

Fractured Object Reassembly via Robust Surface Registration

P. Mavridis, A. Andreadis and G. Papaioannou

Athens University of Economics and Business




Figure 1: Automatic reassembly of fractured cultural heritage objects from their parts using the proposed pipeline.

Abstract

The reassembly of fractured 3D objects from their parts is an important problem in cultural heritage and other domains. We approach reassembly from a geometric matching perspective and propose a pipeline for the automatic solution of the problem, where an efficient and generic three-level coarse-to-fine search strategy is used for the underlying global optimization. Key to the efficiency of our approach is the use of a discretized approximation of the surfaces' distance field, which significantly reduces the cost of distance queries and allows our method to systematically search the global parameter space with minimal cost. The resulting reassembly pipeline provides highly reliable alignment, as demonstrated through the reassembly of fractured objects from their fragments and the reconstruction of 3D objects from partial scans, showcasing the wide applicability of our methodology.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Geometric algorithms, languages, and systems

1. Introduction and Related Work

The problem of fractured object reassembly can be approached from a surface matching perspective. Under the assumption that two matching fragments share a common contact surface, finding the optimal aligning transformation of the fragments is equivalent to finding the optimal alignment of the corresponding contact surfaces.

A recent survey on surface alignment methods can be found in Tam et al. [TCL*13]. However, in the context of Cultural Heritage, this problem poses significantly more challenges than common geometric alignment, since physical erosion of the contact surfaces, along with small missing parts, create a large number of outliers. For this reason,

specialized approaches have been proposed. A first solution for the reassembly of general 3D solids was proposed by Papaioannou et al. [PKT01], however the method assumes that fracture facets are nearly planar and match each other to a wide extent. Huang et al. [HFG*06] proposed a feature-based alignment method that properly handles these cases and provides remarkable results, but the proposed system is rather complicated, consisting of several specialized algorithms for segmentation, multi-scale feature extraction, correspondence determination, registration, collision detection and supervised learning, something that makes the implementation and adoption of the system difficult. Winkelbach et al. [WW08] also proposed a reassembly method based on

a branch-and-bound search heuristic, although they focus on the global pairwise matching, while other parts of a reassembly pipeline, like multi-part matching, are not addressed.

Contributions. Our reassembly pipeline avoids the complexity of previous approaches by using general, robust and well-understood geometric matching methods, with minimal application-specific additions. To avoid negatively affecting the efficiency of the system by the generality of the algorithms, we provide several insights on how to improve the overall performance by using highly efficient state-of-the-art data structures and by introducing a three-level coarse-to-fine search strategy for the geometric matching that combines the desirable characteristics of several algorithms.

Overview. Our method first extracts a set of potentially fractured facets from each one of the input fragments (Sec. 2), then calculates their pairwise matching score (Sec. 3) and finally computes the set of pairwise matches that lead to the desired reassembly (Sec. 4).

2. Fracture Facet Extraction

For the facet extraction our method uses a standard region-growing procedure, which produces regions separated by sharp edges, using as the stopping criterion the deviation of averaged normals. Regions with a high variance in the *bending energy* [HFG*06] are then classified as potential fracture facets and are tagged for matching against fractured surfaces of other fragments. Unlike previous methods, our approach does not rely on sophisticated segmentation procedures, since the subsequent geometric matching step is robust-enough to handle potential inconsistencies and partial overlap between the extracted facets.

3. Pairwise Rigid Geometric Registration

The next stage of the pipeline performs a pairwise matching operation for all pairs of fracture facets that belong to distinct fragments. The matching score is based on the residual distance between two fragments after their rigid registration, which is performed using a three-level coarse-to-fine search strategy that we present in the remainder of this section.

