

Geometry Based Faceting of 3D Digitized Archaeological Fragments

Hanan ElNaghy, Leo Dorst

Computer Vision Group, Informatics Institute, University of Amsterdam
The Netherlands

{hanan.elnaghy, l.dorst}@uva.nl

Abstract

We present a robust pipeline for segmenting digital cultural heritage fragments into distinct facets, with few tunable yet archaeologically meaningful parameters.

Given a terracotta broken artifact, digitally scanned in the form of irregularly sampled 3D mesh, our method first estimates the local angles of fractures by applying weighted eigenanalysis of the local neighborhoods. Using 3D fit of a quadratic polynomial, we estimate the directional derivative of the angle function along the maximum bending direction for accurate localization of the fracture lines across the mesh. Then, the salient fracture lines are detected and incidental possible gaps between them are closed in order to extract a set of closed facets. Finally, the facets are categorized into fracture and skin. The method is tested on two different datasets of the GRAVITATE project.

1. Introduction

Our research is part of the H2020 GRAVITATE project¹, which stands for Geometric Reconstruction And noVel semantic reunification of cultural heritage objects. GRAVITATE [9] is intended to create a software platform that will allow archaeologists and curators to digitally reassemble fragmented cultural heritage objects, to identify and reunify parts of a cultural object that have been separated across collections and to recognize associations between cultural artifacts that will allow new knowledge and understanding of past societies to be inferred.

The initial GRAVITATE data collection consists of terracotta artifacts that have been broken long ago with highly abraded, faded and cracked surfaces and partially eroded fracture edges. The terracotta fragments are digitally scanned in the form of irregularly sampled 3D meshes through different archiving and scanning practices.

During the breaking of an archaeological artifact, its surface is incidentally divided into two main regions: skin

(original surface) and fracture. Each of these regions may be subdivided into distinct facets, where each facet is characterized by its own geometrical properties of roughness and sharpness of its boundary. A facet is considered a semantically relevant subpart of a given fragment within the GRAVITATE pipeline. We will refer to the process of producing these facets as 'faceting'.

One of the most significant delineations of the facets is the fact that they are bordered by sharp fracture lines of characteristic breaking angles. Detecting such lines and measuring the angles forming them are essential for later partitioning of the fragment into distinct facets. However in case of highly abraded fragments, such fracture lines are not as contiguous as one might hope and may not reflect the actual angle of the real fracture fold when the original archaeological object broke.

Within the context of GRAVITATE, faceting is a core preprocessing step towards much analysis by different modules. The fracture facets are used for complementarity mating, while the skin facets are used for local continuity of color and patterns belonging to the original surface. Moreover, the skin facets could be further analyzed for enriching the semantic annotation of the archaeological fragments through detecting different decorative patterns, semantic features and manufacturing styles.

In this paper, a complete faceting pipeline is proposed for segmenting 3D archaeological fragments into distinct contiguous facets. It may also be useful for other applications beyond GRAVITATE. Only geometrical properties of the 3D archaeological fragments are considered in our work rather than also colorimetry or material. First, we define a bi-plane fracture model based on least square fitting of two eigen-planes and measurement of their internally encompassed angle. Then, locally salient fracture lines are detected for contouring the facet curves and the gaps between them are closed. Finally, each facet is classified as either being fracture or skin facet depending on its local geometrical surface properties. Figure 1 shows the flow of the presented faceting pipeline.

¹<http://gravitate-project.eu>

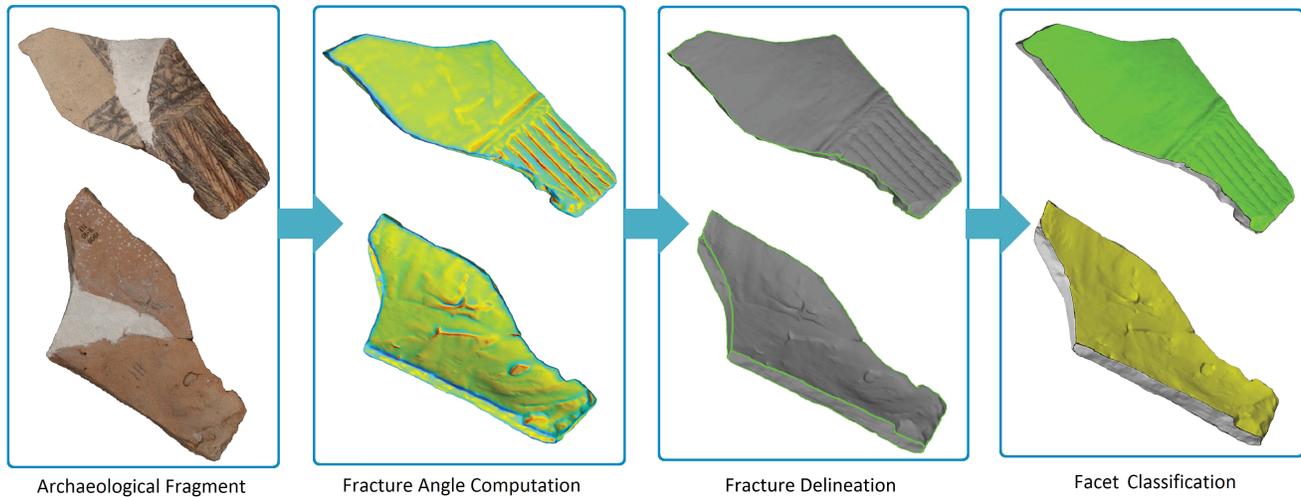


Figure 1. Faceting Pipeline

2. Previous work

A wide range of applications have recently applied 3D segmentation as a preprocessing step to divide a mesh into meaningful components based on different requirements and goals. Beside the cultural heritage community, segmentation of 3D meshes have been addressed by CAD, medical and geological applications and 3D face analysis [10].

