



Optimal, Non-Rigid Alignment for Feature-Preserving Mesh Denoising

Gawrilowicz, Florian; Bærentzen, Jakob Andreas

Published in:
Proceedings of the International Conference on 3D Vision

Link to article, DOI:
[10.1109/3DV.2019.00053](https://doi.org/10.1109/3DV.2019.00053)

Publication date:
2019

Document Version
Peer reviewed version

[Link back to DTU Orbit](#)

Citation (APA):
Gawrilowicz, F., & Bærentzen, J. A. (2019). Optimal, Non-Rigid Alignment for Feature-Preserving Mesh Denoising. In *Proceedings of the International Conference on 3D Vision IEEE*.
<https://doi.org/10.1109/3DV.2019.00053>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Optimal, Non-Rigid Alignment for Feature-Preserving Mesh Denoising

Florian Gawrilowicz

flgw@dtu.dk

J. Andreas Bærentzen

<http://people.compute.dtu.dk/janba>

Technical University of Denmark, Dept. of Applied Mathematics and Computer Science
Kgs. Lyngby, Denmark

Abstract

We present a simple, yet effective method for non-rigid alignment of point clouds. Our focus lies on developing a practical approach that allows us to do efficient, multi-way alignment of millions of points such as those produced by structured light scanners. Starting from an initial displacement field over the combined point cloud, our solution relies on an iterative smoothing scheme on the neighborhood graph of each sub-scan, reducing the Dirichlet energy of the displacement field. We compare a number of schemes for computing the initial displacement field, ranging from estimating the Laplacian of the combined point clouds to more traditional measures such as the point to point distance or point to plane distance.

1. Introduction

Optical surface scanning technologies ranging from Lidar over time-of-flight cameras to structured light scanners allow us to acquire digital models of many types of objects for a vast range of applications. However, almost all of these applications require us to scan each object from several directions since each scan only covers the object as seen from a particular point. Often, we are interested in a closed surface model for applications such as additive manufacturing, and even in cases where we do not need a complete 3D model, concavities and self-occlusion make it impossible to capture the entire region of interest from a single vantage point. For these reasons, we invariably have to deal with point clouds that consist of multiple sub-scans, each covering a part of the object. It is crucial these are aligned precisely.

Substantial progress has been made when it comes to rigid alignment of such point clouds. In a typical processing pipeline, global, rigid registration is followed by a local, rigid alignment. However, rigid alignment cannot account

for errors caused by faulty intrinsic calibration or due to motion artifacts during the scanning process. These problems cause a non-rigid distortion of the scanned surfaces. The misalignment results in high-frequency artifacts when the combined point cloud is later on triangulated to produce a surface model. Clearly, we can remove the noise using smoothing methods, but while feature preserving smoothing methods are available, features and noise might well overlap in the frequency domain making it effectively impossible to remove the noise without harming features. The problem is illustrated in Figure 2 where three overlapping sub-scans are shown as triangle meshes in 2a, 2b, and 2c. In Figure 2d, a triangulation of the combined point cloud is seen. It is clear that the combined point cloud contains high-frequency noise, which is not present in the original sub-scans. These high-frequency components (as already observed in [5]) are caused by a misalignment of the sub-scans that could not be corrected through rigid alignment.

We approach this problem via a scheme that specifically removes the high-frequency noise that arises when sub-scans are combined into a single point-cloud. Our approach is to apply a combinatorial reconstruction algorithm, e.g., [4], to obtain a (noisy) triangle mesh of the combined scans. We then cast the non-rigid alignment as a feature preserving smoothing of the triangle mesh performed after the reconstruction.

We know that much of the noise in the reconstruction is introduced by the misalignment of the sub-scans. Thus, we constrain the smoothing to be an almost rigid deformation of each sub-scan. To be more specific, we first compute displacement vectors for each vertex in the combined point-cloud. This can be done in a number of ways, and the natural approach is to compute the Laplacian – assuming that we have a triangulation of the point set. In the next step, we apply the smoothing based on the computed displacement vectors. However, we constrain the displacement of each vertex to be coherent within its respective sub-scan.

