

Image-Based Registration of 3D-Range Data Using Feature Surface Elements

Gerhard Heinrich Bendels, Patrick Degener, Roland Wahl, Marcel Körtgen, and Reinhard Klein

Institute for Computer Science II – Computer Graphics, University of Bonn, Germany

Abstract

Digitizing real-life objects via range scanners, stereo vision or tactile sensors usually requires the composition of multiple range images. In this paper we exploit intensity images often recorded with the range data and propose a fully automatic registration technique using 2D-image features with intrinsic scale information for finding corresponding points on the 3D-views. In our approach, the fine registration of two range images is performed by first aligning the feature points themselves, followed by a so-called constrained-domain alignment step. In the latter, rather than feature points, we consider feature surface elements that are derived using the scale information inherently established with the 2D-features. The global registration error is minimized using graph relaxation techniques to mediate the transformations required to align the multiple range images. We demonstrate the power and feasibility of our method by a case-study in the cultural heritage domain.

Categories and Subject Descriptors (according to ACM CCS): I.3.8 [Computer Graphics]: Applications I.4.1 [Image Processing And Computer Vision]: Digitization and Image Capture

1. Introduction

Due to its accuracy, inexpensiveness, and non-intrusiveness, digitizing 3D-Objects with Laser-Range Scanners is the method of choice for many applications, ranging from the automotive over the entertainment industries to creative design and cultural heritage applications. However, to produce a complete surface of the object to be digitized, the measurement of a single view seldom provides sufficient data, such that multiple, often dozens of views have to be registered. Registering two views of an object is usually a two-stage process: First, an initial transformation is estimated, which, in turn, is used as a starting point for the second stage, the fine registration.

The fully automatic registration of multiple range images is still an area of active research in computer graphics. State-of-the-art systems often still rely on user-interaction to determine the initial transformation [CCG*03], making the pre-registration a tedious and time-consuming task. To overcome this drawback, in some applications additional information available from the scanning process can be exploited to derive the initial transformation: For instance, the relative viewpoint position might be known, e.g. from tracking the scanner position or by using a turntable on which the object




Figure 1: Range and colour image acquired with a laser range scanner (in this case a Minolta Vivid 900)

to be digitized was situated. Although direct and convenient, this is not always feasible due to the nature of the object, its dimensions or location. Therefore, a common approach is to derive an initial transformation by aligning a small set of corresponding feature points in the range images. These feature points are either found as local geometric features on the surface of the object or by placing additional markers on or in the surrounding of the object. In the former case, robustness of the feature detection is of vital importance, whereas in the latter, special care has to be taken in the placement of the markers [Akc03], as markers should be visible from as many viewpoints as possible whilst casting preferably no shadows on the object. Aside from the inconvenience, the placement

of markers on the object is *infeasible* in cultural heritage applications, where artifacts to be digitized often must not be touched at all. The need for close-up scans for detailed and spacious objects also diminishes the use of markers placed in the surrounding (see figs 1 and 8).

On the other hand, scanning devices commonly capture not only geometry but also color information or light intensities for the scene (cf. figure 1). These intensity images are far less subject to noise and as opposed to range images do not exhibit missing values. As a consequence, feature points extracted from these images are more robust than those extracted from range images, making them more suitable for correspondence computation. In addition to the robustness, expressiveness and mere number of the features available in the 2D-image information, the key to the ensuing registration steps lies in the fact that the features used in this paper provide *scales* – an indication of in how far is the surrounding of the feature also part of the feature. It is this that allows us to define the feature surface elements and thereby efficiently derive a high-quality registration.

After solving the pairwise registration procedure, the registration problem has to be solved for the full set of available range images. This becomes necessary as the range scans usually overlap with a number of neighbouring range images. In real-world data sets, the range images will be noisy and erroneous due to *material properties* (colour, shininess, transparency, etc.), *lighting situation*, and *object dimensions* (due to a limited depth of focus in the optical system of the scanner). For each neighbour the bilateral registration will therefore result in more or less differing minimizing positions. This non-conformity necessitates mediation among the respective, bilaterally optimal, transformations. We consider the multi-view registration as a directed cost graph, where the range images constitute the nodes and two nodes are connected by an edge iff the corresponding range images overlap *well* (cf. sections 3 and 5 for details).

