Facet Extraction and Classification for the Reassembly of Fractured 3D Objects

A. Andreadis P. Mavridis G. Papaioannou

AUEB Computer Graphics Group

PRESIOUS http://presious.eu/

Problem statement: The reassembly of fractured 3D objects is a critical problem in computational archaeology. An essential part of this problem is to identify which facets of a fragment are fractured. A general strategy to solve this region classification problem is to first divide the geometry into regions and then classify each one as intact or fractured, based on statistical properties.

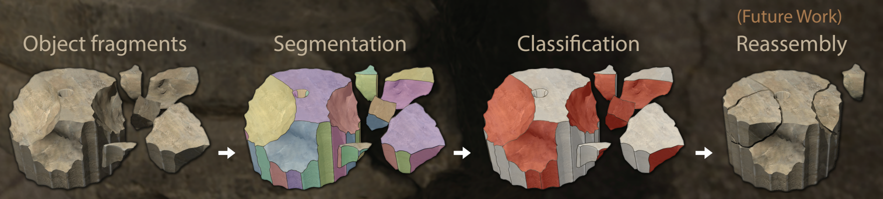

Contributions:
1. Comparative evaluation of some well-known segmentation strategies in the context of reassembly, in terms of performance and quality of segmentation.
2. A novel method for the classification of the segmented regions into intact and fractured ones, based on their statistical properties.

Classification Results

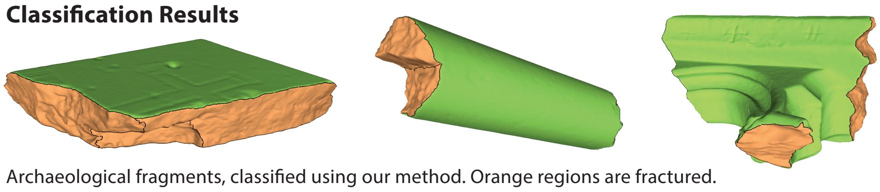

Archaeological fragments, classified using our method. Orange regions are fractured.

1. Segmentation - Distance Metrics

Global: The angle between the average normals of two segments.

Local: The angle between average normals of two segments computed on the local neighbourhood at their common border.

Comments: The global metric performs well on planar surfaces, but it results in over-segmentation on curved ones. This can be aleviated using the local metric.

Region Growing - Naive

- Grows one cluster at a time
- Merges a random neighboring element with distance below a threshold
- Creates a new cluster, when no compatible neighbouring element can

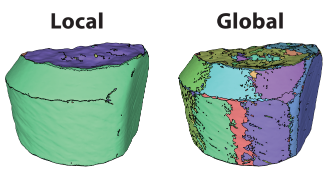

1600 clusters 2847 clusters

Region Growing - Best First

- Grows one cluster at a time
- Merges the neighbouring element with the closest distance
- Creates a new cluster, when no compatible neighbouring element can be found

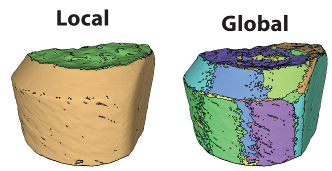

1863 clusters 1863 clusters

Hierarchical Agglomerative

- Starts with every element as a cluster
- Merges the two nearest clusters with respect to the metric used
- Stops when the minimum distance is higher than a threshold

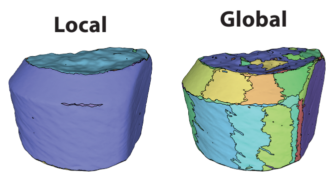

521 clusters 933 clusters

2. Post-Processing

The greedy nature of the merging algorithms can lead to severe over-segmentation. This is fixed by a custom post-processing step that first decomposes small regions into single elements, which are subsequently merged to the nearest neighbouring segments.

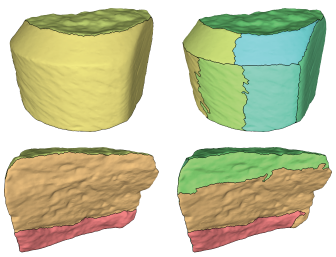

5 clusters 11 clusters

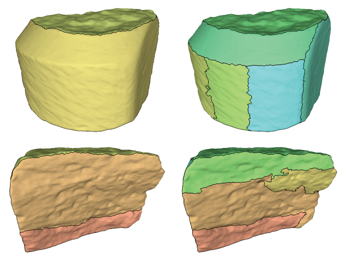

5 clusters 11 clusters

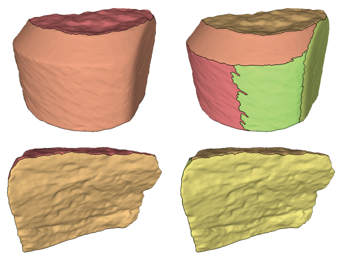

4 clusters 7 clusters

3. Classification

In order to discriminate regions to fractured and intact ones, we estimate the surface roughness using the Sphere Volume Integral Invariant. A semi-automatic machine learning approach is used to classify segments as fractured or intact.

Classifier: AdaBoost

TP	FN
89.3%	10.7%
FP	TN
21.7%	78.7%

Cross-Validation Performance

Orange: Fractured
Green: Intact

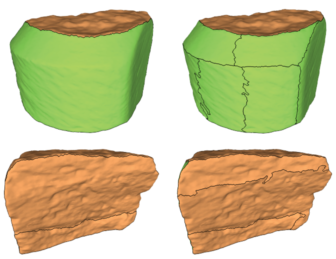

3.123sec 3 Fractured 2 Intact — 1.582sec 4 Fractured 7 Intact

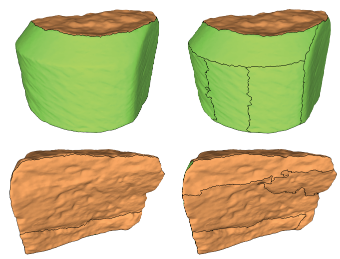

26.91sec 3 Fractured 4 Intact — 2.310sec 5 Fractured 6 Intact

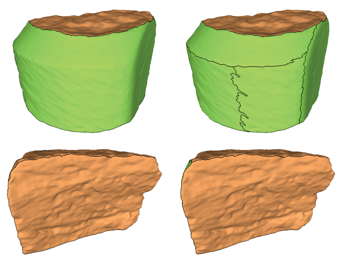

177.1sec 2 Fractured 2 Intact — 4.001sec 2 Fractured 5 Intact

Conclusions

Our results indicate that the choice of a distance metric has a far greater impact on the segmentation quality than choosing an optimal order of operations. A robust post-processing is essential for making region growing practical, since omitting this step leads to a large number of segments.

Acknowledgments

The datasets used in this poster are from the PRESIOUS project data collection. This work was supported by EC FP7 STREP Project PRESIOUS, grant no. 600533.

http://presious.eu/resources/3d-data-sets

Contact information: A. Andreadis, anthousis@aueb.gr