



Pairwise Alignment of Archaeological Fragments Through Morphological Characterization of Fracture Surfaces

Hanan ElNaghy¹ · Leo Dorst¹

Received: 14 March 2021 / Accepted: 30 May 2022 / Published online: 9 July 2022
© The Author(s) 2022

Abstract

We design a computational method to align pairs of counter-fitting fracture surfaces of digitized archaeological artefacts. The challenge is to achieve an accurate fit, even though the data is inherently lacking material through abrasion, missing geometry of the counterparts, and may have been acquired by different scanning practices. We propose to use the non-linear complementarity-preserving properties of Mathematical Morphology to guide the pairwise fitting in a manner inherently insensitive to these aspects. In our approach, the fracture surface is tightly bounded by a concise set of characteristic multi-local morphological features. Such features and their descriptors are computed by analysing the discrete distance transform and its causal scale-space information. This compact morphological representation provides the information required for accurately aligning the fracture surfaces through applying a RANSAC-based algorithm incorporating weighted Procrustes to the morphological features, followed by ICP on morphologically selected ‘flank’ regions. We propose new criteria for evaluating the resulting pairwise alignment quality, taking into consideration both penetration and gap regions. Careful quantitative evaluation on real terracotta fragments confirms the accuracy of our method under the expected archaeological noise. We show that our morphological method outperforms a recent linear pairwise alignment method and briefly discuss our limitations and the effects of variations in digitization and abrasion on our proposed alignment technique.

Keywords Archaeological reconstruction · Fragment alignment · Scale space · Fracture representation · Boundary morphology · Morphological features · Lipschitz · ICP

1 Introduction

1.1 Alignment in Virtual Archaeology

Virtual archaeology, the study of artefacts by means of their digital representations, should enable an archaeologist to perform common workflow tasks remotely. This may involve the virtual fitting or reassembly of broken objects, dispersed over different collections, or simply too heavy to be manipulated physically. Even a single accurate alignment may considerably affect the cultural interpretation of artefacts (Sommella Mura, 2011), so accuracy is required. Typically, archaeological artefacts are not only broken, but their fractures may be incomplete (e.g., if the artefact was broken again

and the other piece is missing), abraded, or even deformed. Moreover, their digital mesh representation may contain digitization noise and point cloud registration errors from the 3D scanning process. These demands, paired with a sensible time behaviour, make pairwise alignment a daunting computational task (as other studies such as (Toler-Franklin et al., 2010; McBride and Kimia, 2003; Palmas et al., 2013) confirm).

Our method is part of a larger system for reunification of broken archaeological artefacts in the GRAVITATE project (Phillips et al., 2016). Our role was to focus on pairwise fitting of fracture surfaces belonging to fragments, nominated by other selection modules in the system, by aligning the fragments to an accuracy compatible with the digitization noise expected within the scanning process (or report failure to the selection module). For specific sets of fragments (such as in FORMA URBIS ROMAE (Koller et al., 2006), or fresco restoration (Vendrell-Vidal and Sánchez-Belenguer, 2014)), information on the ‘skin’ patterns on the original outside of the artefact would very likely be highly beneficial to improve

Communicated by Takeshi Oishi.

✉ Hanan ElNaghy
hanan.elnaghy@uva.nl

¹ Computer Vision Lab, Informatics Institute, University of Amsterdam, Amsterdam, The Netherlands

the computation of a good alignment, in speed or accuracy. A choice was made early on in GRAVITATE to separate the purely geometrical from semantic or colorimetric pattern information as was also done in PRESIOUS (Andreadis et al., 2015).

Our geometrical method addresses *pairwise alignment*, we specifically do not perform global reassembly or optimized global alignment of multiple pieces, although we will see in Sect. 8.3 that our results may form a good basis for such an extension.

1.2 Pairwise Alignment in Literature

The main drive of our work is to address the computational problem of pairwise alignment of archaeological artefacts through investigating the geometric characterization of their involved fracture surfaces. The actual fractures are typically close to flat and fractal, and the variations that should lead to alignment subtend a small range of orientations. This makes the archaeological pairwise alignment problem distinctly different from global alignment methods such as (Mellado et al., 2014).

The pairwise alignment of artefacts has been studied for quite a while. The early work started with adapting 2D techniques to 3D alignment (Kalvin et al., 1986; Kong and Kimia, 2001; Papaioannou et al., 2002; Leitao and Stolfi, 2002; Sagirolglu and Ercil, 2005; Koller et al., 2006; Huang et al., 2013). This may involve the 3D fracture surfaces, but is often partly based on the continuity of the carving or painting depicted on the skin facets. Truly 3D fitting may use freshly broken and complete artefacts (Huang et al., 2006; Brown et al., 2008; Winkelbach and Wahl, 2007; Son et al., 2017), and combine the geometrical and combinatorial aspects, with the aim to perform a total reassembly.

The high resolution at which archaeological artefacts are scanned makes their brute force alignment on the 3D raw data a computationally expensive process (as other studies such as (Toler-Franklin et al., 2010; McBride and Kimia, 2003; Palmas et al., 2013) confirm). Therefore, most recent methods employ feature-based alignment techniques by matching significant corresponding features of counter-fitting fragments. This reduces the pairwise alignment complexity while maintaining the intended performance. These features are mostly detected and extracted using differential geometric techniques which are linearly sensitive to missing information. Such features are not expected to be stable under abrasion and many of them are not descriptive enough to allow the effective hierarchical scale-based fitting.

In Table 1, we have collated some representative methods that contain purely geometrical pairwise alignment, either as their goal or as part of their total procedure. We focus on how global reassembly methods address the pairwise alignment as part of their whole pipeline without going into the details of

Table 1 Representative geometry-based pairwise alignment methods in literature, focusing on the types of features used. Where the pairwise alignment was part of a global reassembly, we denote this by an asterisk on the reference. PC denotes Point Cloud, M denotes Mesh. ‘Fresh’ denotes non-archaeological newly broken fragments

| Year | References | Data | Dataset | Features | Alignment method |
|------|---|-------|------------------------|------------------------------|------------------------|
| 2002 | Papaioannou et al. (2002)* | M | Real fragments | Featureless | Simulated annealing |
| 2006 | Huang et al. (2006)* | PC | Fresh 3D | Multi-scale curvature | Forward search |
| 2007 | Winkelbach and Wahl (2007) | PC | Fresh 3D | Cluster tree | Depth-first search |
| 2014 | Altantsetseg et al. (2014) | PC | Flint tools | Fourier on curves | Spectral matching |
| 2014 | Vendrell-Vidal and Sánchez-Belenguier (2014) | M | Flat fresco | Contour only | Hierarchical search |
| 2015 | Mavridis et al. (2015)* | PC | Real fragments | Featureless | Rigid registration |
| 2015 | Zhang et al. (2015)* | PC | Thin fragments, skulls | Histograms of orientations | RANSAC, template based |
| 2017 | Andreadis et al. (2015), Papaioannou et al. (2017)* | PC, M | Real fragments | Skin feature curves | Rigid registration |
| 2017 | Son et al. (2017)* | M | Fresh 3D | Signature curves, spin image | RANSAC |
| 2018 | Zhang et al. (2018)* | PC | Thin fragments | Guiding contour points | GAN, ICP |
| 2019 | Alzaid and Dogramadzi (2019)* | M | Fresh 3D, simulated | Fourier on contours | ICP |
| 2020 | Li et al. (2020) | PC, M | Fresh 3D, terracotta | Convexity clusters | ICP |
| 2020 | Liao et al. (2020)* | PC, M | Fresh 3D, bone | Cluster tree | ICP |
| 2021 | Our method | M | fresh 3D, terracotta | MM scale space | RANSAC, ICP |

their overall strategy of multi-fragment reassembly. This covers most related work dealing merely with pairwise matching and alignment of fractured archaeological surfaces. We have indicated what objects were used, and how the problem was solved in terms of the alignment method on the features that were computed as its input. Although some authors apply their method on other datasets besides their own (notably that of (Huang et al., 2006)), the field has no generally agreed benchmark datasets and no agreed-upon evaluation criteria.

As observed from Table 1, the linear feature-based techniques can be mainly categorized into curvature based, Fourier based, cluster tree based or contour based techniques.

The *curvature based techniques* typically analyse the local curvature information of the fracture surfaces (concave and convex regions) to perform pairwise matching and alignment. One of these methods is the early well-known method by (Huang et al., 2006) for automatic reassembly of 3D digital models of broken fragments. They extract clusters of multi-scale integral invariant features for each fracture surface and use a forward search algorithm to identify patches of similar features for pairwise matching. (Li et al., 2020) is one of the recent extended methods of Huang et al.'s approach. They present a simple well-evaluated method for pairwise matching and alignment of 3D fragments. They conduct an initial boundary-curve based comparison for excluding different fracture surfaces. Then they perform fine alignment based on matching concave and convex patches of the complementary fracture surfaces. They extract such concave and convex regions using Gaussian and mean curvature. (Son et al., 2017) also propose reassembly of fractured objects using a surface signature descriptor based on whether a point on the fractured surface is convex or concave. The similarity between two fracture surfaces is calculated based on the spin images of feature curve points. The feature curve is the boundary between convex and concave regions and describes the geometric characteristics of a fracture surface. For matching fracture surfaces, distance and normal deviation between two feature curves are used.

Fourier based techniques compute the Fourier coefficients of salient curves extracted from the fracture surfaces. For pairwise alignment, the descriptor curves are usually clustered and matched according to their Fourier representation. One of these methods is (Altantsetseg et al., 2014). They characterize the fracture surfaces by clusters of feature points and extract descriptor curves along the principal directions of these clusters. The Fourier coefficients of each descriptor curve are computed by using Fast Fourier Transform. For the alignment of fracture surfaces represented as point clouds, the total power of the extracted descriptor curves of each cluster are compared. (Alzaid and Dogramadzi, 2019) present another method for automatic reassembly of fractured objects based on Fourier representation. Their method

uses 2D boundary curve features associated with Fourier descriptors for pairwise matching and alignment.

(Winkelbach and Wahl, 2007) propose a *cluster-tree based method* for pairwise matching and alignment of 3D fragments. Their method follows a cluster tree matching strategy to find optimum alignment. The fragments are matched pairwise, in a hierarchical manner, by descending the cluster trees simultaneously in a depth-first fashion. The use of a cluster tree representation for object reconstruction and bone fracture reduction has also been studied by (Liao et al., 2020). They introduce a region-pair invariant feature (based on geodesic distance) combined with a balanced cluster tree to select the best region contact pairs from a valid potential pool for matching.

The *contour based techniques* perform matching and aligning of the fracture surfaces focusing on their outer border (outline), not their inner fracture region. These techniques are more well-suited for flat fresco fragments or non-abraded fracture surfaces. (Vendrell-Vidal and Sánchez-Belenguer, 2014) present a discrete pairwise approach for registering flat fresco fragments. Their method performs a hierarchical search strategy that minimizes a cost function which only focuses on the contour of fragments, instead of the inner local patches. Papaioannou et al. (2017) also propose a semi-automatic pipeline for reassembly of 3D archaeological objects and restoration of missing parts. For pairwise matching and alignment, they combine the featureless rigid geometric registration of (Mavridis et al., 2015) for the fractured surface of the facets with the alignment of the extracted feature curve point sets of the intact object's surface, as proposed by (Andreadis et al., 2015). Their method is an extension of (Mavridis et al., 2015), where feature curves from the non-fracture surfaces of the objects are also introduced in order to address the matching of adjacent fragments with a minimal or even absent matching surface.