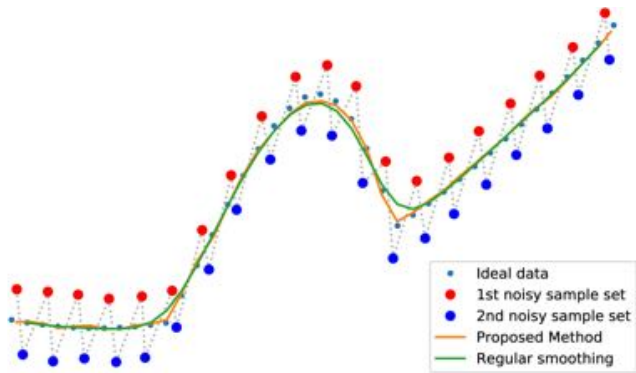


Figure 1: To illustrate the principle of our method the red and blue samples are translated, and noise is added. Simple Laplacian smoothing yields the green curve. The smoothing removes the translations and reduces the noise, but it also blurs the sharp features. In comparison our proposed method (orange curve) enforces coherence of the displacement vectors within each sample set. This preserves features much better at the expense of less noise reduction.

The process is illustrated in 2D in Figure 1. The light-blue dots show samples of the original signal. Another set of samples of this signal has been corrupted by noise and then translated to emulate misalignment: the red dots have been moved up and the blue dots have been moved down. Consequently, the curve connecting the samples (red and blue) is very erratic. We can remove the noise and jaggedness by regular smoothing, but this leads to significant blurring of the sharp features as illustrated by the green curve. The orange curve illustrates our method which more effectively preserves the features of the original signal while also removing noise.

In the following, we assume the individual sub-scans to be already rigidly aligned - globally and locally. Besides the aforementioned approaches to global alignment, scans are often already coarsely aligned - either through the structure-from-motion (SfM) step necessary in a multi-view-stereo (MVS) pipeline or directly measured while capturing, for example through a turntable or robotic arm. Furthermore, we assume our sub-scans also to be locally aligned, for example, by classical rigid ICP.

2. Related Work

There exist a plethora of algorithms for registration or alignment of 3D data. Giving a comprehensive overview is clearly beyond the scope of this paper, but we refer to [27] for a comparative survey on rigid and non-rigid alignment techniques for point clouds and meshes.

The continued work in the field shows the lack of a general solution to the various use cases of this fundamental idea. The literature on non-rigid registration differs vastly

in the underlying motivation and application. Early work focused on fitting a single scan and a predefined template, e.g., for parametrization of human bodies [1]. This continues to be a challenging task as finding correct correspondences is crucial [17, 30]. Commonly, only a small number of feature points is therefore considered and based on those a deformation model is fitted, e.g., using thin-plate splines [9] or isometric deformations [15]. The approaches in [6] and [23] divide the point cloud into clusters, which are then individually aligned via rigid ICP. Spatial smoothing of the resulting transformations results in a global transformation function, which aligns the scans with a template or base model.

Feature-point based methods are popular for global alignment, allowing for arbitrarily translations and rotations [13, 26, 19, 34]. If multiple scans need to be aligned, the transformations need to be consistent [8]. For the combined task of global, non-rigid alignment of multiple scans, feature-points are used as well [5]. In our work, we assume the transformations to be small, i.e., we are only concerned with the final step of the alignment and assume global alignment as a separate task.

A wide range of feature preserving smoothing techniques exist, e.g. [18, 22, 12, 14] and generally use prior assumptions to constrain the smoothing. The specific properties to differentiate noise from surface features either has to be carefully picked and tuned or estimated from ground truth data, like it has been done in *Mesh Denoising via Cascaded Normal Regression* (CNR) [31].

Our approach to feature preserving smoothing is similar to the deformation in the optimal step framework of [2], but here we apply the technique to multi-way alignment. It is also closely related to *High Fidelity Scan Merging* (HiFi) [10] – in the sense that our approach shares the idea of combining smoothing on the merged scans as well as on the partial sub-scans. However, their approach does not explicitly optimize an error function.

2.1. Contribution

Our method allows for feature-preserving smoothing of multiple partial scans. When a mesh is reconstructed from a set of sub-scans using a combinatorial method, i.e., a method which preserves the input points as vertices of the output mesh, information about the scanning process is retained. As mentioned, any misalignment between the sub-scans then manifests itself as high-frequency noise in the output mesh. We propose to formulate feature preserving smoothing after reconstruction on the output mesh as non-rigid alignment. Feature preservation is based on the observation that the displacements due to smoothing should vary smoothly over the sub-scans since the error that requires a non-rigid alignment is assumed to result in a coherent deformation of the sub-scans.




