Visualizing and Contextualizing Outliers in Aegean Seal Collections

Bartosz Bogacz[1], Sarah Finlayson[2],
Diamantis Panagiotopolous[2], and Hubert Mara[1]

[1] FCGL – Forensic Computational Geometry Laboratory
bartosz.bogacz@iwr.uni-heidelberg.de
hubert.mara@hs-mainz.de
Heidelberg University & University of Applied Sciences Mainz
[2] CMS – Corpus of Minoan and Mycenaean Seals
Institut für Klassische Archäologie und Byzantinische Archäologie
Heidelberg University
{sarah.finlayson,diamantis.panagiotopolous}@zaw.uni-heidelberg.de

Abstract. The *Corpus of Minoan and Mycenaean Seals* (CMS) in Heidelberg contains records of approximately 12.000 ancient seals and seal impressions. The study of the seals, their engraved motifs, and sealing practices gives valuable insights into the social, political and economic organization of Aegean Bronze Age societies. A key research question is whether a seal is always used by a single individual. Current archaeological practice is to manually compare sealings and qualitatively assess their similarity or difference. With large collections of seal impressions made by the same seal, this process quickly becomes prohibitive if every detail is to be considered. Our dataset consists of rasterized images of structured-light 3D scanned seal impressions on plasticine casts. We improve upon our previous approach to alignment and introduce methods visualizing and summarizing differences in a collection of highly similar seal impressions. We overlay binarized images of seal impressions to easily detect variations in the motifs. We enrich those with quiver plots displaying only the non-rigid contribution to the deformation between seal impression pairs. By comparing the visualizations of historic seal impressions to our experimentally created modern variants we gather evidence of their authorship.

Keywords: Machine Learning · 3D Computer Vision · Aegean Seals

1 Introduction

The CMS project (Corpus of Minoan and Mycenaean Seals) in Heidelberg, Germany, records approximately 12.000 Aegean Bronze Age seals and sealings. It consists of impressions of seals and sealings in plasticine, silicon and gypsum, photographs, drawings, and a digital database with associated metadata. Seals

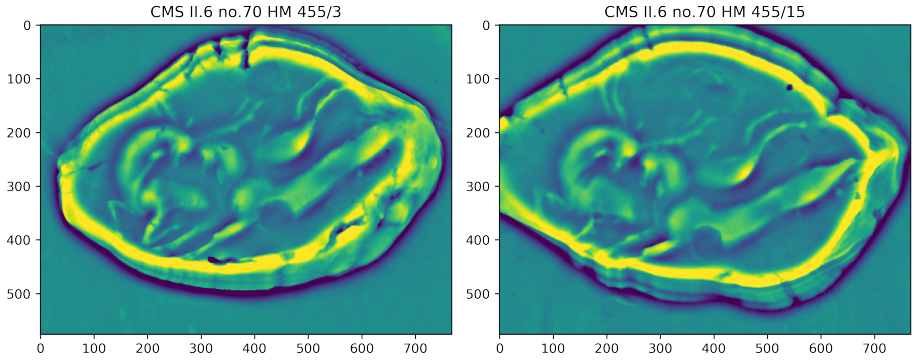

Fig. 1. Rasterized 3D scans of seal impressions CMS II.6 no.70 HM 455/3 and CMS II.6 no.70 HM 455/15. The color indicates the MSII computed surface curvature of the 3D model with a sphere radius of 1.5mm. Dark blue colors denote concave (recessed) regions, bright yellow colors denote convex (exposed) regions.

are typically made from hard and soft stones, bone or ivory, metal, and occasionally man-made materials. They display engraved motifs, from simple geometrical patterns to complex figurative scenes. Beyond their significance as prestige objects, insignia and amulets, their primary purpose is administrative: seals are impressed on clay sealings to secure objects, and to make statements of ownership or responsibility, whether corporate or individual. Therefore, it is key for archaeologists to have certainty that a set of impressions with seemingly identical motifs originates from the same seal, and then to clarify whether or not each impression was made by the same person. Current manual approaches focus on a detailed study of casts, photographs and drawings of similar seal impressions. The detailed study of minuscule visual differences in the motifs of 10 or more seal impressions at once becomes prohibitively time consuming very quickly. We approach this challenge by computing an automated alignment of seal impressions, highlighting visual differences in motifs, and summarizing views to quickly determine outliers in the depiction of motifs across a large collection of mostly identical impressions.