Machine Learning methods are only now beginning to be applied; (Zhang et al., 2018) propose a pipeline for the reassembly of thin-shell and thick fragments with mostly flat fracture surfaces represented as point clouds. They convert the problem of fragment matching into a GAN matching. The initial pairwise alignment is performed by comparing guiding points on the boundary contours and the alignment is then refined using ICP.

None of the methods mentioned so far employ the fundamentally *non-linear* nature of bringing objects into contact. There is an asymmetry in the relative alignment of two fracture surfaces: in principle penetration is not allowed (and should not be treated as some minor squared error), missing elements are fine (and not a punishable as a large squared error), abrasion noise is one-sided and will naturally lead to acceptable gaps in optimally aligned fragments.

We thus observe that the complementary nature of a fit between fracture surfaces, and the one-sided 'noise'

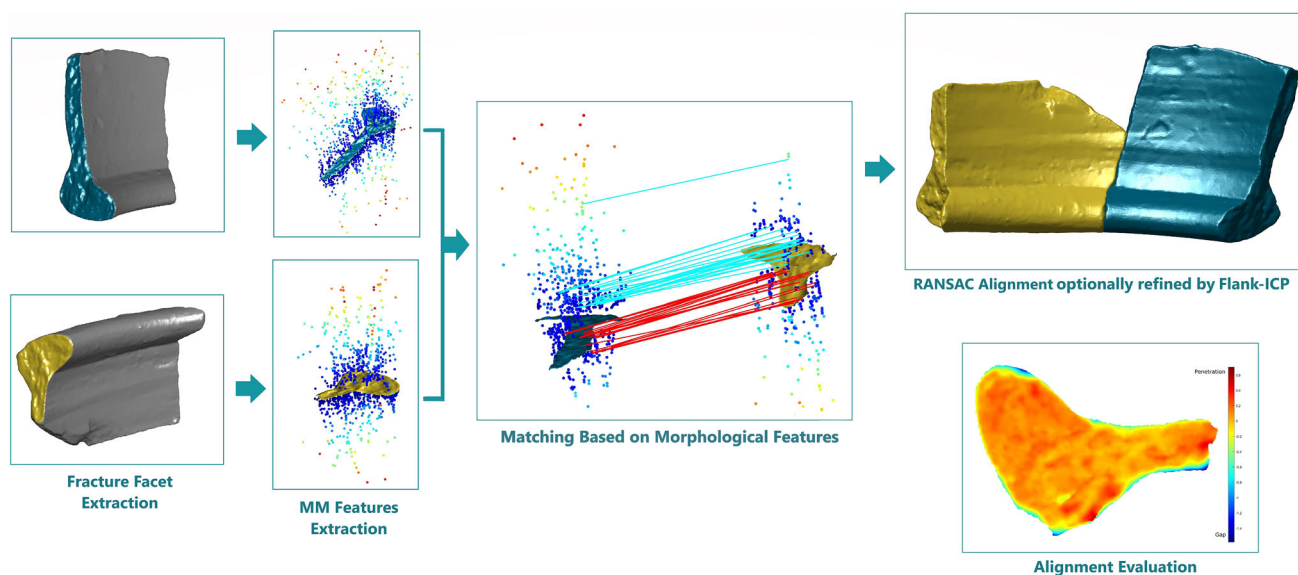


Fig. 1 The pairwise alignment workflow pipeline

(abrasion) that may naturally occur while not affecting our confidence in a match are ill represented by classical linear methods: *Archaeological complementarity is essentially non-linear*. We show in this paper that the framework of Mathematical Morphology (MM) (Serra, 1986) (with its opening and closing scale spaces) enables the proper complementarity-preserving representation of fracture surfaces, while being insensitive to missing information in exactly the correct non-linear way. Mathematical Morphology was originally developed for the study and analysis of geometric and topological structures especially on digital images (Haralick et al., 1987; Dougherty, 2018; Serra and Soille, 2012). It has also been applied for studying the geometry of contact (Dorst and van den Boomgaard, 2000). Our use for alignment of complementary 3D archaeological fracture surfaces is new.

Other people have observed that MM can be useful in quantitative shape characterization tasks due to its insensitivity to missing information (e.g., (Attali et al., 2008)). Specifically, (Chang et al., 2004) used an MM-based method for approximate global point cloud registration of models with distinctive global characteristics, extracting features from a carefully simplified medial axis. Their method would not work on the fractal and rather planar fracture surfaces we encounter; as we will show, a different way of extracting morphological features from a medial axis is then useful to achieve accurate alignment.

1.3 Contributions

The main contributions of our method are as follows:

- We propose a complete pipeline for accurate pairwise alignment of abraded archaeological fracture surfaces using the non-linear complementarity preserving properties of Mathematical Morphology.
- We characterize fracture surfaces through a concise set of robust morphological feature points by directly analysing computed distance transforms and tracing their associated provenance information.
- We perform a perturbation analysis of the morphological features to quantify their reliability.
- The MM features are endowed with descriptors to facilitate the heuristics of establishing correspondences for a RANSAC/Procrustes registration which produces an accurate alignment with little penetration.
- We additionally develop a refined version of ICP that works exclusively on the flank regions of the fracture surface, which were found to be more important than the peaks and valleys for stable alignment.
- We suggest new evaluation criteria for the fitting quality, taking into consideration both penetration and gap regions.

2 The Workflow Pipeline

Our workflow pipeline of Fig. 1 provides an overview of the steps required to perform the pairwise alignment of a suggested pair of complementary archaeological fragments:

1. Extract the meshes A and B of the two fracture facets to be aligned, using our Faceting technique from (ElNaghy and Dorst, 2017).

2. Characterize A and B , each by our novel morphological multi-local features: the centre point of the critical spheres that bound the fracture facet from above (3PC) and below (3PO). Based on the local geometry, establish characteristic descriptors and a measure of the reliability of each MM feature.
3. Establish possible correspondences between the 3PC features of A and 3PO features of B , and *vice versa*, using the morphological descriptors.
4. Employ RANSAC: using a set of random corresponding matched MM features, apply a reliability-weighted Procrustes method to construct a candidate rigid body motion for alignment. Continue sampling till an alignment is supported sufficiently by enough inliers.
5. To refine the alignment suggested by RANSAC, use ICP on a set of morphologically selected mesh vertices (which we call ‘flanks’).
6. Compute and display quality measures on the fit suitable for interpretation by a trained archaeologist.

The paper is organized along the lines of these workflow steps. We apply our method to align real artefacts, with a focus on accuracy of the resulting alignment and the conciseness of the morphological characterization relative to the original fracture surface.

3 Boundary Morphology on Fractures

Our method processes only fracture facets, and does so in a morphological manner. This preliminary section briefly justifies that methodology.

3.1 Fracture Facets from Artefact Meshes

Scans of artefacts are processed when their data is ingested into the GRAVITATE system as described in (Scalas et al., 2020). For each fragment, this ingestion involves standardization of its metadata, but also producing regularized mesh versions of the geometry, in different resolutions suitable for further processing.

In our geometrical approach to fitting and pairwise alignment, we employ only the fracture parts of the artefacts. The terracotta fragments that form the GRAVITATE data are uniform and brittle in consistency. Such artefacts break in a simple way, leading to rather straight fractures. This permits the subdivision of the total fracture regions of the digitized artefact mesh into ‘facets’. The fracture facets are extracted by the Faceting preprocessing procedure we developed in (ElNaghy and Dorst, 2017).

The resulting submeshes are the mesh representation of Monge patches (i.e., local height functions in their own coordinate system), and our faceting should guarantee them to be

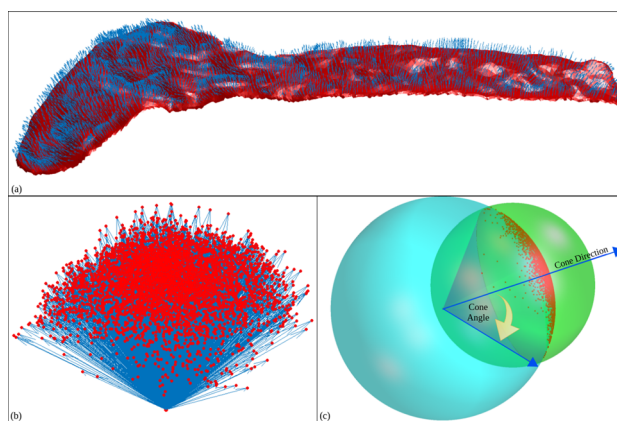


Fig. 2 **a** Fracture surface with blue local normal directions. **b** Mapping of the normals to a unit sphere of directions. **c** The minimum bounding sphere (in green) determines the minimum bounding cone on the unit sphere of directions (in cyan). From (ElNaghy and Dorst, 2020) (Color figure online)

Lipschitz (i.e., slope-limited) surfaces. Under those assumptions, a terracotta fracture facet has a limited and balanced distribution of normals, as illustrated in Fig. 2. We use the Lipschitz principal direction as the normal axis of the fracture facet mesh in all further processing.

3.2 Complementarity Preservation under Simplification

As we noted, Mathematical Morphology (MM) is fundamentally the mathematics of contact (Dorst and van den Boomgaard, 2000). Its operations allow one to establish quantitatively how sets overlap or are complementary to each other. When used on sets with noise, the non-linearity of Mathematical Morphology enables the one-sided treatment we need when performing alignment of broken objects: penetration is undesirable, a small separation explicable by abrasion is allowed; missing parts do not affect the fitting of what did survive. Mathematical morphology is a pertinent framework that preserves these contact properties under its simplifications (ElNaghy and Dorst, 2020).

There are two dual simplification operations in mathematical morphology: *opening* and *closing* (‘dual’ means: in a binary situation, performing an opening on an object is identical to performing a closing on its background, and *vice versa*). Because of isotropy of our given artefacts, we perform those operations with a ball (‘spherical structuring element’), of a variable radius ρ .

In intuitive terms, consider a fracture facet, outside facing upwards. The result of closing a fracture facet at scale ρ is to produce the surface formed by a ball rolling over the fracture (see Fig. 3). Protrusions (‘peaks’) will still be present at the original locations, but valleys will have been filled by spherical caps and hence ‘smoothed away’. Similarly, the




Fig. 3 A fracture facet, with a rolling ball on top forming its *closing*, and a rolling ball below tracing out its *opening*




Fig. 4 A fracture facet, and its closing (above) and opening (below) surfaces at a range of scales. Adapted from (ElNaghy and Dorst, 2020)

opening of a fracture surface is the surface formed by rolling a ball underneath it: this mostly preserves the pits and valleys, but rounds the peaks spherically from below. Together, these opening and closing surfaces bound the actual fracture surface from above and below, with an accuracy parametrized by the spherical scale ρ , see Fig. 4.

In (ElNaghy and Dorst, 2020), we proved that when these dual morphological operations are applied to the common fracture of an object A aligned in exact contact with its counterpart B , the resulting simplified surfaces obtained are still aligned in exact contact: $\text{open}(A)$ matches $\text{close}(B)$ and $\text{close}(A)$ matches $\text{open}(B)$ at any scale ρ_i . Moreover, the coarser scale simplifications subsume the finer scales. More precisely phrased, for different scales R and ρ with $R > \rho$, the complementarity and containment relationships for two masked fracture facets A_M and B_M are:

$$\begin{aligned} \gamma_R(A_M) \subseteq \gamma_\rho(A_M) \subseteq A_M \subseteq \phi_\rho(A_M) \subseteq \phi_R(A_M) \\ \Downarrow c_R \quad \Downarrow c_\rho \quad \Downarrow c \quad \Downarrow c_\rho \quad \Downarrow c_R \\ \phi_R(B_M) \supseteq \phi_\rho(B_M) \supseteq B_M \supseteq \gamma_\rho(B_M) \supseteq \gamma_R(B_M) \end{aligned}$$

Here we denoted a closing by a sphere of scale ρ by ϕ_ρ , and the opening at scale ρ by γ_ρ . The c_ρ denotes the morphological complementarity of the (simplified) fracture surfaces; its

subscript indicates a scale dependence. Our analysis showed that complementarity only holds exactly at a reduced area of the fractures, shrunk by 2ρ .

