Non-Rigid Global Alignment Using Thin-Plate Splines

Benedict J. Brown
Princeton University
bjbrown@cs.princeton.edu

Szymon Rusinkiewicz
Princeton University
smr@cs.princeton.edu

Abstract

A key challenge in reconstructing high-quality 3D scans and image mosaics is registration of data from different viewpoints. Existing global alignment algorithms are restricted to rigid-body transformations, and cannot adequately handle non-rigid warps in the data. Algorithms that can compensate for such warps between pairs of scans do not easily generalize to the multiview case. We present an algorithm for obtaining a globally-optimal alignment of multiple overlapping datasets in the presence of low-frequency non-rigid deformations, such as those caused by device nonlinearities or calibration error. The process first obtains sparse correspondences between views using a locally-weighted stability-guaranteeing variant of iterative closest points (ICP). Global positions for feature points are found using a relaxation method, and the scans are warped to their final positions using thin-plate splines. The technique is efficient, with only minimal overhead beyond comparable rigid-body global alignment techniques. We demonstrate that, relative to rigid-body registration, it improves the quality of alignment and better preserves detail in large 3D-scanned meshes, range images obtained using photometric stereo, and image sequences obtained with an uncalibrated camcorder.

1. Introduction

Range scanning is a common method for acquiring three-dimensional models of real objects. Because a range scanner can see only a single viewpoint at a time, several scans must be merged to obtain a complete model. In the case of small objects either the scanner or the object can be moved using only calibrated motions—by placing the object on a motorized turntable for example—and the relative alignments of each range scan are known. However, this is not possible for large objects such as Michelangelo's statues, nor when the top and bottom of an object need to be scanned as well. In these cases, the range scans must be aligned after acquisition. Because they must be merged and retriangulated to obtain a final mesh, even small misalignment leads to smoothing of high-frequency details in the merged mesh. Larger misalignments result in corrupt reconstructions.

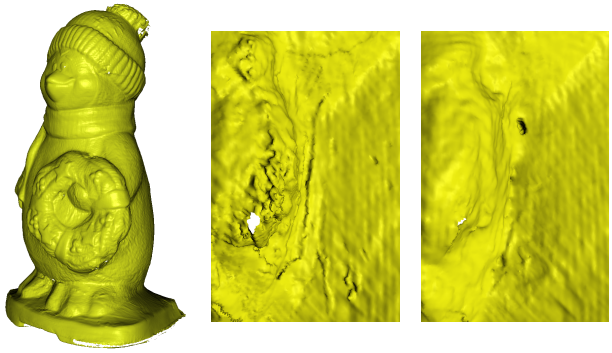

Figure 1. (Left) A penguin model, scanned using photometric stereo. The 15 individual range scans are warped due to normal integration error. (Center) A closeup of the area behind the wreath, aligned and merged using rigid-body ICP alignment. (Right) Aligned and merged using our non-rigid global alignment.

Although existing global registration techniques assume that the scans can be aligned using only rigid-body transformations, there are several practical situations in which this assumption is invalid. For example, in the case of data acquired by the Digital Michelangelo Project [16], the desired precision of approximately .25 mm combined with viewing volumes up to 7.5 m high made perfect scanner calibration unattainable, leading to warping on the order of several millimeters. In the case of range data obtained using shape-from-shading or photometric stereo techniques [25], limitations on the modeling of materials and lighting cause bias in estimating normals, which in turn leads to warping when the normals are integrated to obtain geometry.

We present two major contributions in our paper. The first is a novel framework for global alignment which can align simultaneously align many data sets based on feature correspondences in the presence of low-frequency warp. In the resulting alignment, we minimize the amount of warp applied to each data set, and distribute this warp evenly

across the sets. The second contribution is a feature correspondence algorithms for range scans which robustly handles low frequency warp while retaining most of the speed of rigid-body alignment techniques.

Algorithm Overview We introduce a framework for global alignment that compensates for low frequency warp using thin-plate splines. A set of global feature points are selected, and their positions are found on each range scan, using a new weighted ICP algorithm. The final positions of each feature point are computed using an efficient global optimization process that operates on the feature points themselves. These are then used to compute a non-rigid thin-plate spline warp for each range scan, which preserves high-frequency detail while deforming the range scan exactly enough to bring it into alignment with the other scans. Performance is proportional to the product of the number of range scans with the degree of overlap, and the algorithm is highly parallelizable. For datasets in which the degree of overlap is high, we propose an alternative approach that reduces running time by selecting feature points only on the minimum set of scans necessary to cover the object.

Further Applications The global alignment framework is general, and applies to any problem where multiple data sets must be simultaneously aligned, such as panoramic image construction and alignment of repeated scientific measurements. Our weighted ICP is also a flexible feature correspondence algorithm, and works well for images as well as range scans. As an example, we show our algorithm applied to a video sequence to produce a panoramic image. With the same framework as we use for range scans, we obtain panoramic mosaics from inexpensive, uncalibrated cameras, in the presence of unknown lens distortion, auto-focus, auto-exposure, and variation in perspective and depth of field.

