Vision, Modeling, and Visualization (2020)
J. Krüger, M. Niessner, and J. Stückler (Eds.)

Segmenting Computer-Tomographic Scans of Ancient Clay Artefacts for Visual Analysis of Cuneiform Inscriptions

T. Rolff[1], M. Rautenhaus[2] S. Olbrich[2] and S. Frintrop[1]

[1]Department of Informatics, Universität Hamburg, Germany
[2]Regional Computing Center, Universität Hamburg, Germany

Abstract

We address the automatic segmentation of computer tomographic scans of ancient clay tablets with cuneiform inscriptions enclosed inside a clay envelope. Such separation of parts of similar material properties in the scan enables domain scientists to virtually investigate the historically valuable artefacts by means of 3D visualization without physical destruction. We investigate two segmentation methods, the Priority-Flood algorithm and the Compact Watershed algorithm, the latter being modified by employing a distance metric that takes the ellipsoidal shape of the artefacts into account. Additionally, we propose a novel pre-segmentation method that suppresses the intensity values of the distance transform at contact points between clay envelope and tablet. We apply all methods to volumetric scans of a replicated clay tablet and analyze their performance under varying noise distributions. Evaluation by comparison to a manually segmented ground truth shows best results for the novel suppression-based approach.

CCS Concepts

• ***Human-centered computing*** → *Scientific visualization;* • ***Computing methodologies*** → *Image processing; Volumetric models;*

1 Introduction

In Kültepe, Turkey, approximately 22500 archaeological clay tablets mainly dating back to the nineteenth century BC have been discovered, with more tablets found at other sites [Mic08; Mic16]. Cuneiform inscriptions on these tablets range from legal documents, including contracts, to private letters. They hence provide valuable insight into the customs and rites at the time of their creation. As shown in Fig. 1, many of these tablets are still sealed within a protective envelope which is most likely manufactured from the same material [TC12]. The tablets were handcrafted, and a variety of different shapes in various conditions has been discovered. While some tablets are completely sealed in their envelopes, for others the envelope has been lost or has been accidentally shattered. Both the clay tablets and the envelopes feature inscriptions and illustrations. Therefore, both are of historical value and it is no option to destroy the envelope to gain access to the tablet content. A viable alternative solution is to perform computer tomographic (CT) scans of the discovered artefacts and to analyze the inscriptions using interactive 3D visualization. This raises the need for a robust segmentation algorithm capable of separating the clay tablet from its envelope. In this study, we address the challenge of performing such segmentation on volumetric data that was acquired using a CT scanner.

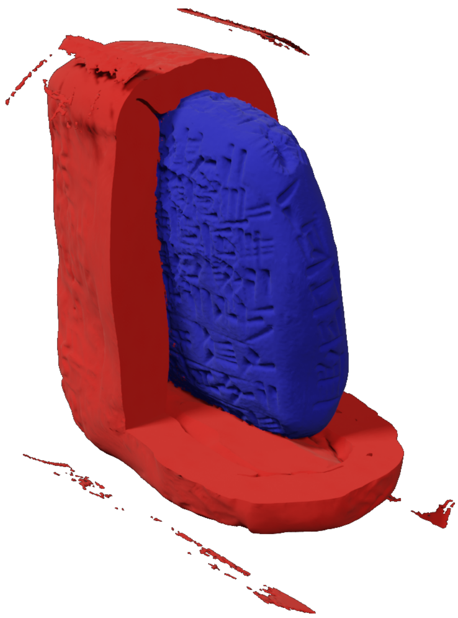

Figure 1: *3D visualization of the partly cut protective clay envelope (red) and the enclosed clay tablet (blue) of the artefact available for our work, segmented using our proposed method. Both parts of the artefact, envelope and tablet, feature inscriptions that domain scientists are interested in analyzing. Best viewed in color.*

Segmentation of the scanned volumetric data faces two major challenges: difficult to distinguish material properties and noise introduced by the scanning process. Considering the first issue, most of the clay tablet and the envelope consists of silt from a riverbank [TC12]. The silt is mostly homogeneous in density, except for impurities such as stones or shells, implying that it is not practical to distinguish tablet from envelope by merely comparing density values. The techniques used in former times to enclose the tablet in the envelope make the situation even more challenging. Some kind of separation material including plant or fabric-based layers may or may not have been used during construction to separate the tablet from the envelope [CT11; TC12]. But even if separation material was used, the material may not be discernible in the scanned data. Reasons are manifold and include scanning parameters such as spatial resolution and noise. Also, since the artefacts were handcrafted, separation material may have been incomplete, resulting in contact points and thus physical connections of tablet and envelope.

With respect to the second issue, the noise distribution of the scanned data heavily depends on the exposure time of the scans [Buz08; PL15]. For the current study, data acquired using a micro CT scanner capturing volumes at micrometer resolution was provided by the Deutsches Elektronen-Synchroton (DESY), with exposure times ranging from 2 hours to 20 hours. While the 20 h scans exhibit very low noise levels, the 2 h scans show a pronounced noise distribution that heavily impacts the otherwise homogeneously distributed density values. Fig. 2 shows an example, highlighting an area in which it is challenging to distinguish between tablet and envelope. The application of our proposed method is envisaged to be applied to a larger number of discovered artefacts in the future. For such activities, scan times of 20 h are infeasible, and having a segmentation method available that can deal with the noise levels exhibited by low-exposure time scans is essential.