In this paper we present a fully automatic registration approach based on 2D-image feature correspondences incorporating the following key features:

- No need for special markers
- Robustness with respect to noise and missing geometry data
- Automatic incorporation of additional markers if available

Our registration algorithm is incremental in the sense that additional range images can be incorporated into a set of already registered range images very efficiently. The feature detection is performed unilaterally (constant time), whereas the feature matching has to be done with respect to each of the 2D-images in the given set (linear). Finally, the graph relaxation procedure is performed on the full set of range images. Results from previous range image integration can nonetheless be exploited, as extending an already relaxed graph with additional range images converges very fast.

2. Previous Work

2.1. Pairwise Registration

One of the most popular registration methods in literature is the iterative closest pair algorithm (ICP) by Besl and McKay [BM92]. It iteratively searches for closest point pairs in two surface patches and optimizes the transformation to minimize the distances between these points. However, since this algorithm implicitly assumes that closest points on different patches correspond to each other, it only converges toward a reasonable solution if the patches are roughly pre-aligned. In order to overcome this drawback, various improvements and variants of the original ICP were proposed. This includes verification of closest point pairs by additional attributes like colour or surface normal which is sometimes referred to as the iterative closest compatible point algorithm (ICCP). Furthermore, more sophisticated optimization schemes were proposed as for example simulated annealing or evolutionary algorithms. [RFL02] and [RL01] provide good surveys over these ICP variants. Although these measures improve the convergence properties of the original ICP algorithms and achieve high registration accuracy, they still do not allow for a registration of several completely unaligned surface patches in reasonable time.

To automatize the registration process, several authors proposed to detect special surface feature points on the surface patches [FH86] [SM92] [YF02] [SLW02] [JH97] [SA01] [HHI01] [KPH02] [FA96] [AF97] [WG02] [TB99] [KPJR91]. Constraining the search for correspondences to these features can accelerate the registration process drastically and automatic registration becomes possible. Feature-based approaches primarily differ in their definition of feature points and in the way they are matched. A common drawback of these approaches is that they rely on a sufficient number of prominent or salient features in the geometry. Especially in the presence of noise or missing values this is often problematic.

To circumvent this problem Chen et al. [CHC99] developed a different approach: for pairwise registration they propose a randomized selection of control points on one of the surface patches followed by an exhaustive rigidly constrained search for corresponding points on the other surface. Robertson and Fisher [RF02] also proposed an exhaustive search for automatic registration. Instead of searching for correspondences, they use a parallel search in pose space based on evolutionary algorithms. While the method of Chen et al. is sensitive to noise, the method of Robertson and Fisher requires relatively large overlaps in the surface patches in order to converge to the correct solution. Furthermore, both methods require substantial computational efforts.

Considering the desirable properties of image feature detection, it is not surprising that the idea of exploiting 2D-features for 3D-registration problems is not new. In [Rot99]



Figure 2: Photograph of a medieval rood-screen that was scanned and reconstructed using the approach presented in this paper

Roth uses the popular Harris feature detector [HS88] to extract features from a intensity image that is aligned with a range image. Because of the large number of detected feature points, the author refrained from considering all possible feature point pairs for matching. Instead, the feature points of each surface in three space are tetrahedrized individually using a Delaunay tetrahedrization and the search for correspondences is restricted to the faces of these tetrahedrizations. Two triangles are considered a match if their edge lengths match. However, due to occlusion and missing values in the range images, feature points might be present in only one of the two range images and the Delaunay tetrahedrizations become inconsistent. Therefore, the method is limited to relatively small view point changes and range images with only few missing values.

Another approach related to our method was presented by DePiero in [DeP03]. While his method is not based on image features, it detects KLT features [LK81] in range images and maintains these features together with a graph structure in a database. Targeting at the fast registration of range image sequences, the method predicts the sensor movement from the previous images and uses this prediction to project a subgraph from the database into the next range image in the sequence. This predicted subgraph is then fitted against the detected features, and corresponding features are identified by a graph matching algorithm. While this approach can register a range image sequence at rates of up to 10Hz on current PC hardware, it relies on the viewpoint changes between subsequent images to be comparatively small.

2.2. Multiview Registration

If more than two range images are to be registered a simple solution is the incremental approach taken in [BM92]

[MSY96] and [SG00]: From the set of unregistered patches U two patches are chosen and registered using a pairwise registration method. The two registered patches are then merged into a single patch which is put back into U . This process is repeated until the set U contains only a single surface patch. This incremental approach suffers from the accumulation of local registration errors leading to possibly large global registration errors.