To compute such complementarity directly on fracture facets, in (ElNaghy and Dorst, 2020) we extended Mathematical Morphology to a ‘Boundary Morphology’, specializing the usual morphological operators to be valid on boundary surfaces rather than volumes. We will employ that framework in this paper.

4 Pairwise Alignment by Morphological Features

4.1 Morphological Features

As we demonstrated, the opening and closing surfaces bound an original fracture surface by spheres, and thus simplify its representation. However, not all parts of these bounding surfaces at all scales (i.e., spherical radii) are equally descriptive. We propose to use a small set of the rolling balls at characteristic locations and scales as morphological features to encode the surface, and then use only those features (rather than the surfaces themselves) for alignment.

To produce the features required for alignment, we recall the observation from Sect. 3.2 that the morphological closing process results in a surface of which large parts are spherical caps (see Fig. 4, especially at the coarser scales, and Fig. 5). Moreover, as the scale increases, the surface tends to be hierarchically composed of fewer, larger and more clearly separated spherical caps. The centres of the balls to which these spherical caps belong can be extracted from the provenance map of the distance transform computation (which specifies which point on the original fracture generated a given point on the MM scale space surface). In fact, these centres are medial axis points of the complement of the extruded fracture mesh volume.

With increasing scale, the ongoing morphological simplification can be characterized in terms of fewer yet more distinctive medial axis feature points, which can therefore serve as 3D morphological features. Moreover, these points can be tagged by their significance for the MM simplification (see Fig. 5a). Medial axes are notoriously sensitive to discretization (though they can be made robust (Attali et al., 2008)) but with our provenance approach, we can retrieve their key features effectively. Fig. 5b is an exploration of this idea. It shows the 30 most distinctive sphere centres which are capable of representing 30% of the simplified (closed) fracture facet at scale 30 mm. The fracture facet is the same as in Fig. 4, which is about 6×2.5 cm size.

At a fixed scale ρ , the ball of that radius will roll mostly on one or two contact points of the surface; but at certain positions, it touches three contact locations, see Fig. 6. Such

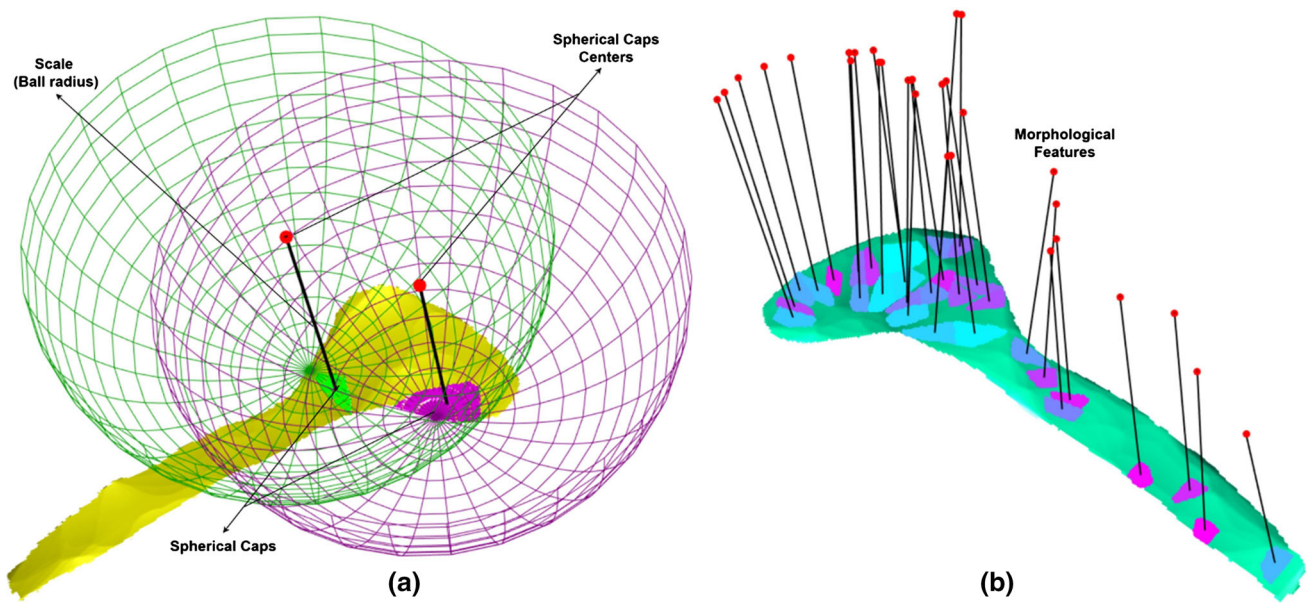


Fig. 5 Morphological characterization of fracture surfaces through 3D sphere centres points. **a** Each medial axis point locally simplifies a part of the surface by a spherical cap at the chosen scale. **b** The 30 most relevant feature points and their spherical caps. From (ElNaghy and Dorst, 2020)




Fig. 6 A 3-point closing (3PC) feature (the bounding sphere centre and radius) and some of its descriptors based on the 3-point contact geometry. We also show the peak components (in gold) of the facet at the scale of the sphere (i.e., the vertices that can be touched by a rolling ball with the shown radius)

3-point contacts are very localized, and therefore supply useful features for accurate alignment. Moreover, the same three points (or points very close to them) will support a range of contact sphere sizes, so they are relatively stable. As we grow the scale, the local contact sphere will become so large that it touches the surface at a fourth point; increase the radius slightly more, and that contact will have taken over from one of the three original points, and support the new contact sphere. The radius at which this happens thus localizes the multi-local morphology not only in location, but also in scale.

Such ‘3-point contact spheres’ are the preferred scale space events to use as our morphological features, since they remove the arbitrariness of the scale parameter from our characterization. The features’ scale is locally determined intrinsically by the multi-local morphology of the surface itself; and the same is likely to happen in the complementary counterpart, since that fracture should have a similar, but dual, morphology.

We thus simplify both the opening and closing surfaces as being bounded by these characteristic spheres, and store their location and radius as our *MM features* of the fracture facet. One feature point is thus a crystallization of much multi-local information on how to tightly bound the surface around that spatial neighbourhood and that scale. One can visualize the fracture surface in our morphological feature representation as having a set of spatial points suspended over, but rigidly tied to, the surface (as in Fig. 7). The 3-point contact features of closing we denote as 3PC, of opening as 3PO. As we will demonstrate in Sect. 7, aligning a small subset of such features is enough to align the original fracture facets accurately.

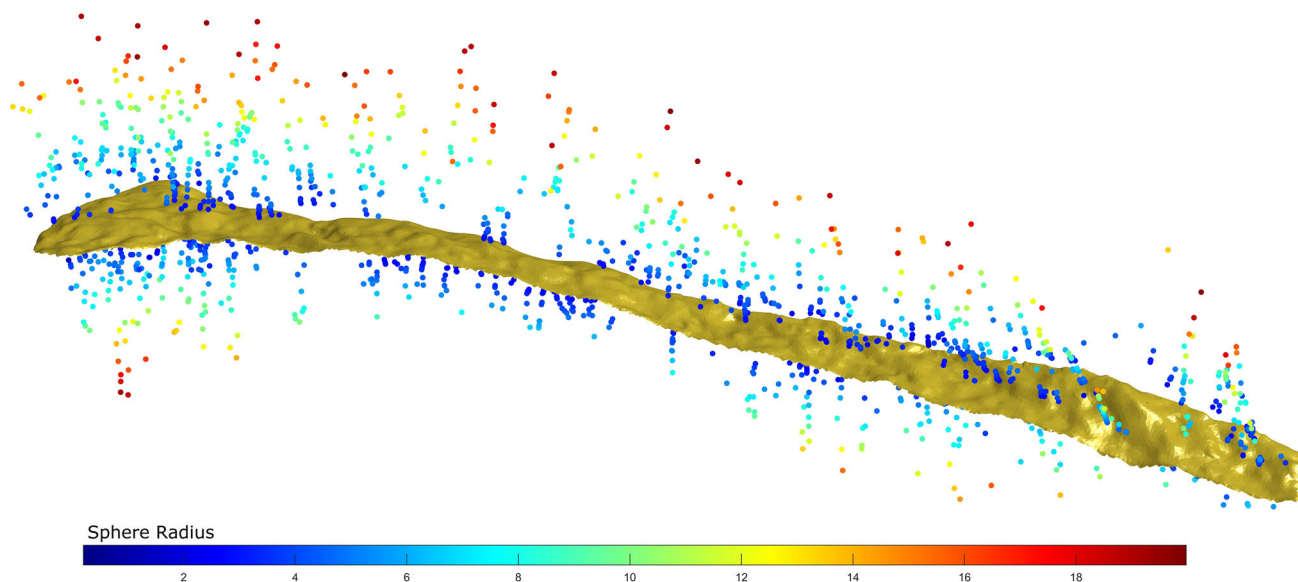


Fig. 7 Closing (3PC, above) and Opening (3PO, below) morphological features bounding a fracture surface from both sides. Each MM feature point is colour coded by the scale (sphere radius) it represents

4.2 Distance Transform Based Computation of MM Features

Though our MM features are inspired by the bounding through opening and closing, they can be computed from a study of the simpler operations of *dilation* (spherical thickening) and *erosion* (spherical thinning) surfaces at different scales. The closing features are based on the (dilation) *distance function*, which denotes for each point above the surface its distance to the closest point of the surface. To determine whether a point is a 3-point contact sphere, we analyse the local differential structure of the distance function: these feature events have a discontinuity along the iso-distance-surface they reside on, the branching feature points have moreover a discontinuity in the scale direction.

We follow the Boundary Morphology framework presented in (ElNaghy and Dorst, 2020) and compute the distance function on a discrete grid, with one axis aligned with the Lipschitz principal direction of the facet (see Fig. 2). Thus the features we compute are invariant under rigid body motions of the facet. We employ the linear-time ‘distance transform’ algorithm of (Maurer et al., 2003), but we compute the Euclidean distance directly to the fracture mesh vertices rather than on the binary representation of the surface. For each grid point, we not only maintain the actual distance ρ , but also a pointer to the actual mesh vertex that achieved this distance (one of the contact points). The whole 3D grid of these vertex indicators we call the *provenance map*.

We use the neighbourhood structure of the provenance map to extract the 3PC and the 3PO feature points by

analysing the type of contact at the disappearance event of every point in the discrete distance field. We specify this procedure in detail in Algorithm 1 (see Fig. 8). Consider the distance transform grid points at increasing distance values of their scales. When a facet mesh vertex V last appears at scale ρ in the dilation distance field, this implies that V must have been a part of the closing surfaces at all scales up to ρ ; but at larger scales, it no longer is. We label each vertex of the mesh with the value of ρ at which it disappears; a heat map of this function is shown in Fig. 9. By thresholding this function, we can determine the connected components of the mountain peaks (connected along the original mesh edges) at different scales. Those ‘peak components’ are then used to determine whether scale space events are possibly characteristic, for at characteristic scales a peak stops contributing to the provenance map. We thus need to keep track of the mountain peaks in order to detect the 3PC and the 3PO feature points. This is done by exploring the distance field going from large scales to smaller ones. This order of exploration enables the detection of all 3-point contact medial axis points resting on distinct mountain peaks. Later on, a subsequent RANSAC retains the most significant ones which are mostly the branching points of the medial axis. Going the other way around (from small to larger scales) would not allow identifying the mountain peaks since they would always be merged into larger contiguous components, which might lead to the detection of non-characteristic points as features.

