

Efficient Global Registration for Nominal/Actual Comparisons

Sarah Berkei¹, Max Limper¹, Christian Hörr², and Arjan Kuijper¹

¹TU Darmstadt & Fraunhofer IGD ²Carl Zeiss Optotechnik GmbH

Abstract

We investigate global registration methods for Nominal/Actual comparisons, using precise, high-resolution 3D scans. First we summarize existing approaches and requirements for this field of application. We then demonstrate that a basic RANSAC strategy, along with a slightly modified version of basic building blocks, can lead to a high global registration performance at moderate registration times. Specifically, we introduce a simple feedback loop that exploits the fast convergence of the ICP algorithm to efficiently speed up the search for a valid global alignment. Using the example of 3D printed parts and range images acquired by two different high-precision 3D scanners for quality control, we show that our method can be efficiently used for Nominal/Actual comparison. For this scenario, the proposed algorithm significantly outperforms the current state of the art, with regards to registration time and success rate.

CCS Concepts

• **Computing methodologies** → *Reconstruction; Matching; Mesh models;*

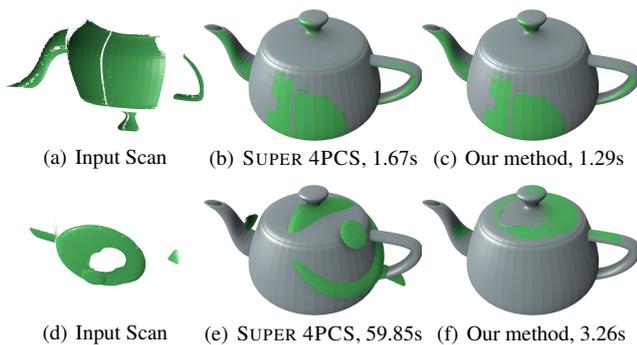


Figure 1: Registration of range images for Nominal/Actual comparison. Using existing building blocks, we propose a new method for global registration that achieves high success rates and fast registration, even for challenging input data. Our study uses real-world data from different high-end industrial scanners (top row: structured light scanner; bottom row: handheld laser scanner).

1. Introduction

Registration of 3D scans against a reference model is a topic which has been extensively studied within the past [BM92, RL01, PLH04]. One important use case is *Nominal/Actual* comparison. The aim is to assess deviations from the original CAD model, which might have occurred during manufacturing. Being able to use only a single range image for this purpose, or at least just a few ones, is an important requirement in practice. Often, deviations should be measured on just one specific part of an object during the evaluation process. To keep processing times at a reasonable level, it is crucial to avoid the necessity to reconstruct the whole surface, as 3D scanning is usually a time-consuming step. By aligning a single range image and the CAD model in a common coordinate system,

distances between both data sets can be efficiently measured and visualized. However, computing an alignment, also referred to as *registration*, reliably and in a fully automatic fashion is a hard task. Typical problems include parameters that need to be tweaked by hand, excessive computation times, as well as the automatic assessment of the quality of the resulting registration. Moreover, many global alignment methods cannot deal with partial data. For these reasons, professional software packages do usually not provide automatic alignment methods, but provide user-assisted methods as a first choice. One possible workflow consists in the manual selection of three or more roughly corresponding points in each of the data sets (as available in open source software such as *Meshlab* or *CloudCompare*, for example). Other possibilities are the selection of two points with normals, or the selection of a single corresponding point under similar viewing angles onto both models. Given such global constraints from the user, the alignment problem reduces to a local registration, which can be solved fast and efficiently using one of the variants of the well-known Iterative Closest Point (ICP) method [RL01].

Within this paper, we first summarize relevant approaches for *Range Image to CAD* registration in the context of *Nominal/Actual* comparison. We then demonstrate that a basic RANSAC strategy with some modifications is able to significantly outperform the state of the art with respect to global registration time and success rate. Using data from two different high-precision industrial scanners, we demonstrate the utility of the proposed method for the use case of fully-automatic alignment of 3D scans and CAD data for *Nominal/Actual* comparison, using the example of 3D printing.

Our main contribution is an efficient and robust solution of the global registration problem based on a basic RANSAC strategy, which even works well with partial scans and structureless surfaces. In contrast to previous work, our approach only asks the user for a single input parameter (δ_H), which is needed in principle to distinguish between possible registration artifacts from true devi-

ations that occurred during manufacturing. Furthermore, we introduce a simple feedback loop for global alignment by incorporating the ICP algorithm directly into the computation of a global alignment. Thanks to the quick error minimization of the ICP method, this approach allows us to efficiently discard global alignments that would otherwise result in wrong transformations. As one of our key findings, we show that this is especially important in order to cope with shortcomings of the commonly used Largest Common Point Set (LCP) measure. We demonstrate the applicability of our approach and its high efficiency by comparing it against the state-of-the-art methods, using high-precision range images that have been acquired from 3D-printed models as a possible use case for nominal/actual comparison.

2. Local Alignment Methods

The de facto standard *local* method for computing a precise alignment is the Iterative Closest Point (ICP) algorithm [CM92, BM92]. Rusinkiewicz and Levoy presented variants of the original method and optimized configurations for fast convergence [RL01]. Pottmann et al. presented a method that can be regarded as an alternative to ICP which is based on instantaneous kinematics and quadratic approximants of the squared distance function to a surface [PLH04]. Low presented a fast point-to-plane distance minimization method, which allows for a significant speedup of the ICP method [IL04]. Pottmann et al. analyzed the convergence properties of ICP more in detail [PHYH06]. Pointing out that ICP converges only slowly if the optimization runs through tangential moves along the surface, they propose alternatives with better convergence behavior. All of the mentioned methods for local registration have in common that they need a good rough alignment as initialization. If this alignment is not of sufficient quality, the local alignment step might converge into a wrong local minimum, producing a wrong result. Another common difficulty is the sensitivity to missing data or outliers. Most of the methods have several parameters to tune, which may be difficult with varying data quality. Therefore, Bouaziz et al. presented a new variant of the ICP algorithm, called Sparse ICP [BTP13]. Within they describe an algorithm that optimizes the registration by using sparsity inducing norms for noisy data sets.