However, these methods are not immediately suitable to handle the quantitative needs of archaeological faceting, which has some rather specific properties and desiderata:

- Relatively sharp transition between disparate smooth and fractal regions.
- Variable abrasion and accidental surface perturbations due to long archiving periods and environmental effect.
- Irregular 3D mesh sampling due to different scanning procedures and post processing stages.
- Minimum non-expert user interaction using archaeologically meaningful and easily tunable parameters.

Nevertheless, some methods have been dedicated for segmenting 3D archaeological fragments, usually as a sub-stage of a larger reassembly or reconstruction framework. Huang et al. [4] have presented a segmentation and classification method of fractured objects as a first stage of their robust reassembly pipeline. Each fragment is initially segmented into faces by using a multi-scale edge extraction. Then, a graph cut algorithm is applied to partition the initial set of faces into original and fracture faces. Finally, the over-segmentation is handled by merging adjacent fracture faces that are more likely to belong to one larger fracture

face. Local variation of normals is used to discriminate original and fracture faces. However, the method was not tested on real abraded fragments with a variety of archaeological challenging conditions and it was only developed for fragments represented as point cloud surfaces.

Within the PRESIOUS project ², extracting archaeological facets has also been addressed as a preprocessing step within their scope of predictive digitization, restoration and degradation assessment of cultural heritage objects. Their facet extraction and classification work is based on [6], where a simple region growing algorithm is adopted by exploiting the local mesh polygons' normal variation to extract an initial set of facets. A region merging stage is applied to refine the initially extracted facets. Finally, an elevation map for each facet is constructed to pick the facets nominated for potential matching. This work has been also extended by [1, 2] and further applied as pre-processing stage within the latest work presented by Papaioannou et al. [7],

When we applied those methods to our data collection, we found that reasonable results might only be obtained by fragment-specific fine-tuning of non-intuitive parameters and not for all the fragments we have. The segmentation quality deteriorates or even fails in case of fragments with curved skin surfaces or with highly detailed decoration patterns. Moreover, geometrically similar fragments may produce considerably different segmentation results. Neither of these are therefore the robust and efficient faceting technique running with minimum set of easily-perceivable and tunable parameters that GRAVITATE requires.

²<http://www.presious.eu/>

3. Eigen Analysis for Plane Fitting

Before we start, some general background needs to be covered. In order to define the unit normal vector n of the best fitting plane to a set of points N , eigenanalysis of the points' covariance matrix is commonly performed [3]. The symmetric positive semi-definite covariance matrix C (3×3) defined over a set of points N is defined by:

$$C = \begin{bmatrix} p_1 - c \\ \dots \\ p_k - c \end{bmatrix}^T \begin{bmatrix} p_1 - c \\ \dots \\ p_k - c \end{bmatrix}, \quad (1)$$

where c is the average position centroid of the set of points N and $p_1 \dots p_k \in N$.

If $\lambda_1 \leq \lambda_2 \leq \lambda_3$ represent the eigenvalues of C associated with unit eigenvectors e_1, e_2 and e_3 respectively, then the smallest eigenvalue λ_1 is the sum of squared distances to the best fitting plane of the set of points contained in N . Accordingly, the unit normal vector of the best fitting plane to N can be defined by e_1 , which is the unit eigenvector associated with λ_1 .

From the covariance matrix, we can compute the *Surface Variation*, which was first introduced by Pauly et al.[8] as:

$$\sigma = \frac{\lambda_1}{\lambda_1 + \lambda_2 + \lambda_3}. \quad (2)$$

Thus σ defines the percentage of variance in the direction of the smallest extent of the data (minimum variance direction) compared to the overall variance. It ranges from 0 (for all points lying in the same plane) to 1/3 (for completely isotropically distributed points).

The above standard analysis acts on point sets treating all points equally. However, our data points are organized in a mesh, giving additional information on neighborhood connectivity and irregular density. Accordingly, we include the surface area a_i (A_{Mixed}) [5] occupied by each point p_i in N for weighting the eigenanalysis (plane fitting) of C . The A_{Mixed} is an extension of the Voronoi finite-volume area to apply to both obtuse and non-obtuse triangles in 3D surface meshes. Therefore, analogous to Eq. (1), we can now write for the 3D arbitrary mesh the covariance matrix C as follows:

$$C = \begin{bmatrix} p_1 - c \\ \dots \\ p_k - c \end{bmatrix}^T \begin{bmatrix} a_1 & & \\ & \ddots & \\ & & a_k \end{bmatrix} \begin{bmatrix} p_1 - c \\ \dots \\ p_k - c \end{bmatrix}, \quad (3)$$

where a_i is the local surface area occupied by point p_i and c is redefined as:

$$c = \frac{\sum_{i=1}^N a_i * p_i}{\sum_{i=1}^N a_i}, \quad (4)$$

since for consistency we should use the area weighted centroid instead of the average position of the points.

