Research Article

Open Access

Hanan ElNaghy* and Leo Dorst

Boundary Morphology for Hierarchical Simplification of Archaeological Fragments

<https://doi.org/10.1515/mathm-2020-0101>

Received November 21, 2019; accepted April 20, 2020

Abstract: When fitting archaeological artifacts, one would like to have a representation that simplifies fragments while preserving their complementarity. In this paper, we propose to employ the scale-spaces of mathematical morphology to hierarchically simplify potentially fitting fracture surfaces. We study the masking effect when morphological operations are applied to selected subsets of objects. Since fitting locally depends on the complementarity of fractures only, we introduce ‘*Boundary Morphology*’ on surfaces rather than volumes. Moreover, demonstrating the Lipschitz nature of the terracotta fractures informs our novel extrusion method to compute both closing and opening operations simultaneously. We also show that in this proposed representation the effects of abrasion and uncertainty are naturally bounded, justifying the morphological approach. This work is an extension of our contribution earlier published in the proceedings of ISMM2019 [10].

Keywords: Boundary Morphology, Scale Space, Fracture Representation, Abrasion, Lipschitz, Extrusion.

1 Introduction

The GRAVITATE H2020 project aims at providing archaeologists with the virtual tools to analyse digital artifacts, distributed across several collections (<https://gravitate-project.eu/>). The project involves archaeological institutes and research partners with experience in 3D geometry processing, semantic analysis and ICT integration. The artifacts are digitally scanned by standard scanning techniques in the form of 3D meshes, capturing the geometrical properties of each object and some of its photometric properties. Reassembly of broken artifacts is one of the core objectives of the project.

Figure 1: Morphological simplification pipeline for archaeological fracture facets.

***Corresponding Author: Hanan ElNaghy:** Computer Vision Lab, Informatics Institute, University of Amsterdam, Amsterdam, The Netherlands, E-mail: hanan.elnaghy@uva.nl

Leo Dorst: Computer Vision Lab, Informatics Institute, University of Amsterdam, Amsterdam, The Netherlands, E-mail: l.dorst@uva.nl

The test case is on terracotta, which does not deform; but typical fragments are lacking material through abrasion and chipping. Each fragment undergoes a preliminary preprocessing step [9] which partitions its surface into significant sub-parts called ‘*facets*’. Each facet is characterized by its own geometrical properties of the roughness of its surface and the sharpness of its boundary, which in turn guides its categorization as either belonging to the fracture region or outside skin region of the fragment.

Our contribution to the project focuses on structuring the pair-wise alignment of promising fragments nominated for fitting by other selection modules. The computational approach is based purely on the geometrical properties of the fracture facets, ignoring other clues (such as possible continuity of decorative patterns on the object facets).

Optimal pair-wise alignment of potentially counter-fitting fragments is computationally expensive. We therefore take a hierarchical approach, considering the fragments in increasing resolution from simplified to detailed. We could then explore the possible fit and approximate alignment at the coarse level, and refine this when descending to finer resolutions, all at approximately the same computational cost per level. But to be relevant to alignment, such a simplification should be contact preserving and sufficiently specific to structure and guide the pair-wise alignment.

The common way of simplifying shapes represented as meshes is by various forms of linear filtering of their vertices; letting them move with a differentiable flow (such as based on the Poisson equation) to produce a smoother model [16]. Structure-aware mesh decimation, to preserve the structure of the mesh by representing it as a set of pre-computed canonical proxies, has been studied by [22]. Such simplifications are strongly related to characterizing a surface by differential geometric features, such as those computed by some means of discrete differential geometry [7]. Indeed most methods to reassemble broken objects [12, 14, 18, 20, 26] are based on such features. However, a small chip, resulting in a missing part of an object, will affect such differential measures in a *linear* fashion: the deeper the missing bit, the stronger the simplified shape is affected. This is undesirable: the actual match between part and counterpart is simply lacking some local evidence of a more binary nature. The effect of this gap should not depend on the depth of the hole (a gap with the same surface outline but twice as deep is not twice as bad); so our representation method should not be too sensitive to this either. For the same reason, applying ICP (Iterative Closest Point) algorithm to find a matching counterpart based on squared errors of corresponding points is not appropriate; some extensions of ICP that include distance transforms [11] are more in line of what we might employ.

In our hierarchical approach, we need a representation where a coarser, simplified version of the fracture surface maintains its potential fit with a coarser, simplified version of the counterpart, under abrasion-type noise. Mathematical morphology with its opening and closing scale spaces, of increasing size, is the natural choice for just that kind of representation. Contrary to linear methods, it is not affected by the depth of missing parts, but merely by their spatial scale. Such scale spaces have been studied in [2, 3, 15], and it has been demonstrated that they simplify the shape, when measured in terms of certain descriptive features such as local maxima [15], or zero-crossings of the curvature [6]. But our intended use for alignment and complementarity checking appears to be new.

In this paper, we therefore hone standard Mathematical Morphology (MM) to this novel application. First, we define scale-based morphology of complementary objects (rather than of one object and its background). Then, we define morphological operations working only on the fracture region of the fragments. We implement this ‘Boundary Morphology’ by taking advantage of the Lipschitz nature of the individual fracture facets. Our new extrusion method uses a single distance transform to simultaneously produce both scaled dilations and scaled erosions of a fracture facet (as in Figure 1), leading to a simplified representation at every scale. We briefly discuss how the fracture Lipschitz property also naturally gives quantitative bounds on the detectable complementarity and collision-free alignment at each scale of our representation, even in the presence of abrasion. This paper does not perform the actual fitting, but Section 6 outlines how the proposed morphological representation of the fracture surfaces can be employed in optimal pair-wise alignment.

This paper extends and supersedes our work published in [10]. We now give full technical detail in the form of lemmas and proofs of the validity of complementarity properties within a masked region in Section 3.1 and 3.2, preparing this by precise notation in Section 2. Section 6 is also new contribution relative to [10].

2 Basics of Mathematical Morphology

Mathematical morphology (MM) was developed in the 1960s as a non-linear framework to describe shape in its qualitative (morphological) aspects, with structuring elements acting as probes to describe the shape hierarchically at different scales. MM flourished mostly in image (pre)processing, in which its set of non-linear operations produces results that cannot be obtained by linear, convolutional filtering. This effect-oriented use of MM has obscured the fact that it is actually a powerful language for describing shapes in contact [8]. In this paper, we employ it as such.

Since it is always clear which structuring element we use, we adopt the following notational shorthand. Let X be a connected set in $\mathbb{R}^3$ and S be our structuring element. Then:

$$\text{Dilation: } \delta(X) = X \oplus S = \{x + s \mid \forall x \in X, \forall s \in S\}$$

$$\text{Erosion: } \varepsilon(X) = X \ominus S = \{x \mid \forall s \in S : x + s \in X\}$$

$$\text{Closing: } \phi(X) = \varepsilon\delta(X) = (X \oplus S) \ominus S$$

$$\text{Opening: } \gamma(X) = \delta\varepsilon(X) = (X \ominus S) \oplus S$$

$$\text{Reflection: } \check{S} = \{s \mid -s \in S\}$$

$$\text{Reflected dilation: } \check{\delta}(X) = X \oplus \check{S}$$

$$\text{Reflected erosion: } \check{\varepsilon}(X) = X \ominus \check{S}$$

$$\text{Reflected closing: } \check{\phi}(X) = \check{\varepsilon}\check{\delta}(X) = (X \oplus \check{S}) \ominus \check{S}$$

$$\text{Reflected opening: } \check{\gamma}(X) = \check{\delta}\check{\varepsilon}(X) = (X \ominus \check{S}) \oplus \check{S}$$