2. Previous Work

2.1. Iterative Closest Points

The most common form of range-scan alignment is Iterative Closest Points (ICP), originally introduced by Besl and McKay [4] and Chen and Medioni [8]. ICP begins with an initial estimate of the alignment, obtained by a method such as tracking scanner position, matching local surface descriptors [15, 11], exhaustive search for corresponding points [7], or user input. The alignment is refined by repeatedly selecting points on one or both models, finding closest points in the other mesh, and computing the rigid body transform that minimizes the least squares error between the two point sets. Gelfand *et al.* [12] and Ikemoto *et al.* [14] extend this to select feature points that will constrain all degrees of freedom in the rigid-body transform, improving alignment performance and quality. Because rigid-body transforms have only six parameters, they are efficient to compute and apply, and do not overfit the (possibly) poor correspondences in early iterations. On the other hand, the algorithm relies on these properties, and is not easily extended to more complex transforms.

Global Rigid Body Alignment Several extensions to rigid-body ICP for global alignment have been proposed. Chen and Medioni [8] perform global alignment by incrementally aligning new range scans to all previous ones. Since error accumulates with successive scans, this generally proves unsatisfactory, and more recent algorithms focus on distributing the error more evenly across scans. Bergevin *et al.* [2] iteratively perform incremental alignment, until the individual alignment errors stabilize and are well-distributed. Pulli [20] selects a stable set of global feature correspondences, then repeatedly aligns single scans to all others using these features until the system converges. Neugebauer [19], Lu and Milios [17] and Williams and Bennamoun [24] all solve for the alignment of all scans simultaneously.

2.2. Non-Rigid Alignment Techniques

Ikemoto *et al.* [14] address the rigid-body constraint of ICP by dicing each range scan into overlapping pieces, and performing global rigid-body alignment on the pieces. This can address low-frequency warps, but causes smoothing and has running time quadratic in the amount of dicing performed. The Florentine Pietà Project [3] uses ICP alignment but also incorporates image-based alignment and conformance smoothing of overlapping scans along scanner lines-of-sight to improve the registration further. Hänel *et al.* [13] present an extension to ICP that allows deformable objects to be aligned by computing rigid-body warps on different parts of a jointed skeleton. Allen *et al.* [1] use an affine transformation at each vertex of the source mesh to allow non-rigid registrations of full-body scans to a high-resolution template. In contrast, we assume a low-frequency deformation which allows us to more efficiently compute a warp using many fewer parameters.

Thin-Plate Splines The thin-plate spline that we use to represent our warp has been used extensively in medical imaging applications, where its use was first proposed in Bookstein [5]. Chui and Rangarajan [9] use a softassign/deterministic annealing framework to iteratively compute point correspondences and align both medical and non-medical data in 2-D and 3-D using thin-plate splines. Rather than assigning features based on closest points as

in ICP, each pair of points is assigned a probability of corresponding based on a Gaussian function of their distance from each other (softassign). This framework works well in many situations, but we have found it to be less successful for range scan alignment. Our range scans typically contains hundreds of thousands of points, so—just as in ICP—a small subset of points from the source scan must be used. As a result, the set of potential correspondences (which is *every* point in the target scan) consists almost entirely of outliers, and the process is too unconstrained to converge on correct correspondences (ICP avoids this problem precisely by using a heavily constrained transformation). Also, because the thin-plate spline must be recomputed and applied at every iteration, a pairwise alignment requires minutes at best. This is fine in the pairwise case, but is impractically slow to incorporate in a global alignment framework.

More recently, the thin-plate spline has been used for *pairwise* registration of large range scans [6]. Feature correspondences are obtained by dicing one range scan hierarchically, and aligning each piece separately to the other range scan using ICP. The method provides improved pairwise alignments, but includes no mechanism for computing any globally consistent warping among many range scans. Also, the discontinuities between piecewise alignments in the feature correspondence computation are apparent in the structure of the spline warp.

3. Global Alignment Framework

The goal of global alignment is to find alignments of each range scan, such that overall alignment error is both minmized and evenly distributed among scans. This misalignment error generally results in smoothing and/or artifacts in the final mesh. Since low frequency error in the individual range scans yields substantial rigid-body misalignment, it effectively causes error at *all* frequencies in the merged result. In the non-rigid case, the undesirable high frequency misalignment can be traded for a low frequency warp in the range data, restricting the final error to the frequencies where error already existed. On the other hand, it is possible to produce a badly warped global alignment which nevertheless has low alignment error, so the amount of warp must be minimized and evenly distributed as well.

We solve the problem of minimizing and distributing warp by first selecting a set of representative feature points on each range scan. Correspondences are found to all overlapping range scans using the weighted ICP method described in section 4, then the positions of these feature points are global optimized so that total warp will be minimized and properly distributed. (Just as in ICP, an approximate initial alignment of the scans is assumed to exist; any method used for initial ICP alignment is sufficient.) Finally, using these features as control points, a thin-plate spline warp is computed for each range scan which brings it into alignment with the global feature point positions, and hence with all other range scans.

Figure 2. The head of Michelangelo's David, aligned using our non-rigid, global alignment method. The green points are the global feature points used.

3.1. Feature Point Selection

Feature points must be selected in a manner which both cover the final mesh, and yields stable correspondence computations between range scans. Half the features are selected uniformly on each range scan, while the other half are selected using the stable sampling technique described in section 4. Since features are selected at a constant sampling rate, this results in denser sampling in areas of greater overlap, which has not proved to be a problem in practice.