Figure 2: Three separate sub-scans from different viewing directions (a-c) with individual meshes resulting from 2D Delaunay triangulation in the corresponding image plane. When combining all points (aligned with rigid ICP) into a global mesh the non-rigid misalignment results in visual noise as seen in sub-figure (d)

Many methods for non-rigid alignment, e.g., [1, 6, 17, 23], address the case of aligning a single set of source points to a target shape. In order to apply them to the case where there are multiple scans, an incremental approach would be necessary. However, this has the inherent drawback that the result depends on the order in which the sub-scans are added. This arbitrary choice would lead to accumulating errors. Furthermore, it bears the peril of initializing with a scan with particularly pronounced distortion.

Thus, our goal and main contribution is a method that aligns all sub-scans jointly and symmetrically in order to reduce high-frequency noise in the final mesh. We build on the framework introduced in [2] and show how to apply it to multiple scans jointly. Furthermore, we extend this framework with an error function focused on denoising the resulting output mesh with a regularization using the sub-scan information. The resulting algorithm is able to correct the misalignment of multiple scans comprised of millions of points in only a few minutes.

3. Method

For non-rigid alignment, we are concerned with two driving forces. We want to achieve proper alignment of the positions of the points resulting in a smooth global mesh when triangulated. Nevertheless, we also want to move the points coherently, i.e., we likewise want the displacement vector field during this process to be smooth. Intuitively the procedure applies Laplacian smoothing to the final mesh, but before actually moving the vertices, we smooth the resulting displacements across the each sub-scan.

3.1. Continuous Energy

At the core of our approach lies the idea of reducing the local variations measured by the Dirichlet energy on a domain Ω , which, in this context, are 2-manifolds embedded in \mathbb{R}^3 . For the two energies in non-rigid alignment, we have two distinct domains. One being the manifold Θ given by the surface of the scanned object, which is approximated by the surface reconstruction step. This constructs a global mesh of the merged scans. The other one, Ξ , is given by the

individual sub-scans and the triangulation of them.

Driving the alignment of the point cloud is a smoothing process, which seeks to reduce the Dirichlet energy on the global manifold Θ

$$E_a[p] = \int_{\Theta} \|\nabla p(x)\|^2 dA, \quad (1)$$

where function value $p(x)$ is the position in the embedding space. We are not interested in the actual minimum of this energy, but rather take gradient steps

$$\frac{\partial}{\partial p} E_a = -2\Delta p \quad (2)$$

to yield a smoother surface, giving the classical Laplacian smoothing [28].

Our displacement energy then measures how close our actual displacement $d(x)$ is to a given target displacement $t(x)$, obtained from, but not limited to, the gradient mentioned above eq.(2).

$$E_d[t] = \int_{\Theta} \|d(x) - t(x)\|^2 dA. \quad (3)$$

As mentioned earlier, we also want to enforce a coherent movement of the points in each of the sub-scans in order to limit distortion. To that end, we minimize the Dirichlet energy of the displacement vector field $t \in \mathbb{R}^3$ associated with each point x on the k manifolds seen by each of the k sub-scans

$$E_c[t] = \sum_k \int_{\Xi_k} \|\nabla d(x)\|^2 dA. \quad (4)$$

Our combined energy is now simply a weighted sum of the energies

$$E_t = E_d + \alpha E_c, \quad (5)$$

and its gradient is

$$\frac{\partial}{\partial d} E_t = d(x) - t(x) + \alpha \sum_i \Delta_{\Xi_i} d(x). \quad (6)$$

3.2. Discrete Formulation

A triangulated surface S , with vertices $V = \{v_i\}$, $i = 1, \dots, N$ and edges $E = \{e_{ij}\}$, defines a graph $G = (V, E)$. The discrete Laplace operator on that graph can be expressed as a matrix L such that

$$(Lu)_i = \sum_j w_{ij}(u_i - u_j), \quad (7)$$

where u_i is the value at vertex v_i . There are multiple variants of the discrete Laplace operator discerned by their weights ω_{ij} [33]. We use the uniformly weighted Laplacian given by $\omega_{ij} = \frac{1}{d_i}$, where d_i is the degree of vertex i .