Algorithm 1 Extraction of 3PC MM feature Points *

Input: **V**: list of $n \times 3$ vertex coordinates of fracture facet mesh **A**.
C: $n \times n$ adjacency matrix of the vertices fracture facet mesh **A**.

Output: list of 3PC MM feature points of fracture facet mesh **A**.

1. Compute the dilation distance field **D** of fracture facet mesh **A**.
2. Compute the provenance map **P** associated to the propagation of mesh points in the dilation distance field **D** (i.e., **P** is therefore the list of vertices v_i that caused each distance).
3. **for each** $v_i \in \mathbf{V}$ **compute the following**:
 - (a) $\rho_i \leftarrow$ the value of the scale ρ at which v_i last appears in the dilation distance field **D**.
 - (b) $x_i \leftarrow$ the grid location at which v_i last appears in the provenance map **P**.
 - (c) $nbr_{v_i} \leftarrow$ a list of provenance points in grid locations directly adjacent to x_i in the provenance map **P**, within a $3 \times 3 \times 3$ neighbourhood.
4. Sort all mesh vertices v_i into a list **V_{sorted}** in descending order of their associated ρ_i .
5. **for** $k = 1$ **to** n :
 - (a) $v_k \leftarrow \mathbf{V}_{\text{sorted}}(k)$, $\mathbf{NbrComps}_k \leftarrow \{\}$.
 - (b) Remove the points v_j satisfying: $\rho_j < \rho_k$ from the neighbourhood list nbr_{v_k} .
 - (c) Using the adjacency mesh information of **C**, extract the connected mesh components **MaxComps** constructed by vertices v_l such that: $\rho_l \geq \rho_k$.
 - (d) **for each** $n_i \in nbr_{v_k}$
 - i. find the mesh component $comp_i$ such that:
 $n_i \in comp_i$ and $comp_i \in \mathbf{MaxComps}$
 - ii. add $comp_i$ to $\mathbf{NbrComps}_k$ (the list of neighbouring components).
 - (e) Compute the number of unique components **NComps_k** of all mesh components in **NbrComps_k**.
 - (f) **if** $\mathbf{NComps}_k == 3$ **then**
 add v_k to the list of 3PC MM feature points of fracture facet **A**.
endif

Fig. 8 Algorithm for the extraction of 3PC MM feature points. The same algorithm is followed to extract the 3PO MM feature points using the dilation distance field of the upside down fracture surface facet (with inverted normals)

4.3 Description and Reliability of MM Features

To perform pairwise alignment, we could treat the locations of the set of feature points of the two fractures as unlabelled point clouds, and use a standard 3D point cloud alignment algorithm such as ICP. Instead, we choose to increase effectiveness and precision by employing the multi-local structure of the individual 3-point contacts that led to those features by encoding it in well-chosen *morphological descriptors*. We illustrate those descriptors in Fig. 6: the sphere radius, the area of the 3-point contact triangle and the area of the circumscribed circle of the triangle. These descriptors are the outcome of a detailed stability analysis performed in the Appendix of this paper.

The Appendix shows that not all 3-point contacts are equally stable: we perform a geometrical analysis on how much the sphere of an MM feature is affected by small perturbations of the fracture surface. The result is that spheres resting on peaks and whose 3-point contacts form an obtuse triangle are very sensitive to small displacements caused by




Fig. 9 A function on a facet mesh that shows, for each vertex V , the largest scale at which V can be touched by a rolling ball of that scale as its radius. At larger scales, the vertex can no longer appear in the closing surface of the facet. We show both closing radius (dilation distance function value) and opening radius (erosion distance function value). The colour bar is logarithmic, the value 2 indicates a radius of $\exp(2) \approx 7.5$ mm; the maximum computed radius was $\exp(4) = 60$ mm. This fracture facet is approximately 60×10 mm long

abrasions; such spheres are most stable when resting on a more-or-less equilateral triangle of peaks. Therefore, as a *reliability measure* of each MM feature we use ‘equilateralness’ (defined as the ratio of the areas of the contact triangle and its circumscribed circle, see Fig. 6).

We also demonstrate in the Appendix that 3-point contacts with very small triangles should be excluded as features, as they tend to be very unstable under discretization (not only of the scanning process, but also of the discretized computations we perform). Mid-range spheres are generally found to be quite stable. This involves the ‘flanks’ of the fracture surface, which will be treated later in Sect. 5.2.

5 Pairwise Alignment

Our preparations thus characterize each fracture facet by two sets of morphological feature points. The 3PC features bound the surface from above at characteristic locations, and are derived from the dilation distance field; the 3PO features bound from below, and are computed from the erosion distance field.¹ The MM features, augmented by their descriptors and the reliability measures, provide the information to perform the pairwise alignment.

¹ We in fact compute the 3PO features as 3PC features of the dilation distance transform of the upside-down facet.

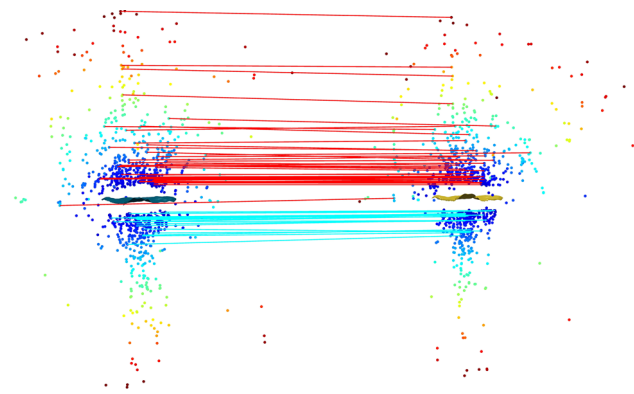


Fig. 10 Morphological feature correspondences: RANSAC inliers of a pair of complementary matching fracture facets. Note how RANSAC adaptively chooses among MM features at many scales, selecting correspondences for their effectiveness in the actual pairwise alignment

5.1 RANSAC

Our alignment method employs standard RANSAC (Derpanis, 2010), but on the MM features rather than on the original mesh (as illustrated in Fig. 10). To provide RANSAC with the required initial subset of hypothetical inliers, we compute a set of potential corresponding pairs between the MM features of the two fractures, based on similarity of their descriptors. The actual alignment of the fracture surfaces is based on the spatial locations (the sphere centres) of the MM features. Thus, the feature descriptors are only used to guide RANSAC.

For RANSAC to determine a 3D rigid body transformation (translation and rotation), we need a minimum of three corresponding pairs of feature points. To employ both upper and lower bounding of the surface in a balanced manner, we actually use two pairs of each (for a total of four pairs), enforcing two of the preliminary corresponding pairs at each iteration to be selected from the initial correspondences between the 3PC points of the first surface and the 3PO points of the second surface, and the other two from the initial correspondences between the 3PC points of the second surface and the 3PO points of the first surface. It employs a weighted version of the Procrustes method (Gower, 1975) (using the reliability measures of Sect. 4.3 as weights) to compute a hypothetical transformation. We have not found the need for initial approximate alignment, so the initial transformation is fully arbitrary. The RANSAC algorithm terminates after reaching the best possible transformation between a set of outlier-free feature correspondences.

RANSAC may fail to find an acceptable 3D rigid body transformation (based on the corresponding MM features), this may be an indication that the fractures suggested for pairwise alignment are actually not matching. Though enough matching corresponding pairs of features can usually be




Fig. 11 The flanks of a fracture in blue, non-flanks in yellow (and in cyan the transitions). The fracture surface size is about 30×25 mm (Color figure online)

found (based on the similarity of the associated feature descriptors), the failure occurs when either a transformation could not be found or the RANSAC feature-based alignment error is unacceptably high. This suggests that our method could also sensibly be used in the fragment selection phase of the archaeological workflow, when fragments are nominated for pairwise alignment.

5.2 ICP on Flanks

As we will see in the experiments of Sect. 7, RANSAC on MM feature points provides good alignments. But it is quite customary to follow RANSAC's best effort on inlier features with a further refinement step on the result, using more of the actual data. A standard ICP (Besl and McKay, 1992; Chen and Medioni, 1992; Fitzgibbon, 2003) on all facet points would try to minimize an MSE, and thus potentially destroy our carefully maintained morphological complementarity (with its distinction between penetration and missing parts).

We therefore propose to preserve the morphological nature of our complete pairwise alignment method by applying ICP only on the mesh points in the 'flank regions' of the fracture surfaces. These are the relatively flat sides of the mountainous fracture landscape, rather than its peaks or valleys ('inflection points' (Park and Lee, 1996) is a correct intuition). Our analysis in the Appendix suggests that the 3-point contacts resting on the flanks are especially stable. Archaeologically such contact points are less likely to abrade; and when they do, the displacement of the exact contact point tends to be tangential to the sphere, which the Appendix shows has no effect to first order.

Figure 11 shows such flanks (in blue) of a fracture, morphologically computed as the surface areas at which the maximum balls touching from above and below have approx-

imately equal sizes. Such places must exist contiguously around each mountain and valley by the intermediate value theorem: the sizes of the opening spheres (small at the peaks and big at the valleys) and closing spheres (big at the peaks and small at the valleys) have to cross over somewhere. The flank regions can thus be completely identified morphologically, through computing the difference between the dilation and erosion distance functions of each fracture surface (see Fig. 9). To the best of our knowledge, this way of identifying flank points morphologically is new.

We use ICP on this morphologically determined subset of the fracture mesh to refine the final alignment that RANSAC produces and we will study its effect in the experiments of Sect. 7.

6 Measuring Alignment Quality

We carefully consider how to evaluate pairwise alignment results in general, before proceeding with conducting our experiments and comparison. Our evaluation measures should reflect the asymmetrical nature of fitting. We agree with (Liao et al., 2020) that quantifying the regions of penetration is much more meaningful for evaluating the quality of the fit than just measuring the minimized total MSE error (as authors of linear methods typically do). This fails to measure the important distinction between gap and penetration areas. We propose a more relevant overall measure of the alignment quality, obtained by separately computing the weighted mean of the sum of the distances (Weighted Mean Error) of the areas of the facets with gaps (WMEG) and penetration (WMEP) between the aligned fragments. We define these as:

$$\text{WMEG} = \frac{\sum_{i \in I_G} a_i * d_i}{\sum_{i \in I_G} a_i},$$

and

$$\text{WMEP} = \frac{\sum_{i \in I_P} a_i * d_i}{\sum_{i \in I_P} a_i}.$$

Here d_i are the distances, which are counted negative for the gaps, and positive for the penetrations. I_G is the set of mesh point indices for which $d_i < 0$ (the gap part) and I_P is the set of mesh point indices for which $d_i > 0$ (the penetration part). The local surface area a_i (Meyer et al., 2003) occupied by each mesh vertex is used for weighting the alignment measurement. We also compute the absolute weighted mean error:

$$\text{AWME} = \frac{\sum_{i=1}^N a_i * |d_i|}{\sum_{i=1}^N a_i}.$$

Therefore, all measures are volumes over areas.