In the present paper, we propose a segmentation and visualization pipeline capable of segmenting clay tablets with the described characteristics. Our approach aims at minimal user input, allowing domain scientists with little knowledge in data processing to use the system. While previous studies [MKJB10; FWMC13; BHM16] have focused on the detection and extraction of cuneiform letters, there is to the best of our knowledge no research on the segmentation of clay tablets which are sealed inside an envelope. However, there are multiple general approaches for the segmentation of volumetric data especially in the medical sciences. An overview of segmentation methods is provided by ELNAKIB, GIMEL'FARB, SURI, and EL-BAZ [EGSE11] and NOSRATI and HAMARNEH [NH16]. As baseline for our work we choose a variant of the *watershed* algorithm first introduced by BEUCHER and LANTUÉJOUL [BL79]. It segments an image by interpreting it as a heightmap and then "flooding" it from its minima. This allows to segment the image into multiple regions by constructing "barriers" where assumed "water" in two basins would flow into each other [BL79; GW17]. In the final result, these barriers mark the borders of the generated segments. However, one disadvantage of the watershed algorithm is that it tends to generate irregular shapes [NP14]. NEUBERT and PROTZEL [NP14] addressed this issue by modifying pixel intensity values by adding the Euclidean distance of each pixel to its seed. In this work, we further adapt this *Compact Watershed* algorithm to the problem at hand by supporting ellipsoidal geometry reflecting the shape of the artefacts. Additionally, we introduce a novel pre-segmentation method that suppresses the physical contact points between envelope and tablet. We compare the performance of our approach to the original formulation of the algorithms.

In short, we contribute the following:

- We propose a pipeline that has the ability to segment high-resolution low-exposure volumetric scans (i.e., exhibiting high noise levels) of clay tablets sealed inside their envelope.
- We generalize the Compact Watershed algorithm by adapting the distance metric to segment spherical geometry relfecting the shape of the artefacts under investigation.
- We introduce a novel pre-segmentation method to suppress physical connections between tablet and envelope in order to improve the performance of the Watershed algorithm.

The paper is structured as follows. Sect. 2 gives an overview of the most important mathematical concepts used in the study. Sect. 3 presents an overview of the segmentation pipeline. Sect. 4 introduces the proposed changes to the Compact Watershed algorithm and presents the novel suppression algorithm. In Sect. 5, we introduce the sample dataset and the annotation process. Sect. 6 presents the evaluation results by comparing the novel method against the existing algorithms. The paper is concluded in Sect. 7.

2 Preliminaries

In this section we define some mathematical concepts used in later sections.

2.1 Superellipsoids

Superellipsoids as defined by BARR [Bar81] are a set of shapes which can be described through

$$f(x,y,z) = \left(\left(\frac{x}{a_1}\right)^{\frac{2}{\varepsilon_2}} + \left(\frac{y}{a_2}\right)^{\frac{2}{\varepsilon_2}} \right)^{\frac{\varepsilon_2}{\varepsilon_1}} + \left(\frac{z}{a_3}\right)^{\frac{2}{\varepsilon_1}}, \qquad (1)$$

where a_1, a_2, a_3 describe the size and $\varepsilon_1, \varepsilon_2$ the shape parameters of the superellipsoid. In addition to these parameters, BARR [Bar81] assigns each superellipsoid a rotation and a translation.

The shape itself can be described through the implicit function

$$f(x,y,z) = 1. \qquad (2)$$

Hence, a point $p = (x,y,z)^\intercal$ is considered on the superellipsoid if it satisfies the condition above. Otherwise it is either considered inside $f(x,y,z) < 1$ or outside $f(x,y,z) > 1$.

2.2 Orthogonal Distance Fitting

To fit an implicit surface AHN, RAUH, and RECKNAGEL [ARR01] proposed an algorithm that minimizes the orthogonal distance of all points to the surface. Considering that there might be no analytical solution to calculate the orthogonal distance from a point to an arbitrary implicit surface, they propose calculating the distance numerically. Hence, they defined the orthogonal contact point $p' = x_\Lambda(p)$

on the surface described by the parameters Λ and which follows by minimizing

$$g(p', p_0, \Lambda) = \begin{pmatrix} f(p', \Lambda) \\ \nabla f(p', \Lambda) \times (p' - p_0) \end{pmatrix}, \tag{3}$$

using a non-linear least squares solver. In this instance, ∇f describes the gradient of the function f which defines the implicit surface. Furthermore, p_0 defines the starting parameter of the non-linear least squares problem and $\times$ is the cross product. With the definition of a contact point it is possible to define the error function which is used in another least squares minimization process as

$$\sigma^2(X', X) = (X' - X)^\intercal \cdot P^\intercal P \cdot (X - X'). \tag{4}$$

Here X defines the input points used to fit the implicit surface and $P^\intercal P$ the covariance matrix of the distances calculated using the orthogonal contact point algorithm.

3 Method Overview

Fig. 3 illustrates our approach. We base our system model on a source, filter, and mapper architecture as discussed, e.g., by SCHROEDER, MARTIN, and LORENSEN [SML18] and SCHUMANN and MÜLLER [SM00]. SCHROEDER, MARTIN, and LORENSEN [SML18] define the source as a processing step that interfaces with external data sources or generates data from parameters, whereas the mapper terminates the visualization pipeline by generating a visual representation of the data. The filter describes the transformation process of the data. It generally requires one or multiple inputs, processes the data and then passes the transformed data on to one or multiple outputs. Using this framework, we in the following describe how the data is pre-processed by denoising, and how the segmentation of the tablet is approached.

Source: Input to our pipeline are 3D density volumes reconstructed from captured x-ray projections. We input these into the pipeline as raw volumetric data.

Denoising: To reduce noise, the data are pre-processed by means of an image denoising algorithm. To avoid smearing cuneiform letters or closing small gaps that might be part of a symbol we employed an iterative bilateral filter because of its edge preserving properties [PKTD09]. To speedup the computation process, we restricted the denoising using a box window of size $2\sigma + 1$, with σ being the spatial parameter of the bilteral filter.

Segmentation of Foreground & Background: We refer to the *background* as parts of the volume that do not contain any information about the actual artefact, such as the air that got captured in the data acquisition process or the sample holder used to hold the artefact in place. We define the *foreground* as all voxels that belong to the artefact. This includes the materials used for construction as well as impurities including stones. By utilizing the fact that the density of air and of the sample holder is lower than the density of the clay we are able to employ a thresholding method such as Otsu's method [Ots79] to differentiate between foreground and background. To perform the foreground-background segmentation we binarize the volume and utilize the 2-dimensional Otsu's method [JWY91], which has the advantage of being robust to the noise that might still be left after the denoising process.