For the alignment of multiple surfaces S_k each corresponding to one individual scan, we have disjoint edge sets E_k . Indexing all points of all sub-scans jointly we can gather the corresponding Laplacian operators in one matrix L_Ξ . The edges are generated by a Delaunay triangulation of the 2D positions in the image plane of each sub-scan. We remove edges that are longer than ten times the average edge length to avoid separate pieces of the scanned surface to be dependent on each other. The discrete formulation of eq. (5)

$$E_t = \|D - T\|^2 + \alpha \|M_\Xi D\|^2, \quad (8)$$

where M_Ξ is the incidence matrix with respect to the sub-scans. The gradient is then expressed as

$$\nabla E_t = D - T + \alpha M_\Xi^T M_\Xi D \quad (9)$$

$$= D - T + \alpha L_\Xi D, \quad (10)$$

where $T \in \mathbb{R}^{N \times 3}$ is the matrix of target displacements, and $D \in \mathbb{R}^{N \times 3}$ the displacements actually applied after optimization, akin to [2] but formulated for the multi-way alignment case. The second equality in eq.(9) only holds for the unweighted discrete Laplacian.

After applying the calculated translations D , the target displacements change with the updated positions. For that reason we repeat this process for a fixed, small number of iterations or until the smallest eigenvalues of the local covariance matrices (see Figure 8a) fall below a certain threshold, which is a parameter to our algorithm.

If we were to do simple gradient descent, starting with an initial guess of using the target displacements, i.e., $D_0 = T$, we would obtain a gradient equal to $\alpha L_\Xi T + 2T$ and a resulting translation of

$$T_1 = 3T + \alpha L_\Xi T \quad (11)$$

This effectively defines a smoothing process of the target translations over the individual sub scans. Although we used a conjugate gradient solver in practice, it can thus be seen that the non-rigid alignment is conceptually just smoothing the displacements, resulting from a smoothing operator, with respect to the proximity in the sub-scans.

3.3. Target Displacements

As input to the optimization framework, we need a target displacement T . As our goal is denoising, we evaluate displacement vector fields generated by smoothing operators for meshes as well as point clouds.

GLP Graph-Laplacian smoothing only uses the connectivity information of the mesh discarding any spatial information. As a simple approximation to the Laplace-Beltrami operator, it is the natural choice to approximate eq (1) in case a global mesh is given.

MCM Mean Curvature Motion on point clouds [11], is computationally more demanding as it fits a total least squares (TLS) regression plane to the neighborhood around each point and projects the point onto it. The plane passes through the centroid of the k -nearest neighbors and its normal points in the direction of least variance.

Centroid Instead of estimating a regression plane, simply moving each point to the centroid of the k -nearest neighbors directly is another way to induce a displacement vector field. It also relates to the Graph-Laplacian, but instead of using the edges of a mesh as edges in the graph, the edges connect to the k -nearest neighbors around each point.

P2P The classical point-to-point correspondence [3] is essentially equivalent to the centroid of the k -nearest neighbors, where $k = 1$. Although some additional heuristics have been incorporated to reject false or misleading correspondences [25]. For non-rigid alignment, it is not well-suited as it lumps points together (see Figure 6a), but it is included for completeness.

P2PI Point-to-plane (P2PI) [7] correspondence is widely used for ICP. It is similar to MCM in that it also restricts the displacements to the normal directions. Furthermore, the standard way of normal estimation via Eigen decomposition of the covariance matrix is equivalent to TLS regression [20]. The only difference is the point the regression plane passes through, which in the case of P2PI is the centroid of the k -nearest neighbors with $k = 1$, and not n .

4. Results and Comparison

We evaluated our method on a set of four different scans depicted in Figure 3. The famous *Stanford Bunny* was taken from the *The Stanford 3D Scanning Repository* [29]. The three other models (Owl, Facade, and Skull) were scanned with a structured light scanner while placed on a turntable to obtain pre-aligned sub-scans.

In Figure 3 each of the sub-scans is shown in a distinct color. All of the scans show substantial non-rigid misalignment, albeit being carefully calibrated. This becomes obvious in the top row showing the combined sub-scans after rigid alignment via ICP. The bottom row shows our non-rigid alignment process based on the Graph-Laplacian substantially reduces the error for all models. In case of the Skull, the inside of the nose has only been scanned from a single direction. This is visible as an orange patch before and after the non-rigid alignment.

Figure 4 and the top row of 5 show that the non-rigid alignment error is substantially reduced. Laplacian smoothing [21] alone introduces far more blurriness before the high-frequency noise caused by the misalignment vanishes as shown in the bottom row of 5. It also becomes apparent that *Bilateral mesh denoising* [12] models the high level of noise as features and even after 50 iterations taking 48 minutes cannot fully remove the misalignment noise. *High Fidelity Scan Merging* (HiFi) [10] is able to align the sub-scans well, taking 390 seconds for five iterations, resulting in less remaining noise. However, it still is not able to align the sub-scans as well as ours, and a more visible pattern can be seen in Figure 7 indicating larger deviations in between the sub-scans. The final image in Figure 4 shows the result of *Cascaded Normal Regression* (CNR) [31]. There is also some remaining misalignment visible after 412 seconds of processing, as the misalignment noise gets transformed to lower frequencies, resulting in vertical streaks. It should also be noted that using point to plane (P2Pl) displacements boils down to the method proposed in [2].