3. Global Alignment Methods

Algorithms. Finding a rough, global alignment of two 3D objects for later local alignment via ICP or similar methods is a problem that has been extensively studied within the past. A basic approach is to simply align object centroids of both models, which solves the translational part of the global alignment. The rotation can then be determined in various ways, for example by performing a Principal Component Analysis (PCA) and trying different alignments of the resulting axes of both data sets [LR09]. However, such approaches do typically not work with partial scans because the principal component may differ between two scans. Drost et al. presented a global description based on oriented point pairs [DUNI10]. Similar features on the model are grouped together, which allows to use a sparse representation of the point data. More recently, Birdal and Ilic presented a revised version of the framework, enhancing the pose retrieval by using segmentation and per-segment pose estimation, supported by an occlusion-aware ranking scheme. [BI15]. Another solution is provided by RANSAC-based methods. One example is the DARCES approach of Chen et al. [CHC99]. Their method searches similar triangles in both point data sets and computes a corresponding rigid transformation to obtain a candidate alignment, which is then ranked using the LCP metric. As will be shown within Section 4.2, our basic approach for global registration is closely related to this method, but contains an additional stage of quick

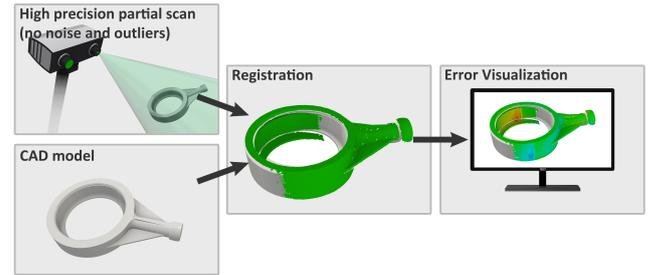


Figure 2: General overview over Nominal/Actual comparison. A part, manufactured from CAD, will be scanned and compared against its CAD by aligning the scanned data without noise and outliers onto the CAD data. Thus the deviation can be visualized.

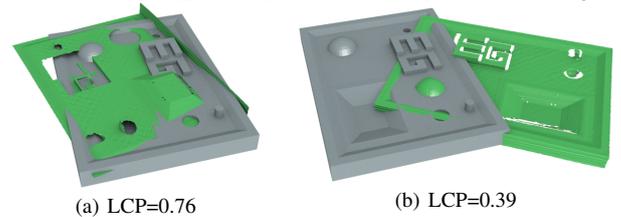


Figure 3: Problems of using a pure LCP-based approach. The configuration on the left has a relatively high LCP value. However, it will converge to an incorrect result after local refinement (ICP). In contrast, the result on the right converges to the global optimum in just 14 iterations, although the initial LCP value is rather low.

intermediate validation via ICP. Zhou et al. presented Fast Global Registration (FGR) as a method for efficient alignment of partially overlapping 3D surfaces [ZPK16]. This algorithm optimizes between a candidate set on each surface, based on Fast Point Feature Histograms (FPFH) [RBB09] for each candidate, without updating these correspondences. Therefore, the algorithm is able to deal with noisy correspondences and so with noisy datasets. Aiger et al. presented 4-Point Congruent Sets (4PCS), a method for efficient global alignment of partial 3D scans [AMCO08]. The algorithm exploits geometric properties which are invariant against transformation, described as congruent points. The method is robust against noise and outliers and has shown good results on a variety of data sets. However, in the worst case, it has a quadratic complexity in the number of sampled points. Mellado et al. have significantly improved the 4PCS approach by proposing an advanced variant, entitled SUPER 4PCS, that uses a smart indexing scheme [MAM14]. Yang et al. presented the Go-ICP method, which combines local and global alignment strategies in order to find a global optimum of the L_2 error metric used by ICP [YLCJ16]. Since the standard ICP algorithm may get stuck converging into local optima, their method uses a global branch-and-bound strategy over the domain of rigid transformations in order to compute different global alignments, which are then locally refined via ICP. Alternating between these stages, a global minimum can be found, and, as the authors state, the algorithm can hence be used as an optimality benchmark. For case studies like ours, it could be interesting to use Go-ICP to generate ground truth data, where no such data is available. Within our evaluation, we compare the results of our algorithm for global alignment to the ones generated by SUPER 4PCS and FGR, using semi-manually generated ground truth alignments.

Measuring Metrics. The most popular metric for measuring which candidate transformation should be considered the best one is the Largest Common Point Set (LCP) [ATT97]. The value of this metric will always be a real number within the interval $[0, 1]$,

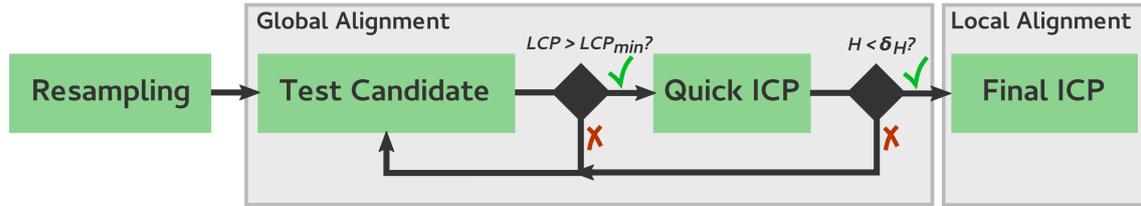


Figure 4: Our framework. Our global alignment algorithm includes a quick ICP stage in order to quickly evaluate whether a given candidate transformation should be regarded as an acceptable result. The quick ICP stage is only executed if the result already has a minimum quality.

measuring the amount of overlap that was detected between both data sets. To measure the LCP value for two roughly aligned sets of points P and Q , a threshold δ_{LCP} is used for the distance between corresponding points p_i and q_i , where q_i is the nearest neighbor of p_i . The threshold δ_{LCP} distinguishes points that are regarded as close points on the one hand, meaning that a local overlap exists, and distant points on the other hand, meaning that there is no local overlap between the data sets. The overall LCP is then simply the ratio between the amount of close points and the overall amount of points:

$$LCP(P, Q) = \frac{|\{p_i \in P : \|p_i - q_i\| < \delta_{LCP}\}|}{|P|} \quad (1)$$

While this measure is easy to compute and produces good results for many cases, we found that a pure LCP-based optimization of the global alignment may also lead to some undesired results. Problems can arise in both directions: On the one hand, a high LCP value does not always guarantee that a good solution has been found. On the other hand, a low LCP may produce better results than another initialization with higher LCP. Two examples illustrating this problem are shown in Fig. 3. Because of such issues, our method does not solely rely on LCP to judge whether a candidate transformation should be considered as the best one.