The performance of this matching is largely determined by the number of pairs of range scans which must be considered, which in the above feature selection algorithm is $O(nd)$ where n is the total number of range scans and d is the average degree of overlap between scans. It is also parallelizable—our implementation splits the computation into d independent processes which can all run on separate computers. With small modifications, every pair of range scans could be considered in parallel.

In situations such as image sequences, where the degree of overlap is very high, d approaches n, yielding an $O(n^2)$ algorithm. In such cases features can be selected only on scans in a covering set of the final model. (On image sequences, a covering set is trivially computed by incremental alignments.) Since a covering set will be small when the degree of overlap is high, this scheme requires only $O(n)$

scan pairs to be considered.

3.2. Global Point Positioning

In order to minimize warp, feature points should be positioned such that rigid body alignment alignment of scans to the global feature positions will result in minimal, well-distributed error. In that case, the warp after non-rigid alignment will also be minimal and well distributed.

We achieve this by setting soft constraints on the distances between pairs of feature points which we derive from the range scans themselves and from the computed correspondences. Because we assume that warps are low frequency the distance between two feature points p_A and q_A which were both generated on scan A should not change much. We therefore constrain the distance between them to remain $|p_A - q_A|$. Note that because rigid-body transformations preserve distance, any rigid-body transformation applied to A will maintain all such internal constraints exactly. Now let $c(p_A)$ be the corresponding point to p_A on scan B, and let p_B be a feature point on B. If A and B are well aligned with respect to each other, p_A and $c(p_A)$ will coincide, so the desired distance between p_A and p_B is $|c(p_A) - p_B|$. This set of constraints enforces the alignment between range scans determined by the feature correspondences. Again, if there exists a rigid-body transformation that correctly aligns A and B, this transformation will simultaneously satisfy all constraints between A and B as well as the internal constraints on A.

We model all these constraints as a set of springs between pairs of points whose rest length is the desired distance between them, with the spring energy being the square of the deviation from its rest length. We use a simple iterated gradient descent method to minimize the system energy, finding optimal target positions for each feature point. Since poor alignments correspond to high energy on springs between scans, and large warps to any given scan correspond to high energy on the springs within a scan, we obtain our goal of accurate alignment with minimal warp.

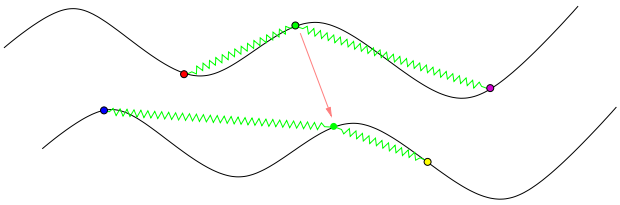

Figure 3. Global feature points are attached to each other by springs whose desired length is based on distances along range scans. Point positions are adjusted to minimize the spring energy.

For a point p which is attached to points $q_1, \ldots, q_n$ by springs $s_1, \ldots, s_n$, the error function is

$$\sum_{i=1}^{n} \text{tension}_i \left(|q_i - p| - \text{desired\_length}(s_i) \right)^2 .$$

We minimize this function one point at a time, using a gradient descent method. We have found this to be efficient and robust in practice, despite its simplicity.

The minimization procedure has the virtues of simplicity and efficiency, but is not guaranteed to converge to a global minimum. In fact, we strongly suspect that some feature points converge to incorrect local minima. Furthermore, some correspondences may be incorrect, especially in flat areas, leading to incorrect springs, and poor feature positioning. We resolve both these issues by performing outlier rejection and thinning prior to warping. A point is considered an outlier if optimization causes it to move too much from its initial position relative to its neighbors. Thinning is used to enforce a minimum distance between points—since the warp is assumed to be low frequency, there is no need for two points on top of each other.

3.3. The Thin-Plate Spline

Once the global features are positioned, all range scans must be warped to align to them. The nature of the warp can be very complex, difficult to model, and vary between data sets. For instance, calibration error, lens distortion, combining input from different sensors, parallax, and data movement all cause different kinds of warp. All of these are difficult to model on their own, and we would like to handle all of them seamlessly. We therefore turn to the thin-plate spline.

The thin-plate spline, introduced by Duchon [10] is a non-rigid spline function with several desirable properties for our application. It is globally smooth, easily computable, separable into affine and non-affine components, and contains the least possible non-affine warping component to achieve the mapping. By the last statement, we mean that the sum of squares of all second order partial derivatives is minimized. So, if $f : \mathbf{R}^n \to \mathbf{R}$ is an n-dimensional thin-plate spline, the *bending energy*,

$$J = \int \left(\sum_{i,j} f_{x_i x_j}^2 \right) \mathrm{d}x_1 \ldots \mathrm{d}x_n \qquad (1)$$

is minimal. Note that since affine transformations are linear, they contribute no error under this metric.

Duchon [10] proves that, for two corresponding point sets $X = \{x_1, \ldots, x_m\}$ and $Y = \{y_1, \ldots, y_m\}$, there is a unique function f such that $f(x_i) = y_i$ and whose bending energy is minimal. Furthermore, this function takes the form $xd + Kw$, where x is a point written in homogeneous

coordinates, d is an affine transformation, w is a fixed m-dimensional column vector of non-affine warping parameters constrained to $X^t w = 0$, and K is an m-dimensional row vector where K_i is the Green's function $U(|x - x_i|)$. In our case (minimizing second order partials in $\mathbf{R}^3$), U is simply $|x - x_i|$; the constant factor is implicitly folded into w.