This work builds upon and improves on our previous research in seal impression alignment [1] in two key areas: (i) Our previous two step process of an initial rigid fit with RANSAC [4] and a subsequent fine-tuning of residuals with TPS-RPM [2], both requiring optimization till convergence, are now replaced by a single direct estimation of the alignment with a SVR [3], requiring only a single parameter to control its smoothness. However, the dense visual descriptor sampling with DAISY [10] remains, making the quality of our alignments comparable to our previous work. (ii) We introduce visualization and summarization techniques developed in interdisciplinary collaboration to maximize their legibility and trustworthiness. We focus on disentangling and thus enabling the investigation of impression motifs and their deformations in separation.

2 Related Work

Keypoint registration methods are typically used for matching images from different viewpoints, e.g. stereo vision and 3D reconstruction, or for matching a template in a target image. In any case, a set of keypoints from the source and target are extracted and correspondences between them are established to fit an underlaying model.

The quality of the matching can be improved by using better image descriptors [10], or by optimizing image descriptors for a specific task. In their work [12] Verdie et al. train a regressor to predict optimal locations for keypoints of handcrafted image descriptors. The authors use a set of images taken with identical camera position and parameters but with changing illuminations due to weather conditions. The resulting regressor focuses on unchanging large structures, e.g. buildings, while ignoring foliage and changing weather. Papadaki et al. in [9] propose to reduce the amount of outliers in the correspondence computation by training a random forest (RF) classifier to reject keypoint not contributing to a successful match. The classifier is trained on a representative set of images for a specific task. For this task, the matching performance in computational speed and robustness is increased.

Keypoint registration is also used for non-rigid matching of images. In [11] Tran et al. show that consensus based model fitting such as random sample consensus (RANSAC) [4] can be used to estimate correspondences for a model without bounding the degrees of freedom. The authors exploit the property that correspondence four-tuples, from source to target (x, y, x', y'), form a two-dimensional manifold embedded in four dimensional assignment space. For a limited amount of deformation the manifold is a hyperplane with inlier correspondences tightly clustering around it while outliers are further apart. A different approach to removing outliers is proposed by Li et al. in [6]. The authors train a support vector regressor (SVR) [3] on the assignment space to match the expected manifold while being robust to outliers. Li et al. repeat the procedure, each time peeling of an increasing number of outliers.

Our approach is also based on manipulating the correspondence manifold. However, we estimate it with a manifold embedding technique, such as multi dimensional scaling (MDS) [8], and remove outliers with high embedding stress. A subsequent SVR is used for its regularization capability to smooth the correspondence manifold.

3 Dataset and 3D-Acquisition

The work presented in this article is done within the *3D forensic analysis and contextualisation of Aegean seals and sealings* (ErKon3D) project, which has already acquired high-resolution 3D-models i.e. triangular meshes of a selection of sealings. These are modern casts of the original ancient sealings; a sealing is a piece of clay on which a seal has been impressed. In this work, we use 3D scans of the modern casts. All the seals and sealings have been published, and

photographs, drawings, textual descriptions and meta-data are available in print and online resources[3]. The latter are connected to online databases of metadata such as *ARACHNE*[4], the central object database of the *German Archaeological Institute* (DAI) and the Archaeological Institute of the University of Cologne.

We focus on renderings of the high-resolution 3D-datasets as those are free from interpretation as compared to manually created tracings and drawings. Further, motifs on seals and their impressions only become fully visible under a changing light-source. A single photograph does not reproduce its full three-dimensional structure and is subject to occlusion effects. Additionally the color information of the impressed material is distracting for experts and machine learning algorithms alike. Therefore, we face a similar challenge as in previous work on cuneiform tablets and apply *Multi-Scale Integral Invariant* (MSII) [7] filtering to the 3D models. The filtering and rasterization of the seal impression models was done using the Open Source *GigaMesh Software Framework*[5].