Of course, other methods such as the one of Huttenlocher and Kedam can be used, which calculates the minimum Hausdorff distance between two point sets under translation [HK90]. The Hausdorff distance is defined as the largest value out of the closest distance values of all points of one data set, as measured towards the other surface [AScE02].

4. Overview of our Nominal/Actual Comparison Framework

During manufacturing processes it may be essential to compare the manufactured parts against their CAD data for quality control, as shown in Fig. 2. In order to detect subtle real-world deviations, it is not sufficient to use low-end 3D reconstruction pipelines (such as the Microsoft Kinect, for example). Instead, a precision in the sub-millimeter range is required, and outliers may not exist in the data that is used for Nominal/Actual comparison. Therefore, only high-precision scan systems are able to produce the expected data quality. Apart from assuming that we operate on this kind of data, we furthermore assume that the partial scans do not contain any other objects than the ones we are having inside the CAD data (i.e., no surroundings have been captured during scanning).

A basic overview of our approach is shown in Fig. 4. It mainly consists of four stages: First, during the initial setup stage, we perform a sampling of the input data sets and estimate parameters of the global alignment algorithm. The main RANSAC loop for global alignment then samples the space of all possible transformations. We do so using random triangle pair alignments on both point data sets, in order to obtain candidate transformations for a coarse alignment. This frequently reoccurring step is supported by intermediate executions of the ICP algorithm (the *Quick ICP* stage), in cases

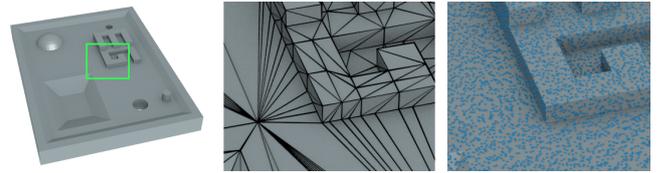


Figure 5: Random samples on the specimen block CAD model.

where a candidate transformation has been rated as a potential success candidate. The RANSAC loop terminates as soon as the quick ICP validation of a candidate transformation was successful, or as soon as the given time budget has been fully exhausted. During the final stage, we perform a more precise local alignment via ICP and classify the result, using the Hausdorff distance calculated from the scan to the CAD model.

The following sections describe the mentioned stages of the pipeline more in detail.

4.1. Sampling and Parameter Estimation

As high-precision range scans can be complex data sets, consisting of many thousands or even millions of points, we have to reduce the amount of data in order to perform an efficient registration. Existing algorithms typically reduce the amount of points for global alignment and local alignment stages to just a few hundred, or, at maximum, a very few thousand samples (cf. [RL01, MAM14]). For the high-resolution input range images, which are typically very dense, we have found that a straightforward random sampling already gave results of sufficient quality - although other strategies, such as Poisson disc sampling, are clearly possible [Bri07]. For CAD reference models, a bit more effort is necessary to obtain reasonable results. Since the tessellation might be highly irregular, simply subsampling the existing vertices is not sufficient. Therefore, we first employ a stratified triangle sampling to create a uniform, high-resolution point cloud on the CAD surface. To do so, we subdivide each triangle temporarily into smaller ones and then generate random samples within the resulting triangles, using the simple scheme of Turk [Tur90]. Fig. 5 shows an example. Although our CAD models always consisted of triangulated meshes, it is worth noting at this point that similar point sampling strategies are possible for parametric surface data. From this high-resolution point sampling, we can then simply draw subsamples in the same fashion as for the range image data. Note that this process allows us to work with two point data sets, which is a more common input for most algorithms and, especially, allows for faster sampling (and hence more candidate evaluations) than more sophisticated strategies, such as computing precise point-to-mesh distances.

In order to efficiently decide whether we should perform a quick run of the ICP method, we have to be able to judge the quality of a candidate transformation that was computed during the RANSAC loop. Like most other methods, we use the popular LCP metric for this purpose, which can be efficiently computed and

leads to a good rough classification of the candidate transformations (cf. [CHC99, AMCO08, MAM14]). However, we have found that the LCP metric might, in general, not always lead to a consistent ranking of candidates, regarding their usability as a final solution (cf. Fig. 3). We therefore used a fixed, low LCP threshold LCP_{min} to decide for each candidate whether the quick ICP stage should be entered or not (the final value used for our experiments was $LCP_{min} = 0.3$). To compute the LCP value for a given candidate transformation, we need to find the closest sample point on the CAD reference model, for each sample within the transformed range image, which we do using a k-d tree. We then measure the Euclidean distance between both points and check whether it falls below the LCP delta value δ_{LCP} (see Sec. 3). As the concrete value of δ_{LCP} should depend on the scale and density of the input data, we compute this value during the setup stage, using the average distance of each point to the closest neighbor within the sampled point data sets of the scans.

Besides δ_{LCP} , which is computed automatically, another parameter is needed to determine whether a good result has been obtained through the quick ICP stage, meaning that the global alignment procedure was successful and can be terminated. We use the Hausdorff distance for this purpose, comparing against a threshold δ_H that specifies the maximum allowed distance between both surfaces. This is the only parameter of our method that must be provided by the user. In contrast to other parameters, such as ICP, it is intuitively easy to grasp: δ_H simply measures which real-world deviations should be expected and tolerated, and which ones should be regarded as registration errors. It is worth noting at this point that this distinction is in every case necessary, since any registration algorithm used for quality control must be able to tolerate real-world deviations (which we want to detect after registration), but should not produce registration errors.

4.2. RANSAC-Based Global Alignment

Our global alignment procedure consists of two stages: the computation and testing of a candidate transformation, and a quick ICP evaluation of the result (if the result already had an LCP value exceeding LCP_{min}). To obtain a candidate transformation, we use a straightforward approach, which is similar in spirit to the RANSAC-based DARCES method of Chen et al. [CHC99]. We search three different points from the sampled range image which form a roughly equilateral triangle. The desired side length, which also determines the search radius, is computed in a similar fashion as the LCP delta value δ_{LCP} . This triangle, also referred to as the *base*, is then matched against several *candidate* triangles on the sparse point samples of the CAD reference model. Instead of using all possible candidates, we limit the search to a fixed maximum number of n_c candidates that are evaluated before a new base is selected (for our evaluation, we used $n_c = 100$). The transformation between the base triangle and each of the candidate triangles is then computed using an SVD. During this step, we also recognize and fix mirroring transformations as described by Eggert et al., using the determinant of the rotation component [ELF97].