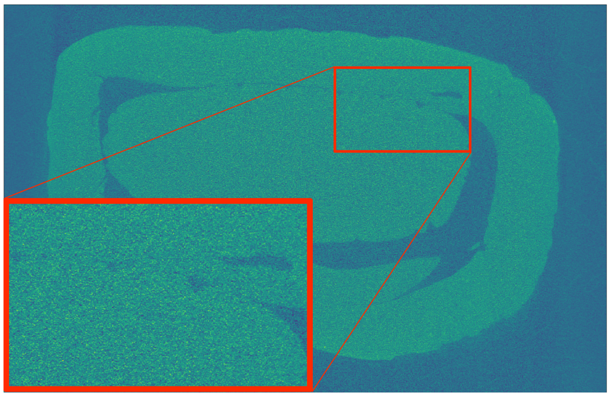

Figure 2: *Slice through the noisy volume data. While when viewing the entire slice it is possible to differentiate between tablet and envelope, it is in some regions challenging to distinguish both when only local information is available. Such a region is magnified in the red rectangle. Here, tablet and envelope are physically touching each other and appear as one part. Best viewed in color.*

Segmentation of Tablet & Envelope: Using the binarized volume, we calculate a distance transform [GW17] by computing the Euclidean distance for all foreground voxels to their nearest background voxel. This distance-transformed data is subsequently input to the different segmentation algorithms discussed in Sect. 4. For the proposed novel *suppressed watershed* algorithm, an additional suppression step is applied to the distance-transformed data. One manual step is included in our approach: All segmentation algorithms require initial markers that act as seeds for the individual segments to define their labels. These markers have to be hand-placed by the user inside the volume at voxels known to belong to either envelope or tablet; thus we require the user to label at least one voxel of the clay tablet and one voxel of the envelope prior to the segmentation process. Parts of our proposed approach depend on superellipsoid geometry for orthogonal distance fitting of the segmented regions; here we additionally require the user to select points located on the border between tablet and envelope. We utlize these user specified points to fit a superellipsoid using orthogonal distance fitting as described in Sec. 2.2.

Mapper: Polygonal isosurfaces of the labelled volume are computed by means of the Marching Cubes algorithm [LC87]. Isosurface contruction as well as 3D display are implemented using the Visualization Toolkit [SML18].

4 Compact & Suppressed Watershed

In this section, we discuss the basics of two watershed algorithm variants; we analyze, the Priority-Flood algorithm as proposed by BARNES, LEHMAN, and MULLA [BLM14] and the Compact Watershed algorithm by NEUBERT and PROTZEL [NP14]. We present our modification to the Compact Watershed algorithm and introduce our novel suppression-based pre-segmentation method that suppresses physical connections between clay tablet and envelope.

As introduced in Sect. 1, watershed algorithms interpret an image as a height-map, which is "flooded" from its low elevations. In

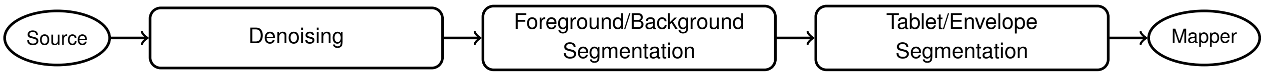

Figure 3: *Method overview illustrating the proposed pipeline for clay tablet segmentation. The volumetric data from the CT scan is denoised using a bilateral filter, then foreground and background voxels are segmented using Otsu's method [Ots79]. Clay tablet and envelope are segmented in multiple processing steps, including a distance transform followed by contact point suppression (cf. Sect. 4) and application of either the Priority-Flood or the Compact Watershed algorithm. The segmented data is mapped to a polgonal isosurface for 3D rendering. For further details see Sect. 3.*

general Priority-Flood algorithms are characterized by employing a priority queue to order cells by their elevation [BLM14], starting with the lowest elevation, in order to achieve improved runtime performance. For this study, we consider a Priority-Flood variant proposed by BARNES, LEHMAN, and MULLA [BLM14]. It is not constraint to a specific connectedness of the data and, particularly important for our application, generalizes to 3-dimensional input grids. Additionally, BARNES, LEHMAN, and MULLA are able to achieve an optimal runtime of $\mathcal{O}(n)$ for integer and a runtime of $\mathcal{O}(n \cdot \log_2(n))$ for float-point data, important when processing large volumes like ours.

The Compact Watershed algorithm proposed by NEUBERT and PROTZEL [NP14] adresses the issue that watershed algorithms tend to generate irregular shapes of the segmented regions (cf. Sect. 1). NEUBERT and PROTZEL proposed penalizing pixels that are farther away from their seed (initial marker, cf. Sect. 3) by adding the Euclidean distance between the currently considered pixel and the corresponding segment seed to the pixel's intensity (here the value of the distance transform), thereby decreasing the pixel's priority in the priority queue. Furthermore, they introduced a compactness factor that is multiplied with the Euclidean distance to control how much the shape of the segments is influenced.

Fig. 4a shows one particular failure point when computing a segmentation with the Priority-Flood or the Compact Watershed algorithm: high-intensity ridges in the distance-transformed volume. Such ridges emerge from the foreground-from-background segmentation at locations where tablet and envelop touch or where background voxels are erroneously classified as foreground. At regions where the contact areas between tablet and envelope are relatively small, the distance-transform value for the affected voxels is relatively small as well, as the nearest background voxel is closeby (Fig. 4d). However, when the contact areas become larger (Fig. 4c), distances to the nearest background voxel increase and so the distance-tranform values increase as well. These high-intensity values at the rigdes cause higher prioritisation of the ridge voxels in the segmentation algorithm's priority queue, hence the larger the ridge's intensity the more likely the algorithm will connect tablet and envelope (the ridges will be "flooded" early in the segmentation process). Once such a connection has been established, it is likely that the algorithm continues to wrongly label parts of the tablet as envelope or vice versa.