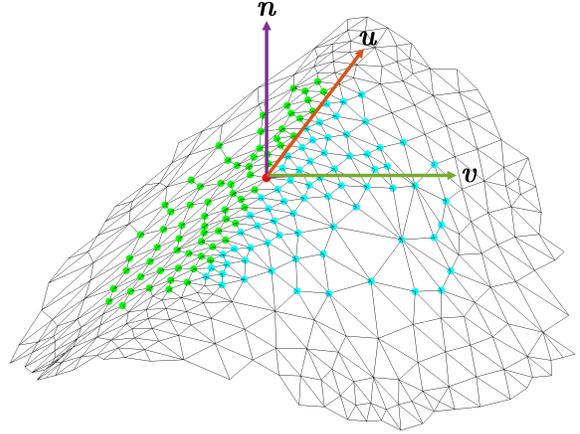


Figure 2. Local neighborhood of a convex point with specific radius ρ . The vectors (u, v, n) form an orthonormal basis: n is the unit normal vector, u is the vector along the local minimum bending direction and v is the vector along the local maximum bending direction.

4. Bi-plane Fracture Model

The expected high level of abrasion leads us to design a specific model of the fracture boundary in the mesh, rather than using general measures of curvature to detect it. We consider a fracture line locally as the meeting of two planes: usually one along the skin, the second along the fracture. Finding such a 'bi-plane fracture model' is subdivided into two consecutive phases: fracture angle computation and fracture line localization. The first phase accurately measures the fracture angle through area-weighted eigenanalysis over the local mesh neighborhood. The second phase localizes the fracture lines across the surface by local minimum detection of the fracture angle computed by the first phase. The fracture angle and the neighborhood over which it is computed are archaeologically meaningful parameters that can be easily interpreted and controlled by a non-expert user of the faceting pipeline.

4.1. Fracture Angle Computation

Let p with unit normal vector n be the point of interest where we would like to accurately measure the fracture angle function α . The local bending of the mesh surface at point p is measured by the two orthonormal directional vectors u and v [5], where u is directed along the minimum bending direction while v is directed along the maximum bending direction.

The principal curvature plane at p , perpendicular to the maximum bending direction v , is used for clipping the local neighborhood with specific radius ρ into two set of points $N_1(p)$ and $N_2(p)$, occupying approximately equal areas. Let c be the area weighted centroid of the points in $N_1 \cup N_2$, then the relative orientation between $(c - p)$ and n is used

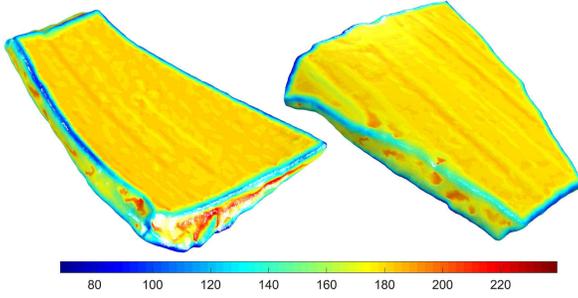


Figure 3. Color-coded Fracture Angle α computed for one of the GRAVITATE fragments, where the blue corresponds to low values and red corresponds to high values. The fragment is 17K vertices and 35K triangles.

to differentiate between convex and concave cases. In case of locally convex surfaces, where the angle α is less than π , the clipping plane is along the direction of the minimum curvature. In case of locally concave surfaces, where the angle is greater than π , the clipping plane is along the direction of the maximum curvature. Figure 2 shows the local neighborhood of a fracture line convex point, where the two directional vectors u and v define the local bending of the surface at this point and the unit normal vector is defined by n . The blue and the green points show how the local neighborhood is divided using the clipping plane perpendicular to the direction of v .

For each group of points $N_1(p)$ and $N_2(p)$, the corresponding covariance matrix is constructed and area-weighted eigenanalysis is conducted to identify the unit normal vectors n_1 and n_2 of the best fitting planes to $N_1(p)$ and $N_2(p)$ respectively. The average local outward pointing unit normal for the two group of points is used as a reference to ensure that the best fitting planes' unit normal vectors are pointing in the right direction (outward) and consistent with the points' local normal vectors.