Therefore, several authors proposed to solve for the position and orientation of all patches simultaneously [BSGL96] [EFF98] [BM94] [SB97]. All of these approaches minimize the sum of squared distances between closest point pairs or the distance between a point and the tangent plane to the corresponding point as suggested in [CM92]. As correspondences are iteratively recomputed during the optimization, these methods are computationally expensive. To tackle this drawback, [Pul99] proposes using a generalization of the so-called *concrete-mate* approach, where point-point correspondences remain fixed during the multiview alignment. Also, [CS99] discusses methods that solve the multiview registration problem in case of known point correspondences. In combination with a feature point detection and matching scheme, these approaches can also be used for automatic multiview registration. However, their sensitivity to noise especially in cases, where only a small number of feature points can be found and matched, lead us to propose a hybrid approach incorporating both feature point and closest point correspondences.

3. Feature Detection and Matching

Finding geometric features in range images is a hard task for several reasons. While 3D-feature detection is already a difficult task in closed object representations, situation worsens




Figure 3: Two range images (green and blue) with matching feature point and scale. Only inside the scale-induced feature surface element (red circle) the two range images can robustly be expected to contain corresponding parts of the object.

in case of surface patches acquired digitizing real-life objects as only parts of the object's surface are visible due to occlusion and limited field-of-view. Moreover, the fact that 3D-descriptors are incapable of distinguishing local regions on surfaces of constant curvature (e.g. on planes, cylinders and balls) makes this approach infeasible for many objects, in particular if they are geometrically highly self similar or rotationally symmetric.

On the other hand, finding and matching features in 2D-images is a well-researched topic, and algorithms robustly detecting features that are insensitive to brightness changes, scaling or local occlusions exist.

In a recent survey [MS03] Mikolajczyk and Schmid compared the performance of several local feature descriptors. In particular they examined the robustness of the features with respect to noise, lighting and view point changes up to 60 degrees. They found the Scale Invariant Feature Transform (SIFT) which was developed by Lowe [Low99] (see also [Low04]) based on earlier work by Lindeberg [Lin93] to perform best. As a Scale-Space based method SIFT detects features with a scale parameter that reflects the spatial extension of its defining image neighbourhood. This scale property is of vital importance for our method since it allows to robustly estimate a 3D-position for each detected image feature.

While a 3D-feature position from a 2D-feature could easily be derived using the one-to-one correspondence between the pixels in the intensity image and the depth values in the range image usually established during the data acquisition process, this is not advisable, as the resulting 3D-point is sensitive to noise and small feature deviations. Furthermore,




Figure 4: 3D-feature surface elements are derived from scale-equipped 2D-features.

feature points in 2D-images might correspond to places on the 3D-object where no geometry data is available e.g. holes, dark or reflective spots on the object's surface. Therefore instead of using a single 3D-point (the direct corresponding point to the 2D-feature point) as feature, we use the set of all points corresponding to the image area determined by the position and scale of the feature (see 4). We call these sets *feature surface elements* to accent that they are indeed a surface realization of the scale-equipped feature points. Please note, that the similarity to the notion of *surfels*, i.e. surface points equipped with normals, is not accidentally: Surfels implicitly store a local first-order approximation of the neighbouring surface. Analogously, feature surface elements represent a sampling of the neighbourhood. Unlike surfels though, the feature surface elements represent a region on the surface with a well-defined size known from the 2D-image features.

According to the above definition, we define a *feature point* as the center of gravity of the respective feature surface element and denote by $C_{\mathbf{t}\kappa}$ for any pair (\mathbf{t}, κ) of range images the set of corresponding feature points (see fig. 3).

Although the SIFT method already provides good matching results, false positive matches are nevertheless possible. Since the subsequent registration steps are sensitive to such false correspondences, we apply an additional filtering to the matches based on the RANSAC method [FB81]. A set of matching features in a pair of images can be validated as soon as the 3D-positions of the features have been de-