Ellipsoid-based Compact Watershed: Fig. 5a-c shows the effect of increasing the compactness factor c of the Compact Watershed algorithm. For low c, flooding across the ridges occurs, connecting tablet with envelope. When c is large, the penalty factor dominates a pixel's (or in 3D a voxel's) priority in the priority queue.

Figure 5: *Slice of the segmented volume, using Compact Watershed with Eucleadian-distance-based compactness factors c of (a) $c = 0$, (b) $c = 0.1$, (c) $c = 10$, and (d) with our ellisoid-based modification and $c = 0.1$. Seed points are marked with a red cross. With Eucleadian-distance-based compactness (a-c), flooding across the high-intensity ridges occurs; higher values of c result in the distance between seed and voxel dominating over the distance transform value (intensity). Ellipsoid-based compactness (d) successfully capture the shape of the tablet.*

In these cases and for our special case of two seeds, the segmentation simply cuts the region along a planar surface. To counteract, we modify the algorithm as follows. We assume that the clay tablet can be approximated by a superellipsoid (or in the simplest case a sphere) defined by parameters Λ with center of mass p_m. Based on this assumption, a function $G_\Lambda(p)$ is defined for the penalization of a voxel p depending on its distance to the surface of the superellipsoid (compared to its Euclidean distance to the seed as in the original NEUBERT and PROTZEL algorithm). Recall from Sect. 3 that the user manually places seeds into tablet and envelope. To define the shape of the superellipsoid, we require multiple such seeds $\{p_l^1, \ldots, p_l^n\}$ for each label l and assume that the seeds are all placed on the boundary of the superellipsoid. To fit a superellipsoid we utilize orthogonal distance fitting as described in Sect. 2.2 by minimizing the total orthogonal distance of all seed to the respective superellipsoid. We solve this non-linear least squares problem using the Levenberg-Marquardt algorithm [Lev44; Mar63] resulting in superellipsoids that approximate the user specified shapes. As the initial state of the Levenberg-Marquardt algorithm we set the center of the superellipsoid equal to the center of mass of the seeds and derive the initial rotation via Principle Component Analysis (PCA) [Pea01]. The other parameters a, b, c are set equal to the distance of the seed furthest from the center of mass and $\varepsilon_1 = \varepsilon_2 = 1$ thus using a sphere as our initial configuration to fit the superellipsoid. Hence, for the case of a simple sphere, a single seed for tablet and envelope each suffices. In this spherical case we define the center of the sphere as the center of mass of the clay tablet and envelope. Further, it allows to express G_Λ as the orthogonal distance between a voxel at coordinate p and the sphere's surface spanned up by the seed p_l^1 as $G_\Lambda = |||p - p_m||_2 - ||p_l^1 - p_m||_2|$ with $\Lambda = \{p_l^1, p_m\}$.

In the case of a superellipsoid, G_Λ becomes $G_\Lambda = ||p - x_\Lambda(p)||$,

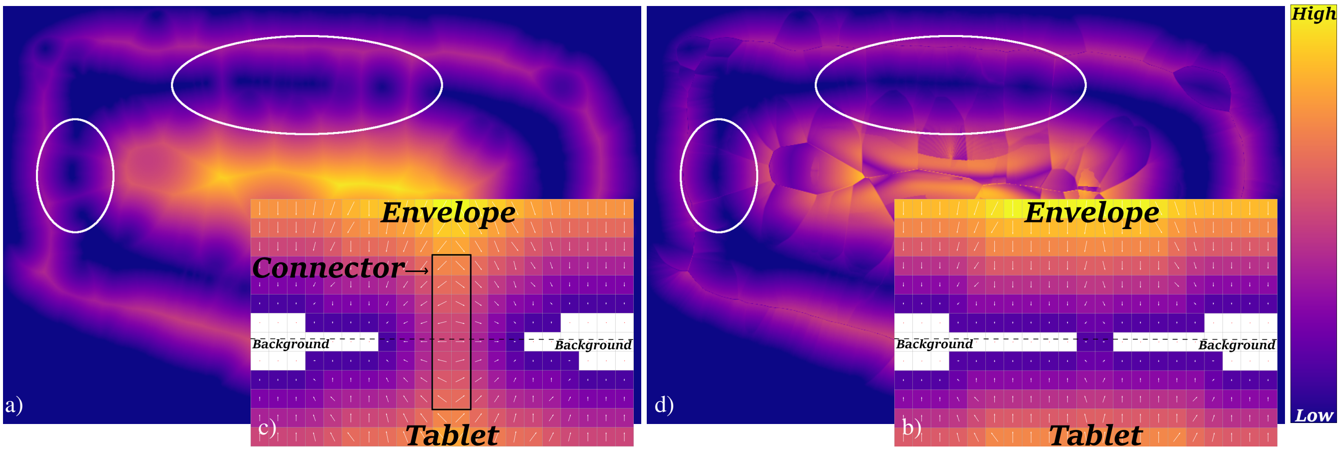

Figure 4: *(a) Original and (b) suppressed distance-transformed data of the slice shown in Fig. 2. The original distance transform contains strong connections between clay tablet and envelope (white ellipses), making segmentation challenging. The suppressed distance transform largely improves upon this situation. Panels (c) and (d) zoom in on two selected ridges, see text for details.*

where Λ specifies the superellipsoid parameters (cf. Sect. 2.1) and $x_\Lambda(p)$ is the orthogonal contact point of p (cf. Sect. 2.2).