We compute the distances d_i with some care, to accommodate locally varying mesh density. Calling the larger fracture surface of the counter-fitting pairs ‘target mesh’ and the smaller ‘source mesh’, we compute d_i for every source vertex as the minimum of three kinds of distances: to the nearest target vertex, to the nearest target edge, and to the nearest target triangle.

These error measures are averages, though. A much more revealing impression of the quality of the alignment is the spatial distribution of gaps and penetrations through employing heat maps, as already suggested by (Brown et al., 2012). We will show an example of this visual evaluation in Experiment 1 below.

7 Experimental Results

To test and evaluate our pairwise alignment pipeline, we performed two experiments. The first experiment involves pairwise alignment of real archaeological fragments that have been scanned by an archaeological institute. The second is a comparison with a state of the art linear method for pairwise alignment on non-archaeological fragments.

7.1 Experiment 1: Archaeological Fragments

We first test our pairwise alignment pipeline on different pairs of real archaeological artefacts belonging to the same collection. The artefacts are some fragmented Iron Age terracotta vessels from the excavation at Tell Es-Safi/Gath in Palestine² made available to the GRAVITATE project by Bar-Ilan University. The fragments are about 5 to 12 cm in size.

The fragments are highly abraded, where the potential pairs of complementary fracture facets have been indicated by an expert. We show the results on six different pairs of fragments (see Fig. 12 and Table 2). The first two pairs (17 - 19) and (23 - 17) are challenging examples because the fracture surfaces are only partially fitting with some missing geometry, while the remaining four pairs have narrow fracture surfaces.

We applied the steps of our morphological pairwise alignment workflow to the Gath fragments. Figure 12 shows visually the alignment results. In many cases, the results of the RANSAC alignments are already convincing. Figure 13 shows an example of the best and the worst cases of

² Information about Gath excavation can be found at [https://en.wikipedia.org/wiki/Gath_\(city\)](https://en.wikipedia.org/wiki/Gath_(city)) and <https://gath.wordpress.com>.




Fig. 12 Experiment 1: Morphological feature-based alignment of archaeological fragment pairs. First and third rows: the pairs of fragments with their extracted facets outlined in red colour. Second and fourth rows: Alignment after RANSAC and ICP on flanks

RANSAC-based alignment and their refinement after applying ICP on flanks only. The result columns (column 3–5) of Table 2 shows how much we have reduced the data required to achieve these results: the MM features are a factor 3 to 10 less than the original facet mesh vertices; and the inliers that suffice to produce an accurate alignment are a factor 8 to 10 less in number still.

The results of the RANSAC alignment confirm our geometrical intuition: we are indeed capable of computing the alignment based on producing a simplified bounded representation of the fractures, at sufficiently characteristic locations. The small number of features determining the final alignment demonstrates the power of the MM representation. A more detailed analysis of which features end

up as RANSAC inliers confirms that the abraded fracture borders are not significant for fitting noisy fragments (they are apparently too sensitive to abrasion). The results support our geometrical analysis of the Appendix, that the centres of spheres leaning on 3-point contacts forming nearly equilateral triangles should be expected as the most useful for fitting purposes. The *inliers* are apparently the most stable and representative features for our pair of fracture facets.

The WMEx measures of Table 2 show that the additional ICP refinement on flanks indeed slightly refines the already well-aligned results of RANSAC. Table 2 shows (in its last column) that only a few iterations are required for this refinement. Since the flank regions contain only about 40% of the fracture surface, this extra refinement is relatively fast. It

Table 2 Evaluation of our alignment method on different pairs of complementary archaeological fragments of Fig. 12. The Size and Vertices are those of the fracture facets of interest

| Pairs | Fracture Size (mm) | Fracture Vertices | MM Features | Inliers | Inlier RMSE (mm) | RANSAC (mm) | | | ICP on flanks (mm) | | | ICP Iterations |
|-------|--------------------|-------------------|-------------|---------|------------------|-------------|------|-------------|--------------------|------|-------------|----------------|
| | | | | | | WMEG | WMEP | AWME | WMEG | WMEP | AWME | |
| 17 | 60 × 25 | 5840 | 705 | 67 | 0.55 | -0.23 | 0.11 | 0.17 | -0.27 | 0.11 | 0.15 | 3 |
| 19 | 30 × 25 | 3767 | 415 | | | | | | | | | |
| 23 | 128 × 25 | 11,878 | 1451 | 32 | 0.96 | -0.23 | 0.15 | 0.19 | -0.16 | 0.12 | 0.13 | 5 |
| 17 | 58 × 25 | 5499 | 358 | | | | | | | | | |
| 01 | 34 × 4.5 | 2157 | 530 | 63 | 1.32 | -0.64 | 0.29 | 0.55 | -0.24 | 0.16 | 0.18 | 10 |
| 09 | 34 × 4.5 | 1871 | 492 | | | | | | | | | |
| 11 | 50 × 4 | 2638 | 669 | 53 | 1.02 | -0.24 | 0.25 | 0.25 | -0.11 | 0.12 | 0.12 | 6 |
| 06 | 50 × 4 | 2391 | 691 | | | | | | | | | |
| 10 | 69 × 5 | 2972 | 307 | 70 | 0.97 | -0.9 | 1.65 | 1.32 | -0.28 | 0.33 | 0.31 | 12 |
| 13 | 60 × 5 | 2736 | 250 | | | | | | | | | |
| 05 | 35 × 4 | 1728 | 772 | 81 | 1.34 | -0.58 | 0.09 | 0.54 | -0.14 | 0.16 | 0.16 | 8 |
| 07 | 37 × 4 | 1349 | 428 | | | | | | | | | |

The bold indicates the best results as usually performed in most Computer Vision Literature




Fig. 13 Refinement achieved by applying ICP on flanks following RANSAC alignment for two different pairs of fragments. The left column shows alignment after RANSAC, while the right column shows the refinement after applying ICP. The refinement in alignment achieved by ICP on flanks over RANSAC is more noticeable for the lower pair of fragments




Fig. 14 Colour-coded signed alignment error computed for a fracture facet of about 30 × 25 mm, after the ICP refinement on flanks. The blue corresponds to gap regions, and red corresponds to penetration regions. Note that zero error is indicated by orange and that the alignment error is measured in mm. In our example, the digitization noise is estimated to be maximum about 0.6 mm. (This happens to be the overlap region resulting from the pairwise alignment of fragments 17 and 19 after application of flank-based ICP in the top right of Fig. 13.) (Color figure online)

achieves a fitting with small misalignment error as shown quantitatively in the AWME column of Table 2 and visually in Fig. 12.

The final alignment can be further visually analysed through a graphical heat map depiction of the local distance map as shown in Fig. 14. Complementary visual results like these give the archaeological user the most easily interpretable indication of the quality of the pairwise alignment by seeing the distribution of the gap and penetration regions across the fracture surface. Coupled with WMEs for penetration and gap, this provides the archaeologist with visual and quantitative measures that could help him/her better evaluate




Fig. 15 Pairwise alignment through the concave-convex patches characterization in the CCP method of Li et al. (Li et al., 2020)

the quality of the fit from a broader archaeological perspective.

7.2 Experiment 2: Comparison with Linear Method

The field of digital archaeological reassembly has no generally agreed benchmark datasets or criteria for effective evaluation and comparison. This severely hampers the comparison process with other proposed techniques. In addition, it is not easy to obtain the scanned 3D digitized versions of real archaeological fragments due to different provenance and proprietary considerations.

(Huang et al., 2006) made their dataset of freshly broken artefacts available³, and it begins to play the role of a benchmark. However, since the pipeline in (Huang et al., 2006) focuses mainly on multi-piece alignment “global reassembly” rather than pairwise alignment, we can not directly compare our results to theirs.

Fortunately, (Li et al., 2020) proposed a recent method for pairwise alignment to which we can compare. It is inspired by Huang et al.’s approach and applied on Huang’s public 3D data of brick fragments. This data consists of 6 fragments made of stone and scanned in the form of dense point set surfaces (not meshes).

Li et al. used the boundary curves and concave-convex patches of fracture surfaces to perform pairwise matching (see Fig. 15). Let us refer to their method as ‘CCP’. In the CCP method, the fracture surface pairs are aligned using a modified version of ICP which focuses only on the points of concave-convex matched patches to optimize the fine alignment results. Before the extraction of concave and convex patches, the fragments undergo a preprocessing step of adaptive linear smoothing (Lavoué, 2007) to reduce their noisy bumpiness. (Li et al., 2020) offer evaluation results of their method using both the smoothed and the non-smoothed versions of the brick fragments.

We compare the performance of our method to theirs on pairs of the non-smoothed brick fragments, since providing the best alignment for those is what matters archaeologically. As a preprocessing stage, we first convert the dense point set

surfaces into 3D meshes using (Cignoni et al., 2008). We then extract the fracture surfaces using our Faceting technique (ElNaghy and Dorst, 2017). Figure 16 shows the alignment results of our method, using both RANSAC and subsequent flank-based ICP on the eight different pairs of the brick fragments.

Table 3 shows the comparison between the performances of Li et al.’s CCP method and our method on the brick fragments in terms of the RMS error, alignment time and number of extracted features. Note that the CCP method fails to align pair (2 - 5) while our method fails to align pair (2 - 4)⁴. Comparing the quantitative measures, our method outperforms the CCP method in all respects with a smaller number of features used, a faster alignment speed and a higher accuracy of alignment measured in terms of RMS error. We use RMS error for evaluation and comparison, since this is the measure used by Li et al. for evaluating their own results. Unfortunately, their implementation is not publicly available so we could not apply our own evaluation measures to their alignment results.

In a sense, the two methods are using complementary sets of mesh points in their final ICP alignment: the linear method works by aligning the points of the concave-convex patches (see Fig. 15), while our non-linear morphological method mainly aligns the points of the flank regions (see Fig. 11). As we reasoned in the Appendix, the flanks are the more stable regions to use for alignment; and this comparative study nicely confirms that insight.

7.3 Parameter Settings and Time Analysis

For experiments 1 and 2 we used the same parameter settings for our method, since the fragments are of similar size. We provide those details now.

Using a distance transform 3D grid of resolution 0.25 mm with a maximum allowed morphological scale of 60 mm, we effectively achieve a bounded simplified representation of each of the counter-fitting fracture surfaces. The resulting alignment is exact within the accuracy of the sampling we employ for the distance transform computation (see results of Experiments 1 and 2). This resolution of 0.25 mm was initially chosen for storage and computational reasons. RANSAC aligns fractures using MM features based on this grid resolution; the fact that subsequent refinement using flank-based ICP (which works on the original fracture surface) hardly improves on the RANSAC results shows that the chosen grid resolution was sufficient. We performed an additional experiment to study the effect of going to half the grid resolution twice (0.5 mm and 1 mm). This considerably reduces the time and storage required for the distance trans-

³ The 3D Digital Models of the dataset are available at <https://www.geometrie.tuwien.ac.at/ig/3dpuzzles.html>.

⁴ This is actually the first example of a failed alignment in all our experiments; a possible explanation is given in the caption of Fig. 17.