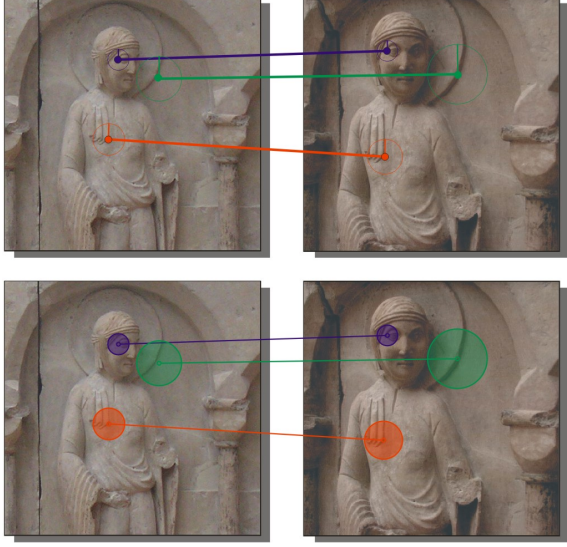


Figure 5: In the first registration stage, only the centres of the feature surface elements are aligned (top). The next stage aligns all available corresponding points pairs contained in the feature surface elements (bottom).

terminated by checking their conformity with respect to rigid transformations. Since it is computationally expensive to actually compute the largest conformal set of matching features (maximum clique), the RANSAC method selects a set of three feature pairs randomly and computes its support, i.e. the set of all feature pairs conforming to the implied transformation. A support set is rejected if it is below a certain size (for our results we used a value of 6). This allows us to remove unreliable correspondences, since large sets of false, yet conforming matches are extremely improbable.

Although the 3D-feature point positions are stable with respect to noise, the sampling of a feature surface element in different images is usually not consistent. In addition to missing range values due to reflective spots, shadowing etc., this might lead to slight deviations in their 3D-positions. While such deviated features can be filtered out using the RANSAC approach to improve the registration accuracy, we tolerate these deviations to a certain extent to increase the number of conformal matches. This constitutes a trade-off between the connectivity in the registration graph (see section 5) and the accuracy. An additional constrained domain alignment step described in section 4 compensates for this tolerated feature deviation.

4. Pairwise Registration

From the algorithm described in the previous section, we have for any range image \mathfrak{r} a set $P^{\mathfrak{r}}$ of scale-equipped feature points $\mathbf{p}_i^{\mathfrak{r}}$, $i = 1, \dots, n_{\mathfrak{r}}$. Moreover, for any pair (\mathfrak{r}, κ) of range images we have a (possibly empty) set of correspondences

$$C_{\mathfrak{r}\kappa} = \{(i, j) \mid \mathbf{p}_i^{\mathfrak{r}} \in P^{\mathfrak{r}} \text{ and } \mathbf{p}_j^{\kappa} \in P^{\kappa} \text{ corresponding}\}.$$



Figure 6: Detail View of the reconstructed angel using 17 range images. Registration was performed on the feature points and the feature surface elements only.

In this section, we describe a two-stage registration procedure for a pair (\mathfrak{r}, κ) with non-empty correspondence set $C_{\mathfrak{r}\kappa}$ (see figure 5).

Coarse Registration

The first registration step consists simply of aligning the point sets $P^{\mathfrak{r}}$ and P^{κ} in a least squares sense, i.e. among the set of all rigid transformations we're looking for the solution to the local minimization problem

$$T_{\mathfrak{r}\kappa} = \underset{T}{\operatorname{argmin}} \quad \varepsilon(T \cdot \mathfrak{r}, \kappa), \quad (1)$$

where the registration error ε is defined as

$$\varepsilon(\mathfrak{r}, \kappa) = \sum_{(i,j) \in C_{\mathfrak{r}\kappa}} d^2(\mathbf{p}_i^{\mathfrak{r}}, \mathbf{p}_j^{\kappa}). \quad (2)$$

Since correspondences are known and fixed, this is a non-iterative procedure (in our implementation solved using the method described in [Hor87]), leading fast and efficiently to an initial registration for \mathfrak{r} and κ . However, the alignment based solely on the feature points accounts only for a fraction of the information available in the range images. (Typically, the number of feature points is in the order of dozens compared to the several hundred thousands of data points.) To compensate for the errors induced in the feature point computation as described in the previous section a second registration step is performed.

Fine Registration

Basically, it would be possible now to register the pre-aligned pair of range images applying one of the many vari-




Figure 7: Detail of the registered rood-screen before and after relaxation. In the left picture, registration errors are noticeable in the area of the chin, the cheek and in the neck.

ants of the ICP-algorithm. They have proven to lead to excellent registration results for good starting positions. Unfortunately, by our experience, they are computationally non-trivial and imperilled of false correspondence computation, which might lead to slow convergence and, more importantly, is susceptible to run into local minima. We tackle these problems by restricting the domain for the correspondence computation to regions of the object that are known to correspond: From the feature detection in the 2D-images, we know that the feature surface elements introduced in section 3 constitute corresponding parts of the surface.