With this notation, the underlying principles of MM [23] that we will employ can be summarized as follows:

(B₁) **Duality:** Erosion (opening) and dilation (closing) are *dual* operations relative to complementation (denoted $X \rightarrow X^c$) of the sets they act on. This means that they satisfy:

$$\varepsilon(X)^c = \check{\delta}(X^c) \text{ and } \delta(X)^c = \check{\varepsilon}(X^c)$$

$$\gamma(X)^c = \check{\phi}(X^c) \text{ and } \phi(X)^c = \check{\gamma}(X^c)$$

(B₂) **Distributivity:** The dilation is distributive over union and the erosion is distributive over intersection:

$$\delta(X \cup Y) = \delta(X) \cup \delta(Y)$$

$$\varepsilon(X \cap Y) = \varepsilon(X) \cap \varepsilon(Y)$$

but the converse statements are weaker:

$$\delta(X \cap Y) \subseteq \delta(X) \cap \delta(Y)$$

$$\varepsilon(X \cup Y) \supseteq \varepsilon(X) \cup \varepsilon(Y)$$

(B₃) **Local Knowledge Principle :** Morphological operations satisfy the "local knowledge principle". Shapes are usually seen through a window (e.g, image frame), within which MM is restricted:

$$\forall \text{ bounded } M, \exists \text{ bounded } M^* : [\Psi(X \cap M)] \cap M^* = \Psi(X) \cap M^*$$

In words, for any bounded window M in which we want to define a morphological transformation $\Psi(X)$, there exists a bounded window M^* , in which the knowledge of X is sufficient to locally perform the transformation Ψ .

3 Complementarity Preserving Scale Space

Consider a fracture between two fragments X and Y in perfect alignment (or, if you prefer, consider a newly developed crack in an object). Within a well-chosen mask M (such as the part of space the object volume occupies), the two fragments X and Y , viewed as sets, are almost complementary in the usual MM sense, with the fracture area being ambiguous (does it belong to X or Y ?). In an implementation such an infinitely thin layer is not going to make a difference, so we arbitrarily choose X to be a closed set and Y to be open. Then the M -restricted parts of X and Y , $X_M (= X \cap M)$ and $Y_M (= Y \cap M)$ are two complementary sets within M ($X_M^c = Y_M$). The properties of *exact complementarity* (no overlap, no gaps) can then be algebraically formulated as:

$$\text{No Overlap} \iff \text{Non-Intersection: } X_M \cap Y_M = \emptyset \quad (1)$$

$$\text{No Gaps} \iff \text{Completeness: } X_M \cup Y_M = M \quad (2)$$

In our work, one does not acquire the fragments in alignment, or in perfect condition, or even together with their counterpart. We choose MM scale spaces to process the fracture region of each of the fragments separately, producing a simplified representation for each fracture facet of each object at a range of scales ρ . Morphological scale spaces are well suited to this, since they can simplify objects while closely maintaining local complementarity, even when objects are slightly damaged (as we will discuss in Section 5). Performing erosions or dilations would change the objects considerably, so we prefer to use *openings* and *closings* to process the fragments (even though much essential structure is already contained in eroded and dilated versions).

The complementarity of two exactly fitting objects (within a Mask M) and their morphological opening (γ) and closing (ϕ) may thus be summarized with the following schema:

$$\begin{aligned} \check{\gamma}(X_M) &\subseteq X_M \subseteq \phi(X_M) \\ \Downarrow c_{M^*} &\quad \Downarrow c_{M^*} \quad \Downarrow c_{M^*} \\ \phi(Y_M) &\supseteq Y_M \supseteq \check{\gamma}(Y_M) \end{aligned} \quad (3)$$

Here c_{M^*} refers to complementarity within a more restricted mask M^* containing the common fracture of X and Y . This ‘well-chosen’ actually hides some essential details, since masking does not commute with morphology (e.g., $\phi_\rho(X_M) \neq (\phi_\rho(X))_M$). We refine this procedural sketch in the following section.

3.1 Masked MM

Before we move to the scale space complementarity, we need intermediate results on the effects of masking morphological operations when applied to one object only (independent of its complementary counterpart). The setup will be that within a mask M we have only $X \cap M$ available. We are interested in the largest possible region M^* in which the results of our operations are identical to these operations applied to the unmasked X . So our results will have the form:

$$\alpha(X \cap M) \cap M^* = \alpha(X) \cap M^* \quad (4)$$

for each operation α , with the post-masking set M^* to be determined. Therefore the lemmas will all be specific instances of the principle of local knowledge (B_3).

For many specific examples of X and M , the masking by M does not really interact with operations on X . For instance, if X is a point at the origin, and M is a large set enclosing the origin, then small dilations will satisfy $\delta(X \cap M) = \delta(X)$. In this case $M^* = M$ would suffice. However, if just outside M there are elements of the set X , these will dilate into the mask M , and the identity then only holds in a shrunken version of M ; in fact, optimally shrunken to $\check{\epsilon}(M)$, as we will show below in Lemma 3.2. In our typical use of masking, M will select a part of X (it is meant to focus on a fracture facet of an artifact), so the precise delimiting of the validity matters. We are therefore interested in lemmas that hold universally, for all X and M . To formulate this more precisely, let us introduce the weak notion of ‘necessity’ required to formulate the lemmas.

We will call a parametrized set a ‘*generally necessary lower bound*’ for the validity of a theorem involving sets X and M , if for all X and all M , there is no smaller set for which the theorem holds. In order to demonstrate that a particular set is a generally necessary lower bound, it suffices to show a specific choice of X and M in which masking of the result by a larger set would make the equality invalid.

Lemma 3.1. *For masked erosion, the generally necessary lower bound is $\varepsilon(M)$, i.e. $\varepsilon(X \cap M) \cap \varepsilon(M) = \varepsilon(X) \cap \varepsilon(M)$, for all X and M .*

Proof. Sufficiency is straightforward from the commutation and distributivity properties of the operations of erosion and intersection:

$$\varepsilon(X) \cap \varepsilon(M) = \varepsilon(X) \cap \varepsilon(M) \cap \varepsilon(M) = \varepsilon(X \cap M) \cap \varepsilon(M).$$

We give a critical example to demonstrate the general necessity in 1D. Consider a mask $M = (-\infty, 0)$, the half line left of zero, and an object $X = (-\infty, 0]$, an interval that includes the point at zero. Picking as our structuring element the asymmetric 1-sphere $B = [-\rho, 0]$, we find that $\varepsilon(X)$ equals $(-\infty, -\rho]$, so it includes the point at $-\rho$. Yet $\varepsilon(X \cap M)$ equals $(-\infty, -\rho)$, and therefore does not include this point. We can precisely eliminate this point by limiting the post-masking of the results by $M^* = \varepsilon(M)$, and no more severe limitation than this is needed.