Given the two unit normal vectors n_1 and n_2 of the two best fitting planes to N_1 and N_2 respectively, the fracture angle α incorporated between these two planes is measured using the following two equations:

In case of convex fracture angle ($\alpha \in [0, \pi]$):

$$\alpha = \pi - \cos^{-1}(n_1 \cdot n_2) \quad (5)$$

In case of concave fracture angle $\alpha \in [\pi, 2\pi]$:

$$\alpha = \pi + \cos^{-1}(n_1 \cdot n_2) \quad (6)$$

Moreover, the rugosity level γ of a point p can now be sensibly expressed in terms of how well the bi-plane model is fitted to the two set of points N_1 and N_2 , as follows:

$$\gamma(p) = \frac{\sigma_1 + \sigma_2}{2} \quad (7)$$

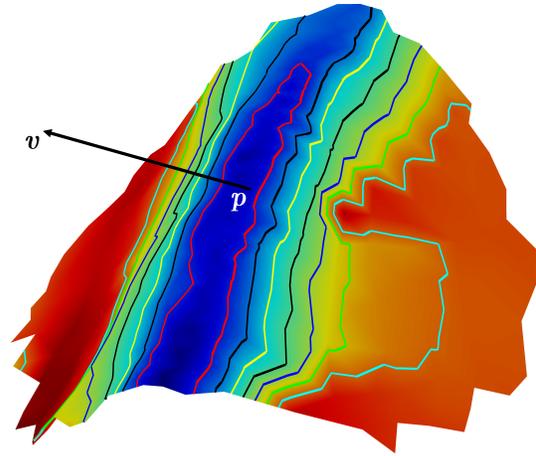


Figure 4. Fracture angle function for a point p along the fracture line. The iso-lines with different colors show how the angle function exhibits a local minimum along the fracture line in the direction of maximum local surface bending v . Color code is the same as Figure 3.

where σ_1 and σ_2 are the surface variation of plane P_1 and P_2 respectively (see Eq. (2)).

Figure 3 shows the angle function computed for an archaeological 3D mesh of a fragment that is about 8 mm thick. The angle function ranges from 80° to 220° and 2 mm is the size of the neighborhood radius ρ over which the angle is computed. It is obvious from the figure that, at convex fracture lines, the angle function exhibits a local minimum, while at concave regions the function exhibits a local maximum.

4.2. Towards Fracture Line Localization

To localize the fracture line near a point p , we need to find the local minimum of the angle function α in the direction v of maximum curvature, since this is approximately perpendicular to the nearby fracture line (see Figure 4).

Mathematically, we need the zero crossing of the directional derivative $\nabla_{\hat{v}}\alpha$ of the angle function, since this represents the rate at which α will change while the input moves with velocity vector \vec{v} . This directional derivative is expressed in 3D coordinates as:

$$\nabla_{\hat{v}}\alpha(x, y, z) = \frac{\partial\alpha}{\partial x}v_x + \frac{\partial\alpha}{\partial y}v_y + \frac{\partial\alpha}{\partial z}v_z, \quad (8)$$

Of course the angle function is strictly only defined on the mesh surface, since it is an estimate of the value of the local angle between two best fitting planes passing through a point. However, we see in Figure 3 (with a sensible choice of neighborhood radius ρ) this function is smooth, and it appears mostly dependent on the in-plane variation of the points at which it is evaluated as they vary across the local surface. Moreover, in fracture regions, we would like to

consider the perpendicular variation as incidental (due to scanning noise, or the fractal nature of the fracture) and not affect our determination of the local minimum much.

For these reasons, and for ease of computation, we have decided to base the required partial derivatives of the angle function not directly on $\alpha(x, y, z)$, but on a *second order 3D fit* to the angle function, in the local neighborhood of known values.

Applying this idea, performing a least-square fit of a function $\hat{\alpha}$ of the form:

$$\hat{\alpha}(x, y, z) = b_0x^2 + b_1y^2 + b_2z^2 + b_3xy + b_4xz + b_5yz + b_6x + b_7y + b_8z + b_9 \quad (9)$$

yields the local 10 coefficients parameters $b_0...b_9$ of this quadratic polynomial function approximation to α .

Those in turn determine the required local derivatives in Eq. (8) as:

$$\begin{aligned} \frac{\partial \alpha}{\partial x} &\approx \frac{\partial \hat{\alpha}}{\partial x} = 2b_0x + b_3y + b_4z + b_6, \\ \frac{\partial \alpha}{\partial y} &\approx \frac{\partial \hat{\alpha}}{\partial y} = 2b_1y + b_3x + b_5z + b_7, \\ \frac{\partial \alpha}{\partial z} &\approx \frac{\partial \hat{\alpha}}{\partial z} = 2b_2z + b_4x + b_5y + b_8 \end{aligned} \quad (10)$$

and the location of the minimum of the directional derivative is then easily determined.

5. Fracture Delineation

5.1. Fracture Line Detection

For our terracotta fragments, the fracture lines are most likely to pass through regions where the angle function is less than approximately 145° . At these regions, guiding markers are picked by looking for the local minimum of the directional derivative function over circular local patches.

Figure 5 shows the initial markers selection process: (a) shows the discontinuous fracture line regions bordered with white lines where the angle is less than approximately 145° . The selected markers are colored in red in (b) where the local patches for selecting the markers are equal to the neighborhood radius ρ .