To align the feature surface elements, we perform an ICP on constrained domains: For all pairs $(i, j) \in \mathcal{C}_{\mathbf{ik}}$ we find new correspondences as closest point pairs in the according sets of 3D-points. These enhanced correspondence sets are then aligned using standard ICP-techniques. Figure 6 shows a detail of the reconstructed rood-screen after the two-stage registration process.

We would like to stress, that the 2D-feature matching procedure does not take into account the distribution of the feature points over the range images. In cases where the bounding box of the feature surface elements is very small compared to the bounding box of the range image itself, the two registration steps presented above might leave a registration error noticeable in regions far from the feature surface elements. In these cases, as a consequence of the high-quality pre-registration, a final ICP stage performed on the full data will resolve the remaining inconsistency. In all our experiments, though, (and in all the pictures presented in this pa-

per), the fine registration by feature surface element alignment proved to be sufficient.

5. Multiview Registration

For real-life, erroneous data, the bilaterally optimal transformations will be non-conforming, i.e. the optimal transformation of a range image with respect to one other range image will not be optimal with respect to the remaining range images. To mediate between the competing transformations, we solve in this section the multiview registration problem with a graph relaxation algorithm.

5.1. Graph Setup

Let \mathcal{G} be a directed graph $(\mathcal{N}, \mathcal{E})$. The nodes \mathcal{N} represent the set of range images. An edge $e = (\mathbf{t}, \mathbf{k})$ is element of \mathcal{E} iff the correspondence set $\mathcal{C}_{\mathbf{ik}}$ is non-empty. To every edge $e = (\mathbf{t}, \mathbf{k})$, we assign a rigid transformation $T(e) = T(\mathbf{t}, \mathbf{k})$ that is initialized to be the solution of the bilateral alignment process of the two adjacent range images. Additionally, we store with every edge the registration error $\varepsilon(e) = \varepsilon(\mathbf{t}, \mathbf{k})$ induced by this initial registration. The antisymmetry $T(\mathbf{k}, \mathbf{t}) = T(\mathbf{t}, \mathbf{k})^{-1}$ in the edge attributes is the reason why \mathcal{G} needs to be a *directed* graph – in all other respects \mathcal{G} can be treated as undirected.

The task is now to find for every node \mathbf{t} a transformation $T_{\mathbf{t}}$ such that the global registration error

$$\Sigma := \sum_{e \in \mathcal{E}} \varepsilon(e) \quad (3)$$

is minimal. In other words: Let \mathbf{T} be the vector (T_1, \dots, T_n)




Figure 8: Registered range images. 84 Patches, 20 million points. Note that many patches cover exclusively the interior of the object, a fact that would make the exploitation of synthetic marker points attached in the surrounding infeasible.

of rigid transformations, then we're looking for the solution to the global minimization problem

$$\mathbf{T} = \operatorname{argmin}_{(\mathbf{t}, \kappa) \in \mathcal{E}} \varepsilon(T_{\mathbf{t}}, T_{\kappa}). \quad (4)$$

5.2. Graph Collapse

Clearly, problem (4) has a degeneracy in the sense that the error Σ is invariant under any rigid Transformation Q :

$$\Sigma(\mathbf{T}) = \Sigma(Q\mathbf{T}) = \Sigma(QT_1, \dots, QT_N)$$

Therefore, we choose an arbitrary node t_0 s.t. T_{t_0} is the identity transformation. An initialization $T_{\mathbf{t}}$ for all nodes \mathbf{t} can then be found by computing a minimal spanning tree of \mathcal{G} and combining the transformations from t_0 to \mathbf{t} along the paths in the spanning tree (cf. fig. 9). For numerical reasons




Figure 9: The registration graph and a corresponding spanning tree. Setting T_3 to be the identity would give, e.g., $T_1 = T(1, 3)$, and $T_5 = T(4, 5) \circ T(4, 3)$ as initial transformations.

it is beneficial to choose the root node t_0 s.t. the average path length from t_0 to all remaining nodes is minimal, otherwise the choice is arbitrary.