Our counterexample to prove the general necessity of the lower bound was taken from 1D, for clarity. This is more universal than it may seem. It has been shown that the boundary points that can possibly be present in the result of erosions (or dilations) emanate along straight lines from the original set boundaries, as the structuring element scales (see [27], Section IV). In morphology in n dimensions, there will always be at least one critical 1-dimensional propagation direction that determines the lower bound of validity, through the effect most purely exposed in this 1D counterexample. ■

Lemma 3.2. *For masked dilation, the generally necessary lower bound is $\check{\varepsilon}(M)$, i.e. $\delta(X \cap M) \cap \check{\varepsilon}(M) = \delta(X) \cap \check{\varepsilon}(M)$, for all X and M .*

Proof. We first show sufficiency of the bound:

$$\begin{aligned} \delta(X) \cap \check{\varepsilon}(M) &= \delta((X \cap M) \cup (X \cap M^c)) \cap \check{\varepsilon}(M) && (\text{by } X = (X \cap M) \cup (X \cap M^c)) \\ &= (\delta(X \cap M) \cup \delta(X \cap M^c)) \cap \check{\varepsilon}(M) \\ &= (\delta(X \cap M) \cap \check{\varepsilon}(M)) \cup (\delta(X \cap M^c) \cap \check{\varepsilon}(M)) \end{aligned}$$

The second term is empty by duality between dilation and erosion:

$$\begin{aligned} \delta(X \cap M^c) \cap \check{\varepsilon}(M) &\subseteq (\delta(X) \cap \delta(M^c)) \cap \check{\varepsilon}(M) \\ &= \delta(X) \cap (\check{\varepsilon}(M))^c \cap \check{\varepsilon}(M) \\ &= \delta(X) \cap \emptyset \\ &= \emptyset \end{aligned}$$

So it follows that

$$\delta(X \cap M) \cap \check{\varepsilon}(M) = \delta(X) \cap \check{\varepsilon}(M)$$

The general necessity of the masking by $\check{\varepsilon}(M)$ can be established by giving an example in which no larger mask (i.e., smaller eroded mask) would work to preserve the identity. In 1D, consider a mask $M = (-\infty, 0)$, the half line left of zero and an object $X = \{0\}$ (the point at the origin). Clearly $X \cap M = \emptyset$, so to preserve the identity $\delta(X \cap M) \cap M^* = \delta(X) \cap M^*$, the right hand side should be empty as well. Picking as our structuring element the asymmetric 1-sphere $B = [-\rho, 0]$, we find that $\delta(X)$ extends to the point $-\rho$ into M ; therefore M^* needs to extend at least as far, and this is indeed what the choice $M^* = \check{\varepsilon}(M)$ minimally does. This argument can be extended to n dimensions by considering any 1-D direction in those spaces. ■

Lemma 3.3. *The universal equivalence of closing before and after masking by M is limited by the generally necessary lower bound $\varepsilon\check{\varepsilon}(M)$, i.e. $\phi(X \cap M) \cap \varepsilon\check{\varepsilon}(M) = \phi(X) \cap \varepsilon\check{\varepsilon}(M)$ for all X and M .*

Proof. Sufficiency:

$$\begin{aligned}\phi(X) \cap \varepsilon\check{\varepsilon}(M) &= \varepsilon\delta(X) \cap \varepsilon\check{\varepsilon}(M) \\ &= \varepsilon(\delta(X) \cap \check{\varepsilon}(M)) \\ &= \varepsilon(\delta(X \cap M) \cap \check{\varepsilon}(M)) && \text{(by Lemma 3.2)} \\ &= \varepsilon\delta(X \cap M) \cap \varepsilon\check{\varepsilon}(M) \\ &= \phi(X \cap M) \cap \varepsilon\check{\varepsilon}(M)\end{aligned}$$

Necessity of this post-masking lower bound for general X and M is demonstrated by the following 1D example on the real line. Let $M = (-\infty, 0)$ (the half line), and $X = \{0\} \cup (-\infty, -2\rho)$. Then with $B = [-\rho, \rho]$ (a closed interval), the closing of the two components of X equals the interval $\phi(X) = (-\infty, 0]$, but $\phi(X \cap M) = (-\infty, -2\rho)$. The largest masking in which the non-empty parts of these sets coincide is by the open interval $(-\infty, -2\rho)$, which is indeed $\varepsilon\check{\varepsilon}(M)$. ■

Lemma 3.4. *The universal equivalence of opening before and after masking by M is limited by the generally necessary lower bound $\check{\varepsilon}\varepsilon(M)$, i.e. $\gamma(X \cap M) \cap \check{\varepsilon}\varepsilon(M) = \gamma(X) \cap \check{\varepsilon}\varepsilon(M)$ for all X and M .*

Proof. Sufficiency:

$$\begin{aligned}\gamma(X) \cap \check{\varepsilon}\varepsilon(M) &= \delta\varepsilon(X) \cap \check{\varepsilon}\varepsilon(M) \\ &= \delta(\varepsilon(X)) \cap \check{\varepsilon}(\varepsilon(M)) \\ &= \delta(\varepsilon(X) \cap \varepsilon(M)) \cap \check{\varepsilon}(\varepsilon(M)) \\ &\quad \text{(by applying Lemma 3.2 on } \varepsilon(X) \text{ and } \varepsilon(M) \text{ instead of } X \text{ and } M \text{ respectively)} \\ &= \delta(\varepsilon(X \cap M)) \cap \check{\varepsilon}\varepsilon(M) \\ &= \gamma(X \cap M) \cap \check{\varepsilon}\varepsilon(M)\end{aligned}$$

General necessity for the opening follows the same principle as for closing, but now with $X = \{0, -2\rho\}^c$ (a set of two point holes), for $M = (-\infty, 0)$ and $B = [-\rho, \rho]$. ■

Actually, the post-masking validity regions of opening and closing are identical, since we have:

Lemma 3.5. $\varepsilon\check{\varepsilon} = \check{\varepsilon}\varepsilon$

Proof.

$$\begin{aligned}\check{\varepsilon}\varepsilon(A) &= (A \ominus S) \ominus \check{S} \\ &= A \ominus (S \oplus \check{S}) \\ &= A \ominus (\check{S} \oplus S) \\ &= (A \ominus \check{S}) \ominus S \\ &= \varepsilon\check{\varepsilon}(A)\end{aligned}$$

For notation convenience, we will denote both $\varepsilon\check{\varepsilon}$ and $\check{\varepsilon}\varepsilon$ by ε^2 , which better shows their equality. In fact the resulting structuring element has interesting properties:

Lemma 3.6. *The structuring element $S \oplus \check{S}$ involved in ε^2 is symmetrical and contains the origin.*

Proof. $S \oplus \check{S} = \{x = b_1 - b_2 \mid b_1, b_2 \in S\}$ so choosing $b_1 = b_2$ gives $x = 0$, and swapping b_1 and b_2 shows that $-x$ also belongs when x does. ■

This implies that the post-masking region for opening and closing only shrinks the sets symmetrically but does not displace them, even for general S .

3.2 Masked Complementarity

In this section we prove that when originally there was exact complementarity between X and Y within a mask M (i.e., $X_M^c = Y_M$), then exact complementarity still holds between $\phi(X)$ and $\check{\gamma}(Y)$, or between $\gamma(X)$ and $\check{\phi}(Y)$, but is only guaranteed within a doubly eroded mask $M^* = \varepsilon^2(M)$. This is important to establish before we move into MM scale space, where the operations will be performed with scaled version of the structuring element S . It shows precisely how the validity region depends on S and its scale.