Thin-plate splines need not be interpolating. Instead, they can minimize the energy functional [10, 23]

$$E_\lambda = \frac{1}{m} \sum |y_i - f(x_i)| + \lambda J \tag{2}$$

The spline will not be interpolating in this case, but for any fixed λ, there will still be a unique minimum, of the form described above.

For the interpolating case, the thin-plate spline specification provides a linear system of equations, which [5] solves directly:

$$\left(\begin{array}{c} w \\ \hline d \end{array}\right) = \left(\begin{array}{c|c} K & X \\ \hline X^t & 0 \end{array}\right)^{-1} \left(\begin{array}{c} Y \\ \hline 0 \end{array}\right) \tag{3}$$

where $K_{ij} = U(|x_i - x_j|)$. For the approximating case [23, eqs. 2.4.23 and 2.4.24] derives a similar system of equations by rewriting equation 2 in matrix form, performing a QR decomposition on X, and simplifying:

$$\begin{array}{rcl} Xd + (K + m\lambda I)w & = & Y \\ X^t w & = & 0 \end{array} \tag{4}$$

An analogous derivation yields similar equations when confidence values or covariance matrices are used [22, 21].

Because we would like very precise alignments (and the data multiple measurements of rigid objects should align precisely), the spline must be heavily weighted toward interpolation, and any effects of covariance matrices or more exotic restrictions are minimal. For this reason, we use no covariance constraints, and rely on accurate correspondences to produce a good alignment. We nevertheless set λ to a very small value to account for measurement noise, and because Equation 4 becomes unstable when $\lambda = 0$ and many point pairs are used.

4. Weighted ICP for Correspondences

Our algorithm exploits the ability of ICP to provide good local alignment, but focuses on aligning individual "feature points" rather than dicing the mesh into small patches as in Ikemoto *et al.* [14]. To do this, we propose a weighted "local ICP" algorithm that selects a set of feature points P on each range scan, then registers the part of the mesh around each point to overlapping scans. This is done by modifying the point selection stage of ICP: instead of selecting points randomly, we instead select them according to a probability distribution that makes it more likely that points near P are used. The major component of the probability function is a decreasing function of distance from the feature point P under consideration:

$$p_{feature}(x) = \frac{1}{\varepsilon + \|x - P\|^2}.$$

If, however, the weight were based entirely on this, there would be a danger of selecting only points that lie on a region of the mesh without sufficient geometric variation to constrain all six degrees of freedom of a rigid-body transformation. Therefore, we augment the probability function with a term that assigns high probability to locations on the mesh that do a good job of constraining the alignment:

$$p_{stability}(x) = \begin{pmatrix} x \times n_x & n_x \end{pmatrix} C^{-1} \begin{pmatrix} x \times n_x \\ n_x \end{pmatrix}$$

where C is the 6×6 point-to-plane ICP covariance matrix computed over the entire region of overlap of the scans (see Gelfand *et al.* [12] for a derivation of the expression for ICP stability). Intuitively, the covariance matrix has small eigenvalues corresponding to transformations that are not well-constrained, hence using its inverse will assign higher probability to points that are important for these under-constrained degrees of freedom.

Our final probability function therefore consists of the product of $p_{feature}$ and $p_{stability}$. This is normalized, integrated into a CDF, then numerically inverted to transform a uniform random variable into samples to be selected. A separate weighted ICP is run for each feature point P, then the nearest point on the target range scan is selected as the correspondence to P. Since thin-plate splines are then used to warp the data so that all corresponding feature points are aligned to each other, accurate global alignment is obtained with no additional smoothing.

Our feature correspondence algorithm adds little to the total computation time. Because an initial overall alignment of each pair of range scans is performed first, it takes only a small number of iterations of ICP to perform each weighted alignment. The most expensive parts of the ICP, computing kd-trees of points and the overlap between each pair of meshes, are computed only once per pair. Another long process, smoothing the normals, is precomputed and cached beforehand, and is therefore performed only once per range scan.

5. Image Alignment

Although we have presented, and indeed developed, our alignment algorithm primarily with range scanning in mind, the framework is quite general, and applies to any area where multiple data measurements must be aligned in the presence of low frequency warp. One obvious, and

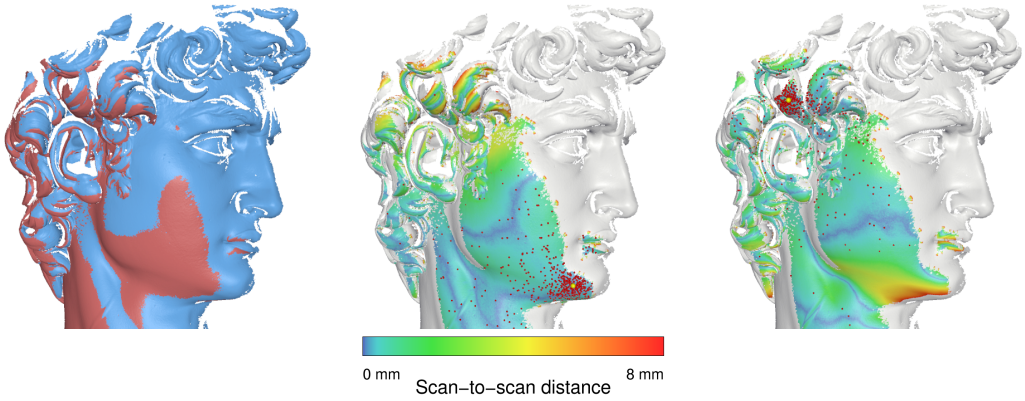

Figure 4. Two range scans from David's head (left). Two weighted ICP alignments centered on the chin (middle) and above the ear (right), color-coded by alignment quality. The yellow point on each scan indicates the center of interest; the red points show the samples select for weighted ICP.