4.3. Quick ICP Evaluation

Having a candidate transformation at hand, computed using RANSAC-based global alignment, we compute the LCP and execute the quick ICP stage, if the LCP is equal or greater than the threshold LCP_{min} . The purpose of the quick ICP evaluation is to detect whether the local alignment algorithm is converging against a local minimum. For this process, just a very few iterations of ICP are necessary, as we do not aim for perfect alignment at this point. Our ICP implementation is based on the fast variant of Low, but any fast local alignment algorithm could be used at this point [IL04].

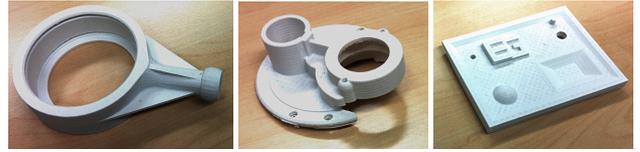


Figure 6: Three of the printed objects used for evaluation. From left to right: Mechanical Part 1, Mechanical Part 2, Specimen Block.

During the quick ICP stage, we check if the ICP algorithm successfully converged. If this was not the case, we regard the result of this stage as not successful and proceed with the testing of the next candidate transformation.

The global alignment and quick ICP stages are frequently executed until the quick ICP stage detects that a result has a sufficient quality, measured via the Hausdorff distance (see previous section).

4.4. Local Alignment and Validation

As soon as the global alignment procedure has terminated (either because a good transformation was found, or because the time budget was fully exhausted), we perform a final local alignment via ICP. At this stage, we allow for a higher number of iterations and use a stricter convergence criterion. The final result is then again checked against the maximum allowed Hausdorff distance δ_H , trying to classify whether the registration procedure was successful or not.

5. Results

Within this section, we present experimental results of our algorithm, measured on a variety of real-world data sets. We also discuss Nominal/Actual comparison as an application and highlight current limitations.

5.1. Experimental Setup

Range images used for our evaluation have been obtained from 3D printed models, which contain several typical artifacts that may arise during the layered manufacturing process. In order to detect subtle real-world deviations, a precision in the sub-millimeter range is required, and outliers may not exist in the data that is used for Nominal/Actual comparison. Therefore, we have used two different, high-precision scanning devices for quality control: a ZEISS T-SCAN hand-held laser scanner and a ZEISS COMET structured light scanner. The results do not contain significant noise or outliers. However, we performed an automatic initial cleanup in order to remove tiny, isolated pieces, which were present in the results from the COMET and resulted from parts of the turntable being scanned along with the model.

The objects used for evaluation are a CAD model of a block-shaped specimen with different geometrical features, two different mechanical parts, the well-known model of the Utah Teapot, and two famous models from the Stanford archive (Bunny and Armadillo). Some of the printed objects are shown in Fig. 6. Each object has been scanned from several directions, with both of the mentioned scanning devices. Fig. 8 shows some examples of the resulting range images.

Printers used were a *Makerbot Replicator* (Bunny, Armadillo and Teapot) and a *Big Builder Dual-Feed* (Mech. Parts and Specimen Block). For all objects, white PLA was used as a material, and the size of the printed objects was up to 20cm within each spatial dimension. The specimen block shown in Fig. 6, for example, has a size of $12\text{cm} \times 8.3\text{cm} \times 0.5\text{cm}$

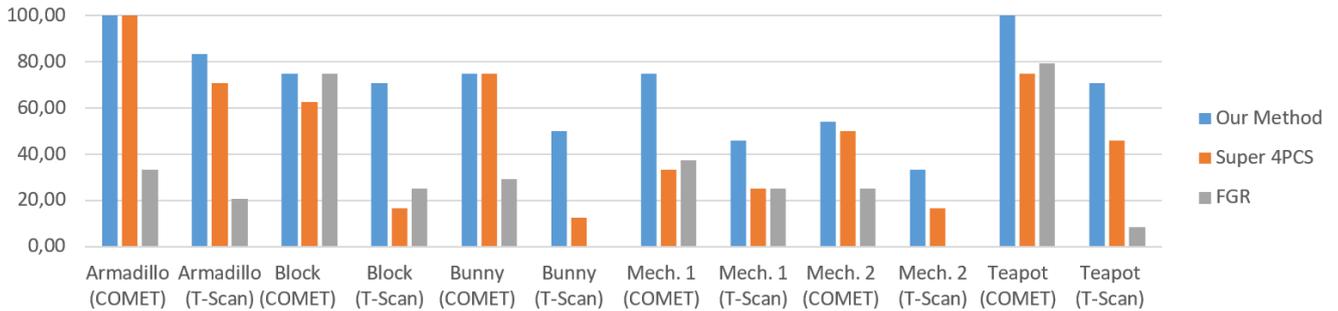


Figure 7: Registration success rate, measured as percentage of runs that achieved the correct result (the manually generated ground truth).

5.2. Performance Evaluation

For our experiments, we used 6 different models, scanned with 2 different devices, resulting in 12 different data sets, each consisting of multiple scans. We first selected a representative subset of 4 scans for each of those data sets, which resulted in 48 different range images. For each of those range images, we created a correct alignment, using our automated method and additional manual alignment steps where it was necessary. We used the user-guided global alignment functionality of the open-source *Cloud Compare* software for this purpose. Based on this ground truth, we are able to evaluate the success rate of any global registration method by comparing the residuum of the respective transformations after the final local alignment stage, using the Frobenius norm of the difference matrix.

In order to obtain representative results, we executed, for each of the methods tested, each registration test on each of the 48 range images 6 times. From the 6×4 test runs for each data set, we then computed the minimum, maximum and average success rate and running time.