Mathematical Formulation. Pairwise rigid registration methods search for the transformation that optimally aligns a *source* surface \mathcal{X} to a *target* one \mathcal{Y} . When the two surfaces are discretized this problem can be formulated as:

$$\arg \min_{\mathbf{R}, \mathbf{t}} \sum_{i=1}^n \phi(\mathbf{R}\mathbf{x}_i + \mathbf{t}, \mathcal{Y}), \quad (1)$$

where n is the number of vertices in the source surface, $\mathbf{R} \in \mathbb{R}^{3 \times 3}$ denotes a rigid rotation matrix and $\mathbf{t} \in \mathbb{R}^3$ is a translation vector. The function $\phi(\mathbf{x}, \mathcal{Y})$ measures the distance of an arbitrary point $\mathbf{x} \in \mathbb{R}^3$ to the surface \mathcal{Y} and is defined as:

$$\phi(\mathbf{x}, \mathcal{Y}) = \min_{\mathbf{y} \in \mathcal{Y}} \|\mathbf{x} - \mathbf{y}\|, \quad (2)$$




Figure 2: Alignment of two fragments using ℓ_p -norm minimization on the fractured surfaces. Missing parts on the left fragment create outliers, resulting in sub-optimal registration when using the ℓ_2 -norm. The contribution of large distances is reduced with ℓ_p -norms, making the method robust to outliers for sufficiently low values of the parameter p .

where the metric $\phi(\mathbf{x}, \mathcal{Y})$ measures the distance between two points in space. This equation is often referred to as the *distance transform* or the *distance field* of surface \mathcal{Y} .

Many methods use the squared Euclidian norm as the distance metric and optimize Equation 1 using a least squares formulation. However, in the presence of outliers, an optimiser will skew the solution in order to reduce the large penalty associated with distant points, as shown in Figure 2. To avoid this problem, similar to recent work in the field [BTP13], we define the distance metric in Equation 2 as:

$$\phi(\mathbf{x}, \mathcal{Y}) = \mu_p(\|\mathbf{x} - \mathbf{y}\|_2), \quad \mu_p(x) = |x|^p.$$

Using this ℓ_p -norm formulation of the problem, our method becomes more robust to outliers. To further improve the robustness, we also *trim* a small percentage, 5%, of the elements with the largest residual distances.

Discretized Distance Field. The minimization of Equation 1 requires many evaluations of the distance function (Equation 2), thus an efficient data structure is required to minimize the associated cost. To this end, the distance function is discretely sampled on a 3D grid that extends over the narrow band of the target surface and is stored using a sparse hierarchical volumetric data structure, the VDB [Mus13]. Subsequent distance queries are simply resolved by fetching the pre-computed values from this structure, which is a key to the efficiency of our approach. For manifold objects we compute a *signed* distance field representation, where interior points have a negative distance. The memory consumption of this data structure depends on the resolution of the distance field volume. In our tests we have used volumes that consume roughly 40 MBytes of memory.

Coarse Initial Alignment. When the surface area of two facets is roughly equal, we perform a coarse initial alignment of the centroids and the average normals of the facets. Otherwise, we compute the alignment using a standard RANSAC-based alignment procedure that is based on 3-point con-

gruent sets [CHC99]. The 4-point congruent sets method [AMCO08] provided similar results and might be preferable in more noisy datasets. Because of this initial alignment, the original pose of the fragments does not influence the accuracy of the final alignment.

Simulated Annealing. In the next stage, the alignment transformation is parameterized using three free variables for the translation and three (euler angles) for the rotation. These variables are initialized with the alignment of the previous stage and a *Simulated Annealing* method is used to further minimize the residual distance, as measured in Equation 1. The range of the search is restricted around the alignment of the previous stage. Our implementation uses the *Enhanced Simulated Annealing* [SBDH97] method, configured to perturb one variable at a time. To increase the efficiency, the residual distance evaluation of this stage uses a uniformly sampled subset of points from the source surface.

Local Refinement. In the last geometric alignment stage, the computed registration is locally refined using the sparse ICP algorithm [BTP13]. The choice of Sparse ICP is an important one, because unlike previous approaches, it can handle the outliers created from missing parts or eroded surfaces without introducing any application-specific heuristics or weighting functions. Furthermore, this stage of the pipeline uses accurate distance queries, instead of the discretized distance field, to avoid compromising the quality of the final alignment and to allow the use of low resolution distance field volumes in the previous stages.