Table 1 lists timing for the alignment on the different models. A Intel® Xeon® E5-2660v3 CPU with 10 cores was used during processing. The adjacency matrix for the sub-scans had to be assembled for all approaches in order to compute L_{Ξ} . This adds a pre-processing time listed in the fourth column of 1 to all variants. The meshing step is of course only mandatory for the GLP displacements. We use the implementation of the *Co3Ne* algorithm [4] available in Geogram [16]. Timings for reconstructing the global mesh are listed in the fifth column. All of the variants for generating the displacement vector fields, except for P2P which produces lumped points, reduce the misalignment. A clear advantage of the GLP based approach is a substantial reduction in computation time.

The same set of parameters is employed for all models shown in Figure 3. We use the 16 nearest neighbors for computing the MCM, as well as for the centroid. P2P and P2Pl use an upper bound on the euclidean distance of 18-times the average closest point distance. Running ten iterations with a stiffness factor α between 15 and 30 gave good results on all models.

In contrast to the other approaches for calculating the displacements, GLP does not require to recompute the

neighborhood relation in each iteration. Furthermore, using the uniform Graph-Laplacian alleviates us from recomputing the weights in each iteration. For P2P and Centroid, it is necessary to update the kD-tree used for nearest neighbors search adding computation time. When using MCM or P2Pl, an additional Eigen decomposition has to be computed for each point. These factors accumulate to 20% to 60% reduced iteration time for GLP compared to the other methods.

Although the setting and the data is different in [23] as they are concerned with aligning a single scan to a template and the data is not publicly available, our implementation in Python achieves timings on the same order. Comparing, for instance, their results for aligning a point cloud of 8.7 million points in 221 seconds with the 279 seconds for our Skull model with 7.2 million points. The more related case discussed in [6] of aligning the 1406 scans of the *The Digital Michelangelo Project* containing 28 million vertices took 1.5 hours on a cluster of 60 nodes.

One thing to note is that when a global mesh is initially given, using the GLP displacements is not only faster but also avoids folding triangles as can be seen in Figure 6. It inherits the mesh fairing properties of the Graph-Laplacian [21] and is the only approach of the ones evaluated that directly results in a proper mesh.

Figure 8 gives a quantitative comparison between the classical energy terms based on Point to Point (P2P) and Point to Plane (P2Pl), the proposed methods Centroid, Mean Curvature Motion (MCM), Graph Laplacian (GLP), and simple Laplacian smoothing. Figure 8a shows the quality of the alignment measured by the smallest eigenvalue of the local covariance matrix of the merged points. All methods show a similar result indicating a successful reduction in noise. Laplacian smoothing shows the most rapid decrease as there is no regularizing term. The Dirichlet energy in Figure 8b is computed with respect to the global mesh. For the Point to Point based method this energy term notably increases after an initial decline as a result of drastic tangential drift of the points. The method is oblivious to the mesh as correspondences are only based on proximity. On the other hand, the Graph Laplacian is explicitly accounting for this and thereby able to reduce the energy beyond the other regularized methods.

5. Discussion

In this paper, we proposed a method bridging non-rigid registration of partial scans and feature preserving mesh denoising. While coherent displacements preserve features much better, it is still effective at removing the noise due to alignment errors.

We have compared several procedures for computing the displacements that drive the smoothing process, and we find that the Laplacians tend to give the best quality, yet also the

Table 1: Processing times for setup & 10 iterations of the alignment process in seconds for the models shown in 3.

	#Scans	#Points	Preprocessing		Alignment				
			Adjacency	Meshing	MCM	GLP	Centr.	P2P	P2PI
Bunny	10	362k	3.32	2.36	28.47	15.40	24.00	19.27	28.38
Owl	3	692k	15.73	7.47	71.49	27.89	58.92	55.25	69.87
Facade	3	3.8M	77.43	40.53	284.77	150.77	229.08	189.83	262.76
Skull	18	7.2M	40.50	58.53	566.59	279.00	443.95	378.18	526.25



Figure 3: Color coded triangulation of the sub-scans aligned with rigid ICP (top row), and non-rigid ICP with Graph-Laplacian (GLP) based displacements (bottom row).

best run-time performance.