For the experiments and validation of our methods we chose sets of sealings from Neopalatial Crete (1750/1700 - 1500/1450 BCE). The 46 sealings form 8 groups, each containing between 3 and 7 sealings impressed by the same seal. The first 4 groups all come from the same building at the site of Haghia Triada; a great deal of administrative activity took place in this building, including the storage of sealed goods, but we do not understand the complex sealing pattern there, in which a few seals are impressed frequently and the remainder only once or twice. The second 4 groups contain sealings of a specific form, used to seal folded parchment documents; here, sealings impressed with the same gold sealing ring are found at different sites around Crete. The archaeological research question is fundamentally the same for both sets - did the same person always use the same seal? - but the socio-political significance of the answers differs greatly.

4 Alignment

We build our approach upon the insights of Tran et al. [11] and Li et al. [6] that describe the alignment process between two images as fitting a two-dimensional manifold onto a four-dimensional set of points. Our main concern is then the generation, filtering, and smoothing of these correspondences that finally leads to an alignment manifold used to warp the source image onto the target image.

4.1 Descriptor Transform

Our image alignment process depends on determining keypoints in the source and target images that share the same visual patterns. We extract visual patterns from the images by means of a local visual feature descriptor such DAISY [10]. The descriptors are extracted densely, that is, for each location of a regular grid

[3] https://www.uni-heidelberg.de/fakultaeten/philosophie/zaw/cms/
[4] https://arachne.uni-koeln.de
[5] https://gigamesh.eu

defined on the image pair. For this particular dataset we found that a grid spacing of 5 pixels with a DAISY kernel size of 30 pixels and 5 rings of 8 histograms yields the best visualization results. We reduce the count of dimensions from 328 to 16 by a principal component analysis (PCA) to save computational resources. A higher count of dimensions did not improve the fidelity of our visualizations.

4.2 Descriptor Filtering

Contrary to typical image registration, we are only interested in aligning the central motif of a seal impression. The aligning process is required to ignore any material deficits, damage from weathering, and material deformation that is unlikely to be equal between impressions. Shared damage between seal impressions, however, is indicative that the seal itself was already damaged.

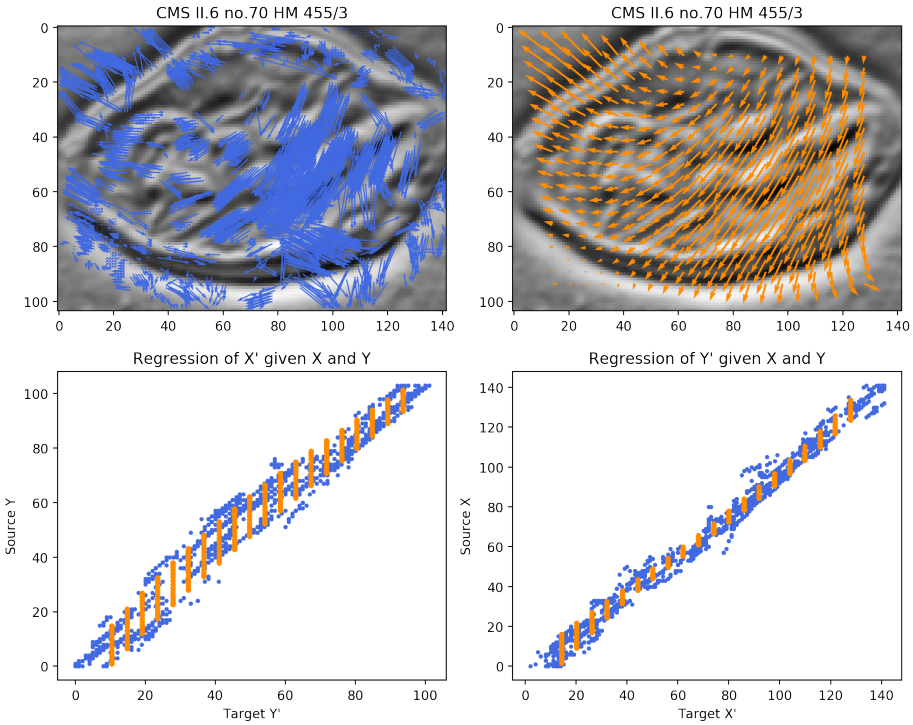

Fig. 2. Both on the left and right column the same seal impression CMS II.6 no. 70 HM 455/3 depcited grayscale by its DAISY filter response. Colors denote position on the one-dimensional PCA embedding. Overlaid is a quiver of correspondences to CMS II.6 no. 70 HM 455/15, not depicted here. The left blue quiver is before SVR smoothing, the right orange after SVR smoothing. The bottom graphs show two 2D projections of the 4D regressed assignment manifold on top of the initial assignment manifold.

