEEE Trans. Pattern Anal. Mach. Intell. 14*, 2 (1992), 239–256.

[BMR*99] BERNARDINI F., MITTLEMAN J., RUSHMEIER H., SILVA C., TAUBIN G.: The ball-pivoting algorithm for surface reconstruction. *IEEE TVCG 5* (1999), 349–359.

[BR04] BROWN B., RUSINKIEWICZ S.: Non-rigid range-scan alignment using thin-plate splines. In *3DPVT'04* (2004). printed.

[BR07] BROWN B., RUSINKIEWICZ S.: Global non-rigid alignment of 3-D scans. *ACM Transactions on Graphics (Proc. SIGGRAPH) 26*, 3 (Aug. 2007).

[BS97] BENJEMAA R., SCHMITT F.: Fast global registration of 3d sampled surfaces using a multi-z-buffer technique. In *Image and Vision Computing* (1997), pp. 113–120.

[BSW09] BELKIN M., SUN J., WANG Y.: Constructing laplace operator from point clouds in rd. In *Proc. SODA '09* (Philadelphia, PA, USA, 2009), SIAM, pp. 1031–1040.

[CL96] CURLESS B., LEVOY M.: A volumetric method for building complex models from range images. In *SIGGRAPH '96* (New York, NY, USA, 1996), ACM Press, pp. 303–312.

[CM92] CHEN Y., MEDIONI G.: Object modeling by registration of multiple range images. *Image Vision Comput. 10*, 3 (1992), 145–155.

[CR03] CHUI H., RANGARAJAN A.: A new point matching algorithm for non-rigid registration. *Comput. Vis. Image Underst. 89* (2003), 114–141.

[DMMSL09] DIGNE J., MOREL J.-M., MEHDI-SOUZANI C., LARTIGUE C.: Scale space meshing of raw data point sets. preprint CMLA 2009-30 - ENS Cachan, October 2009.

[FDCO03] FLEISHMAN S., DRORI I., COHEN-OR D.: Bilateral mesh denoising. *ACM Trans. Graph. 22*, 3 (2003), 950–953.

[GIRL03] GELFAND N., IKEMOTO L., RUSINKIEWICZ S., LEVOY M.: Geometrically stable sampling for the icp algorithm. In *Proc. 3DIM 2003* (2003), pp. 260–267.

[Hor87] HORN B. K. P.: Closed-form solution of absolute orientation using unit quaternions. *Journal of the Optical Society of America A 4*, 4 (1987), 629–642.

[JH99] JOHNSON A. E., HEBERT M.: Using spin images for efficient object recognition in cluttered 3d scenes. *IEEE PAMI 21* (1999), 433–449.

[Kaz05] KAZHDAN M.: Reconstruction of solid models from oriented point sets. In *SGP '05* (Aire-la-Ville, Switzerland, 2005), Eurographics Association, p. 73.

[KBH06] KAZHDAN M., BOLITHO M., HOPPE H.: Poisson surface reconstruction. In *SGP '06* (Aire-la-Ville, Switzerland, 2006), Eurographics Association, pp. 61–70.

[KFR03] KAZHDAN M., FUNKHOUSER T., RUSINKIEWICZ S.: Rotation invariant spherical harmonic representation of 3d shape descriptors. In *SGP '03* (Aire-la-Ville, Switzerland, 2003), Eurographics Association, pp. 156–164.

[KMA06] KIL Y., MEDEROS B., AMENTA N.: Laser scanner super-resolution. Point Based Graphics.

[KST09] KOLOMENKIN M., SHIMSHONI I., TAL A.: On edge detection on surfaces. In *CVPR* (2009), pp. 2767–2774.

[LC87] LORENSEN W. E., CLINE H. E.: Marching cubes: A high resolution 3d surface construction algorithm. In *SIGGRAPH '87* (New York, NY, USA, 1987), ACM Press, pp. 163–169.

[NRDR05] NEHAB D., RUSINKIEWICZ S., DAVIS J., RAMAMOORTHI R.: Efficiently combining positions and normals for precise 3d geometry. *ACM Trans. Graph. 24*, 3 (2005), 536–543.

[PKG06] PAULY M., KOBBELT L. P., GROSS M.: Point-based multiscale surface representation. *ACM Trans. Graph. 25*, 2 (2006), 177–193.

[RL01] RUSINKIEWICZ S., LEVOY M.: Efficient variants of the icp algorithm. In *Proc. 3DIM 2001* (2001), pp. 145–152.

[SA01] SUN Y., ABIDI M. A.: Surface matching by 3d point's fingerprint. In *ICCV 2001. Proceedings.* (2001), vol. 2, pp. 263–269 vol.2.

[Tau95] TAUBIN G.: A signal processing approach to fair surface design. In *SIGGRAPH '95* (New York, NY, USA, 1995), ACM Press, pp. 351–358.

[VSR01] VRANIĆ D. V., SAUPE D., RICHTER J.: Tools for 3d-object retrieval: Karhunen-loeve transform and spherical harmonics. In *Multimedia Signal Processing, 2001 IEEE Fourth Workshop on* (2001), pp. 293–298.

[YF99] YAMANY S. M., FARAG A. A.: Free-form surface registration using surface signatures. In *ICCV '99* (Washington, DC, USA, 1999), IEEE, p. 1098.

[ZH99] ZHANG D., HEBERT M.: Harmonic maps and their applications in surface matching. In *IEEE (CVPR '99)* (1999), vol. 2.

[ZTS09] ZATZARINNI R., TAL A., SHAMIR A.: Relief analysis and extraction. *ACM Trans. Graph. 28*, 5 (2009), 1–9.