5.3. Relaxation

To resolve the non-conforming transformations we iterate over the graph and re-align each node with respect to the adjacent nodes. Again, this is a two-stage procedure: First, the relaxation is performed taking into account the feature points only, whereas in the second stage, the correspondences in the feature surface elements are accounted for.

In the literature, different approaches have been discussed concerning the recomputation of correspondences in-between iterations. Recomputing the correspondences between two iterations is not only computationally expensive, it might also exhibit slow convergence speed. This is due to the fact that changing the correspondences actually constantly changes the function to minimize. Moreover, since thresholding is applied during correspondence computation, the registration graph might even get disconnected in cases where subgraphs of the graph are connected only by very few cross-edges. Keeping the correspondences fixed during the whole relaxation, on the other hand, is sensitive to noise and prone to run into local minima. Hence, we pursue a hybrid approach that keeps correspondences fixed during the relaxation and afterwards repeats the process with recomputed correspondences. In pseudo-code the relaxation reads:



Figure 10: Right: Detail photograph of the rood-screen. Left: Reconstruction of the detail; 17 range images were used, no global fine registration step applied. Note that the images were taken from a slightly different viewpoint. The colour difference mainly results from using a flashlight for the photography

```

relax( $\mathcal{G}, stage$ )
while  $\Sigma$  improves do
  if stage > 1 then
    recompute correspondences;
  end if
  while  $\Sigma$  improves do
    for all  $t \in \mathcal{N}$  do
      align  $t$  with adjacent nodes;
    end for
    for  $t \in \mathcal{N}$  do
      evaluate  $\Sigma$ ;
    end for
    evaluate  $\Sigma$ ;
  end while
end while

```

Finding corresponding pairs as closest points results in asymmetric correspondence sets, i.e. $\mathcal{C}_{1\kappa} \neq \mathcal{C}_{\kappa 1}$. This is appropriate if one range image has to be aligned to another (since this relationship, too, is asymmetric). In multiview-registration, however, range images have to be aligned mutually. Otherwise, for an edge (t, κ) , a next relaxation step (where t is the current node to be re-aligned) might simply try to undo the transformation just achieved in the last step (where κ was re-aligned), leading to slow convergence. Hence, we define the correspondence set for all edges $(t, \kappa) \in \mathcal{E}$ to be the union of the *one-sided* correspondence sets $\mathcal{C}_{1\kappa}$ and $\mathcal{C}_{\kappa 1}$. Obviously, this is not necessary in the first relaxation stage, as here, the correspondence sets consists only of the feature points themselves and, therefore, is symmetric by construction.

Also, in all our experiments we found it sufficient to perform during the pairwise registration the first stage only, i.e. we omit the alignment of the feature surface elements in the bilateral case and apply both stages not until the relaxation of the registration graph.

6. Results and Conclusions

Figure 8 shows the 84 patches that were registered to reconstruct the rood-screen depicted in figure 2 using the two registration steps described in sections 4 and 5. Figures 10, and 11 show point renderings of the the reconstructed rood-screen. For the given examples, the complete registration process from feature detection and matching to the graph relaxation based on the feature surface elements took less than an hour on standard PC hardware and was performed without any user-interaction.

In this paper we presented a novel fully automatic registration algorithm for multiple range images. The key to our approach is the use of robust and expressive image features that additionally contain scale information. This extensive feature information allows us to perform a two-stage registration process in which a feature-point alignment precedes an alignment of feature surface elements. The latter is basically a constrained-domain ICP where the domains are consistently derived from the scales established in the 2D-feature detection and matching process. This approach scales well to large data sets and avoids local minima. The thresholds for the correspondence computation in the second



Figure 11: Reconstruction of the complete rood-screen, point rendered with per vertex-colours

registration stage are naturally derived from the registration error of the foregoing stage.

Our approach is very simple in concept, but profits naturally from robust feature point correspondences. In particular, feature detection and matching on basis of 2D-images gives access to 3D-feature points at places infeasible using only the 3D-data, e.g. holes in the object, or spots on the object that do not deliver a 3D-point, but can easily and robustly be identified on the corresponding 2D-image. Thus our approach is robust with respect to missing data in the range images due to the object geometry, material properties, or the scanning process itself, that were a major challenge in previous registration approaches. Another important benefit of exploiting image-based features for our registration procedure is that even surface patches that are geometrically indistinguishable can be robustly registered. Thus, rotationally symmetric objects can be reconstructed as well as objects that are highly self-similar if there is image information that can be evaluated.