somewhat related example, is that of image alignment, for which we present a proof of concept based on our range scan alignment pipeline.

A standard method of performing image alignment is with optical flow [18], in which approximate image alignment is refined by following image gradients. If a grayscale image is treated as a height field, surface normals correspond to image gradients, and ICP can be directly used for alignment and feature correspondence computation. Changes in image brightness are accommodated as the third dimension. Image-based feature correspondence algorithms will likely result in more robust correspondences. Also, although we do not handle color in our current implementation because this would require 5 dimensions, global point positioning and thin-plate splines both extend elegantly to arbitrary dimensions and could easily handle this.

The global alignment phase has several nice properties for image alignment. First, the thin-plate spline subsumes all linear transformations, and reduces to a linear transformation where that is sufficient, so it will not warp images unnecessarily. At the same time, it can absorb lens distortion effects such as uneven lighting and radial distortion, and small-scale parallax which is not easily filtered out using dominant motion computations. Finally, global point positioning provides a global viewpoint of the scene, rather than favoring any particular frame's viewpoint. The latter is still easily accomplished by pairwise warping each frame to the desired viewpoint before performing global alignment.

6. Results

Figure 5 shows closeups of 1 mm model of the head of Michelangelo's David, created with global, rigid-body ICP alignment, and with our non-rigid alignment. Figure 7(a) shows a histogram of alignment error under rigid and non-rigid alignments, while figures 7(b) and 7(c) show a cross section of the range scan positions for each range scan at the right edge of the eye, after ICP alignment and after our non-rigid alignment. This model contains 1392 range scans at .25 mm precision, comprising approximately 8 GB of data. Feature correspondences required about 1.5 days to compute on 6 Pentium 4 processors. Warping required approximately 3 hours (on a single processor), of which approximately 2:14 minutes were spent on global feature positioning.

The most obvious difference between the models is the elimination of edge smoothing caused by poor alignment perpendicular to the surface (see figure 5(e)). In addition, our new head model has substantially increased surface detail compared to the ICP model, which is due to the improved alignment along the surface (figures 5(e) and 5(g)). Finally, misalignment in highly detailed areas such as the hair can cause confusion between surfaces. This leads to surface noise and "bridging" between the two surfaces. Our improved alignment not only eliminates these artifacts, but reveals interesting surface detail they obscure (note particular the chisel marks revealed between locks of hair in figure 5(g) as compared to 5(f)). Using the non-rigid alignments, we have also produced a .25 mm-resolution model of David's head, closeups of which are show in figure 5(h). Sharp edges are resolved more precisely and the rough surface texture is revealed at this resolution. Because the rigid-body alignment error is too high to give meaningful results, no quarter-millimeter model of David's head (or most of the rest of the Digital Michelangelo data) has been produced before this.

Figure 1 shows a model acquired in our lab using a photometric stereo rig. There is a consistent bias in the normal field, which translates into a low frequency warp

in the range scans. Because of this, rigid-body alignment fails, especially on the face and feet, and behind the wreath. Our alignment substantially improves these areas, however. While photometric range scan data can be substantially improved by imposing boundary conditions derived from stereo triangulation, not all scanning systems can provide this information. When it is not available, non-rigid alignment can still preserve the accurate high frequency information. The model consists of 15 range scans, and required 2:05 minutes to compute feature correspondences. Warping took 10 seconds.

Figure 6 shows a panoramic image created from a digital video sequence of Lake Thingvellir, Iceland shot with a consumer DV camera. There is significant depth-of field, and the camera is in autofocus mode. Because the sequence doubles back on itself, non-adjacent frames in the video sequence overlap in space. Due to the autofocus, a linear brightness scaling cannot account for the brightness variation that occurs between non-temporally-adjacent frames. 64 image frames were used.

7. Discussion and Future Work

We have presented an algorithm for global non-rigid alignment of three dimensional range scans. Our algorithm is robust not only to noise, but also to non-linear warp caused by scanner calibration error. We recover accurate feature correspondences using a weighted, stable ICP computation, then obtain a consistent set of range scans using thin-plate spline warps. The algorithm scales well with data size, requiring running time proportional to the number of overlapping pairs of scans.