Suppressed Watershed: Using an ellipsoid-based penalization largely improves the segmentation; Fig. 5d shows the technique applied to the section failing when using Euclidean distances. However, flooding across high-intensity ridges still persists in many locations, an example is shown in Fig. 7. It is hence desirable to be able to detect if a voxel at position p belongs to such an "erronous" connection, and to correct its distance-transform intensity correspondingly. Let I be the original distance-transform intensity field and f be a classification function indicating whether a voxel belongs to an undesired ridge. We define a "suppressed" distance field I' (that suppresses the undesired ridges) by

$$I' = I \cdot f. \tag{5}$$

Ideally, f would be

$$f(p) = \begin{cases} 0 & \text{if the voxel at } p \text{ is part of a connector} \\ 1 & \text{otherwise.} \end{cases} \tag{6}$$

However, such binary classification would require prior knowledge about the resulting segmentation and is therefore not feasible.

Instead, we define f as a probability function that represents the likelihood of a voxel not belonging to an undesired connector. To avoid false positive assignments, which may result in setting the intensity value of a voxel in the distance field to zero, we boost the intensities instead of suppressing them. Therefore, we introduce similar to NEUBERT and PROTZEL [NP14] a factor γ that increases the intensity values via

$$I'(p) = I(p) \cdot (\gamma f(p) + 1). \tag{7}$$

We note that we require the user to choose a value for γ for the Ellipsoid-based Compact Watershed and the Suppressed Watershed algorithm. For a lower γ the algorithm will weight the distance transform more, whereas for a higher γ it will boost the intensity values of voxels that do not belong to a connector. As for f a straightforward formulation can be obtained by noticing that most ridges are parallel to the normal vectors of a superellipsoid representing the tablet. This implies that the gradient of the distance field along the ridges is orthogonal to the superellipsoid normals. We hence define f as the dot product between the superellipsoid normals and the gradient of the distance field:

$$f(p) = \left| \frac{\nabla I(p)}{||\nabla I(p)||} \cdot \frac{R(p)}{||R(p)||} \right|, \tag{8}$$

where ∇I denotes the gradient of the distance field and R represents the superellipsoid normals:

$$R(p) = p - x_\Lambda(p), \tag{9}$$

where $x_\Lambda(p)$ is the contact point for a superellipsoid with parameters Λ, as introduced in Sect. 2.2. As previously mentioned, we fit superellipsoids from initial seeds provided by the user. For our evaluation, we also consider using a simple sphere instead of a superellipsoid:

$$R(p) = p - p_m. \tag{10}$$

Fig. 4b shows the resulting distance field as obtained using superellipsoid suppression. The undesired ridges present in Fig. 4a have largely been eliminated.

5 Data

For the study at hand, we were provided with a dataset consisting of a total of six volumetric CT scans of two replica tablets of approximately 5cm×5cm×3cm in size, with three scans of varying exposure time (ranging from 2 h to 20 h) for each tablet. The tablets were handcrafted as examination objects to mimic tablets found at archaeological sites. They are enclosed in their envelope with both containing cuneiform inscriptions on the envelope and the tablet.

Fig. 1 shows a 3D visualization of the segmentation obtained from the only 2 h low-exposure scan in our dataset. The scans were conducted at German Electron Synchrotron (DESY) (`www.desy.de`) in order to evaluate the quality of the scanned volumes with regards to different scanning parameters, including *exposure time* and *frames per projection*. The size of the data volumes ranges from $2283 \times 2284 \times 1132$ to $2283 \times 2284 \times 2304$ voxels (approx. 5.903 to 12.014 gigavoxels) stored as 16 bit intensity values with a resolution per voxel of 27µm. The scans were conducted such that the clay tablets where kept at a fixed position, making it possible to directly compare the scans to each other.

For the reasons discussed in Sect. 1, our particular interest is in segmenting scans with short exposure times, which exhibit higher noise levels than the scans with long exposure times. Unfortunately the datset only contained one such low-exposure scan, hence, reducing our dataset to a single volume. To increase the size of the test data and to evaluate the algorithm on different grid resolutions, we subsampled this single volume by factors of 2 to different decreased resolutions listed in Tab. 1. Also, to evaluate our approach for data with varying noise levels, we computed derived volumes by applying Gaussian noise as listed in Tab. 2.

To obtain a *ground truth* annotation for evaluation, we utilized the VAST Lite software (Volume Annotation and Segmentation Tool) [BSL18] and manually segmented the data, using three different labels for background, clay envelope and clay tablet. First, we coarsely annotated foreground and background using a threshold, then finely adjusting the initial segmentation. To segment difficult regions as the one shown in Fig. 2, we assumed that a larger space corresponds to a cuneiform letter and that a letter is always fully stamped inside the tablet. Hence, the furthest point outwards from an empty space, when approached from the center of the clay tablet, corresponds to the border of the clay tablet.

6 Evaluation

This section introduces the metrics used for evaluation and presents our experimental results on our dataset described in Sect. 5. We evaluate our results using the *Intersection over Union* and the (average) *Hausdorff* distances to measure the difference between the automatically obtained segmentations to the ground truth segmentation. Furthermore, we redefine the Hausdorff distance to compare volumes of different resolutions.

6.1 Intersection over Union

The Intersection over Union (IoU) is categorized by TAHA and HANBURY [TH15] as an overlapping metric. It calculates the overlap between the ground truth S_t and the generated segmentation S_p via [TH15]

$$\mathrm{JAC} = \frac{|S_t \cap S_p|}{|S_t \cup S_p|}, \tag{11}$$

We choose this metric based on the recommendation by TAHA and HANBURY [TH15] given that they consider this metric particularly useful when outliers exists.

6.2 Hausdorff Distance

As the second metric we chose the Hausdorff distance and the average Hausdorff distance. These metrics are classified by TAHA and HANBURY [TH15] as distance based metrics and measure the distance between the contour of the ground truth and the contour of the prediction. The metric itself is defined as the maximal distance of all points in the set A to their closest point in the set B through [TH15]

$$\mathrm{HD}(A,B) = max(h(A,B), h(B,A)), \tag{12}$$

with h defined as [TH15]

$$h(A,B) = \max_{a \in A} \min_{b \in B} ||a-b||. \tag{13}$$

The metric is considered particularly useful by TAHA and HANBURY [TH15] when comparing the contour between the ground truth and the generated segmentation.