Non-penetrating Registration. When using unconstrained registration, it is expected that the final alignment will exhibit mutual penetrations of the source and target surfaces (see Figure 3 - left). However, for the reassembly of fractured objects a penetration-free alignment might be desired, depending on the accuracy of the scanned data. In our approach, penetrations can be detected very efficiently using the sign of the distance field. When a penetration is detected, the corresponding candidate alignment is rejected, by assigning a sufficiently high value to the cost function. Similarly, arbitrary alignment constraints can be integrated in our pipeline. However, since the inclusion of such constraints is not trivial for ICP methods, when arbitrary constraints are required, we perform the local refinement using the ESA method with a very small search range, instead.

4. Multi-part Matching

During the last stage of our pipeline we determine the set of pair-wise matches that result in a plausible reassembly. Similar to [Hub02], we construct a graph, where each fragment corresponds to a node, and each pair-wise match to an edge between two nodes. The optimal set of pair-wise connections is computed by finding the *Minimum Spanning Tree* of the graph, using the well-known Kruskal’s algorithm, with the additional constraint that, when merging two sub-forests,



Figure 3: Left: Unconstrained registration yields mutual penetrations of the source and target contact surface, which might be undesirable in object reassembly applications. Right: Inclusion of constraints for non-penetrating registration. To better illustrate the mutual penetrations, the fragment in front is rendered with transparency.

Data Structure	T_p	154K	77K	38K	9K
VDB	1.9	25.5	9.3	4.1	1.8
ANN	0.01	78.4	22.8	9.1	2.4
k-D Tree	0.01	890.6	236.8	67.6	8.5

Table 1: Average facet alignment time when using the discretized distance field (VDB), Approximate Nearest Neighbor (ANN) and a k-D tree, w.r.t. the number of source surface points. T_p : precomputation time of the corresponding data structure (Target points: 155K). Times in seconds.

the selected connection should not create geometric penetrations. If this is the case, the corresponding edge is dropped from the graph and the algorithm continues with the next best edge. In some cases, this procedure might not lead to the desired reassembly. For this reason, the user can also manually include or exclude graph edges from the solution.

5. Results

Figures 1, 2 and 3 demonstrate our pipeline on the reassembly of fractured cultural heritage objects. Most fragments are eroded and small parts are missing or damaged, increasing the difficulty of the underlying alignment problem. Even in these challenging cases, our approach provides the desired reassembly. All results use the $\ell_{0.4}$ -norm.

Performance. Our measurements in Table 1 indicate that the use of a precomputed distance field offers large efficiency gains over the more traditional kd-tree or *Approximate Nearest Neighbor* (ANN) data structures. ANN was configured to provide roughly the same accuracy as the VDB structure. The time required to create the discretized distance field is higher than the other data structures, but this cost can be amortized over many alignment queries, which is typical for problems involving more than two fragments.

Sparse ICP can be directly used in order to align the fragments, provided that the coarse initial alignment (Sec 3) is good enough. However, as shown in Table 2, this approach has a very slow convergence rate. Our pipeline vastly

Algorithm	N_s	T_s	T_{total}
Ours (ESA & S-ICP)	2	2.74	4.12
S-ICP only	191	96.9	97.1
ℓ_2 -ICP & S-ICP	57	30.2	31.4

Table 2: The performance of our three-level alignment strategy compared to other approaches for the Figure 2 dataset (Source: 120K points, Target: 150K points). All methods include the same initial alignment stage and provide roughly the same accuracy. N_s : Sparse ICP iterations. T_s : Sparse ICP run time. T_{total} : Total run time. Times in seconds.




Figure 4: Reconstruction of an object from partially overlapping scans using our pipeline. In this case we omit the facet extraction stage, which is specific to the fractured object reassembly. The insets show the input partial scans in their final aligned position. The initial position of the scans does not influence the registration results.

increases the alignment efficiency, by incorporating an intermediate ESA optimization step, which performs larger steps than Sparse ICP, and can approach the optimal solution much faster than Sparse ICP, albeit with less accuracy. Another option to approach the solution faster is to perform a traditional ℓ_2 -ICP before Sparse ICP, but our measurements indicate that this approach is less efficient. All experiments were conducted on an Intel Core i7-3820 CPU at 3.6GHz.