Our work relies on two essential assumptions. The first of these is that we already have a rigid alignment. The energies we minimize are very local and cannot establish correspondences between points far apart. Fortunately, it appears that modern methods for rigid alignment [24, 32] make this assumption reasonably easy to justify. Our second assumption, in case of the GLP based energy, is that we have triangle mesh connectivities for both the combined point cloud and for each sub-scan. Fortunately, the partial scans can generally be triangulated using 2D Delaunay triangulation in the image plane, and several combinatorial reconstruction algorithms are available for reconstruction of mesh connectivity from the combined point cloud, e.g., [4].

If these assumptions are met, our post-reconstruction non-rigid alignment procedure can be applied. Compared to merely smoothing out the noise due to misalignment or (almost equivalently) performing a volumetric reconstruction, the benefit is that we are able to preserve features more accurately. We see the utility of our method predominantly as the last step in the pipeline – somewhat unusually – after meshing of the points.

When comparing to the two related classes of algorithms, i.e., feature-preserving mesh denoising and non-rigid alignment, we arrive at the best of both worlds. Non-rigid alignment operates on point clouds, ignoring the mesh

connectivity. Feature-preserving denoising uses the topology information of the mesh, but ignores the knowledge about the scanning process. Our method uses the fact that we want each input scan to deform almost rigidly as an informal prior, turning an approach for non-rigid registration into feature preserving smoothing. It also should be noted that more sophisticated smoothing algorithms are straightforward to include in our framework. In summary, by using both sources of information, we obtain superior results in the frequent scenario of creating high fidelity meshes from scanned data.

Acknowledgements

The models were rendered with Open3D [35].

References

- [1] B. Allen, B. Curless, and Z. Popović. The space of human body shapes: Reconstruction and parameterization from range scans. *Acm Transactions on Graphics*, 22(3):587–594, 2003. 2, 3
- [2] B. Amberg, S. Romdhani, and T. Vetter. Optimal step non-rigid icp algorithms for surface registration. *Proceedings of the Ieee Computer Society Conference on Computer Vision and Pattern Recognition*, page 4270190, 2007. 2, 3, 4, 5
- [3] P. J. Besl and N. D. McKay. Method for registration of 3-d shapes. In *Sensor Fusion IV: Control Paradigms and Data Structures*, volume 1611, pages 586–607. International Society for Optics and Photonics, 1992. 4
- [4] D. Boltcheva and B. Lévy. Surface reconstruction by computing restricted voronoi cells in parallel. *Computer-Aided Design*, 90:123 – 134, 2017. SI:SPM2017. 1, 5, 6, 8
- [5] B. Brown and S. Rusinkiewicz. Global non-rigid alignment of 3-D scans. *ACM Transactions on Graphics (Proc. SIG-GRAPH)*, 26(3), Aug. 2007. 1, 2
- [6] B. J. Brown and S. Rusinkiewicz. Non-rigid range-scan alignment using thin-plate splines. In *Proceedings. 2nd International Symposium on 3D Data Processing, Visualization and Transmission, 2004. 3DPVT 2004.*, pages 759–765. IEEE, 2004. 2, 3, 5
- [7] Y. Chen and G. Medioni. Object modelling by registration of multiple range images. *Image and vision computing*, 10(3):145–155, 1992. 4
- [8] S. Choi, Q.-Y. Zhou, and V. Koltun. Robust reconstruction of indoor scenes. In *Proceedings of the IEEE Conference*




Figure 4: Color coded triangulations: each sub-scan is triangulated individually by 2D Delaunay triangulation in the image plane and displayed in a distinct color. The actual 3D positions of the vertices are result to the respective method for generating the displacement vectors after 5 iterations, see 3.3.