However, a drawback of the Hausdorff distance is that it is sensitive to outliers; therefore TAHA and HANBURY [TH15] recommend using the average Hausdorff distance instead. This metric calculates the average over all closest points by changing h to

$$h'(A.B) = \frac{1}{|A|} \sum_{a \in A} \min_{b \in B} ||a-b||. \tag{14}$$

6.3 Normalized Hausdorff Distance

We note that the Hausdorff distance decreases when subsampling the volume even if the content is kept the same. To avoid this behavior we normalized the Hausdorff distance by dividing the closest distance by the total size of the volume s. This transforms the volume via scaling into an unit volume and can therefore be interpreted as a simplified approach of [Sua05]. Thus, we redefine h for the Hausdorff distance as

$$h_s(A,B) = \max_{a \in A} \min_{b \in B} \frac{||a-b||}{||s||} = \frac{h(A,B)}{||s||} \tag{15}$$

and for the average Hausdorff distance as

$$h'_s(A,B) = \frac{1}{|A|} \sum_{a \in A} \min_{b \in B} \frac{||a-b||}{||s||} = \frac{h'(A,B)}{||s||}. \tag{16}$$

Assuming that all points in A and B are in the volume of size s then the normalized Hausdorff distance ensures that the error is between zero and one.

6.4 Results

Fig. 6 shows the evaluation metrics for tablet and envelope segmented by the different methods in dependence of subsampling factor and added noise. For all considered volume sizes and noise distributions, our proposed suppression-based approach performs best compared to the alternative algorithms. The top row of Fig. 6 also shows that the suppression-based approach is reasonably stable with respect to changes in volume size (i.e., subsampling), in comparison the performance of the ellipsoid-based Compact Watershed algorithm worsens with increasing subsampling factor. The Priority Flood algorithm performs worst for all volume sizes. Similar results are obtained for varying noise distributions (bottom row of Fig. 6), although results are less stable with respect to increasing noise.

We could confirm these quantitative results by a qualitative visual analysis of the data. As Fig. 7a shows, our approach is able to

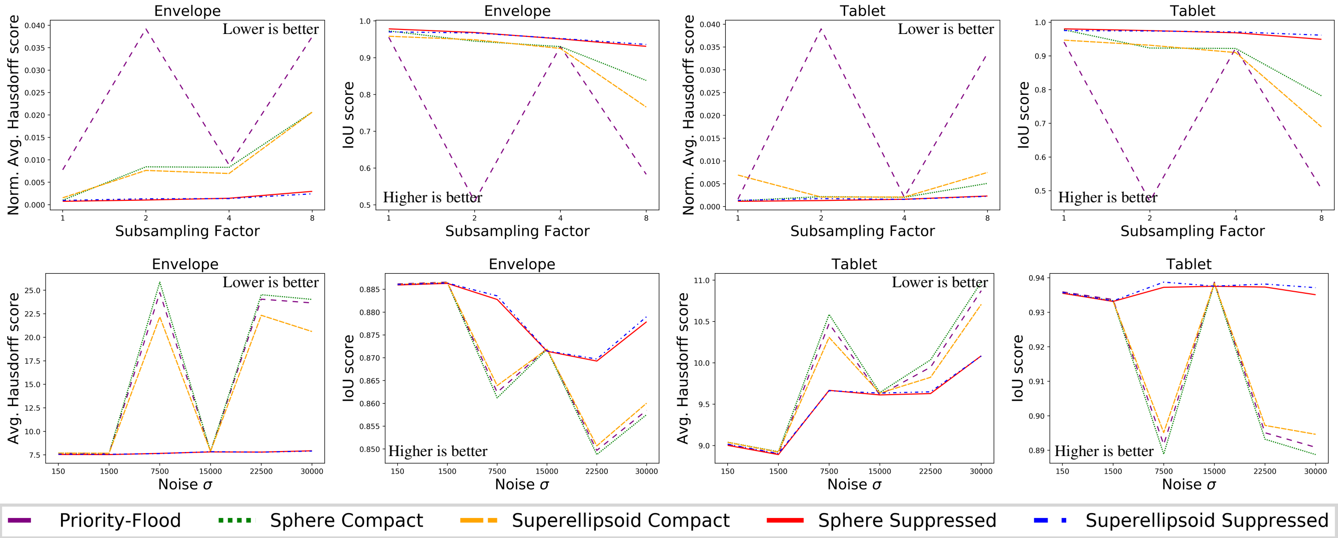

Figure 6: *Evaluation results for the scan described in Sect. 5. Upper row shows evaluation metrics (cf. Sect. 6) comparing automatic segmentation to ground truth segmentation for volumes of varying resolution (cf. Tab. 1); bottom row shows metrics for data with added Gaussian noise (cf. Tab. 2). Our suppression-based method results in more stable segmentations for all datasets, while the Compact Watershed worsens with decreasing resolution and increasing noise. The Priority-Flood algorithm is the least stable and often generates the worst result.*

preserve fine details of the inner clay tablet. For example, it is capable of maintaining a *ruling* on the inner clay tablet that was used to align the handwritten letters, whereas the other algorithms often mislabel parts of the envelope as tablet (e.g., Fig. 7b). While this is less the case for the ellipsoid-based Compact Watershed compared to the Priority-Flood algorithm, it still mislabels part of the envelope, whereas the suppressed watershed only keeps tiny blobs. We note that all algorithms tend to generate erroneous patterns at locations where envelope and tablet exhibit physical connections. For instance, the patterns shown in Fig. 8 manifest themselves as multiple planar surfaces similar to a “crystal” and could potentially obscure letters or decrease the readability.