Lemma 3.7. *For erosion and dilation, the non-intersection property (equation 1) holds within the entire M , i.e. $\delta(X_M) \cap \check{\varepsilon}(Y_M) = \emptyset$ and $\varepsilon(X_M) \cap \check{\delta}(Y_M) = \emptyset$.*

Proof.

$$\begin{aligned}\delta(X_M) \cap \check{\varepsilon}(Y_M) &= \delta(X \cap M) \cap \check{\varepsilon}(X^c \cap M) \\ &= \delta(X \cap M) \cap \check{\varepsilon}(M) \cap \check{\varepsilon}(X^c) \\ &= \delta(X) \cap \check{\varepsilon}(M) \cap \check{\varepsilon}(X^c) && \text{(by Lemma 3.2)} \\ &= \check{\varepsilon}(M) \cap (\delta(X) \cap (\delta(X))^c) \\ &= \check{\varepsilon}(M) \cap \emptyset \\ &= \emptyset\end{aligned}$$

Similarly,

$$\varepsilon(X_M) \cap \check{\delta}(Y_M) = \emptyset$$

■

Lemma 3.8. *For opening and closing, the non-intersection property (equation 1) holds within the entire M , i.e. $\phi(X_M) \cap \check{\gamma}(Y_M) = \emptyset$ and $\gamma(X_M) \cap \check{\phi}(Y_M) = \emptyset$.*

Proof.

$$\begin{aligned}\phi(X_M) \cap \check{\gamma}(Y_M) &= \phi(X \cap M) \cap \check{\gamma}(X^c \cap M) \\ &= \varepsilon(\delta(X \cap M)) \cap \check{\delta}(\check{\varepsilon}(X^c \cap M)) \\ &\subseteq \varepsilon(\delta(X) \cap \delta(M)) \cap \check{\delta}(\check{\varepsilon}(X^c) \cap \check{\varepsilon}(M)) \\ &= \phi(X) \cap \phi(M) \cap \check{\gamma}(X^c) \cap \check{\gamma}(M) \\ &= (\phi(M) \cap \check{\gamma}(M)) \cap (\phi(X) \cap \check{\gamma}(X^c)) \\ &= (\phi(M) \cap \check{\gamma}(M)) \cap (\phi(X) \cap (\phi(X))^c) \\ &= (\phi(M) \cap \check{\gamma}(M)) \cap \emptyset \\ &= \emptyset\end{aligned}$$

Similarly,

$$\gamma(X_M) \cap \check{\phi}(Y_M) = \emptyset$$

■

In contrast to this universal validity of non-intersection within M , the completeness property needs to be restricted, as shown by the following lemmas.

Lemma 3.9. *For dilation and erosion, the completeness property (equation 2) holds only within $\check{\varepsilon}(M)$, i.e. $[\delta(X_M) \cup \check{\varepsilon}(Y_M)]_{\check{\varepsilon}(M)} = \check{\varepsilon}(M)$ and $[\varepsilon(X_M) \cup \check{\delta}(Y_M)]_{\check{\varepsilon}(M)} = \check{\varepsilon}(M)$ and no larger area provides general validity.*

Proof. Sufficiency of limiting by $\check{\varepsilon}(M)$:

$$\begin{aligned}
 [\delta(X_M) \cup \check{\varepsilon}(Y_M)]_{\check{\varepsilon}(M)} &= [\delta(X_M) \cup \check{\varepsilon}(Y_M)] \cap \check{\varepsilon}(M) \\
 &= [\delta(X \cap M) \cup \check{\varepsilon}(X^c \cap M)] \cap \check{\varepsilon}(M) \\
 &= [\delta(X \cap M) \cap \check{\varepsilon}(M)] \cup [\check{\varepsilon}(X^c \cap M) \cap \check{\varepsilon}(M)] \\
 &= [\delta(X) \cap \check{\varepsilon}(M)] \cup [\check{\varepsilon}(X^c) \cap \check{\varepsilon}(M)] \quad (\text{by Lemmas 3.2 and 3.1}) \\
 &= [\delta(X) \cup \check{\varepsilon}(X^c)] \cap \check{\varepsilon}(M) \\
 &= [\delta(X) \cup (\delta(X))^c] \cap \check{\varepsilon}(M) \\
 &= \check{\varepsilon}(M)
 \end{aligned}$$

Similarly,

$$[\varepsilon(X_M) \cup \check{\delta}(Y_M)]_{\check{\varepsilon}(M)} = \check{\varepsilon}(M)$$

Necessity in the sense of ‘*generally necessary lower bound*’ follows from the general necessity of the erosion and dilation masking of Lemmas 3.1 and 3.2. ■

Lemma 3.10. *For closing and opening, the completeness property (equation 2) holds within a further shrunk area $\varepsilon^2(M)$, i.e. $[\phi(X_M) \cup \check{\gamma}(Y_M)]_{\varepsilon^2(M)} = \varepsilon^2(M)$ and $[\gamma(X_M) \cup \check{\phi}(Y_M)]_{\varepsilon^2(M)} = \varepsilon^2(M)$ and no larger area provides general validity.*

Proof. Sufficiency of limiting by $\varepsilon^2(M)$:

$$\begin{aligned}
 [\phi(X_M) \cup \check{\gamma}(Y_M)]_{\varepsilon^2(M)} &= [\phi(X_M) \cup \check{\gamma}(Y_M)] \cap \varepsilon^2(M) \\
 &= [\phi(X \cap M) \cup \check{\gamma}(X^c \cap M)] \cap \varepsilon^2(M) \\
 &= [\phi(X \cap M) \cap \varepsilon^2(M)] \cup [\check{\gamma}(X^c \cap M) \cap \varepsilon^2(M)] \\
 &= [\phi(X) \cap \varepsilon^2(M)] \cup [\check{\gamma}(X^c) \cap \varepsilon^2(M)] \quad (\text{by Lemmas 3.3 and 3.4}) \\
 &= [\phi(X) \cup \check{\gamma}(X^c)] \cap \varepsilon^2(M) \\
 &= [\phi(X) \cup (\phi(X))^c] \cap \varepsilon^2(M) \\
 &= \varepsilon^2(M)
 \end{aligned}$$

Similarly,

$$[\gamma(X_M) \cup \check{\phi}(Y_M)]_{\varepsilon^2(M)} = \varepsilon^2(M)$$

Necessity in the sense of ‘*generally necessary lower bound*’ follows from the general necessity of the opening and closing masking of Lemmas 3.3 and 3.4. ■

3.3 Ball Scale Spaces for Archaeology

We will not need MM in its full generality, which would allow one to choose a structuring element for all basic operations to affect the directions and extent of dilation and erosion (and hence of closing and opening). In our archaeological context, we assume that the objects are placed in arbitrary orientation, and that abrasion (by wear or wind) has no preferential direction. As a consequence, we should perform all our morphological operations with structuring elements that are balls (filled spheres), of a specified radius ρ as the effective scale. We now refine our notation to be specific on the radius and use ρ as a subscript to our δ as in δ_ρ . The balls we use are convex, symmetric and contain the origin; this simplifies our treatment since we no longer need to distinguish between S and $\check{S}$, and consequently ε and $\check{\varepsilon}$ etc.

At a scale ρ within a suitably chosen mask $\varepsilon_\rho^2(M)$, the opened version $\gamma_\rho(X)$ will be complementary to the closed version $\phi_\rho(Y)$, and vice versa. The closing $\phi_\rho(X)$ simplifies the local geometry of the fracture X at its

valleys, and the opening $\gamma_\rho(Y)$ does the same to Y at its peaks. Therefore, the complementarity of $\phi_\rho(X)$ with $\gamma_\rho(Y)$ becomes less specific in those areas, and vice versa. To maintain specificity of both peaks and valleys of the common fracture surface, we must therefore *compute both opening $\gamma_\rho(X)$ and closing $\phi_\rho(X)$ for each fragment X* . Since both opening and closing are increasing, coarser levels of the scale space are guaranteed to contain less detail, thus enabling the hierarchical approach.