We want to connect each adjacent pair of markers by a path of mesh points that is short, and runs along minima of the directional derivative. To achieve this, we define a cost function for each mesh edge as follows:

$$W(i, j) = \frac{|\nabla_{\hat{\alpha}} \alpha_i| + |\nabla_{\hat{\alpha}} \alpha_j|}{2} \|p_i - p_j\|, \quad (11)$$

such that $W(i, j)$ is the edge weight between point i and point j , $\nabla_{\hat{\alpha}} \alpha_i$ and $\nabla_{\hat{\alpha}} \alpha_j$ are the angular directional derivatives computed at points i and j respectively and $\|p_i - p_j\|$ is the length of the edge between the two points.

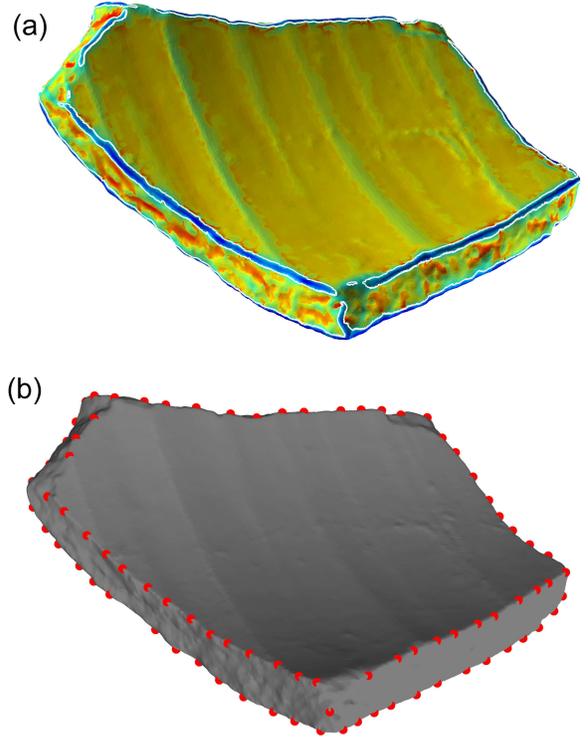


Figure 5. (a) Initial fracture line localization by thresholding of the angle function at 145° . (b) Selected guiding markers at local minima of the directional derivative $\nabla_{\hat{\alpha}} \alpha$.

Accordingly, the fracture stretches are extracted by finding the cheapest path between consecutive guiding markers. The green lines in Figure 6 show an example of the initially extracted fracture stretches for the same fragment in Figure 5. Such weighting scheme has shown robust results even when decreasing the number of initially selected markers.

5.2. Closing Gaps

Due to the observed variable abrasion, even within one fragment, and consequent fracture line discontinuity, not all the points along a fracture line are expected to be found by a globally fixed threshold over the angle function, and thus they may not perfectly enclose a contiguous facet. Using a wider threshold, greater than 145° , would include too many spurious detections away from the potential fracture lines.

Therefore, a post-processing step for closing these gaps is needed. First, the open ends of the initially extracted fracture stretches are detected by inspecting the first-ring neighborhood of each fracture point. A fracture line point with only one member of its one-ring neighborhood connected to the other fracture line points is considered an open end point. Then, we close a gap by the cheapest path from an

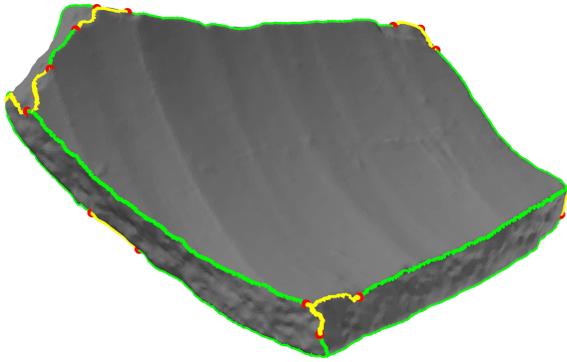


Figure 6. Closing Fracture Gaps

end point p to any other fracture line point in the general direction of the outward pointing vector u of p .

The yellow lines in Figure 6 show gap closure between the open ends indicated in red.

6. Facet Classification

Using the fracture lines, the facets of the artifact can now be extracted as enclosed regions. To enable effective reassembly within GRAVITATE, we subsequently need to distinguish fracture facets from skin facets, which is the last phase in the faceting pipeline in Figure 1. For many fragments (such a broken pottery), one can hope to identify the fractures purely geometrically, based on their characteristic fractal roughness (even when abraded). We may use the rugosity measure of Eq. (7) to compute a 'rugosity density' for each facet:

$$\bar{\gamma}_F = \frac{\sum_{i=1}^M a_i * \gamma_i}{\sum_{i=1}^M a_i} \quad (12)$$

where a_i is the local area occupied by point p_i and $\sum_{i=1}^M a_i$ is the total area of a facet F with M points. A facet F is classified as fracture if the $\bar{\gamma}_F$ is greater than ϵ , while F is classified as skin if the $\bar{\gamma}_F$ is less than ϵ . We have found that the value of ϵ can be automatically determined for fragments with clear distinction in rugosity between skin and fracture facets. Figure 7 shows a color coded example of the rugosity density $\bar{\gamma}_F$ computed for different extracted facets. It is clear that the skin facets (colored in blue) exhibit a rugosity density (less than 0.001) considerably lower than the fracture facets which sensibly allow their correct classification. Alternatively, we might apply the surface roughness measure from [4], but within our context it saves cost to use our precomputed $\gamma(p)$ from Eq. (7).