It is important to be clear about the ways in which our results are "better" than those produced by rigid-body alignment. Indeed, we can not claim to produce results whose overall shape is necessarily closer to the original, since there is just as much uncertainty in the overall position of the scans. However, by avoiding the loss of high-frequency detail caused by low-frequency warps, our final meshes are more *precise* than those produced using rigid-body alignment. Moreover, because of the deformation-minimizing properties of thin-plate splines, our results reflect the minimum non-rigid warp necessary to compensate for deformation in the data. As imagers and 3D scanners become more ubiquitous and inexpensive, and as they are pushed to their resolution and working-volume limits, we believe that in many situations calibration will become the ultimate limit on achievable quality. Our algorithm produces results that maximize the high-frequency fidelity of the result while introducing minimum deviation from the original data.

There are some areas of the algorithm which will benefit from further refinement. In certain areas, such as the top of David's head, the acquired range scans are simply of poor quality. By incorporating existing confidence data into our framework, and detecting outlier scan regions during the merging process, we can improve the final mesh quality in these areas. Doing so would likely reduce high frequency error further at the expense of simplicity.

In the case of image alignment, a scale independent error metric will likely improve results for image sequences containing zoom. We have also not incorporated several standard image alignment techniques such as dominant motion analysis; these are independent of our fundamental algorithms, and incorporating them is not difficult.

We have developed our algorithm with large scale, high-resolution scanning projects, such as the Digital Michelangelo Project, in mind. However we expect it will work well in many situations where sufficiently precise scanner calibration is either impossible or inconvenient, and we plan to explore uses in these areas. For example, we expect our methods to be useful in conjunction with low-cost, do-it-yourself scanning systems used by untrained operators. We have also shown its direct applicability to image alignment, where there are many non-linear effects to account for, and we are exploring its use in other areas such as biomedical measurements.

References

[1] B. Allen, B. Curless, and Z. Popović. The space of human body shapes: reconstruction and parameterization from range scans. *ACM Trans. Graph.*, 22(3):587–594, 2003.

[2] R. Bergevin, M. Soucy, H. Gagnon, and D. Laurendeau. Towards a general multi-view registration technique. *IEEE Trans. Pattern Anal. Mach. Intell.*, 18(5):540–547, 1996.

[3] F. Bernardini, H. Rushmeier, I. M. Martin, J. Mittleman, and G. Taubin. Building a digital model of Michelangelo's Florentine Pietà. *IEEE Computer Graphics and Applications*, 22(1):59–67, 2002.

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

[5] F. L. Bookstein. Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(6):567 – 585, June 1989.

[6] B.J. Brown and S. Rusinkiewicz. Non-rigid range-scan alignment using thin-plate splines. In *Symposium on 3D Data Processing, Visualization, and Transmission*, sept 2004.

[7] C. Chen, Y. Hung, and J. Cheng. RANSAC-based DARCES: A new approach to fast automatic registration of partially overlapping range images. *Trans. PAMI*, 21(11), 1999.

[8] Y. Chen and G. Medioni. Object modelling by registration of multiple range images. *Image Vision Comput.*, 10(3):145–155, 1992.

[9] H. Chui and A. Rangarajan. A new point matching algo-

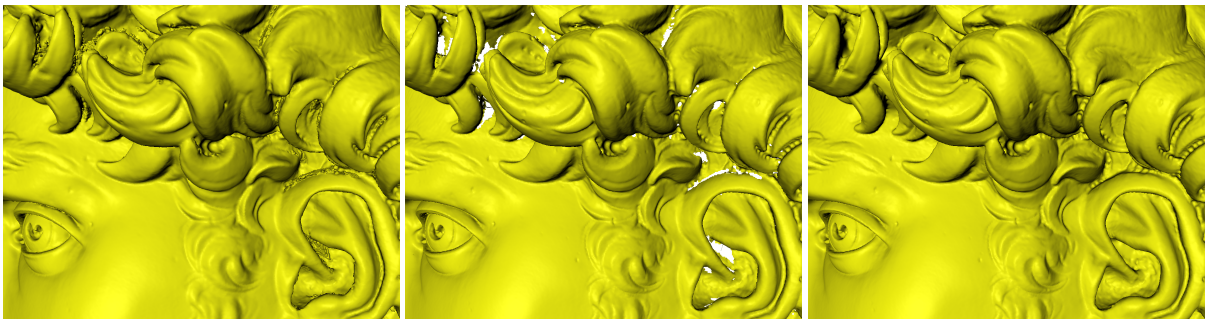

(a) Rigid alignment of David's head (b) TPS alignment of David's head (c) TPS alignment with holes filled

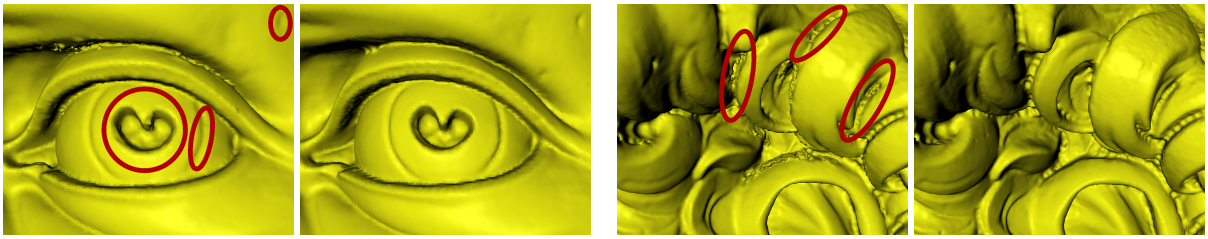

(d) Rigid-body alignment of eye; the pupil is misshapen and the ridges are smooth due to misalignments