For different scales R and ρ with $R > \rho$, the complementarity and containment relationships illustrated by schema 3 can now be extended as follows:

$$\begin{aligned} \gamma_R(X_M) \subseteq \gamma_\rho(X_M) \subseteq X_M \subseteq \phi_\rho(X_M) \subseteq \phi_R(X_M) \\ \Downarrow c_R \quad \Downarrow c_\rho \quad \Downarrow c \quad \Downarrow c_\rho \quad \Downarrow c_R \\ \phi_R(Y_M) \supseteq \phi_\rho(Y_M) \supseteq Y_M \supseteq \gamma_\rho(Y_M) \supseteq \gamma_R(Y_M) \end{aligned} \quad (5)$$

Here c_ρ is a shorthand for $c_{\varepsilon_\rho^2(M)}$ and $\varepsilon_\rho^2(M) = \varepsilon_\rho(\check{\varepsilon}_\rho(M)) = \check{\varepsilon}_\rho(\varepsilon_\rho(M))$, which is equivalent to erosion of M by a ball of twice the radius ($\varepsilon_\rho^2 = \varepsilon_{2\rho}$).

4 Fracture Morphology

Figure 2: Morphological opening of a volumetric broken object at 4 different scales.

Consider a broken vase producing thin sherds as fragments. Volumetric morphological simplification of the fragments would be limited in scale by their thickness: an opening by a ball with a radius ρ of more than half this thickness would result in empty sets, (see Figure 2), which would only be complementary to their counterpart in a trivial way. Since the fractures themselves have informative morphological structure transcending this scale, we design in this section a way to focus the morphological processing purely on the fracture surface (rather than on the fragment volume). This should suffice for our puzzle, since whether two broken fragments can be refit locally depends on the complementary shape of their fractures only.

4.1 Boundary Morphology

In Section 3, we had masked objects X_M and Y_M , with exact complementarity within a chosen mask M (and a slight issue of how we treat their common boundary fracture F). Now consider only one of them, no longer in contact, and investigate how to compute its openings and closings. We are only interested in the effect on F , not on the remainder of the object volume within M .

F is a boundary, not a volume, and the classical mathematical morphology does not apply immediately. We first define what we mean by applying mathematical morphology to a boundary. Let us call it ‘*Boundary Morphology*’ and denote its operations by an over-bar. Consider a general object A of which we take the boundary ∂A (and ignore the masking effects for the moment). As we have seen, for true mathematical set complementarity some objects may be open sets, others closed sets. We need to take the boundary of either, so we will always close the set and then apply the usual boundary operator ∂ to define our boundary operator (which we denote by $\bar{\partial}$):

$$\bar{\partial} A = \partial \bar{A} \quad (6)$$

Thus $\bar{\partial} A$ returns the set of points of $\bar{A}$ that have neighbours in $\bar{A}^c$.

In order to distinguish erosion and dilation, we need to make the boundary oriented, for instance by denoting the outward pointing normal at each location. We assume that the result of the boundary operator ∂ includes this allocation of normal directions (and hence so does $\bar{\partial}$). This orientation should change to its opposite when we consider the boundary of the complement. To more easily denote this, we write the complement as the prefix operator c . We then have:

$$\bar{c} \bar{\partial} A \equiv \bar{\partial} c A (= \bar{\partial} A^c), \quad (7)$$

which defines the complementation $\bar{c}$ of the boundary.

We define what it means to dilate and erode a boundary, defining operators $\bar{\delta}$ and $\bar{\varepsilon}$ in terms of classical volumetric operators by:

$$\bar{\delta} \bar{\partial} A \equiv \bar{\partial} \delta A \quad \text{and} \quad \bar{\varepsilon} \bar{\partial} A \equiv \bar{\partial} \varepsilon A. \quad (8)$$

It now follows that we can perform a reflected boundary erosion as a boundary dilation on the complemented (oppositely oriented) boundary:

$$\bar{\delta} \bar{\partial} A = \bar{\partial} \delta A = \bar{\partial} c \check{\varepsilon} c A = \bar{c} \bar{\partial} \check{\varepsilon} c A = \bar{c} \check{\varepsilon} \bar{\partial} c A = \bar{c} \check{\varepsilon} \bar{c} \bar{\partial} A, \quad (9)$$

which implies that

$$\check{\varepsilon} = \bar{c} \bar{\delta} \bar{c} \quad \text{and} \quad \check{\delta} = \bar{c} \bar{\varepsilon} \bar{c}. \quad (10)$$

In fact, this is merely the duality between dilation and erosion, extended to their boundary versions (with $\bar{c}$ playing the role of complementation). So once we can dilate a boundary, we can use that capability to produce both the dilation $\bar{\delta}$ and the reflected erosion $\check{\varepsilon}$ (by doing dilation on $\bar{c} \bar{\partial} A$, the oppositely oriented boundary, and complementing the result):

This can be easily extended to boundary closing $\bar{\phi}$ and opening $\bar{\gamma}$ defined as:

$$\bar{\phi} \bar{\partial} A \equiv \bar{\varepsilon} \bar{\delta} \bar{\partial} A \quad \text{and} \quad \bar{\gamma} \bar{\partial} A \equiv \bar{\delta} \bar{\varepsilon} \bar{\partial} A, \quad (11)$$

for it follows from eqs. (9) and (10) that we can also perform a reflected boundary opening as a boundary closing on the complemented boundary:

$$\bar{\phi} \bar{\partial} A = \bar{\varepsilon} \bar{\delta} \bar{\partial} A = \bar{c} \check{\delta} \bar{c} \bar{\delta} \bar{\partial} A = \bar{c} \check{\delta} \check{\varepsilon} \bar{c} \bar{\partial} A = \bar{c} \check{\gamma} \bar{c} \bar{\partial} A, \quad (12)$$

which implies that

$$\check{\gamma} = \bar{c} \bar{\phi} \bar{c} \quad \text{and} \quad \check{\phi} = \bar{c} \bar{\gamma} \bar{c}. \quad (13)$$

For symmetric structuring elements like balls, we can ignore the checkmarks within the above MM notations.

Now let us relate this to each of the objects X and Y within a mask M . The mask is chosen to focus on (part of) the common fracture F . With X and Y in perfect contact, the set F is simply $F = \bar{X} \cap \bar{Y} \cap M$ (the overbar denoting set closure). Considering only the object X in isolation, for the contact we focus on the part of its boundary masked by M . We turn this point set into an oriented boundary facet for X , borrowing the outward pointing normals of X at the chosen locations. This defines a facet F_X :

$$F_X \equiv \partial(\bar{X} \cap M). \quad (14)$$

Some care is needed: this definition produces only points of the desired F if M itself does not contribute to the boundary. Therefore, we have to take the mask M as an *open set*. Similarly, one defines $F_Y \equiv \partial(\bar{Y} \cap M)$.

One would then like to define the boundary dilation $\bar{\delta}$ through:

$$\bar{\delta}(\partial(\bar{X} \cap M)) = \partial(\delta(\bar{X}) \cap M), \quad (15)$$

or, denoting the masking by M as an operator,

$$\bar{\delta} \partial M = \partial M \delta, \quad (16)$$

which looks like the natural extension of the unmasked case of eq. (8). But after dilation, the masking by a general M could crop parts of the dilated boundary facet that should have been included, since they were

caused by points of the original boundary within the mask. This issue can be resolved simply, by constructing masks M for boundary MM only as the union of normal line segments emanating from the boundary facet F_X , since the differential equations of morphology transport points along straight lines in the direction of the normals [27].