However, marking facets with high rugosity as fractures is not a universal procedure: for some hollow terracotta statues in our dataset, the rough unfinished original inside is easily confused with a heavily abraded fracture. We may be able to employ a more subtle measure, characterizing the

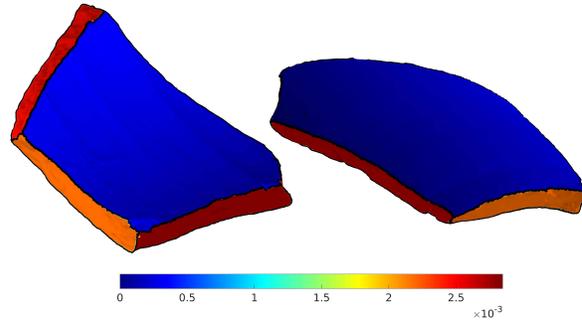


Figure 7. Color-coded rugosity density $\bar{\gamma}_F$ computed for extracted facets of a fragment with 37K vertices and 74K triangles. The blue corresponds to low values and red corresponds to high values.

geometric texture of the roughness, but ultimately the user may be called upon to perform some manual adjustment before approving the faceting and its classification. Within GRAVITATE this is not a problem, since the faceting is part of the once-only inclusion of the fragment data into the system and the vetting of its correctness; limited interaction is a small price to pay for archaeological accuracy of the faceting result on which further analysis depends.

At a later stage, we also need to separate the non-fracture facets into 'external' and 'internal', since they are treated differently in the further GRAVITATE pipeline of processing and detecting artifact similarity. We expect to do this most effectively by combining the geometric, colorimetric and semantic information provided by the remainder of the GRAVITATE system [9], augmented by minimal user interaction.

Figure 8 shows the extracted facets and their classification for the same fragment in Figure 6. The fractured facets are colored in gray, while the skin facets are colored in green and yellow to visually distinguish between external and internal facets.

7. Results and Discussion

The presented faceting algorithm is being applied to two different GRAVITATE datasets from different archives. The first dataset consists of 47 fragments of two broken jars, where the number of vertices per fragment ranges between 5K to 135K and the number of triangles ranges between 11K to 271K (Examples: Figure 9 columns (a-c)). The second dataset consists of 300 fragments rich with both semantic and geometrical properties. A simplified version of the original collection is used for testing, with fixed size of 50K vertices and 100K triangles per fragment (Examples: Figure 9 columns (d) and (e)).

The faceting of various fragments through our pipeline is shown in Figure 9. For each fragment (a-e), the three

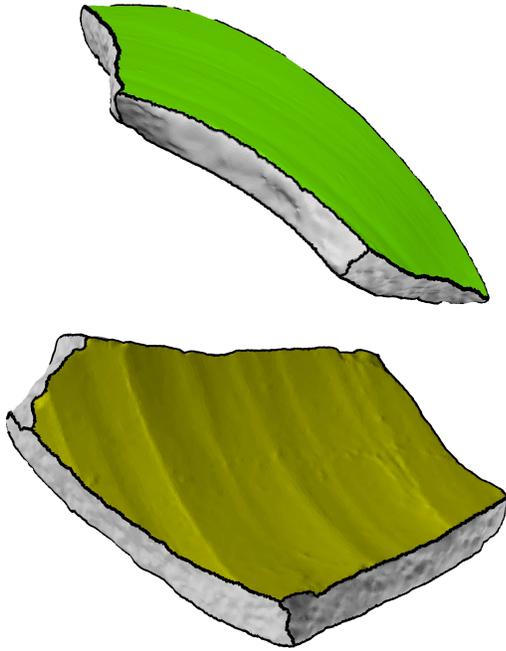


Figure 8. Facet Extraction and Classification

stages of fracture angle computation, fracture delineation and classification are illustrated from top to bottom. The results show the insensitivity of our method to abrasion, irregular mesh sampling and accidental acquisition noise. The skin region is naturally divided into two separate facets (internal/ external) for most of the tested fragments, except for a few rim fragments where the smooth curvature transition between internal and external regions avoid their further decomposition, (see column 9(a)). Some fracture facets are decomposed into smaller yet more distinctive ones due to the existence of sharp fracture lines in the middle of their contained fracture region, (see column 9(b)). This feature consequently demonstrates the potential of reassembling multiple other fragments to the indicated fracture region.

It is interesting to observe that extra semantic features could be further inferred from our results. For example in columns 9(d) and (e), the eyes and the ears were detected as distinct facets. By varying the neighborhood radius ρ when computing the fracture angle, the number of facets being extracted may be affected. This is a useful property which provides the user with enough flexibility on extracting coarse or fine facets. However, in no case should ρ exceed half the thickness of the fragment.