(e) TPS alignment of eye; the pupil is now the correct shape, and the ridges are all sharp

(f) Rigid-body alignment of hair above the ear; misalignment causes smoothing and bridging of gaps

(g) TPS alignment of hair; detail between locks of hair, including chisel marks, is revealed

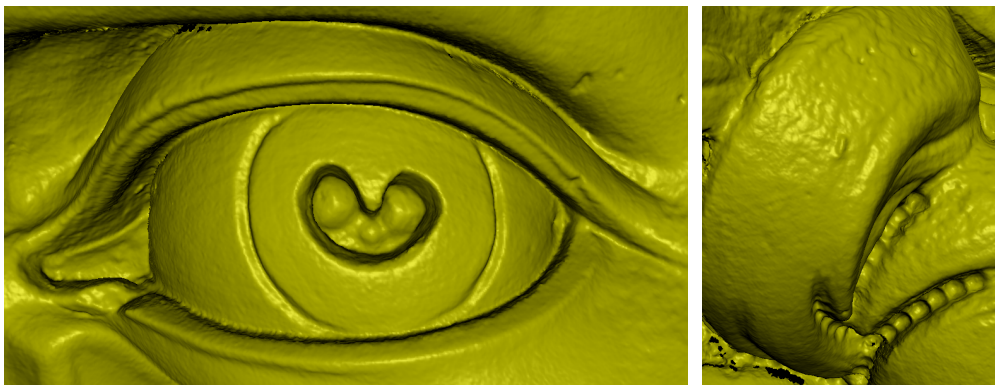

(h) Closeups of the eye and hair merged at .25 mm resolution. At this resolution, finer-scale surface detail is revealed, and edges are sharpened. Only a tps-aligned model of David's head exists at .25 mm, because the rigid-body alignment error is too high to produce meaningful results at this resolution.

Figure 5. Comparison of rigid-body, 1 mm model of David's head with our new, tps-aligned model. .25 mm models created with the tps-aligned data show the further fine-scale detail that these alignments can resolve.

Figure 6. Lake Thingvellir, Iceland. This panorama was created from a camcorder source. The camera is uncalibrated, autofocus, and doubles back on itself (motion from right edge to left edge, and back to center).

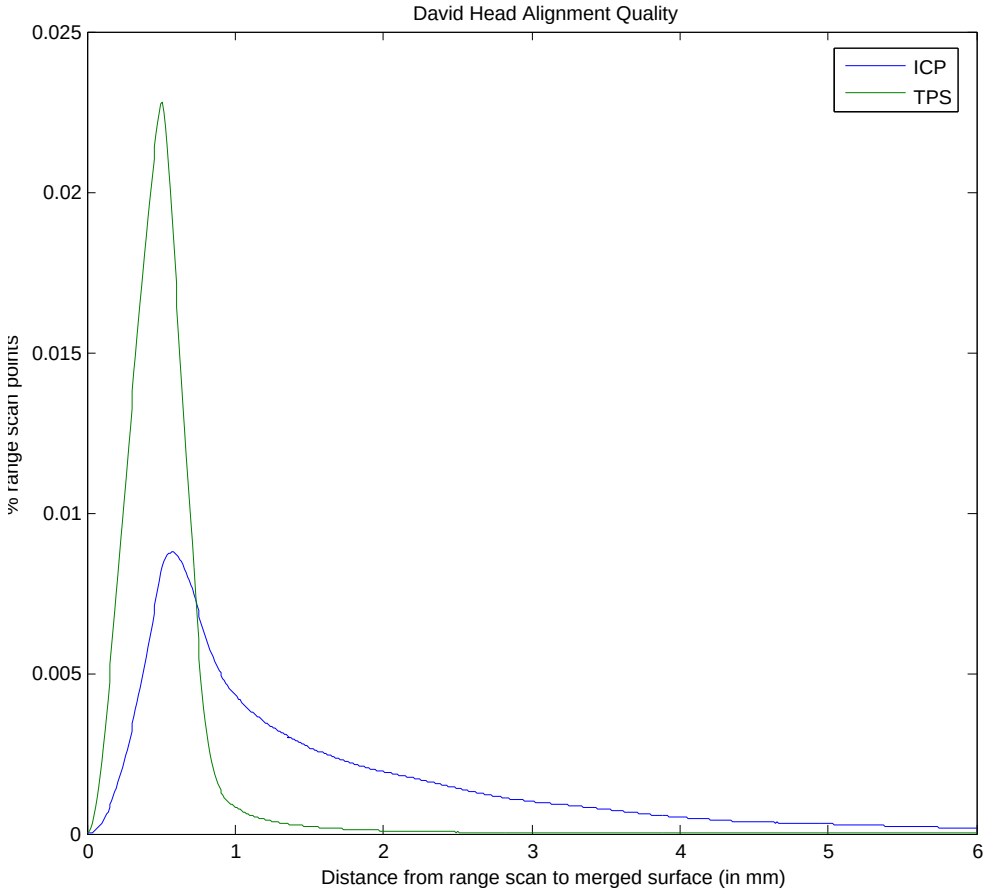

(a) Alignment error histograms for rigid and non-rigid alignments. In each case, the distance from each point on the aligned range scans was measured to the merged model. The rigid-body alignment histogram reveals a long tail of misalignments by several millimeters, which results directly in the observed artifacts and loss of surface detail. The non-rigid alignment, in contrast, has a very small tail, corresponding to sharper features and a reduction in artifacts.