Fig. 16 Experiment 2: Morphology-based alignment of eight brick fragment pairs

Table 3 Comparison between our method and CCP method (Li et al., 2020) for pairwise alignment of brick fragments. CCP* indicates that the number of ICP iterations was computed for aligning the smoothed versions of the brick fragments (no data has been reported by the authors for the original fragments ICP number of Iterations)

| Fragmen Pairs | Fracture Size (mm) | Fracture Vertices | Number of features | | Matching Time (s) | | ICP Iterations | | RMS error (mm) | |
|---------------|--------------------|-------------------|--------------------|------|-------------------|------|----------------|------|----------------|---------------|
| | | | CCP | Ours | CCP | Ours | CCP* | Ours | CCP | Ours |
| 1 | 34.5 × 32.4 | 6157 | 136 | 131 | 15 | 10 | 13 | 6 | 0.2011 | 0.0459 |
| 2 | 33.8 × 33.6 | 6091 | 145 | 60 | | | | | | |
| 1 | 46 × 29.3 | 7096 | 221 | 17 | 21 | 4 | 10 | 6 | 0.1789 | 0.0684 |
| 3 | 45.6 × 29.1 | 7132 | 236 | 18 | | | | | | |
| 1 | 32.4 × 29.5 | 7127 | 186 | 34 | 19 | 7 | 12 | 14 | 0.1267 | 0.0387 |
| 4 | 32.6 × 29.2 | 7056 | 174 | 33 | | | | | | |
| 2 | 29.3 × 21.5 | 7731 | 177 | 91 | 19 | 4 | 9 | 7 | 0.1103 | Fail |
| 4 | 32.6 × 29.2 | 7056 | 174 | 57 | | | | | | |
| 2 | 29.3 × 21.5 | 7731 | 177 | 62 | 15 | 9 | 8 | 13 | Fail | 0.0500 |
| 5 | 28.9 × 17.9 | 7113 | 91 | 30 | | | | | | |
| 4 | 36.6 × 28.6 | 5759 | 190 | 59 | 17 | 5 | 13 | 5 | 0.1267 | 0.0712 |
| 5 | 37.4 × 29.1 | 5654 | 161 | 52 | | | | | | |
| 4 | 29.4 × 11.3 | 6088 | 190 | 45 | 13 | 5 | 10 | 2 | 0.1567 | 0.0518 |
| 6 | 29.7 × 11.9 | 6089 | 56 | 16 | | | | | | |
| 5 | 29.4 × 20.6 | 6178 | 95 | 30 | 11 | 6 | 13 | 5 | 0.1211 | 0.0366 |
| 6 | 29.4 × 20.3 | 5919 | 100 | 23 | | | | | | |

The bold indicates the best results as usually performed in most Computer Vision Literature

form computation. The WMEs measures for the RANSAC step were still accurate enough to perform ICP on flanks, achieving the same accurate alignment results (at the expense of doubling the iterations), though the resolution for computing the MM features has been much reduced.

Each MM feature is characterized by an ‘equilateralness’ area ratio as a reliability measure (see Sect. 4.3). The unreliable feature points with area ratio less than 0.1 are excluded from the feature matching process. The ini-

tial correspondences between the feature points passed to RANSAC for outlier rejection are established by measuring weighted Euclidean distances between their associated descriptors introduced in Sect. 4.3.

Within the RANSAC computation of the best transformation, we weigh the contribution to the Procrustes fit for each pair of corresponding MM features, by the minimum (i.e., worst) of their associated reliability measures. The RANSAC




Fig. 17 The match between fragment 2 and 4 is a partial match, where a relatively small fraction of the counter-fitting fracture surfaces overlaps. The MM features of this overlap region on each of the fragments is heavily affected by the contiguous part of that fragment, especially at larger scales, as explained in Sect. 3.2 (For more detail see Sect. 3.2 in ElNaghy and Dorst (2020)). The failure of our method in this case apparently implies that the overlapping regions are not similar enough in their features to guide the RANSAC alignment

tolerance threshold computed on the MM feature spatial locations was set to about 1.5 mm for our fragments.

In our test cases of Figs. 12 and 16, the yellow fragments have been selected as the source meshes while the blue fragments have been selected as the target. In case of complete surface alignment, it does not make a big difference which one is which. However, in case of partial fracture surface alignment, the fine alignment convergence is much faster when the smaller surface is spatially transformed to the larger fracture surface. The smaller fracture can be easily identified by measuring the overall surface area of each fracture mesh using (Meyer et al., 2003).

In the flank-based ICP fine refinement, the flank regions are computed as having a maximum difference of 3 mm between their opening and closing distance functions; they account for about 40% of the overall fracture surface area in our experiments.

Computation times differ greatly per pipeline step of Fig. 1. Faceting a 3D model with 50K vertices takes less than 5 minutes, where it is mainly dominated by the 3D model size and the scale ρ used for defining the neighbourhood over which the fracture angle is calculated (see Sect. 4 in (ElNaghy and Dorst, 2017)). This faceting processing time is deemed acceptable within GRAVITATE, since the Faceting is part of the once-only inclusion (ingestion) of the fragment data into the system. Most time-consuming is the computation of the distance transform on the voxel grid, step 2. This took on average 3 minutes per fracture mesh. Step 3, the extraction of the morphological feature points and matching them, took maximum 5 seconds. Like the Faceting of step 1, steps 2 and 3 need to be performed only once, when a fragment is entered into the virtual archaeological system. The RANSAC outlier rejection of step 4 took on average about 1 minute and refining




Fig. 18 Abraded fracture facet by 1 mm (on the right) versus non-abraded facet (on the left). The fracture facet measures about 30×25 mm

the alignment by ICP on flanks (step 5) took a few seconds. We performed our experiments on an Intel(R) Xeon(R) CPU 2.9 GHz \times 8 computer with 256 GB RAM.

8 Limitations and Extensions

As the results show, our method works well for a range of fragments. To study its sensitivity, we perturb the data artificially to mimic abrasion and digitization effects. We also obtain a preliminary impression of the robustness of our pairwise alignment method by attempting global reassembly on the dataset of (Li et al., 2020).

8.1 Robustness to Abrasion

Terracotta abrades easily, but also has an inner coherence. The general effect is that abrasion mostly leads to a wear on protrusions like the small peaks and the sharp borders of a fracture. This loss of material can be sensibly modelled by a small morphological opening.

We mimicked the abrasion effect by applying a morphological opening to one of our Gath archaeological fragments by a small ball of radius 1 mm, see Fig. 18. The abraded fragment surface was extracted using the marching cube algorithm and the fracture mesh A' is delineated out of the whole opened fragment using the Faceting method of (ElNaghy and Dorst, 2017). We compared the morphological features of fracture surface A' to the ones characterizing the non-abraded facet A (which has been extracted from the same fragment before applying the opening effect). A is the fracture facet of fragment number 19 in Table 2, involved in its alignment with fragment number 17 (see top left pair of Fig. 12).




Fig. 19 Resampling of a fracture facet mesh. The original fracture surface is first upsampled to increase its resolution. Then, quadratic edge collapse is applied to the upsampled version to retain its original resolution and topology with different sampling. The zoomed-in subfigure shows the difference between the original mesh A (red edges) and its resampled version A' (blue edges) (Color figure online)

Fracture facet A is characterized by 654 3PC and 386 3PO feature points, while A' is characterized by 573 3PC and 345 3PO feature points. The reduction in the number of 3PC points is higher than that in the number of 3PO points, since abrasion (as an opening) affects the morphological closing of a fracture facet more than its opening (see Sect. 5 in (ElNaghy and Dorst, 2020)). The feature points range from scale 5 to 60 mm, as expected the lower scales are more sensitive to abrasion.

Matching the features of A against A' produced 294 3PC and 170 3PO unique matching feature pairs. In an experiment where we made A' play the role of the counterpart of A , RANSAC adaptively selected 67 corresponding MM features as inliers. Since this is comparable to the number of inlier features selected in regular pairwise alignment (see first row of Table 2), we find that a sufficient number of stable features remains after (artificial) abrasion to establish an accurate pairwise alignment.

8.2 Variations in Digitization

In GRAVITATE, we are particularly interested in digitally relating artefacts dispersed over different collections. Such archaeological artefacts are likely to be digitized by different museums or archaeological institutions with different scanners. Non-consistent registration and post-scanning techniques are frequently applied to the raw scanner data. As a consequence, counter-fitting fracture surfaces suggested for alignment across collections are expected to be sampled differently with different digitization. We actually observed in our experiments that the mesh sampling varies even among fragments belonging to the same collection.




Fig. 20 a and b: Two pairwise morphology-based alignments of an artificially broken archaeological artefact. c: Their combination achieves an acceptable global reassembly

To get an impression of the effect of varying digitization on our proposed alignment technique, we performed a virtual experiment (we have no real case to study this effect): we artificially broke a real archaeological fragment into three pieces using another real archaeological fracture surface as a cutter. The fracture surface used for cutting was resampled on either side of the break to simulate different provenance. The applied resampling displaced the spatial locations of the mesh vertices while preserving the original topology (see Fig. 19). In this particular virtual experiment (see Fig. 20), the RANSAC based alignment was highly accurate with no need for further ICP refinement. Even with different sampling, our method managed to automatically identify a small set of stable feature pairs required for pairwise alignment (with an average of WMEG of -0.038 mm, WMEP of 0.036 mm, and AWME of 0.046 mm for an object of about 27×17 mm).

8.3 Towards Global Reassembly

Though we have designed our method to work on pairs of fragments, one might wonder how well it would work when viewed as a step in a global alignment method. Being morphological rather than linear, the reassembly should be insensitive to the incompleteness of fragments themselves, and of the collection they belong to. Figure 21 shows the effect of composing a global restoration from the pairwise results of Experiment 2. The result is a quite satisfactory completed brick, and rather convincingly confirms the accuracy of the individual alignments in our pairwise method. If, as is common in global reassembly, a global optimization step is applied afterwards, it should be based on applying ICP on flanks only to respect the morphological aspects of complementary pairwise alignment. Figure 20c shows a similar result for three artificially broken head fragments. These results confirm that our method could be a reliable subpart of a system for global reassembly of broken archaeological artefacts.




Fig. 21 Global Reassembly by selected pairwise morphology-based alignment of brick fragment pairs. The insert shows in red lines the pairwise alignments that have been involved in the global alignment of the six brick fragments (Color figure online)

9 Conclusions and Future Work

We have shown that features based on the opening and closing scale spaces of mathematical morphology contain sufficient information of the essentials of a fracture surface, for it to be aligned with a counterpart. We demonstrated this on a mesh representation of abraded archaeological fragments and non-archaeological freshly broken fragments.

The morphological features were computed by means of a discrete distance transform on a voxel grid, augmented with a provenance map. This enabled us to determine stable and localizable morphological feature points, determining critical bounding spheres for the fracture surface. In the alignment process, a set of heuristic descriptors was defined to guide preliminary matching correspondence and the weighted Procrustes within RANSAC.

We observed how the alignment actually depended mostly on flank-based MM features, and refined it by applying ICP purely on those flank points. This flank-based ICP slightly improved the RANSAC result and it may be a useful MM-based alignment method by itself. There is acceptable penetration in the final alignment results, within the bounds of the digitization noise of the archaeological fragments. We have also shown how our morphological method performs better than the state of the art linear method of (Li et al., 2020).

The economical characterization by our MM features of essential aspects of fractures, in a scale-space manner, may also make them suitable for an earlier stage in a reconstruc-

tion pipeline: automatically suggesting potential matches. (Actually, in a typical archaeological workflow, such suggestions are based on more than merely the geometrical fracture surfaces; one may employ local curvature, outward designs, or human expertise to do this more effectively.)

The capability of 3D compact morphological feature characterization (Yang et al., 2017) is potentially applicable to other fields like 3D object retrieval (Bustos et al., 2007), 3D object classification (Biasotti et al., 2006) or the adaptive simplification of surfaces for variable resolution rendering in computer graphics (Shaffer and Garland, 2001).