7 Summary and Conclusion

We investigated an unmodified Priority-Flood algorithm and an ellipsoid-based formulation of the Compact Watershed algorithm in terms of their ability to segment volumetric CT scans of clay tablets enclosed inside an envelope and proposed an additional novel pre-segmentation method. It suppresses physical contact points between the clay tablet and the envelope, resulting in a more stable output with regards to noise and resolution of the input. Comparison of automatically obtained segmentation with a ground-truth segmentation by means of both qualitative inspection and quantitative similarity metrics showed that our novel suppression-based approach largely improves upon the existing algorithms.

For the dataset available in this study, the proposed method is able to produce a segmentation that is capable of preserving fine details. However, all methods still produce some erroneous patterns at locations where the tablet is physically connected to the envelope. Also, we noticed that all algorithms tend to mislabel parts of the envelope as clay tablet to some extent. While for the proposed suppression-based watershed algorithm, these incorrectly labelled parts visually are only barely noticeable, further investigation needs to determine whether they can potentially impact the readability or alter the meaning of the inscriptions. In this context, to improve readability, it will also be useful in future work to investigate visual analysis techniques that extract and highlight the inscribed cuneiform letters. Given the stable output on the subsampled dataset it would be worth considering using the algorithm in an hierarchical approach or directly segmenting the polygonal mesh generated from the isosurface to speedup the computation. In conclusion we note, however, that the CT data available to us for this study is too small for generic statements. Hence, further research will be required to validate our results using a larger number of scanned clay tablets with different characteristics.

Scale	SNR	Resolution
1	~0.88	$2124 \times 1411 \times 2191$
2	~0.88	$1061 \times 704 \times 1094$
4	~0.88	$530 \times 351 \times 546$
8	~0.88	$264 \times 175 \times 272$

Table 1: *Subsampling resolutions used in this study.* Scale *denotes subsampling factor,* SNR *denotes signal to noise ratio.*

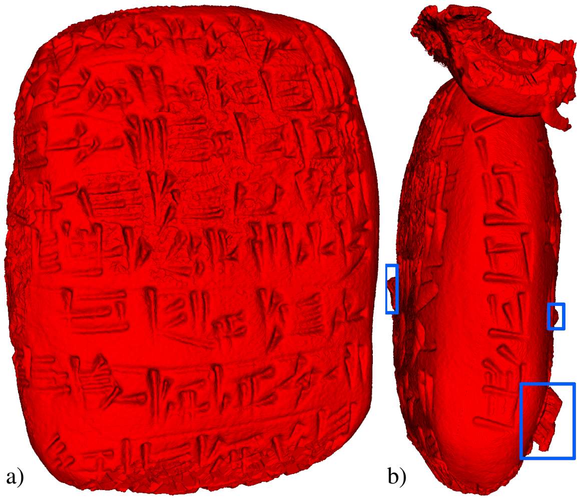

Figure 7: *Left: Inner clay tablet segmented using our novel suppression-based method. Our method preserves many fine details, including a ruling used to align the text. Right: Segmentation obtained with the ellipsoid-based, unsuppressed Compact Watershed algorithm. This algorithm considers parts of the envelope as tablet, resulting in regions of the tablet being obscured. Also, the blobs discussed in Sect. 6 are visible.*

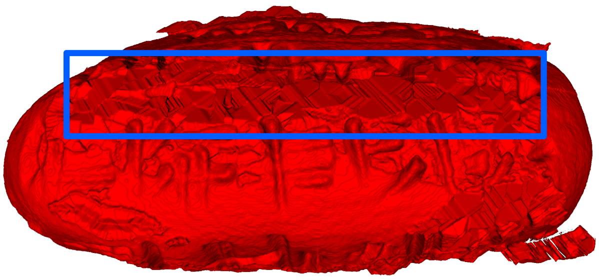

Figure 8: *Example showing the "crystal" pattern discussed in Sect. 6 (blue rectangle). This pattern occurs at locations where the clay tablet physically touches the envelope. This results in the algorithm making an artificial cut between the envelope and the tablet.*

Sigma	SNR	Resolution
150	~49.98	$2124 \times 1411 \times 2191$
1500	~4.99	$2124 \times 1411 \times 2191$
7500	~1.00	$2124 \times 1411 \times 2191$
15000	~0.51	$2124 \times 1411 \times 2191$
22500	~0.40	$2124 \times 1411 \times 2191$
30000	~0.37	$2124 \times 1411 \times 2191$

Table 2: *Gaussian noise distributions added to the original scanned data.* Sigma *denotes Gaussian standard distribution in voxels,* SNR *denotes signal to noise ratio.*

Acknowledgements

We want to thank the DESY research center and especially Prof. Dr. Christian Schroer, Dr. Andreas Schropp and Dr. Maik Kahnt as well as Dr. Imke Greving, Dr. Berit Zeller-Plumhoff from Helmholtz-Zentrum Geesthacht (HZG) for providing us with the data to conduct our research. We further want to thank Prof. Dr. Dr. Cécile Michel for manufacturing the replica and Dr. Wiebke Beyer for providing us with information about the tablets themself. This work is based on the results of the Master's thesis "Segmentation of Noisy Volume Data" by Tim Rolff at Universität Hamburg.

The research for this publication was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy –EXC 2176 'Understanding Written Artefacts: Material, Interaction and Transmission in Manuscript Cultures', project no. 390893796. The research was conducted within the scope of the Centre for the Study of Manuscript Cultures (CSMC) at Universität Hamburg.

References

[ARR01] Ahn, Sung Joon, Rauh, Wolfgang, and Recknagel, Matthias. "Least Squares Orthogonal Distance Fitting of Implicit Curves and Surfaces". Ed. by Radig, B. and Florczyk, S. Lecture Notes in Computer Science, vol. 2191. Springer, 2001 2.

[Bar81] Barr, Alan H. "Superquadrics and Angle-Preserving Transformations". *IEEE Computer Graphics and Applications* 1 (1981) 2.