With F_X defined by eq. (14), we have produced a proper oriented boundary facet, which can take the place of $\bar{\partial} A$ of eq. (6) in the derivations above. We can then define boundary morphology in this manner, with the complementarity relationships like $F_Y = \bar{c} F_X$ indicating that F_Y and F_X are the same set of points, but with opposite orientation. This permits us to translate the volumetric schema 5 to describe the complementarity of two exactly fitting fracture boundaries F_X and F_Y (at scales R and ρ with $R > \rho$), as follows:

$$\begin{aligned} \bar{\gamma}_R F_X &\subseteq \bar{\gamma}_\rho F_X \subseteq F_X \subseteq \bar{\phi}_\rho F_X \subseteq \bar{\phi}_R F_X \\ \Downarrow \bar{c}_R &\quad \Downarrow \bar{c}_\rho \quad \Downarrow \bar{c} \quad \Downarrow \bar{c}_\rho \quad \Downarrow \bar{c}_R \\ \bar{\phi}_R F_Y &\supseteq \bar{\phi}_\rho F_Y \supseteq F_Y \supseteq \bar{\gamma}_\rho F_Y \supseteq \bar{\gamma}_R F_Y \end{aligned} \quad (17)$$

Here $\bar{c}_\rho$ is the boundary complementation within a mask $\varepsilon_\rho^2(M)$.

4.2 The Lipschitz Condition

Figure 3: (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).

Terracotta is uniform and brittle. We assert that it fractures rather simply such that there is at least one 3D direction from which the whole fracture facet is completely visible, i.e. a ray in that direction emanating from each point of the fracture will not encounter any other point of that local fracture facet. Accordingly, a given fracture facet can be represented as a function (a Monge patch) within an orthogonal coordinate frame taking a visibility direction as function axis (while noting that another fracture facet of the same fragment may require a different coordinate frame to be seen as a function). By choosing our frames well, the visibility assumption of the surface can be related to the ‘Lipschitz condition’ of the corresponding function [24].

For any pair of points on the graph of a slope-limited Lipschitz function with slope s , the absolute value of the slope of the line connecting them is always less than s (i.e., the Lipschitz constant). As a consequence,

there exists a double cone whose vertex can be moved along any such Lipschitz-continuous function, so that it always remains entirely visible from the principal direction of the cone while the rest of the function is completely outside the cone. For a local fracture to meet the visibility condition, it should meet the Lipschitz condition over its entire domain.

To compute the maximum Lipschitz slope s , we first collect the local normal vectors of the fracture facet on a unit sphere of directions, and then fit the smallest possible cone that contains this set (see Figure 3). The axis of this cone we call the *cone direction*, and the opening angle we call the *cone angle*. Let us use the cone direction as the principal direction to align the fracture facet, so that it can be described as a Lipschitz function. The corresponding Lipschitz principal direction acts as the ‘average breakage direction’ for the whole fracture facet. The fracture function is then bounded in Lipschitz slope s , which is the cotangent of the cone angle.

In practice, the Lipschitz condition appears to intrinsically hold for each of our terracotta fracture facets with a Lipschitz slope of approximately 1. Even if it does not hold, we can always artificially separate the fracture regions by the *Faceting* preprocessing procedure [9] that delivers the piecewise Lipschitz facets of the fracture region.

4.3 Extrusion

Figure 4: Extrusion of a fracture facet $\bar{\partial} F$ along the Lipschitz principal direction, generating a volume (cylinder) with $\bar{\partial} F$ as its top red surface, and the oppositely oriented $\bar{c}\bar{\partial} F$ as its bottom aquamarine surface.

The fracture surfaces can thus be guaranteed to be locally Lipschitz. Therefore, we can generate a thickened fracture volume by extruding the facet along such visibility direction without self-intersection. Such an extrusion effect is equivalent to making two copies of the fracture surface bounded by a generalized cylinder: one with original (non-inverted) normals $\bar{\partial} F$ and one with inverted normals $\bar{c}\bar{\partial} F$ (see Figure 4). The outward propagation of the copy with non-inverted normals is the fracture surface dilation ($\bar{\delta}_\rho$), while the propagation of the one with inverted normals is the erosion ($\bar{\epsilon}_\rho$), by eq. (10). By a subsequent inward propagation of the resulting expanded surfaces, by the same amount ρ , one then acquires the closing ($\bar{\phi}_\rho$) and the opening ($\bar{\gamma}_\rho$) of the fracture.

Focusing our morphological simplification on the extruded fracture volume has the following advantages:

- Extending the functional morphological scale (ball radius ρ) beyond the minimum thickness of the archaeological object by only needing to perform the closing operation.
- Permitting simultaneous propagation of the fracture surface to both sides, thus allowing the production of MM opened and closed surfaces in one go.
- Avoiding the dilation of elements outside the mask (which could affect the outcome within the mask) by surrounding the extruded surface by sufficient empty space completely isolated from outside influences.
- Avoiding needless processing of the non-fracture regions of the object.

4.4 Distance-Transform Based MM Implementation

Figure 5: Distance field computation, two cross sections are shown.

We obtain our data as 3D meshes, so it would be natural to apply the operations of boundary morphology to fracture facets represented directly as extruded 3D meshes. However, applying morphological operations directly on a mesh representation is not straightforward (the dilation and erosion operations lead to many additional vertices as the mesh faces start intersecting [13], and the ball structuring elements add many more). At this point of our narrative, it would seem that Point Morphology [5] is also a good candidate; however we will see in Section 5 that we need the structure of the distance field to determine sensible morphological features for pair-wise alignment, rather than just the simplified surfaces by themselves.

An alternative implementation of morphology, especially when done in a scale space context, is by means of the isosurfaces of distance functions. Since we do want to perform our operations at different scales of a ball-shaped structuring element, this is especially attractive. Two basic approaches exist: distance transform considered as a numerical level set in a space of one more dimension [25] or classic distance transform on a 3D grid. We adopted the latter approach which: a) suffices as a straightforward demonstration of the validity of our concepts and b) enables us to keep track of the original fracture points propagation in the distance field and their contribution to the generated MM surfaces through what we call *the provenance map*. This map specifies the original boundary point(s) to which each point of the MM surface can be traced back. Our implementation relies on distance transforms in a volumetric representation [28], oriented with the Lipschitz principal direction along its z-direction. We convert the extruded mesh to binary representation by embedding the fracture surface in a 3D voxelized grid of a well chosen resolution, of step size g . The maximum outward and inward displacement that could take place due to discretization effects is no more than $\sqrt{3} g/2$, since each point is displaced by maximum half the grid size along each of its dimensions. This implies that distances (and hence MM surfaces) are affected by no more than this amount. The generated grid is padded with an empty region sufficient to contain the maximally dilated versions of the fracture surfaces. To generate our distance field, the method proposed by [17] is employed, which calculates the euclidean distance transform in linear time on a binary voxelized representation of the object as shown in Figure 5. The closing of the fracture volume with a ball of radius ρ (closing is all we need by virtue of eq. (13)) is performed by first extracting the outward level set at distance ρ , then re-computing the distance field of the background and

extracting the inward level set at distance ρ . This generates a maximum error of $\sqrt{3} g$ for the extracted surface relative to the true closing.

4.5 MM Surfaces

Figure 6: Closing and opening scale spaces simplifying a fracture with increasing scale ρ .

Intuitively, morphologically closing the entire extruded volume is equivalent to rolling a ball (our chosen structuring element) over the fracture surface from both sides. The upper rolling (constrained by the fracture peaks) is equivalent to the closing effect, while the lower rolling (constrained by the fracture valleys) is equivalent to the desired opening effect. At increasingly coarse scales, we compute the closed and the opened versions of the fracture facet. Moving up the scale space, the generated MM surfaces are simpler in shape than the original fracture they represent while maintaining the fitting (i.e., complementarity) essentials.