The results from our experimental evaluation are shown in Fig. 7 and Table 1. For the SUPER4PCS and FGR methods, we have used the publicly available code and provided it with the same time budget as our algorithm (60s). As an LCP threshold for earlier termination of SUPER4PCS, we used 0.9, which provided the best results. While our algorithm used 2000 samples on each data set throughout, we discovered that SUPER 4PCS performs worse in this case, as the computational costs per sample are higher than for our method. We therefore left the initial setting of 200 samples per data set. The chart and table show, for each of the three methods and for each data set, the success rate (measured after a final ICP alignment, using the ground truth transformations as a reference), as well as execution times and LCP value after global alignment.

For the FGR method we used the example application provided, with the suggested radius $r = 0.1$. As the computation time of FGR and FPFH grows linearly with the number of points for each segment, we reduced the amount of data points to 10.000.

As can be seen, our method is able to successfully register, on average, 69% of the input scans. The average time needed for registration was 15 seconds. Compared with the SUPER 4PCS method, which has a success rate of 49% and needs 45 seconds on average, we achieve a significant improvement in registration performance. FGR achieved a success rate of 30% with only 2 seconds on average, so we doubled the success rate with our approach.

While the final LCP value after global alignment is typically higher than for our method, SUPER 4PCS as well as FGR does not check if a result is plausible, leading to output that potentially has a high overlap, but might not be successfully registered (cf. Fig. 3). This illustrates that the correlation between a high LCP value and

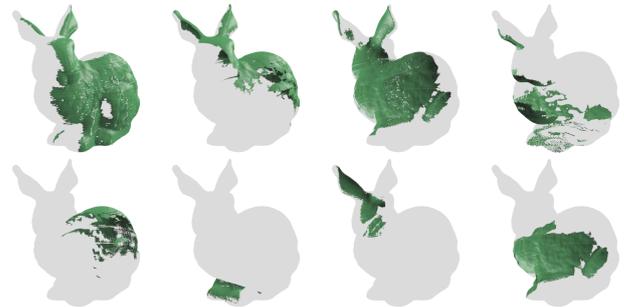


Figure 8: Partial scans from the printed Stanford Bunny, as used for our experiments. Top row: Structured light scanner (ZEISS COMET). Bottom row: Handheld laser scanner (ZEISS T-SCAN).

a good success rate after the final, local alignment is not as strong as one might expect.

It is also interesting to see that the SUPER 4PCS method performs equally well as our method on the Armadillo COMET and Bunny COMET data sets. These scans contain large surface parts and a lot of surface features (cf. Fig. 8). In contrast to the Bunny and Armadillo, the CAD objects have significantly less surface structure, more symmetries, and slippable surface parts. The authors of the original 4PCS method have already reported suboptimal results for such types of data, and we can confirm this finding (cf. also Fig. 1) [AMCO08]. Similarly FGR performs much better on the datasets produced with COMET, for some cases, even better than SUPER 4PCS. However, while the method is very fast, it is the least robust among the three evaluated methods.

Results acquired with the handheld T-Scan system are in most cases more difficult to register for both of the evaluated methods than range images from the COMET structured light scanner. We believe that this is mainly due to the difference regarding the average surface area covered within one scan (cf. Fig. 8). In contrast to the COMET, results produced by the T-SCAN system are harder to match against the CAD model as they potentially contain less significant features, and therefore more ambiguity. Fig. 1 shows examples from the Teapot data sets. Here, the COMET data has much more features and covers a much larger area of the surface. It is therefore easy to match within a very few seconds, for both methods. In contrast, the T-SCAN data consists of a much smaller part of the surface, which is harder to match against the reference model.

We have also evaluated the precision of our classification of results, based on the Hausdorff distance. After the final ICP stage, we added an additional Hausdorff check, comparing the predicted classification of the result (successfully registered / not success-

OUR APPROACH			
Data Set	Success Rate	Avg. LCP (Min./Max LCP)	Avg. Time (Min./Max. Time)
Armadillo (COMET)	100.00%	0.51 (0.41/0.88)	6.61s (0.22s / 25.65s)
Armadillo (T-Scan)	83.33%	0.55 (0.42/0.93)	24.86s (0.16s / 60.00s)
Block (COMET)	75.00%	0.70 (0.45/0.88)	1.99s (0.25s / 6.26s)
Block (T-Scan)	70.83%	0.78 (0.46/0.97)	8.81s (0.62s / 60.00s)
Bunny (COMET)	75.00%	0.52 (0.41/0.76)	18.08s (0.95s / 60.00s)
Bunny (T-Scan)	50.00%	0.59 (0.42/1.00)	23.44s (0.07s / 60.00s)
Mech. 1 (COMET)	75.00%	0.68 (0.46/0.94)	4.66s (0.12s / 25.17s)
Mech. 1 (T-Scan)	45.83%	0.63 (0.40/0.99)	19.54s (0.11s / 60.00s)
Mech. 2 (COMET)	54.17%	0.66 (0.46/0.96)	22.83s (1.56s / 60.00s)
Mech. 2 (T-Scan)	33.33%	0.67 (0.46/0.97)	33.27s (2.40s / 60.00s)
Teapot (COMET)	100.00%	0.71 (0.49/1.00)	3.40s (0.07s / 16.53s)
Teapot (T-Scan)	70.83%	0.72 (0.43/0.99)	15.69s (0.07s / 60.00s)
Total	69.44%	0.64 (0.40/1.00)	15.26s (0.07s / 60.00s)

SUPER 4PCS [MAM14]			
Data Set	Success Rate	Avg. LCP (Min./Max LCP)	Avg. Time (Min./Max. Time)
Armadillo (COMET)	100.00%	0.97 (0.88/1.00)	2.00s (0.32s / 7.51s)
Armadillo (T-Scan)	70.83%	0.72 (0.24/0.99)	45.15s (0.25s / 59.54s)
Block (COMET)	62.50%	0.89 (0.68/1.00)	59.89s (59.71s / 60.00s)
Block (T-Scan)	16.67%	0.62 (0.21/0.99)	59.90s (59.62s / 60.00s)
Bunny (COMET)	75.00%	0.98 (0.95/1.00)	2.15s (0.36s / 6.15s)
Bunny (T-Scan)	12.50%	0.49 (0.23/0.94)	59.53s (59.45s / 59.74s)
Mech. 1 (COMET)	33.33%	0.93 (0.86/1.00)	59.74s (59.62s / 60.00s)
Mech. 1 (T-Scan)	25.00%	0.76 (0.32/1.00)	59.67s (59.48s / 60.00s)
Mech. 2 (COMET)	50.00%	0.93 (0.70/1.00)	59.65s (59.48s / 60.00s)
Mech. 2 (T-Scan)	16.67%	0.51 (0.00/0.87)	59.56s (59.48s / 60.00s)
Teapot (COMET)	75.00%	0.83 (0.27/1.00)	21.85s (0.68s / 60.00s)
Teapot (T-Scan)	45.83%	0.65 (0.00/1.00)	45.01s (0.24s / 59.70s)
Total	48.61%	0.77 (0.00/1.00)	44.51s (0.24s / 60.00s)