Unique Descriptors that are common on the images do not contribute to a good alignment. Constructing a correspondence where either source or target is common leads to ambiguity. There are many good candidates close in feature space yet far apart in image space. This is especially true for our data, as seal impressions have large empty areas and material borders that exhibit very similar visual descriptors yet impair a proper alignment. We compute the prevalence of specific descriptors by estimating the kernel density (KDE) of a Gaussian kernel with bandwidth 0.1 in the joined descriptor space of both images. Then, only descriptors that are less common then the 50 percentile are kept for further processing.

Bidirectional The best target candidate of a source keypoint should also, vice versa, be the best source candidate of the same target keypoint. We enforce each correspondence to point to each other as best candidates. We introduce an acceptable radius of inaccuracy. Correspondences are only kept if the best candidate points back at an area within a radius of 1 descriptor step, here 5 pixels in image space, of the keypoint.

Inlying We make use of the observation of Tran et al. in [11] that for most physical deformations true positive alignment points are distributed compactly on a 2D affine hyperplane in the alignment space. We relax the assumption further and estimate the embedding of an arbitrary correspondence manifold with multi-dimensional scaling (MDS) [8]. We estimate the embedding stress of a correspondence as the sum of squared differences between distances to all other correspondences in the original space and in the embedded space. We only keep correspondences that are in the lower 90 percentile of points in terms of stress.

4.3 Alignment Manifold Regression

Even after the previous steps of filtering, the resultant set of correspondences contains outliers and is irregularly distributed over the image, c.f. the support in Figure 3. We smooth and interpolate correspondences with a radial basis function (RBF) support vector regressor (SVR). Two regression tasks are performed. Target x coordinates and target y coordinates are individually regressed from source x and y. The amount of desired smoothness and rigidity is controlled by the regularization penalty C of the regressor. We set the ϵ-tube where no penalty is applied to a small value of 0.0001. The regressed manifold then closely follows the filtered correspondences. Figure 2 shows correspondance samples and interpolated samples.

4.4 Image Warping

On the basis of the regressed alignment manifold we deform the source image to match the target image. The deformation is computed in two stages. First a finer grid with 30×30 control points resolution is interpolated with thin-plate splines (TPS) [2]. Then, on the basis of this finer grid a piecewise affine

function, quadrilaterals spanned between four grid points, is used to interpolate pixel values of the warped source image.