Figure 6 shows the closed and opened simplified MM surfaces at 6 different scales of a given fracture facet with 5.8K vertices and 11K faces. The grid resolution is $g = 0.2$ mm of size $119 \times 323 \times 29$ and is padded with 151 voxels from each side to permit the outward propagation of the fracture surface up to 30 mm without losing information, thus producing a grid of size $421 \times 625 \times 331$.

The fracture facet Lipschitz computations together with the extrusion and embedding take less than 2 seconds for a mesh of size 10K. The distance field computation time is $O(g^{-3})$: computing the distance transform for $500 \times 500 \times 500$ grid takes at most 1 minute. Extracting the closed and opened MM surfaces takes less than 50 seconds for the same grid size. In the GRAVITATE system, such times are considered acceptable since the morphological characterization of the fracture facets of each artifact is an offline process at ingestion time. We performed our experiments on an Intel(R) Xeon(R) CPU 2.9 GHz $\times$ 8 computer with 256 GB RAM.

5 Inexact Complementarity

In practice, there are many aspects to our data that prevent complementarity from being exact. Fortunately, using the Lipschitz condition for fracture surface characterization together with MM for hierarchical fracture

simplification enables us to establish the bounds on the amount of inexactness. In this section, we briefly discuss the sources causing inexactness and how we expect to treat their associated deviations.

- **Abrasion.** When a fractal surface as brittle as terracotta is rubbed, this will tend to take off the sharp local peaks without affecting the valleys. We propose that a reasonable model of such effect is that the surface undergoes a morphological opening $\bar{\gamma}_\alpha$ of a small size α . Assuming this, the opening scale space will be unaffected for $\rho > \alpha$, since $\bar{\gamma}_\rho(\bar{\gamma}_\alpha(X)) = \bar{\gamma}_\rho(X)$ for $\rho \geq \alpha$. However, the closing is affected by the abrasion. Therefore, exact complementarity of $\bar{\gamma}_\rho(X')$ and $\bar{\phi}_\rho(Y')$ of the schema in eq. (17) no longer holds for the abraded versions $X' = \bar{\gamma}_\alpha(X)$ and $Y' = \bar{\gamma}_\alpha(Y)$, even if it did for the original X and Y . Nevertheless, the difference is bounded for the Lipschitz functions we are considering as the fracture facets. A simple sketch of a local sphere of radius α capped by a Lipschitz cone of slope s shows that the maximal deviation of a possible original surface and its α -opening is $(\sqrt{1+s^2} - 1) \simeq \frac{1}{2}as^2$. Since X' does not differ by more than $\frac{1}{2}as^2$ from X , also $\bar{\phi}_\rho(X')$ will not differ from $\bar{\phi}_\rho(X)$ by more than this amount. With that in mind, we can still test the possibility of complementarity, to within this bound, of the opened and closed abraded fractures.
- **Fracture Adjacency Effects.** Each facet is delineated by an outer border (contour). That border defines the complementarity zone of its processed MM surfaces. In practice, there are two different kinds of facet borders: *fracture-skin* border and *fracture-fracture* border. The former is the contour's segment that is adjacent to a facet which belongs to the exterior skin of the original object. By contrast, a fracture-fracture border occurs when the original fragment is further broken into sub-fragments, causing the adjacency of more than one fracture facet. If the facet has only fracture-skin borders, then complementarity with its counterpart still holds over the entire surfaces of its simplified MM versions across all scales. For facets with fracture-fracture borders, one should expect complementarity to hold within regions away from the border. But there is then a scale-dependent zone adjacent to the fracture-fracture borders where the MM surfaces are no longer guaranteed to be complementary to those of its counterpart. In that zone, the representation of either part may have been affected by the further breaking. The bounds we can obtain on this zone under the general assumptions we made (such as Lipschitz) are weak. We have found that we can much more specifically delineate it through analysing the provenance map of the distance transform (defined in Section 4.4). Our current work is on efficient representation of the simplified scale space surfaces, taking this effect into account.
- **Misalignment.** Complementarity only exists in perfect alignment of the two counterparts; any offset in translation or orientation will destroy it. But this effect is not as boolean as eqs. (1) and (2) makes it appear: the complementarity is quantifiable by means of the MM scale spaces. The distance transform method of producing the openings and closings at different scales can be employed to guide the alignment, from coarser to finer scales, with the separation between one surface's opening and closing at each scale giving an indication of reasonable bounds on the fine alignment at that scale [11].

6 Towards Fitting

The classical way of fitting fracture surfaces is based on matching significant corresponding features characterizing the fracture surfaces of counter-fitting fragments [4, 14, 21]. These features are mostly detected and extracted using differential geometric techniques which are linearly sensitive to missing information, as has been clarified before in Section 1. Such features are not expected to be stable under abrasion and many of them are not compact enough to allow the cheap hierarchical fitting we are anticipating. In this section, we introduce our intended morphological features and how they can be structured in a compact scale space representation that can robustly and cheaply guide our next phase of optimal pairwise alignment.

The key observation is that the closing process results in a surface of which large parts are spherical caps (see Figure 6, especially the coarser resolutions, and Figure 7). Moreover, as the scale increases, the surface

Figure 7: Morphological characterization of fracture surfaces through 3D medial axis points. (a) Each medial axis point representing a part of the surface by a spherical cap at the chosen scale. (b) The 30 most relevant feature points and their spherical caps.

tends to be hierarchically composed of fewer, larger and more clearly separated spherical caps. The centers 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 centers are medial axis points of the complement of the extruded fracture mesh volume. Medial axes are notoriously sensitive to discretization, but they can be made robust [1]. 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 reconstruction (see Figure 7(a)).

Figure 7(b) is an exploration of this idea. It shows the 30 most distinctive MM feature points which are capable of representing 30% of the simplified (closed) fracture facet at scale 30 mm. The fracture facet is the same as in Figure 6, which is about 6×2.5 cm size. Moving up in the hierarchical scale space, these MM feature points are increasingly capable of describing more surface area of the overall simplified fracture surface. It seems natural to use each cap's area for weighting the significance of the corresponding feature point. We thus arrive at the representation of a fracture surface by a set of 3D weighted points, at well-defined positions above it (for closing) and below it (for opening), at a certain scale ρ . Such 'landmarks' may then be used to characterize the morphologically simplified versions of the fracture surface at different scales. We could further reduce the morphological characterization by only keeping the medial axis points at critical scales, as is usual in linear scale spaces. We would then consider replacing our linear time Euclidean distance transform (based on [17]) by the linear time medial axis chamfer approximation of [19].

According to our simplification framework, a morphological graph characterizing the *closing* scale space of a given fracture surface is expected to be very similar to the morphological graph characterizing the *opening* scale space of its counter-fitting fracture surface. Alignment of fracture surfaces using morphological graphs can be done by means of using the well-known Procrustes method. One should then use an extended version of this algorithm which allows for missing point correspondences. But before we perform our alignment in this way, we are currently studying the reproducibility, stability and robustness properties of the morphological feature points and the feasibility of computing them directly from the distance field.