We have shown that the proposed morphological features can guide precise pairwise alignment (and may even be useful in global reconstruction). That capability is not only required in archaeology, but also has uses in other fields: such as bone fracture reduction in medical applications (Liao et al., 2020), skull reassembly in forensic investigation (Zhang et al., 2015) or even in reassembling broken meteorites (de Vet, 2015).

Linear methods prevail as a default approach to alignment (as they do in many fields); we have shown that ‘complementary pairwise alignment’ is better solved by exploiting its particular type of asymmetric non-linearity, which Mathematical Morphology encodes so well.

Acknowledgements This research was funded by the GRAVITATE project (<https://gravitate-project.eu/>) under EU2020 -REFLECTIVE-7-2014 Research and Innovation Action, grant no. 665155. Datasets are from the GRAVITATE data collection. The archaeological artefacts are digitized and kept by the British Museum, Cyprus Institute and the Archaeology laboratory at the Hebrew University of Jerusalem. The 3D models for the comparison study are courtesy of Vienna University of Technology.

Declarations

Conflicts of interest The authors declare that they have no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendix

Let us consider how sensitive the centre of a sphere with radius ρ , passing through three points \mathbf{p}_i (with $i \in \{1, 2, 3\}$) is to perturbations of those points. We consider only small perturbations so that the problem becomes linear. Take the sphere centre as the origin, and consider the perturbation of point \mathbf{p}_1 by a small additive vector \mathbf{t}_1 . The disturbance \mathbf{x}_1 of the sphere centre due to that perturbation alone, keeping the other points in place, is given by:

$$\mathbf{x}_1 = \frac{\mathbf{t}_1 \cdot \mathbf{p}_1}{\mathbf{n}_1 \cdot \mathbf{p}_1} \mathbf{n}_1,$$

where $\mathbf{n}_1 = \mathbf{p}_2 \times \mathbf{p}_3$. For \mathbf{x}_2 and \mathbf{x}_3 the expressions are similar.

Proof We use 3D geometric algebra to analyse this in a coordinate-free manner, employing its wedge product \wedge and its geometric division operation $/$ to construct and manipulate the subspaces involved⁵. Since the radius of the sphere does not change, we have three conditions on \mathbf{x}_1 :

1. Since $\|\mathbf{p}_2 - \mathbf{x}_1\| = \|\mathbf{p}_3 - \mathbf{x}_1\| = \|\mathbf{p}_2\| = \|\mathbf{p}_3\| (= \rho)$ to first order, we must have $\mathbf{x}_1 \cdot \mathbf{p}_2 = \mathbf{x}_1 \cdot \mathbf{p}_3 = 0$; therefore $\mathbf{x}_1 \cdot (\mathbf{p}_2 \wedge \mathbf{p}_3) = 0$. Geometrically, \mathbf{x}_1 is perpendicular to the $(\mathbf{p}_2 \wedge \mathbf{p}_3)$ -plane.
2. The radius remains the same, so $\|\mathbf{p}_1 + \mathbf{t}_1 - \mathbf{x}_1\| = \|\mathbf{p}_1\| (= \rho)$; to first order therefore $\mathbf{t}_1 \cdot \mathbf{p}_1 = \mathbf{x}_1 \cdot \mathbf{p}_1$.
3. Since we are in 3D, we have $\mathbf{x}_1 \wedge \mathbf{p}_1 \wedge \mathbf{p}_2 \wedge \mathbf{p}_3 = 0$.

Combining these conditions into one geometric product, we find $\mathbf{x}_1 (\mathbf{p}_1 \wedge \mathbf{p}_2 \wedge \mathbf{p}_3) = (\mathbf{t}_1 \cdot \mathbf{p}_1) (\mathbf{p}_2 \wedge \mathbf{p}_3)$ so that:

$$\begin{aligned} \mathbf{x}_1 &= (\mathbf{t}_1 \cdot \mathbf{p}_1) (\mathbf{p}_2 \wedge \mathbf{p}_3) / (\mathbf{p}_1 \wedge \mathbf{p}_2 \wedge \mathbf{p}_3) \\ &= (\mathbf{t}_1 \cdot \mathbf{p}_1) / \mathbf{p}_1^\perp, \end{aligned}$$

where $\mathbf{p}_1^\perp \equiv (\mathbf{p}_1 \wedge \mathbf{p}_2 \wedge \mathbf{p}_3) / (\mathbf{p}_2 \wedge \mathbf{p}_3)$ is the component of \mathbf{p}_1 perpendicular to the $(\mathbf{p}_2 \wedge \mathbf{p}_3)$ -plane (see e.g., (Dorst et al., 2009)), the geometric product division by a vector \mathbf{y} is multiplication by $\mathbf{y} / \|\mathbf{y}\|^2$. One may convert this result into the classical linear algebra formula above for \mathbf{x}_1 , by dually characterizing the plane bivector $\mathbf{p}_2 \wedge \mathbf{p}_3$ by its normal vector $\mathbf{n}_1 \equiv \mathbf{p}_2 \times \mathbf{p}_3$, rewriting \mathbf{p}_1^\perp to $\mathbf{p}_1^\perp = (\mathbf{n}_1 \cdot \mathbf{p}_1) \mathbf{n}_1$ and substituting. \square

We make some observations:

- When \mathbf{p}_1^\perp is small relative to \mathbf{p}_1 , which happens when the triangle formed by the three points is acute at \mathbf{p}_1 , the displacement becomes large. Therefore the situations where the three contact points form an acute triangle tend to be less stable.

⁵ https://en.wikipedia.org/wiki/Geometric_algebra.

- On the other hand, if \mathbf{t}_1 is perpendicular to \mathbf{p}_1 , it leads to no displacement of the centre from \mathbf{p}_1 's displacement (to first order). This tends to happen at locations where the sphere rest tangentially on a flat flank of the fracture landscape.
- Note that the magnitude of $\|\mathbf{p}_1\|$, which is the radius ρ , plays no role in the formula: displacements are absolute – though their relative effect given a maximum abrasion displacement $\|\mathbf{t}_1\|$ is smaller for larger spheres.

Small perturbances of the other points propagate in the same manner, so the total displacement \mathbf{x} of the centre is their sum:

$$\mathbf{x} = \sum_i \mathbf{x}_i = \sum_i \frac{\mathbf{t}_i \cdot \mathbf{p}_i}{\mathbf{n}_1 \cdot \mathbf{p}_i} \mathbf{p}_i.$$

Depending on the nature of the contact of the three points, and hence their sensitivity to abrasion \mathbf{t}_i , this equation determines the stability of the MM feature point.

Arguably the worst case is when all displacements of contact points are radial, so that $\mathbf{t}_i \cdot \mathbf{p}_i = \rho \|\mathbf{t}_i\|$. This might happen when all contact points are peaks of the fracture landscape. Assuming these would abrade similarly, we may get a fair impression of the perturbation by taking all $\|\mathbf{t}_i\|$ the same and equal to $\|\mathbf{t}\|$. Then the displacement $\mathbf{x}_* = \|\mathbf{t}\| \rho / \mathbf{c}$, where $\mathbf{c} \equiv 1 / (\sum_i 1 / \mathbf{p}_i^\perp)$ is the vector of the circumcentre of the triangle formed by the contact points.⁶ The magnitude of this worst case \mathbf{x}_* equals $\|\mathbf{x}_*\| = \|\mathbf{t}\| (\rho / \|\mathbf{c}\|)$. Since $\rho / \|\mathbf{c}\|$ is smallest for an equilateral triangle, those are most stable; and the more obtuse the triangle, the larger this worst case error becomes. A well-known measure for the obtuseness of a triangle is the ratio of the area of the triangle relative to the area of its circumscribed circle (the smaller, the more obtuse); this is what we use in the heuristic matching of MM features for RANSAC, see Sect. 5.1.

When we consider the MM features of the original surface relative to those of an abraded surface, the type of contacts we can expect depend on the scale ρ (i.e., the radius of the sphere). Different types of contacts have different expected abrasive perturbances \mathbf{t}_i ; and then the above analysis determines the following effects:

- For tiny scales (radii) ρ , spheres will rest against small protrusions, which are sensitive to abrasion; the $\|\mathbf{t}_i\|$ will be large relative to ρ , and lead to large perturbations. We expect therefore small-radius features to be very unstable; since they are unlikely to occur on the counterpart

⁶ The expression for the circumcentre \mathbf{c} follows from $\|\mathbf{c} - \mathbf{p}_1\| = \|\mathbf{c} - \mathbf{p}_2\| = \|\mathbf{c} - \mathbf{p}_3\|$ so that $\mathbf{c} \cdot \mathbf{p}_1 = \mathbf{c} \cdot \mathbf{p}_2 = \mathbf{c} \cdot \mathbf{p}_3$ and from $(\mathbf{c} - \mathbf{p}_1) \wedge (\mathbf{p}_1 - \mathbf{p}_2) \wedge (\mathbf{p}_2 - \mathbf{p}_3) = 0$, giving $\mathbf{c} \wedge (\mathbf{p}_1 \wedge \mathbf{p}_2 + \mathbf{p}_2 \wedge \mathbf{p}_3 + \mathbf{p}_3 \wedge \mathbf{p}_1) = \mathbf{p}_1 \wedge \mathbf{p}_2 \wedge \mathbf{p}_3$; these combine to $\mathbf{c} = (\mathbf{p}_1 \wedge \mathbf{p}_2 \wedge \mathbf{p}_3) / (\mathbf{p}_1 \wedge \mathbf{p}_2 + \mathbf{p}_2 \wedge \mathbf{p}_3 + \mathbf{p}_3 \wedge \mathbf{p}_1) = 1 / (1 / \mathbf{p}_1^\perp + 1 / \mathbf{p}_2^\perp + 1 / \mathbf{p}_3^\perp)$.

anyway, they will not make good MM features for alignment.

- For large ρ , a sphere is more likely to rest on three peaks. That makes it *absolutely* sensitive to perturbation, since this is similar to the worst case we treated above (though *relatively* less so, since the $\|\mathbf{t}_i\|$ will be small compared to ρ). Especially sensitive would be spheres whose supports deviate strongly from an equilateral triangle, as we showed. We do expect such large scale features to be useful for alignment with the complementary counterpart, so we would like to incorporate a number of them; but we should prefer those from the more equilateral contacts.
- For medium scales, we may find rather stable spheres resting against three ‘flank’ points; those are already more protected from typical archaeological abrasion (which affects the peaks mostly), so the \mathbf{t}_i ’s are small. Moreover, since the sphere rests tangentially at flank points, the \mathbf{t}_i will tend to be perpendicular to the \mathbf{p}_i , leading to almost zero displacement. The acuteness of the triangle is of less importance now. The features based on flanks should moreover be very similar on the complementary fracture surfaces, so they are likely to be both characteristic and stable.