Generality. By omitting the fracture facet extraction, our pipeline can be used to solve more general geometric matching problems, such as the reconstruction of 3D objects from partial scans, as shown in Figure 4. In this case, since the input surfaces are non-manifold, the VDB data structure stores an unsigned distance field, instead of a signed one. The rest of the algorithm remains the same. It is worth noting that we do not explicitly disable the inter-penetration checks, since for unsigned distance fields these tests are always negative.

Limitations. For heavily eroded objects where the fracture surfaces do not touch each other, the problem might be better approached using semi-automatic user-guided methods [PPCS13], instead of automatic ones. Furthermore, our pipeline is designed to handle general 3D objects. For special

cases, such as pottery sherds or planar fragments, specialized algorithms are expected to be more efficient.

6. Conclusions

We have presented a new pipeline for the reassembly of fractured objects. Unlike previous state-of-the-art reassembly methods, our approach computes the desired alignment between two fragments without relying on the computation of features, which can be unreliable on noisy scans or eroded fragments and can also increase the complexity of a method. We have also provided insights on how to improve the geometric matching performance by using efficient data structures and by introducing a three-level coarse-to-fine search strategy for the pair-wise alignment.

Acknowledgements. The authors would like to thank Ioannis Pratikakis from Athena Research Centre for providing the dataset in Figure 4. This work was funded by EC FP7 STREP project “PRESIOUS”, grant no. 600533.

References

- [AMCO08] AIGER D., MITRA N. J., COHEN-OR D.: 4points congruent sets for robust pairwise surface registration. In *ACM SIGGRAPH 2008 Papers* (New York, NY, USA, 2008), SIGGRAPH '08, ACM, pp. 85:1–85:10. 3
- [BTP13] BOUAZIZ S., TAGLIASACCHI A., PAULY M.: Sparse iterative closest point. *Computer Graphics Forum (Symposium on Geometry Processing)* 32, 5 (2013), 1–11. 2, 3
- [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. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 21, 11 (1999). 3
- [HFG*06] HUANG Q.-X., FLÖRY S., GELFAND N., HOFER M., POTTMANN H.: Reassembling fractured objects by geometric matching. In *ACM SIGGRAPH 2006 Papers* (New York, NY, USA, 2006), SIGGRAPH '06, ACM, pp. 569–578. 1, 2
- [Hub02] HUBER D. F.: *Automatic three-dimensional modeling from reality*. PhD thesis, Carnegie Mellon University, 2002. 3
- [Mus13] MUSETH K.: Vdb: High-resolution sparse volumes with dynamic topology. *ACM Trans. Graph.* 32, 3 (July 2013). 2
- [PKT01] PAPAIOANNOU G., KARABASSI E.-A., THEOHARIS T.: Virtual archaeologist: Assembling the past. *IEEE Computer Graphics and Applications* 21 (2001), 53–59. 1
- [PPCS13] PALMAS G., PIETRONI N., CIGNONI P., SCOPIGNO R.: A computer-assisted constraint-based system for assembling fragmented objects. In *Proc. of Digital Heritage 2013 International Congress* (2013), vol. 1, IEEE, pp. 529–536. 4
- [SBDH97] SIARRY P., BERTHIAU G., DURDIN F., HAUSSY J.: Enhanced simulated annealing for globally minimizing functions of many-continuous variables. *ACM Trans. Math. Softw.* 23, 2 (June 1997), 209–228. 3
- [TCL*13] TAM G., CHENG Z.-Q., LAI Y.-K., LANGBEIN F., LIU Y., MARSHALL D., MARTIN R., SUN X.-F., ROSIN P.: Registration of 3d point clouds and meshes: A survey from rigid to nonrigid. *Visualization and Computer Graphics, IEEE Transactions on* 19, 7 (July 2013), 1199–1217. 1
- [WW08] WINKELBACH S., WAHL F. M.: Pairwise matching of 3d fragments using cluster trees. *Int. J. Comput. Vision* 78, 1 (June 2008), 1–13. 1