[BHM16] Bogacz, B., Howe, N., and Mara, H. "Segmentation Free Spotting of Cuneiform Using Part Structured Models". *2016 15th Int'l. Conf. on Frontiers in Handwriting Recognition (ICFHR)*. 2016 2.

[BL79] Beucher, Serge and Lantuéjoul, Christian. "Use of Watersheds in Contour Detection". *In Int'l. Workshop on Image Processing: Real-time Edge and Motion Detection/Estimation*. Vol. 132. Jan. 1979 2.

[BLM14] Barnes, Richard, Lehman, Clarence, and Mulla, David. "Priority-flood: An optimal depression-filling and watershed-labeling algorithm for digital elevation models". *Computers & Geosciences*. Vol. 62. 2014 3, 4.

[BSL18] Berger, Daniel R., Seung, H. Sebastian, and Lichtman, Jeff W. "VAST (Volume Annotation and Segmentation Tool): Efficient Manual and Semi-Automatic Labeling of Large 3D Image Stacks". *Frontiers in Neural Circuits* 12 (2018) 6.

[Buz08] Buzug, Thorsten M. *Computed Tomography. From Photon Statistics to Modern Cone-Beam CT*. Springer, 2008 2.

[CT11] Cartwright, Caroline and Taylor, Jon. "Investigating technological and environmental evidence from plant remains and molluscs in cuneiform tablets". *British Museum technical research bulletin* 5 (2011) 2.

[EGSE11] Elnakib, Ahmed, Gimel'farb, Georgy, Suri, Jasjit S., and El-Baz, Ayman. "Medical Image Segmentation: A Brief Survey". *Multi Modality State-of-the-Art Medical Image Segmentation and Registration Methodologies* (2011) 2.

[FWMC13] Fisseler, Denis, Weichert, Frank, Müller, Gerfrid, and Cammarosano, Michele. "Towards an interactive and automated script feature Analysis of 3D scanned cuneiform tablets". Nov. 2013 2.

[GW17] Gonzales, Rafael C. and Woods, Richard E. *Digital Image Processing*. 2017 2, 3.

[JWY91] Jianzhuang, Liu, Wenqing, Li, and Yupeng, Tian. "Automatic thresholding of gray-level pictures using two-dimension Otsu method". *1991 Int'l. Conf. on Circuits and Systems*. IEEE, 1991 3.

[LC87] LORENSEN, WILLIAM and CLINE, HARVEY. “Marching Cubes: A High Resolution 3D Surface Construction Algorithm”. *ACM SIGGRAPH Computer Graphics* 21 (1987) 3.

[Lev44] LEVENBERG, KENNETH. “A method for the solution of certain non-linear problems in least squares”. *Quarterly of applied mathematics* 2.2 (1944) 4.

[Mar63] MARQUARDT, DONALD W. “An algorithm for least-squares estimation of nonlinear parameters”. *Journal of the society for Industrial and Applied Mathematics* 11.2 (1963), 431–441 4.

[Mic08] MICHEL, CÉCILE. *The Old Assyrian Trade in the light of Recent Kültepe Archives*. 2008 1.

[Mic16] MICHEL, CÉCILE. *Women and Real Estate in the Old Assyrian Texts*. Version 51. Orient, The Society for NearEastern Studies in Japan, 2016 1.

[MKJB10] MARA, HUBERT, KRÖMKER, SUSANNE, JAKOB, STEFAN, and BREUCKMANN, BERND. “GigaMesh and Gilgamesh: –3D Multiscale Integral Invariant Cuneiform Character Extraction”. *Proc. of the 11th Int’l. Conf. on Virtual Reality, Archaeology and Cultural Heritage*. Eurographics Association, 2010 2.

[NH16] NOSRATI, MASOUD S. and HAMARNEH, GHASSAN. “Incorporating prior knowledge in medical image segmentation: a survey”. (2016) 2.

[NP14] NEUBERT, PEER and PROTZEL, PETER. “Compact Watershed and Preemptive SLIC: On improving trade-offs of superpixel segmentation algorithms”. *Int’l. Conf. on Pattern Recognition (ICPR)*. IEEE. 2014 2–5.

[Ots79] OTSU, NOBUYUKI. “A Threshold Selection Method from Gray-Level Histograms”. *IEEE Trans. on Systems, Man, and Cybernetics* 9 (1 1979) 3, 4.

[Pea01] PEARSON, KARL. “On lines and planes of closest fit to systems of points in space”. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 2.11 (1901) 4.

[PKTD09] PARIS, SYLVAIN, KORNPROBST, PIERRE, TUMBLIN, JACK, and DURAND, FRÉDO. “Bilateral Filtering: Theory and Applications”. (2009) 3.

[PL15] PRINCE, JERRY L. and LINKS, JONATHAN M. *Medical Imaging Signals and Systems*. 2nd Edition. Pearson, 2015 2.

[SM00] SCHUMANN, HEIDRUN and MÜLLER, WOLFGANG. *Visualisierung. Grundlagen und allgemeine Methoden*. 2000 3.

[SML18] SCHROEDER, WILL, MARTIN, KEN, and LORENSEN, KEN. *The Visualization Toolkit. An Object-Oriented Approach To 3D Graphics*. Edition 4.1. 2018 3.

[Sua05] SUAU, PABLO. “Adapting Hausdorff metrics to face detection systems: a scale-normalized Hausdorff distance approach”. *Portuguese Conference on Artificial Intelligence*. Springer. 2005, 76–86 6.

[TC12] TAYLOR, JON and CARTWRIGHT, CAROLINE. “THE MAKING AND RE-MAKING OF CLAY TABLETS”. *Scienze dell’Antichità* (17 2012) 1, 2.

[TH15] TAHA, ABDEL AZIZ and HANBURY, ALLAN. “Metrics for evaluating 3D medical image segmentation: analysis, selection, and tool”. *BMC Med Imaging* 15 (2015) 6.