Our registration algorithm is independent of additional user-defined marker points – a point that is vital for cultural heritage applications as artifacts often must not be touched at all. On the other hand, if available, these marker points are naturally included in the registration process.

References

- [AF97] ASHBROOK A. P., FISHER R. B.: Constructing models of articulating objects: range data partitioning. In *Proceedings of the International Conference on Recent Advances in 3-D Digital Imaging and Modeling* (1997), IEEE Computer Society, p. 164. [2](#)
- [AKC03] AKCA D.: Full automatic registration of laser scanner point clouds. In *Optical 3-D Measurement Techniques VI* (2003), Gruen A., Kahmen H., (Eds.), vol. 1, pp. 330–337. [1](#)
- [BM92] BESL P. J., MCKAY N. D.: A method for registration of 3-D shapes. *IEEE Transactions on Pattern Analysis and machine Intelligence* 14, 2 (Feb. 1992), 239–258. [2, 3](#)
- [BM94] BLAIS G., MARTIN L.: *Registering multiview range data to create 3D computer objects*. Tech. Rep. CIM-93-16, McGill Centre for Intelligent Machines, Mar. 1994. [3](#)
- [BSGL96] BERGEVIN R., SOUCY M., GAGNON H., LAURENDEAU D.: Towards a general multi-view registration technique. *IEEE Trans. Pattern Anal. Mach. Intell.* 18, 5 (1996), 540–547. [3](#)
- [CCG*03] CALLIERI M., CIGNONI P., GANOVELLI F., MONTANI C., PINGI P., SCOPIGNO R.: Vclab's tools for 3d range data processing. In *4th International Symposium on Virtual Reality, Archaeology and Intelligent Cultural Heritage (VAST2003) and First EUROGRAPHICS Workshop on Graphics and Cultural Heritage* (Brighton, UK, November 2003). [1](#)
- [CHC99] CHEN C.-S., HUNG Y.-P., CHENG J.-B.: Ransac-based darcs: A new approach to fast automatic registration of partially overlapping range images. *IEEE Trans. Pattern Anal. Mach. Intell.* 21, 11 (1999), 1229–1234. [2](#)
- [CM92] CHEN Y., MEDIONI G.: Object modelling by registration of multiple range images. *Image Vision Comput.* 10, 3 (1992), 145–155. [3](#)
- [CS99] CUNNINGTON S., STODDART A.: N-view point set registration: A comparison. In *British Machine Vision Conference* (Nottingham, UK, 1999), vol. 1, pp. 234–244. [3](#)