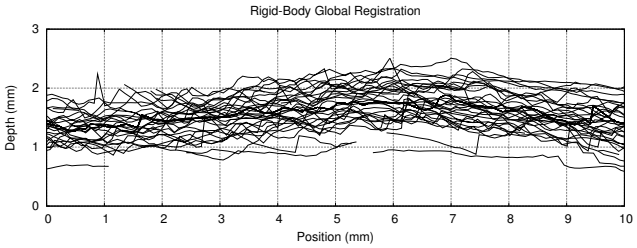

(b) A plot of the depth of each range scan used in rigid-body alignment along a line at the right edge of the eye. Although the alignment here is substantially better than in the hair, misalignment is still greater than 1mm, leading to smoothing and incorrect local shape.

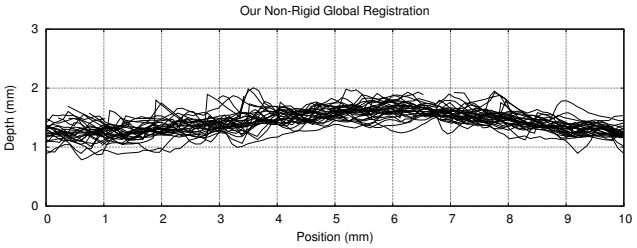

(c) In our non-rigid algorithm, the range scans are very well aligned; there are many fewer outliers, and these are still within 1mm. As a result, the model is merged with sharp features and surface detail preserved..

Figure 7. Graphs of alignment error for the David's head (histogram) and a cross section of the left eye (depth plots) under rigid and non-rigid alignment.

rithm for non-rigid registration. *Computer Vision and Image Understanding*, 89(2-3):114–141, February-March 2003.

[10] J. Duchon. Splines minimizing rotation-invariant seminorms in Sobolev spaces. In *Constructive Theory of Functions of Several Variables*, pages 85–100, Berlin, 1977. Springer-Verlag.

[11] A. Frome, D. Huber, R. Kolluri, T. Bulow, and J. Malik. Recognizing objects in range data using regional point descriptors. In *Proc. ECCV*, 2004.

[12] N. Gelfand, L. Ikemoto, S. Rusinkiewicz, and M. Levoy. Geometrically stable sampling for the ICP algorithm. In *Fourth International Conference on 3D Digital Imaging and Modeling (3DIM 2003)*, October 2003.

[13] D. Hähnel, S. Thrun, and W. Burgard. An extension of the ICP algorithm for modeling nonrigid objects with mobile robots. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI)*, Acapulco, Mexico, 2003. IJCAI.

[14] L. Ikemoto, N. Gelfand, and M. Levoy. A hierarchical method for aligning warped meshes. In *Fourth International Conference on 3D Digital Imaging and Modeling (3DIM 2003)*, October 2003.

[15] A. Johnson and M. Hebert. Surface registration by matching oriented points. In *Proc. Int. Conf. on Recent Advances in 3-D Digital Imaging and Modeling*, pages 121–128, May 1997.

[16] M. Levoy, K. Pulli, B. Curless, S. Rusinkiewicz, D. Koller, L. Pereira, M. Ginzton, S. Anderson, J. Davis, J. Ginsberg, J. Shade, and D. Fulk. The Digital Michelangelo Project: 3-D scanning of large statues. In *Proc. SIGGRAPH*, 2000.

[17] F. Lu and E. Milios. Globally consistent range scan alignment for environment mapping. *Auton. Robots*, 4(4):333–349, 1997.

[18] B. Lucas and T. Kanade. An iterative image registration technique with an application to stereo vision. In *DARPA Image Understanding Workshop*, 1981.

[19] P. J. Neugebauer. Geometrical cloning of 3d objects via simultaneous registration of multiple range images. In *Proceedings of the 1997 International Conference on Shape Modeling and Applications (SMA '97)*, page 130. IEEE Computer Society, 1997.

[20] K. Pulli. *Surface Reconstruction and Display from Range and Color Data*. PhD thesis, University of Washington, 1997.

[21] K. Rohr, M. Fornefett, and H. S. Stiehl. Spline-based elastic image registration: integration of landmark errors and orientation attributes. *Computer Vision and Image Understanding*, 90(2):153–168, May 2003.

[22] K. Rohr, H. S. Stiehl, R. Sprengel, W. Beil, T. M. Buzug, J. Weese, and M. H. Kuhn. Point-based elastic registration of medical image data using approximating thin-plate splines. In *Proc. VBC*, pages 297–306, 1996.

[23] G. Wahba. *Spline Models for Observational Data*, chapter 2.4. Society for Industrial and Applied Mathematics, Philadelphia, PA, 1990.

[24] J.A. Williams and M. Bennamoun. A multiple view 3d registration algorithm with statistical error modeling. *IEICE Transactions on Information and Systems*, E83-D(8):1662–1670, August 2000.

[25] R. Woodham. Photometric method for determining surface orientation from multiple images. *Optical Engineering*, 19(1):139–144, 1980.