References

- Altantsetseg, E., Matsuyama, K., & Konno, K. (2014). Pairwise matching of 3D fragments using fast Fourier transform. *The Visual Computer*, 30, 929–938. <https://doi.org/10.1007/s00371-014-0959-9>
- Alzaid, A., & Dogramadzi, S. (2019). Reassembly of fractured object using fragment topology. In: 10th International Conference on Pattern Recognition Systems (ICPRS-2019), pp 98–105, <https://doi.org/10.1049/cp.2019.0256>
- Andreadis, A., Papaioannou, G., & Mavridis, P. (2015). Generalized digital reassembly using geometric registration. In: 2015 Digital Heritage, vol 2, pp 549–556, <https://doi.org/10.1109/DigitalHeritage.2015.7419572>
- Attali, D., Boissonnat, J.D., & Edelsbrunner, H. (2008). Stability and computation of medial axes—a state-of-the-art report. *Mathematical Foundations of Scientific Visualization, Computer Graphics, and Massive Data Exploration*, pp 109–125
- Besl, P. J., & McKay, N. D. (1992). A method for registration of 3-d shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2), 239–256. <https://doi.org/10.1109/34.121791>
- Biasotti, S., Giorgi, D., Marini, S., Spagnuolo, M., & Falcidieno, B. (2006). A comparison framework for 3D object classification methods. In: International Workshop on Multimedia Content Representation, Classification and Security, Springer, pp 314–321
- Brown, B. J., Toler-Franklin, C., Nehab, D., Burns, M., Dobkin, D., Vlachopoulos, A., Doumas, C., Rusinkiewicz, S., & Weyrich, T. (2008). A system for high-volume acquisition and matching of fresco fragments: Reassembling Thera wall paintings. *ACM Transactions on Graphics*, 27(3), 1–9.
- Brown, B. J., Laken, L., Dutré, P., Van Gool, L., Rusinkiewicz, S., & Weyrich, T. (2012). Tools for virtual reassembly of fresco fragments. *International Journal of Heritage in the Digital Era*, 1(2), 313–329.
- Bustos, B., Keim, D., Saupe, D., & Schreck, T. (2007). Content-based 3D object retrieval. *IEEE Computer Graphics and Applications*, 27(4), 22–27.
- Chang, M.C., Leymarie, F.F., & Kimia, B.B. (2004). 3d shape registration using regularized medial scaffolds. In: Proceedings 2nd International Symposium on 3D Data Processing, Visualization and Transmission, 2004 3DPVT 2004 pp 987–994
- Chen Y, Medioni G (1992) Object modelling by registration of multiple range images. *Image and Vision Computing* 10(3), 145–155, [https://doi.org/10.1016/0262-8856\(92\)90066-C](https://doi.org/10.1016/0262-8856(92)90066-C), <https://www.sciencedirect.com/science/article/pii/026288569290066C>.
- Cignoni, P., Callieri, M., Corsini, M., Dellepiane, M., Ganovelli, F., & Ranzuglia, G. (2008). Meshlab: an open-source mesh processing tool. *Eurographics Italian chapter conference, Salerno, Italy, 2008*, 129–136.
- Derpanis, K. G. (2010). Overview of the RANSAC algorithm. *Image Rochester NY*, 4(1), 2–3.
- Dorst, L., & van den Boomgaard, R. (2000). The systems theory of contact. In: Sommer G, Zeevi Y (eds) Algebraic Frames for the Perception-Action Cycle, Springer, vol LNCS 1888, pp 22–47
- Dorst, L., Fontijne, D., & Mann, S. (2009). Geometric Algebra for Computer Science: An Object-oriented Approach to Geometry. Morgan Kaufman, <http://www.geometricalgebra.net>
- Dougherty, E. (2018). Mathematical morphology in image processing. CRC Press.
- ElNaghy, H., & Dorst, L. (2017). Geometry based faceting of 3D digitized archaeological fragments. In: 2017 IEEE ICCV Workshops, pp 2934–2942, <https://doi.org/10.1109/ICCVW.2017.346>
- ElNaghy, H., & Dorst, L. (2020). Boundary morphology for hierarchical simplification of archaeological fragments. *Mathematical Morphology - Theory and Applications*, 4(1), 46–63. <https://doi.org/10.1515/mathm-2020-0101>
- Fitzgibbon, A. (2003). Robust registration of 2D and 3D point sets. *Image and Vision Computing*, 21, 1145–1153.
- Gower, J. C. (1975). Generalized Procrustes Analysis. *Psychometrika*, 40(1), 33–51.
- Haralick, R. M., Sternberg, S. R., & Zhuang, X. (1987). Image analysis using mathematical morphology. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 9(4), 532–550.
- Huang, H., Yin, K., Gong, M., Lischinski, D., Cohen-Or, D., Ascher, U., & Chen, B. (2013). *Mind the Gap: Tele-registration for Structure-Driven Image Completion*. New York: NY, USA.
- Huang, QX., Flöry, S., Gelfand, N., Hofer, M., & Pottmann, H. (2006). Reassembling fractured objects by geometric matching. In: ACM SIGGRAPH Papers, ACM, New York, NY, USA, SIGGRAPH’06, pp 569–578, <https://doi.org/10.1145/1141911.1141925>
- Kalvin, A. D., Schonberg, E., Schwartz, J. T., & Sharir, M. (1986). Two-dimensional, model-based, boundary matching using footprints. *The International Journal of Robotics Research*, 5, 38–55.
- Koller, D., Trimble, J., Najbjerg, T., Gelfand, N., & Levoy, M. (2006). Fragments of the city: Stanford’s digital Forma Urbis Romae project. In: Proceedings of the Third Williams Symposium on Classical Architecture, Journal of Roman Archaeology Suppl., 61, 237–252.
- Kong, W., & Kimia, B. (2001). On solving 2D and 3D puzzles using curve matching. In: Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001, vol 2, pp II–II
- Lavoué, G. (2007). A roughness measure for 3D mesh visual masking. In: Proceedings of the 4th Symposium on Applied Perception in Graphics and Visualization, Association for Computing Machinery, New York, NY, USA, APGV ’07, pp 57–60, <https://doi.org/10.1145/1272582.1272593>
- Leitao, H., & Stolfi, J. (2002). A multiscale method for the reassembly of two-dimensional fragmented objects. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(9), 1239–1251.

- Li, Q., Geng, G., & Zhou, M. (2020). Pairwise matching for 3D fragment reassembly based on boundary curves and concave-convex patches. *IEEE Access*, 8, 6153–6161. <https://doi.org/10.1109/ACCESS.2019.2961391>
- Liao, S., Xiong, C., Liu, S., Zhang, Y., & Peng, C. (2020). 3D object reassembly using region-pair-relation and balanced cluster tree. *Computer Methods and Programs in Biomedicine*, 197, 105756. <https://doi.org/10.1016/j.cmpb.2020.105756> <https://www.sciencedirect.com/science/article/pii/S0169260720315893>.
- Maurer, C. R., Qi, R., & Raghavan, V. (2003). A linear time algorithm for computing exact Euclidean distance transforms of binary images in arbitrary dimensions. *IEEE Transactions on PAMI*, 25(2), 265–270. <https://doi.org/10.1109/TPAMI.2003.1177156>
- Mavridis, P., Andreadis, A., & Papaioannou, G. (2015). Fractured object reassembly via robust surface registration. In: EG 2015 - Short Papers, The Eurographics Association, <https://doi.org/10.2312/egsh.20151005>
- McBride, J. C., & Kimia, B. B. (2003). Archaeological fragment reconstruction using curve-matching. In: IEEE CVPR Workshops, vol 1, pp 3, <https://doi.org/10.1109/CVPRW.2003.10008>
- Mellado, N., Aiger, D., & Mitra, N. J. (2014). Super 4pcs fast global pointcloud registration via smart indexing. In: Proceedings of the Symposium on Geometry Processing, Eurographics Association, Goslar, DEU, SGP '14, pp 205–215
- Meyer, M., Desbrun, M., Schröder, P., & Barr, A. H. (2003). Discrete differential-geometry operators for triangulated 2-manifolds. In: Hege HC, Polthier K (Eds.) Visualization and Mathematics III, Berlin, Heidelberg, pp 35–57
- Palmas, G., Pietroni, N., Cignoni, P., & Scopigno, R. (2013). A computer-assisted constraint-based system for assembling fragmented objects. In: 2013 Digital Heritage International Congress (DigitalHeritage), vol 1, pp 529–536, <https://doi.org/10.1109/DigitalHeritage.2013.6743793>
- Papaioannou, G., Karabassi, E. A., & Theoharis, T. (2002). Reconstruction of three-dimensional objects through matching of their parts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(1), 114–124.
- Papaioannou, G., Schreck, T., Andreadis, A., Mavridis, P., Gregor, R., Sipiran, I., & Vardis, K. (2017). From reassembly to object completion: A complete systems pipeline. *Journal on Computing and Cultural Heritage*, 10(2), 1–22.
- Park, K., & Lee, C. (1996). Scale-space using mathematical morphology. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(11), 1121–1126.
- Phillips, S. C., Walland, P. W., Modafferi, S., Dorst, L., Spagnuolo, M., Catalano, C. E., Oldman, D., Tal, A., Shimshoni, I., & Hermon, S. (2016). GRAVITATE: Geometric and semantic matching for cultural heritage artefacts. In: Catalano CE, Luca LD (Eds.) Eurographics Workshop on Graphics and Cultural Heritage, The Eurographics Association, <https://doi.org/10.2312/gch.20161407>
- Sagiroglu, M., & Ercil, A. (2005). A texture based approach to reconstruction of archaeological finds. In: Proceedings of the 6th International Conference on Virtual Reality, Archaeology and Intelligent Cultural Heritage, Goslar, DEU, VAST'05, pp 137–142
- Scalas, A., Mortara, M., & Spagnuolo, M. (2020). A pipeline for the preparation of artefacts that provides annotations persistence. *Journal of Cultural Heritage*, 41, 113–124.
- Serra, J. (1986). Introduction to mathematical morphology. *Computer Vision, Graphics, and Image Processing*, 35(3), 283–305.
- Serra, J., & Soille, P. (2012). Mathematical morphology and its applications to image processing, vol 2. Springer Science & Business Media
- Shaffer, E., & Garland, M. (2001). Efficient adaptive simplification of massive meshes. In: Proceedings of the Conference on Visualization '01, IEEE, pp 127–551
- Sommella Mura, A. (2011). La dea col titolo dal tempio arcaico del foro boario. *Deliciae Fictiles IV* pp 177–178
- Son, T. G., Lee, J., Lim, J., & Lee, K. (2017). Reassembly of fractured objects using surface signature. *The Visual Computer*, 34, 1371–1381. <https://doi.org/10.1007/s00371-017-1419-0>
- Toler-Franklin, C., Brown, B., Weyrich, T., Funkhouser, T., & Rusinkiewicz, S. (2010). Multi-feature matching of fresco fragments. *ACM Transactions on Graphics*, 29(6), 1–12.
- Vendrell-Vidal, E., & Sánchez-Belenguer, C. (2014). A discrete approach for pairwise matching of archaeological fragments. *Journal on Computing and Cultural Heritage*, 7, 15. <https://doi.org/10.1145/2597178>
- de Vet, S. J. (2015). The 1925 meteorite fall near Ellemeet and Serooskerke The Netherlands. *Meteoritics & Planetary Science*, 50(11), E1–E17.
- Winkelbach, S., & Wahl, F. M. (2007). Pairwise matching of 3D fragments using cluster trees. *International Journal of Computer Vision*, 78, 1–13.
- Yang, J., Zhang, Q., & Cao, Z. (2017). The effect of spatial information characterization on 3D local feature descriptors: A quantitative evaluation. *Pattern Recognition*, 66, 375–391.
- Zhang, K., Yu, W., Manhein, M., Waggenpack, W., & Li, X. (2015). 3D fragment reassembly using integrated template guidance and fracture-region matching. In: 2015 IEEE International Conference on Computer Vision (ICCV), pp 2138–2146, <https://doi.org/10.1109/ICCV.2015.247>
- Zhang, Y., Li, K., Chen, X., Zhang, S., & Geng, G. (2018). A multi feature fusion method for reassembly of 3D cultural heritage artifacts. *Journal of Cultural Heritage*, 33, 191–200.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.