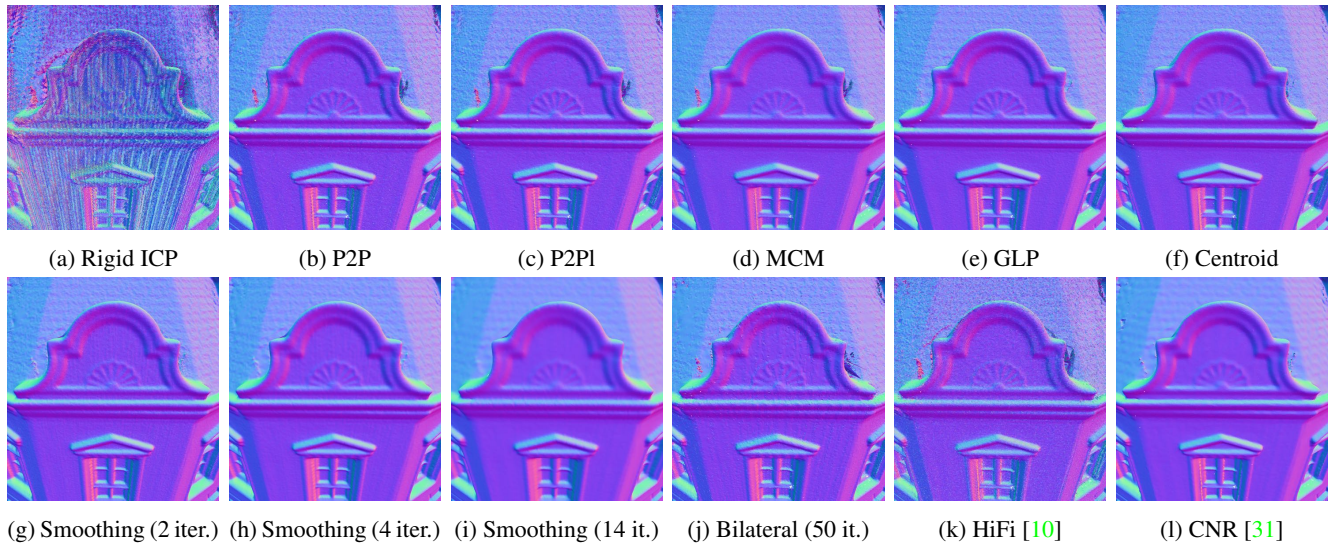


Figure 5: Color coded normals: close-up on a detail of the facade model comparing rigid ICP to non-rigid alignment with different displacement and Laplacian smoothing. Also results of *Bilateral mesh denoising* [12], *High Fidelity Scan Merging* (HiFi) [10], and *Cascaded Normal Regression* (CNR) [31] are shown.

on *Computer Vision and Pattern Recognition*, pages 5556–5565, 2015. 2

- [9] H. Chui and A. Rangarajan. A new point matching algorithm for non-rigid registration. *Computer Vision and Image Understanding*, 89(2):114 – 141, 2003. Nonrigid Image Regis-

tration. 2

- [10] J. Digne, J.-M. Morel, N. Audfray, and C. Lartigue. High fidelity scan merging. *Computer Graphics Forum*, 29(5):1643–1651, 2010. 2, 5, 7, 8
- [11] J. Digne, J.-M. Morel, C.-M. Souzani, and C. Lartigue. Scale




Figure 6: Effect of non-rigid alignment with different displacements on the initial global mesh reconstructed with *Co3Ne* [4] from the original pointcloud. The color encodes the direction of the interpolated face normals.


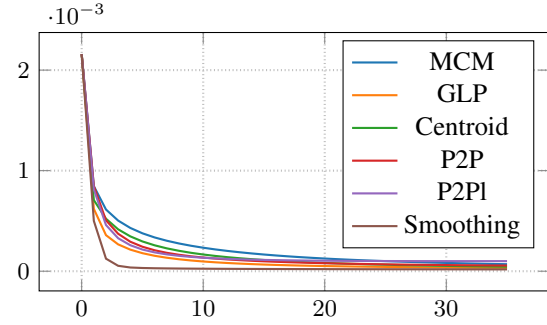



Figure 7: Comparing ours (left) to [10] (right) after 5 iterations each, it can be seen that we achieve a better alignment visible in smaller colored patches representing different partial scans.

space meshing of raw data point sets. *Computer Graphics Forum*, 30(6):1630–1642, 2011. 4