The presented faceting pipeline can be extended for different 3D applications other than its intended use of 3D archaeological segmentation. For example, the angular function (presented in section 4.1) can be used separately as 3D

descriptor for 3D meshes that intrinsically encodes the local mesh structure. Such descriptors are useful for similarity search and partial matching. Moreover, the derivative computation method (presented in section 4.2) could be independently adopted for computing the derivative of any sufficiently smooth scalar function computed over 3D meshes. Accordingly, feature lines known as *ridges and valleys* can be extracted following an approximately similar approach as the one introduced in section 5.1 for fracture line detection. The rugosity density (introduced in section 6) for facet classification can be naturally applied to encode the 3D surface roughness, which has an extensive use in geometric modeling, noise detection and 3D mesh watermarking applications.

We tested our method on Intel Xeon 3.4GHZ $\times 4$ computer and 16GB RAM. Calculating the fracture angle function and its derivative for a 3D model with typical 50K vertices takes less than 2 minutes, where it is mainly dominated by the 3D model size and the scale ρ used for defining the neighborhood over which the angle is calculated. For fracture delineation, the initial stretches extraction takes insignificant additional time of less than 5 seconds. Localizing the fracture line regions by defining the guiding markers significantly improves the time for the overall fracture line extraction. The delineation time is more affected by the number of detected gaps that need to be closed: closing 10 gaps for a 3D model with 50K vertices takes at most 1 minute. These processing times deemed acceptable for GRAVITATE.

8. Conclusion and Future Work

We have presented a complete pipeline for faceting archaeological fragments digitally represented in the form of arbitrary 3D meshes. Our method is robust against abrasion and noise incorporated during the acquisition process. The facets that result are controlled by two main parameters: the radius ρ determining the size of the neighborhood over which the angle function is evaluated and the threshold to constrain the potential fracture line regions. Both are clearly related to the pure geometrical properties of the fragment and can usually be set automatically or easily fine-tuned by an archaeological user.

Future work will focus on how to automate the local selection of a neighborhood radius ρ for the fracture angle computation and the classification of the skin facets into 'internal' and 'external' using the angle function. We will also work on integrating the faceting pipeline with its user-tunable parameters into the GRAVITATE user interface. This includes providing additional flexible functionalities for the user to steer towards more customized faceting results, such as merging over-faceted regions or sub-faceting larger ones.