FAST GLOBAL REGISTRATION [ZPK16]			
Data Set	Success Rate	Avg. LCP (Min./Max LCP)	Avg. Time (Min./Max. Time)
Armadillo (COMET)	33.33%	0.27 (0.04/0.67)	2.49s (1.71s / 3.26s)
Armadillo (T-Scan)	20.83%	0.21 (0.04/0.59)	2.71s (1.80s / 3.34s)
Block (COMET)	75.00%	0.49 (0.36/0.61)	1.48s (1.29s / 1.78s)
Block (T-Scan)	25.00%	0.59 (0.43/0.84)	2.37s (1.38s / 3.35s)
Bunny (COMET)	29.17%	0.50 (0.12/1.00)	1.47s (1.36s / 1.71s)
Bunny (T-Scan)	0.00%	0.15 (0.04/0.34)	1.38s (0.88s / 2.06s)
Mech. 1 (COMET)	37.50%	0.22 (0.08/0.89)	2.23s (1.98s / 2.51s)
Mech. 1 (T-Scan)	25.00%	0.13 (0.09/0.17)	1.54s (0.77s / 2.74s)
Mech. 2 (COMET)	25.00%	0.27 (0.13/0.51)	2.22s (1.62s / 2.95s)
Mech. 2 (T-Scan)	0.00%	0.21 (0.01/0.52)	3.83s (3.03s / 4.72s)
Teapot (COMET)	79.17%	0.42 (0.10/0.99)	1.28s (1.20s / 1.40s)
Teapot (T-Scan)	8.33%	0.08 (0.00/0.27)	0.93s (0.60s / 1.25s)
Total	29.86%	0.29 (0.00/1.00)	1.99s (0.60s / 4.72s)

Table 1: Experimental results for our approach and for the SUPER 4PCS and FGR method. We evaluated all methods on scans of 6 different printed objects. For each model, 8 range images were evaluated (4 for each of the two scanning devices used). To ensure representative results, each of the 48 scans was tested 6 times with each of the algorithms. In addition to success rate and timings, we show LCP values that have been computed after the global alignment stage. Results demonstrate that a high LCP value does not directly imply a high success rate.

fully registered) to the ground truth from manual registration. For our example use case of assessing the quality of 3D printed models, which have been manufactured at a scale of 5cm to 15cm side length (bounding box), we found that a maximum allowed deviation of 3mm worked best. Table 2 shows the results. As can be seen, our classification using the Hausdorff distance is even slightly more reliable than a classification via LCP, although the difference (especially in false negative rate) seems not to be significant.

5.3. Error Visualization

Using the registered range images, we were able to perform Nominal/Actual comparisons on the data sets, in order to assess deviations that occurred during the printing process. To visualize deviations, we have chosen to measure the distance of each data point within the range image to the CAD surface. This was quickly and efficiently done by casting rays following the surface normals, in both directions. Clearly, different methods are possible, all coming with different advantages and disadvantages (consider

Scan Method	succ. rate (gr. truth)	FPR / FNR (Hausdorff)	FPR / FNR (LCP)
Armadillo (COMET)	1.00	0.00 / 0.00	0.00 / 0.00
Armadillo (T-Scan)	0.83	0.00 / 0.00	0.00 / 0.00
Block (COMET)	0.75	0.00 / 0.25	0.00 / 0.00
Block (T-Scan)	0.70	0.57 / 0.00	0.17 / 0.06
Bunny (COMET)	0.75	0.00 / 0.10	0.00 / 0.00
Bunny (T-Scan)	0.50	0.00 / 0.33	0.00 / 0.00
Mech. 1 (COMET)	0.75	0.50 / 0.00	0.75 / 0.10
Mech. 1 (T-Scan)	0.45	0.54 / 0.00	0.64 / 0.15
Mech. 2 (COMET)	0.54	0.17 / 0.28	1.00 / 0.22
Mech. 2 (T-Scan)	0.33	0.25 / 0.00	0.40 / 0.43
Teapot (COMET)	1.00	0.00 / 0.00	0.00 / 0.00
Teapot (T-Scan)	0.70	0.00 / 0.00	0.00 / 0.15
Total	0.69	0.28 / 0.09	0.33 / 0.08

Table 2: Classification performance for our validation method using the Hausdorff distance, compared to an LCP-based approach. A lower false positive rate (FPR) is achieved by the Hausdorff distance, with nearly similar false negative rate (FNR) as LCP.

for example the recent work of Getto et al. for more advanced approach [GKL15]) In principle, any method that produces signed distance values at each data point can be employed at this point. To visualize the signed distance values, we used a common color mapping with a rainbow-like transfer function. While the use of this function has been criticized for various reasons, we found it well-suited in this case, since our deviation is not a unipolar scale (low vs. high), but a bipolar one (cf. [BT107]).

Fig. 9 shows some examples, illustrating different kinds of errors that arised from the layered manufacturing process. As can be seen from these test cases, we can efficiently detect and visualize several forms of deviation from the reference model, such as ringing artifacts, small knobs on the surface or bending of parts of the model. The latter effect can be especially observed for objects that have a large base area, which will be printed as the first few layers. In such cases, the hot, freshly printed material cooling down after its contact with the base plate of the printer leads to bending of those regions. A practical solution to cope with this problem is, for example, to use a printer with a heatable plate instead.