- [DeP03] DEPIERO F. W.: Deterministic surface registration at 10hz based on landmark graphs with prediction. In *14th British Machine Vision Conf. (BMVC2003)* (Norwich, UK, 2003). 3
- [EFF98] EGGERT D. W., FITZGIBBON A. W., FISHER R. B.: Simultaneous registration of multiple range views for use in reverse engineering of cad models. *Comput. Vis. Image Underst.* 69, 3 (1998), 253–272. 3
- [FA96] FELDMAR J., AYACHE N.: Rigid, affine and locally affine registration of free-form surfaces. *Int. J. Comput. Vision* 18, 2 (1996), 99–119. 2
- [FB81] FISCHLER M., BOLLES R.: Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. In *Communications of the ACM* (1981), vol. 24, pp. 381–395. 4
- [FH86] FAUGERAS O. D., HEBERT M.: The representation, recognition, and locating of 3-d objects. *Int. J. Rob. Res.* 5, 3 (1986), 27–52. 2
- [HHI01] HIGUCHI K., HEBERT M., IKEUCHI K.: Building 3-d models from unregistered range images. 41–75. 2
- [Hor87] HORN B.: Closed form solution of absolute orientation using unit quaternions. *Journal of the Optical Society A* 4, 4 (April 1987), 629–642. 5
- [HS88] HARRIS C., STEPHENS M.: A combined corner and edge detector. In *Proceedings of The Fourth Alvey Vision Conference, Manchester* (1988), pp. 147–151. 2
- [JH97] JOHNSON A. E., HEBERT M.: Surface registration by matching oriented points. In *Proceedings of the International Conference on Recent Advances in 3-D Digital Imaging and Modeling* (1997), IEEE Computer Society, p. 121. 2
- [KPH02] KRSEK P., PAJDLA T., HLAVAC V.: Differential invariants as the base of triangulated surface registration. *Comput. Vis. Image Underst.* 87, 1-3 (2002), 27–38. 2
- [KPJR91] KAMGAR-PARSI B., JONES J. L., ROSENFELD A.: Registration of multiple overlapping range images: Scenes without distinctive features. *IEEE Trans. Pattern Anal. Mach. Intell.* 13, 9 (1991), 857–871. 2
- [Lin93] LINDBERG T.: Detecting salient blob-like image structures and their scales with a scale-space primal sketch: a method for focus-of-attention. *Int. J. Comput. Vision* 11, 3 (1993), 283–318. 4
- [LK81] LUCAS B. D., KANADE T.: An iterative image registration technique with an application to stereo vision (darpa). In *Proceedings of the 1981 DARPA Image Understanding Workshop* (April 1981), pp. 121–130. 3
- [Low99] LOWE D. G.: Object recognition from local scale-invariant features. In *Proc. of the International Conference on Computer Vision ICCV, Corfu* (1999), pp. 1150–1157. 4
- [Low04] LOWE D. G.: Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision* 60, 2 (11 2004), 91–110. 4
- [MS03] MIKOLAJCZYK K., SCHMID C.: A performance evaluation of local descriptors. In *Conference on Computer Vision and Pattern Recognition CVPR* (june 2003), vol. 2, IEEE, pp. 257–263. 4
- [MSY96] MASUDA T., SAKAUE K., YOKOYA N.: Registration and integration of multiple range images for 3-d model construction. In *Proceedings of the 1996 International Conference on Pattern Recognition (ICPR '96) Volume I* (1996), IEEE Computer Society, p. 879. 3
- [Pul99] PULLI K.: Multiview registration for large data sets. In *Proceedings of the International Conference on 3D Digital Imaging and Modeling* (1999), pp. 160–168. 3
- [RF02] ROBERTSON C., FISHER R. B.: Parallel evolutionary registration of range data. *Comput. Vis. Image Underst.* 87, 1-3 (2002), 39–50. 2
- [RFL02] RODRIGUES M., FISHER R., LIU Y.: Special issue on registration and fusion of range images. *Comput. Vis. Image Underst.* 87, 1/2/3 (2002), 1–7. 2
- [RL01] RUSINKIEWICZ S., LEVOY M.: Efficient variants of the ICP algorithm. In *Proceedings of the Third Intl. Conf. on 3D Digital Imaging and Modeling* (2001), pp. 145–152. 2
- [Rot99] ROTH G.: Registering two overlapping range images. In *Proceedings of the Second International Conference on 3-D Digital Imaging and Modeling* (1999), pp. 191–200. 2
- [SA01] SUN Y., ABIDI M.: Surface matching by 3d point's fingerprint. In *Proceedings of the Eighth IEEE International Conference on Computer Vision* (2001), vol. 2, pp. 263–269. 2
- [SB97] SCHMITT F., BENJEMAA R.: Fast global registration of 3D sampled surfaces using a multi-Z-buffer technique, Sept. 12 1997. 3
- [SG00] SAPPÀ A., GARCÍA M.: Incremental multiview integration of range images. In *Proceedings of the International Conference on Pattern Recognition* (2000), vol. 1, pp. 546–549. 3
- [SLW02] SHARP G. C., LEE S. W., WEHE D. K.: Icp registration using invariant features. *IEEE Trans. Pattern Anal. Mach. Intell.* 24, 1 (2002), 90–102. 2
- [SM92] STEIN F., MEDIONI G.: Structural indexing: Efficient 2d object recognition. *IEEE Trans. Pattern Anal. Mach. Intell.* 14, 12 (1992), 1198–1204. 2
- [TB99] TAREL J.-P., BOUJEMAA N.: A coarse to fine 3d registration method based on robust fuzzy clustering. *Computer Vision and Image Understanding* 73, 1 (1999), 14–28. <http://www-rocq.inria.fr/tarel/cviu99.html>. 2
- [WG02] WYNGAERD J. V., GOOL L. V.: Automatic crude patch registration: toward automatic 3d model building. *Comput. Vis. Image Underst.* 87, 1-3 (2002), 8–26. 2
- [YF02] YAMANY S. M., FARAG A. A.: Surfacing signatures: An orientation independent free-form surface representation scheme for the purpose of objects registration and matching. *IEEE Trans. Pattern Anal. Mach. Intell.* 24, 8 (2002), 1105–1120. 2