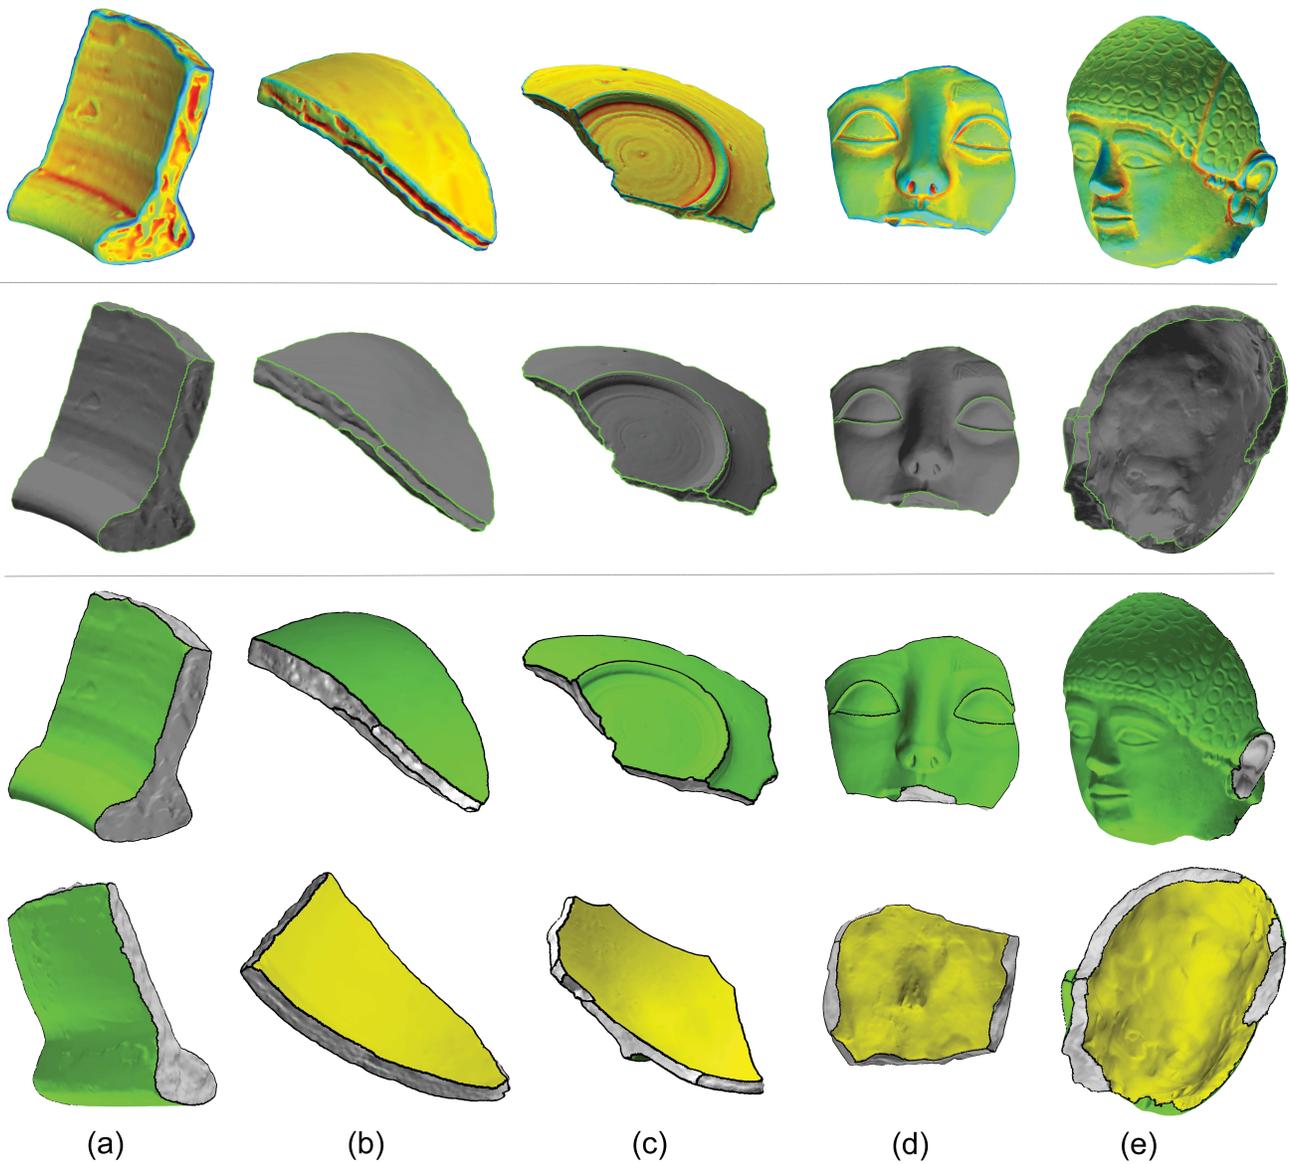


Figure 9. Faceting Results: The top row depicts the fracture angle computed using the bi-plane fracture model (section 4), the second row shows the fracture delineation (section 5) and the two bottom rows show front and back views of the facet classification (section 6). The size of the fragments in columns (a-e) ranges between 27K to 50K vertices and between 65K to 100K triangles. The neighborhood radius ρ used for the angle computation is 2 mm, except for column (a) where ρ is 3 mm.

9. Acknowledgements

This research was funded by GRAVITATE project under EU2020-REFLECTIVE-7-2014 Research and Innovation Action, grant no. 665155. Datasets are from GRAVITATE data collection. The archaeological artifacts are digitized and kept by the British museum, Cyprus Institute and Archaeology laboratory at the Hebrew University of Jerusalem.

References

- [1] A. Andreadis, R. Gregor, I. Sipiran, P. Mavridis, G. Papaioannou, and T. Schreck. Fractured 3d object restoration and completion. In *ACM SIGGRAPH 2015 Posters*, SIGGRAPH '15, pages 74:1–74:1, New York, NY, USA, 2015. ACM.
- [2] A. Andreadis, P. Mavridis, and G. Papaioannou. Facet Extraction and Classification for the Reassembly of Fractured

- 3D Objects. In M. Paulin and C. Dachsbacher, editors, *Eurographics 2014 - Posters*. The Eurographics Association, 2014.
- [3] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, and W. Stuetzle. Surface reconstruction from unorganized points. *SIGGRAPH Comput. Graph.*, 26(2):71–78, July 1992.
- [4] Q.-X. Huang, S. Flöry, N. Gelfand, M. Hofer, and H. Pottmann. Reassembling fractured objects by geometric matching. In *ACM SIGGRAPH 2006 Papers*, SIGGRAPH '06, pages 569–578, New York, NY, USA, 2006. ACM.
- [5] M. Meyer, M. Desbrun, P. Schröder, and A. H. Barr. Discrete differential-geometry operators for triangulated 2-manifolds, 2002.
- [6] G. Papaioannou, E.-A. Karabassi, and T. Theoharis. Virtual archaeologist: Assembling the past. *IEEE Comput. Graph. Appl.*, 21(2):53–59, Mar. 2001.
- [7] G. Papaioannou, T. Schreck, A. Andreadis, P. Mavridis, R. Gregor, I. Sipiran, and K. Vardis. From reassembly to object completion: A complete systems pipeline. *J. Comput. Cult. Herit.*, 10(2):8:1–8:22, Mar. 2017.
- [8] M. Pauly, M. Gross, and L. P. Kobbelt. Efficient simplification of point-sampled surfaces. In *Proceedings of the Conference on Visualization '02*, VIS '02, pages 163–170, Washington, DC, USA, 2002. IEEE Computer Society.
- [9] S. C. Phillips, P. W. Walland, S. Modafferi, L. Dorst, M. Spagnuolo, C. E. Catalano, D. Oldman, A. Tal, I. Shimshoni, and S. Hermon. GRAVITATE: Geometric and Semantic Matching for Cultural Heritage Artefacts. In C. E. Catalano and L. D. Luca, editors, *Eurographics Workshop on Graphics and Cultural Heritage*. The Eurographics Association, 2016.
- [10] A. Shamir. A survey on Mesh Segmentation Techniques. *Computer Graphics Forum*, 2008.