7 Conclusion and Future Work

We have presented morphological simplification of the fracture facets of archaeological fragments based on their scanned 3D mesh representation, as a preparatory phase for optimal pair-wise alignment. This is a problem that lends itself very well to treatment by mathematical morphology, since that framework provides complementarity-preserving simplification of shapes in a manner that is insensitive to the kind of missing information and abrasion that we expect in archaeological fragments. We showed what the effect of masking is on the interpretability of MM results in Sections 3.1 and 3.2 and also how to perform morphology on boundaries in Section 4.1. The assumption that each fracture facet is a Lipschitz function allowed us to set up the extrusion method in Section 4.3. The duplicate internal and external copies of the local fracture facet surface enable computing the opening and closing scale spaces simultaneously, in a volumetric representation in Section 4.5. The resulting surfaces are simpler, but still contain the essentials of complementarity for identification and fitting of counterparts. In Section 5, we discussed the sources of inexact complementarity and how their effects are naturally bounded in our method.

We are currently investigating the compact characterization of the MM scale space surfaces by means of characteristic scale space medial axis points, computed directly from the distance transform and its provenance map (specifying which points are influential at each location) along the lines sketched in Section 6. This should allow the use of standard registration algorithms on those morphologically relevant feature points only. Moreover, we plan to use the provenance map of the distance transform to automatically control the reliability of those representations when parts of the original fracture facet are missing.

Acknowledgements: This research was funded by GRAVITATE project (<https://gravitate-project.eu/>) under EU2020-REFLECTIVE-7-2014 Research and Innovation Action, grant no. 665155. Datasets are from GRAVITATE data collection. The archaeological artifacts are digitized and kept by the British museum, Cyprus Institute and Archaeology laboratory at the Hebrew University of Jerusalem. We thank the reviewers for their helpful comments.

References

- [1] Attali, D., Boissonnat, J.D., Edelsbrunner, H.: Stability and computation of medial axes - a state-of-the-art report pp. 109–125 (2008)
- [2] van den Boomgaard, R., Dorst, L.: The morphological equivalent of Gaussian scale space. In: *Proceedings of Gaussian Scale-Space Theory*. pp. 203–220. Kluwer (1997)
- [3] Bosworth, J.H., Acton, S.T.: Morphological scale-space in image processing. *Digital Signal Processing* **13**(2), 338–367 (2003)
- [4] Brown, B., Laken, L., Dutré, P., Gool, L.V., Rusinkiewicz, S., Weyrich, T.: Tools for virtual reassembly of fresco fragments. *International Journal of Heritage in the Digital Era* **1**(2), 313–329 (2012)
- [5] Calderon, S., Boubekur, T.: Point morphology. *ACM Trans. Graph.* **33**(4), 1–13 (2014)
- [6] Chen, M.h., P.-F., Y.: A multiscaling approach based on morphological filtering. *IEEE Transactions on PAMI* **11**(7), 694–700 (1989)
- [7] Crane, K., de Goes, F., Desbrun, M., Schröder, P.: Digital geometry processing with discrete exterior calculus. In: *ACM SIGGRAPH 2013 courses*. SIGGRAPH '13, ACM, New York, NY, USA (2013)
- [8] Dorst, L., van den Boomgaard, R.: The systems theory of contact. In: Sommer, G., Zeevi, Y. (eds.) *Algebraic Frames for the Perception-Action Cycle*, Lecture Notes in Computer Science, vol. 1888, pp. 22–47. Springer Verlag (2000)
- [9] ElNaghy, H., Dorst, L.: Geometry based faceting of 3D digitized archaeological fragments. In: *2017 IEEE ICCV Workshops*. pp. 2934–2942 (2017)
- [10] ElNaghy, H., Dorst, L.: Complementarity-preserving fracture morphology for archaeological fragments. In: Burgeth, B., Kleefeld, A., Naegel, B., Passat, N., Perret, B. (eds.) *Mathematical Morphology and Its Applications to Signal and Image Processing*. pp. 403–414. Springer International Publishing, Cham (2019)
- [11] Fitzgibbon, A.: Robust registration of 2D and 3D point sets. *Image and Vision Computing* **21**, 1145–1153 (2003)
- [12] Funkhouser, T., Shin, H., Toler-Franklin, C., Castañeda, A.G., Brown, B., Dobkin, D., Rusinkiewicz, S., Weyrich, T.: Learning how to match fresco fragments. *J. Comput. Cult. Herit.* **4**(2), 1–13 (2011)

- [13] Hachenberger, P.: 3D Minkowski sum of polyhedra. In: CGAL User and Reference Manual. CGAL Editorial Board, 4.11.1 edn. (2018)
- [14] Huang, Q.X., Flöry, S., Gelfand, N., Hofer, M., Pottmann, H.: Reassembling fractured objects by geometric matching. In: ACM SIGGRAPH Papers. pp. 569–578. SIGGRAPH'06, ACM, New York, NY, USA (2006)
- [15] Jackway, P.T., Deriche, M.: Scale-space properties of the multiscale morphological dilation-erosion. *IEEE Transactions on PAMI* **18**(1), 38–51 (1996)
- [16] Kazhdan, M., Bolitho, M., Hoppe, H.: Poisson surface reconstruction. In: Proceedings of the Fourth Eurographics Symposium on Geometry Processing. pp. 61–70. SGP '06, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland (2006)
- [17] Maurer, C.R., Qi, R., Raghavan, V.: A linear time algorithm for computing exact Euclidean distance transforms of binary images in arbitrary dimensions. *IEEE Transactions on PAMI* **25**(2), 265–270 (2003)
- [18] McBride, J.C., Kimia, B.B.: Archaeological fragment reconstruction using curve-matching. In: IEEE CVPR Workshops. vol. 1, pp. 3–3 (2003)
- [19] Nacken, P.F.M.: Chamfer metrics, the medial axis and mathematical morphology. *Journal of Mathematical Imaging and Vision* **6**(2), 235–248 (Jun 1996)
- [20] Palmas, G., Pietroni, N., Cignoni, P., Scopigno, R.: A computer-assisted constraint-based system for assembling fragmented objects. In: 2013 Digital Heritage International Congress (DigitalHeritage). vol. 1, pp. 529–536 (2013)
- [21] Papaioannou, G., Schreck, T., Andreadis, A., Mavridis, P., Gregor, R., Sipiran, I., Vardis, K.: From reassembly to object completion: A complete systems pipeline. *J. Comput. Cult. Herit.* **10**(2), 1–22 (2017)
- [22] Salinas, D., Lafarge, F., Alliez, P.: Structure-aware mesh decimation. In: Computer Graphics Forum. vol. 34, pp. 211–227. The Eurographics Association (2015)
- [23] Serra, J.: Image Analysis and Mathematical Morphology. Academic Press, Inc., Orlando, FL, USA (1983)
- [24] Serra, J.C.: Lipschitz lattices and numerical morphology. *SPIE Conf. Image Algebra and Morphological Image Processing II* **1568**, 1568–12 (1991)
- [25] Sethian, J.: Fast marching methods and level set methods for propagating interfaces. *Computational Fluid Dynamics*, 29th Annual Lecture Series (1998)
- [26] Toler-Franklin, C., Brown, B., Weyrich, T., Funkhouser, T., Rusinkiewicz, S.: Multi-feature matching of fresco fragments. *ACM Trans. Graph.* **29**(6), 1–12 (2010)
- [27] van den Boomgaard, R., Smeulders, A.: The morphological structure of images: the differential equations of morphological scale-space. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **16**(11), 1101–1113 (Nov 1994)
- [28] Wu, Z., Song, S., Khosla, A., Yu, F., Zhang, L., Tang, X., Xiao, J.: 3D shapenets: A deep representation for volumetric shapes. In: IEEE CVPR. pp. 1912–1920 (2015)