5.4. Limitations

While our method performs much better for surfaces with less structure, such as the mechanical parts, it should also be noted that we are, on the other hand, not able to robustly handle noisy data. This is especially true for distant outliers, which do not permit to use our validation method based on the Hausdorff distance. We did not find this a problem for our main use case of Nominal/Actual comparison, using high-precision range data: here, any additional noise from the scanning process could not be distinguished from actual deviations (cf. Fig. 9), therefore noisy scans are generally not found to be useful for quality control. However, the ability to deal with noisy scans, as very successfully demonstrated for the 4PCS / SUPER 4PCS and FGR methods, is an important requirement in a lot of other scenarios, such as registration of partially overlapping, noisy range images. We therefore conclude that our algorithm is well-suited for Nominal/Actual comparisons, especially of mechanical parts, while other methods are better at dealing with noisy data.

Another, basic registration problem, which our algorithm inherits from the basic rigid alignment procedure, is the handling of input data with smooth deformations. This problem originates from the minimization of squared distance across the surface during the final, local refinement stage: rather than allowing regions to have larger deviations while leaving other regions in place, ICP distributes the error smoothly, which is - in this case - the wrong solution. Fig. 10 illustrates this problem by example, using the bottom of the specimen block model (viewed from below). Here, heavy bending of the bottom area of the part led to a strong and smooth deformation. During our evaluation, this case was recognized as a failure, since it was found to be too different from the manually obtained ground truth. Similar problems occurred for the mechanical part 2 object, for which the evaluated algorithms showed the smallest success rate. Solving such problem efficiently, fully-automated fashion is a challenging task, and it is unclear how a reliable solution could look like. As we already do locally, using the parameter δ_H , any solution would need to reliably model the expected large, global deviations, in order to distinguish between expected deviations and registration errors.

6. Conclusion

Within this paper, we have summarized some existing approaches for registration of range images against 3D scans. Then we demonstrated, that a basic RANSAC strategy with some modifications

speeds up registration time and success rate for high precise scans without noise. Our main use case is automated registration for Nominal/Actual comparison during quality control. Our key finding is that the common LCP measure is error-prone, regarding the classification of results during global alignment. We have therefore employed an additional quick ICP stage during the main RANSAC procedure, along with a validation based on the Hausdorff distance. For our evaluation, we have used different kinds of 3D-printed models, which have been scanned with two different professional high-precision devices. We demonstrated that our method is able to outperform the current state of the art when considering both primary goals, registration speed and registration performance (success rate).

Different directions are possible for future work. For the sampling step, clearly, different sampling strategies than pure random sampling are possible. We believe that a very regular sparse sampling could lead to less features getting captured in total, compared to random samples. In order to capture important geometric features for matching, we could therefore go one step further and weight the samples by some form of saliency measure, for example using the recent method of Tasse et al. [TKD15]. Such a weighted sampling, however, would also have an impact other stages of the pipeline, for example the search for triangles inside the sampled point data during global alignment. Therefore, this topic needs further investigation. The investigation of alternative methods for local registration, as proposed by Pottmann et al., also seems a promising direction, in order to speed up the reoccurring local alignment task within our algorithm [PHYH06]. The good quality of our results is mainly due to our Hausdorff-based classification method, which allows us to reject bad candidates at early stages of registration, but also to judge the quality of the outcome to a certain extent. In the future, it will be interesting to investigate how a conservative classification of results could look like. Given the possibility of ambiguous registration results, this is a very challenging task. We believe that recognizing such ambiguities, for example by using symmetry detection or sliding window methods during registration, is a key aspect within this challenge.

References

- [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. 2, 4, 5
- [AScE02] ASPERT N., SANTA-CRUZ D., EBRAHIMI T.: Mesh: Measuring errors between surfaces using the hausdorff distance, 02 2002. 3
- [ATT97] AKUTSU T., TAMAKI H., TOKUYAMA T.: Distribution of distances and triangles in a point set and algorithms for computing the largest common point sets. In *Proceedings of the Thirteenth Annual Symposium on Computational Geometry* (1997), SCG '97, ACM, pp. 314–323. 2
- [BI15] BIRDAL T., ILIC S.: Point pair features based object detection and pose estimation revisited. In *Proceedings of the 2015 International Conference on 3D Vision* (2015), 3DV '15, IEEE Computer Society, pp. 527–535. 2
- [BM92] BESL P. J., MCKAY N. D.: A method for registration of 3-d shapes. *IEEE Trans. Pattern Anal. Mach. Intell.* 14, 2 (Feb. 1992), 239–256. 1, 2
- [Bri07] BRIDSON R.: Fast poisson disk sampling in arbitrary dimensions. In *ACM SIGGRAPH 2007 Sketches* (2007), SIGGRAPH '07, ACM. 3
- [BT107] BORLAND D., TAYLOR II R. M.: Rainbow color map (still) considered harmful. *IEEE Comput. Graph. Appl.* 27, 2 (Mar. 2007), 14–17. 7
- [BTP13] BOUAZIZ S., TAGLIASACCHI A., PAULY M.: Sparse iterative closest point. In *Proc. SGP* (2013), SGP '13, Eurographics Association, pp. 113–123. 2
- [CHC99] CHEN C.-S., HUNG Y.-P., CHENG J.-B.: RANSAC-Based DARCES: A new approach to fast automatic registration of partially overlapping range images. *IEEE Trans. Pattern Anal. Mach. Intell.* 21, 11 (Nov. 1999), 1229–1234. 2, 4