- [12] S. Fleishman, I. Drori, and D. Cohen-Or. Bilateral mesh denoising. In *ACM transactions on graphics (TOG)*, volume 22, pages 950–953. ACM, 2003. 2, 5, 7
- [13] N. Gelfand, N. J. Mitra, L. J. Guibas, and H. Pottmann. Robust global registration. In *Symposium on geometry processing*, volume 2, page 5, 2005. 2
- [14] H. Huang, S. Wu, M. Gong, D. Cohen-Or, U. Ascher, and H. R. Zhang. Edge-aware point set resampling. *ACM transactions on graphics (TOG)*, 32(1):9, 2013. 2
- [15] Q. Huang, B. Adams, M. Wicke, and L. J. Guibas. Non-rigid registration under isometric deformations. In *Computer Graphics Forum*, volume 27, pages 1449–1457, 2008. 2
- [16] B. Lévy. Geogram. 5
- [17] H. Li, R. W. Sumner, and M. Pauly. Global correspondence optimization for non-rigid registration of depth scans. In *Computer graphics forum*, volume 27, pages 1421–1430. Wiley Online Library, 2008. 2, 3
- [18] E. Mattei and A. Castrodad. Point cloud denoising via moving rpca. *Computer Graphics Forum*, 36(8):123–137, 2017. 2
- [19] N. Mellado, D. Aiger, and N. J. Mitra. Super 4pcs fast global pointcloud registration via smart indexing. In *Computer Graphics Forum*, volume 33, pages 205–215. Wiley Online Library, 2014. 2
- [20] N. J. Mitra and A. Nguyen. Estimating surface normals in noisy point cloud data. In *Proceedings of the nineteenth*



(a) Mean of the smallest eigenvalue of the local covariance matrices.



(b) Dirichlet energy of the vertex positions E_a eq. (1)

Figure 8: Change in the *Facade* scans across 35 iterations

annual symposium on Computational geometry, pages 322–328. ACM, 2003. 4

- [21] A. Nealen, T. Igarashi, O. Sorkine, and M. Alexa. Laplacian mesh optimization. In *Proceedings of the 4th international conference on Computer graphics and interactive techniques in Australasia and Southeast Asia*, pages 381–389. ACM, 2006. 5
- [22] A. C. Öztireli, G. Guennebaud, and M. Gross. Feature preserving point set surfaces based on non-linear kernel regression. In *Computer Graphics Forum*, volume 28, pages 493–501. Wiley Online Library, 2009. 2
- [23] G. Palma, T. Boubekeur, F. Ganovelli, and P. Cignoni. Scalable non-rigid registration for multi-view stereo data. *ISPRS Journal of Photogrammetry and Remote Sensing*, 142:328–341, 2018. 2, 3, 5
- [24] S. Rusinkiewicz. A symmetric objective function for ICP.

- ACM Transactions on Graphics (Proc. SIGGRAPH)*, 38(4), 7 2019. 6
- [25] S. Rusinkiewicz and M. Levoy. Efficient variants of the icp algorithm. In *3-D Digital Imaging and Modeling, 2001. Proceedings. Third International Conference on*, pages 145–152. IEEE, 2001. 4
 - [26] R. B. Rusu, N. Blodow, and M. Beetz. Fast point feature histograms (fpfh) for 3d registration. In *Robotics and Automation, 2009. ICRA'09. IEEE International Conference on*, pages 3212–3217. Citeseer, 2009. 2
 - [27] G. K. L. Tam, Z.-Q. Cheng, Y.-K. Lai, F. C. Langbein, Y. Liu, D. Marshall, R. R. Martin, X.-F. Sun, and P. L. Rosin. Registration of 3d point clouds and meshes: A survey from rigid to nonrigid. *Ieee Transactions on Visualization and Computer Graphics*, 19(7):1199–1217, 2013. 2
 - [28] G. Taubin. A signal processing approach to fair surface design. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, pages 351–358. ACM, 1995. 3
 - [29] G. Turk and M. Levoy. The stanford 3d scanning repository. 4
 - [30] O. Van Kaick, H. Zhang, G. Hamarneh, and D. Cohen-Or. A survey on shape correspondence. In *Computer Graphics Forum*, volume 30, pages 1681–1707. Wiley Online Library, 2011. 2
 - [31] P.-S. Wang, Y. Liu, and X. Tong. Mesh denoising via cascaded normal regression. *ACM Transactions on Graphics (SIGGRAPH Asia)*, 35(6), 2016. 2, 5, 7
 - [32] Y. Wang and J. M. Solomon. Deep closest point: Learning representations for point cloud registration. *CoRR*, abs/1905.03304, 2019. 6
 - [33] M. Wardetzky, S. Mathur, F. Kälberer, and E. Grinspun. Discrete laplace operators: no free lunch. In *Symposium on Geometry processing*, pages 33–37. Aire-la-Ville, Switzerland, 2007. 4
 - [34] Q.-Y. Zhou, J. Park, and V. Koltun. Fast global registration. In *European Conference on Computer Vision*, pages 766–782. Springer, 2016. 2
 - [35] Q.-Y. Zhou, J. Park, and V. Koltun. Open3D: A Modern Library for 3D Data Processing. *arXiv e-prints*, page arXiv:1801.09847, Jan 2018. 6