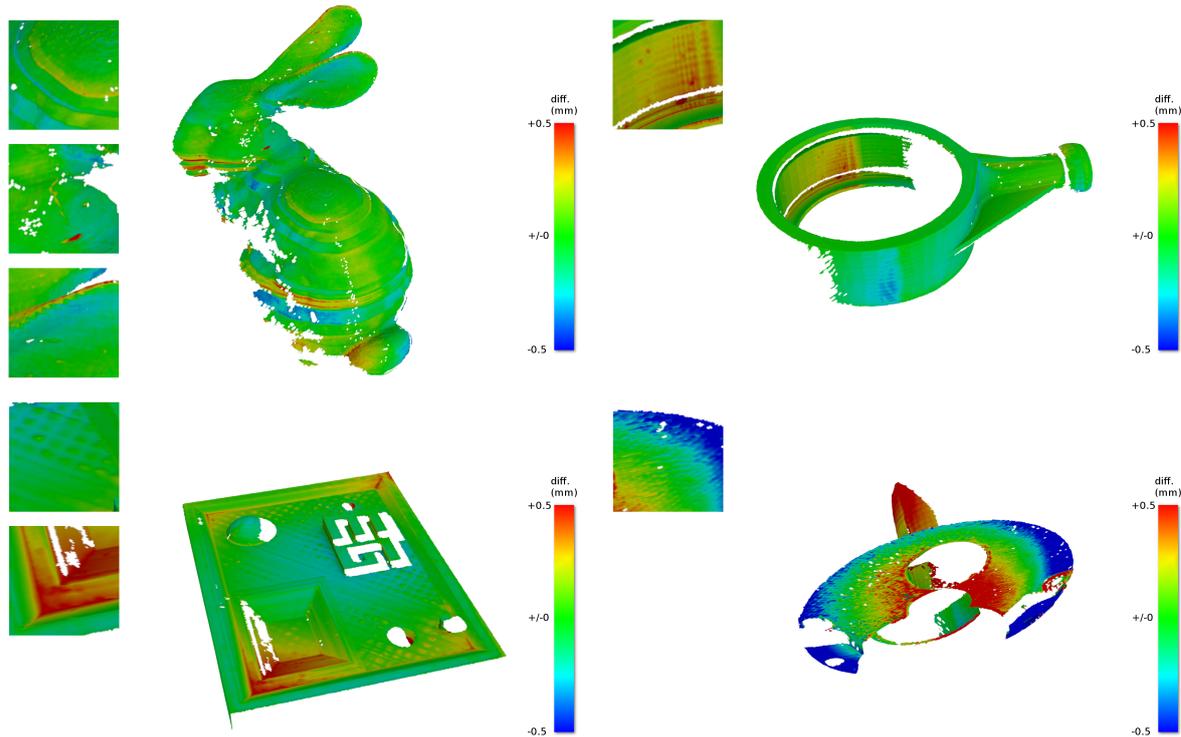


Figure 9: Visualization of different errors, originating from the 3D printing process. Data sets have been registered within a few seconds, using our method. Artifacts that can be seen include (from top left to bottom right): Ringing artifacts, originating from layering, along with small droplets of material and deviations from the desired thickness in steep regions (Bunny), small knobs from temporary congestion of the printing nozzle (Mech. Part 1), peaks on the thin surface, originating from the original support structure, plus slight bending on the edges of the model (Block), as well as strong bending of the base area of the model (Mech. Part 2).

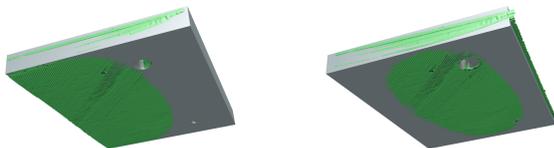


Figure 10: Problems arising from global deformations under rigid matching. Left: manually aligned ground truth. Right: alignment after local optimization via ICP.

[CM92] CHEN Y., MEDIONI G.: Object modelling by registration of multiple range images. *Image Vision Comput.* 10, 3 (Apr. 1992), 145–155. 2

[DUNI10] DROST B., ULRICH M., NAVAB N., ILIC S.: Model globally, match locally: Efficient and robust 3d object recognition. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on* (June 2010), pp. 998–1005. 2

[ELF97] EGGERT D. W., LORUSSO A., FISHER R. B.: Estimating 3-d rigid body transformations: A comparison of four major algorithms. *Mach. Vision Appl.* 9, 5-6 (Mar. 1997), 272–290. 4

[GKL15] GETTO R., KUIJPER A., LANDESBERGER T.: Extended surface distance for local evaluation of 3d medical image segmentations. *Vis. Comput.* 31, 6-8 (June 2015), 989–999. 7

[HK90] HUTTENLOCHER D. P., KEDEM K.: Computing the minimum Hausdorff distance for point sets under translation. In *Proceedings of the Sixth Annual Symposium on Computational Geometry* (1990), SCG '90, ACM, pp. 340–349. 3

[IL04] LIM LOW K.: *Linear least-squares optimization for point-to-plane ICP surface registration*. Tech. rep., 2004. 2, 4

[LR09] LIU Y.-S., RAMANI K.: Robust principal axes determination for

point-based shapes using least median of squares. *Comput. Aided Des.* 41, 4 (Apr. 2009), 293–305. 2

[MAM14] MELLADO N., AIGER D., MITRA N. J.: Super 4pcs fast global pointcloud registration via smart indexing. *Computer Graphics Forum* 33, 5 (2014), 205–215. 2, 3, 4, 6

[PHYH06] POTTMANN H., HUANG Q.-X., YANG Y.-L., HU S.-M.: Geometry and convergence analysis of algorithms for registration of 3d shapes. *Int. J. Comput. Vision* 67, 3 (May 2006), 277–296. 2, 7

[PLH04] POTTMANN H., LEOPOLDSEDER S., HOFER M.: Registration without ICP. *Comput. Vis. Image Underst.* 95, 1 (July 2004), 54–71. 1, 2

[RBB09] RUSU R. B., BLODOW N., BEETZ M.: Fast point feature histograms (fpfh) for 3d registration. In *2009 IEEE International Conference on Robotics and Automation* (May 2009), pp. 3212–3217. 2

[RL01] RUSINKIEWICZ S., LEVOY M.: Efficient variants of the icp algorithm. In *3-D Digital Imaging and Modeling, 2001. Proceedings. Third International Conference on* (2001), pp. 145–152. 1, 2, 3

[TKD15] TASSE F. P., KOSINKA J., DODGSON N.: Cluster-based point set saliency. In *Proceedings of the 2015 IEEE International Conference on Computer Vision (ICCV)* (2015), ICCV '15, IEEE Computer Society, pp. 163–171. 7

[Tur90] TURK G.: *Graphics gems*. Academic Press Professional, Inc., 1990, ch. Generating Random Points in Triangles, pp. 24–28. 3

[YLCJ16] YANG J., LI H., CAMPBELL D., JIA Y.: Go-icp: A globally optimal solution to 3d icp point-set registration. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 38, 11 (Nov 2016), 2241–2254. doi:10.1109/TPAMI.2015.2513405. 2

[ZPK16] ZHOU Q., PARK J., KOLTUN V.: Fast global registration. In *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part II* (2016), pp. 766–782. 2, 6