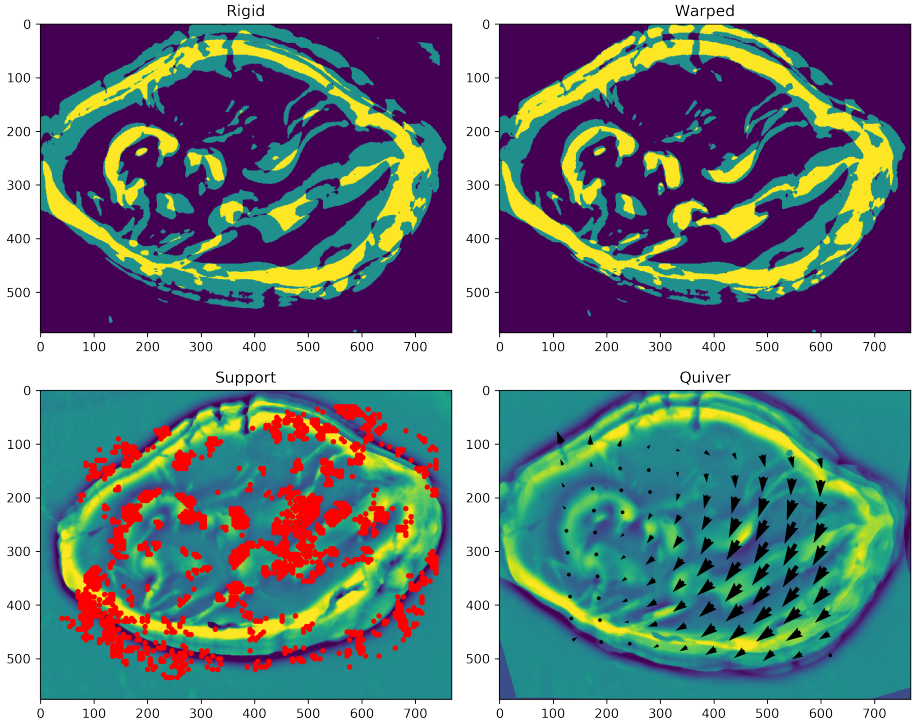

Fig. 3. Alignment of CMS II.6 no.70 HM 455/3 to match CMS II.6 no.70 HM 455/15. The **Rigid** tile depicts an additive overlay of the binarized source and target impressions with the best rigid fit, while **Warped** allows any local deformations. Below, **Support** shows the keypoints (amount of support) on the source used to align to the target. Finally, **Quiver** visualises only the warping transformation necessary to get from the **Rigid** alignment to the **Warped** alignment.

5 Visualization

The purpose of the alignment process is to decompose the images into the deformation of the material induced by the act of impressing a seal, from the deformation of the motifs depicted on the seal impressions. This decomposition enables a detailed study of each aspect while leveling differences of the other aspect. The following section details visualization modes used in the study of differences between seal impressions. Figure 3 shows one tile for each visualization mode. Figure 6 shows a complete comparison matrix for a single mode.

5.1 Binarized Images

All comparison modes make use of overlaying a warped source image on the target image. To increase the legibility of the overlays, we binarize the seal impression images by a single global threshold. The images are zero-centered by substracting their mean standardized by dividing by their standard deviation. Then, a threshold of 0.3 is applied, all pixels below are set to 0 and all above to 1. This binarization results in well-defined semantics for each pixel value. Pixels valued 0 and 2 indicate agreement on the MSII curvature while pixels valued 1 denote disagreement. Examples are shown in Figure 3.

5.2 Rigid and Non-rigid Alignment

Translation and rotation of the seal impression motif images are artifacts of the acquisition process. These transformations are purely dependent on the position and orientation of the motif in the mold and on the virtual embedding into 3D-model space and orthographically projected raster image space. We are intrested only in the non-rigid contribution of transformations needed to deform the source image onto the target image.

We estimate a rigid transformation model with RANSAC on basis of the correspondences of the regressed non-rigid manifold. Then, we compute the movement vectors in the quiver visualization by sampling source image coordinates and transforming them once using the estimate rigid model and once using the regressed non-rigid manifold, by means of the estimated SVRs. In the quiver visualizations this difference is denoted by arrows pointing from the rigidly estimated target position to the non-rigidly estimated target positions. If there is no difference only a point is shown.

5.3 Support Keypoints

An alignment not matching expected features can result from two qualitatively different reasons: i) the images under comparison genuinely do not share any visual features or ii) common visual features have not been properly detected. The first case is a valuable result for experts while the second needs to be reanalysed. Quantifying and visualizing these is crucial to correctly judge the credibility of our visualizations.

We display the keypoints used for regressing the correspondence manifold on top of the source image as shown in Figure 3. These keypoints indicate which image regions were detected in the source that are also present in the target. Regions with a large count of keypoints denote that the alignment of these visual features can be trusted.

6 Experiments and Results

To validate our approach we computed all pairwise alignments within sets of seal impressions containing the same motif. We used 8 groups with respectively

6, 6, 6, 6, 7, 3, 5 and 7 members in each, for a total of $36 + 36 + 36 + 36 + 49 + 9 + 25 + 49 = 276$ comparisons for each visualization mode. For each set we visualized matrices with pairwise comparisons of: (i) binary overlays of only rigid alignment, (ii) binary overlays of warped alignment, and (iii) quiver plots of the warping transformation.

We review here the archaeological significance of the presented visualizations. Within the 4 groups of sealings from Haghia Triada, we draw attention to the pairs of sealings CMS II.6 no. 70 HMs 455/10 and 455/15 in Figure 4, and CMS II.6 no. 11 HMs 441/20 and 441/05 in Figure 4 and Figure 6 which align more closely than the other pairwise comparisons in each set and are especially good targets for other impressions to align to, suggesting that these pairs of sealings could have been impressed by the same person.

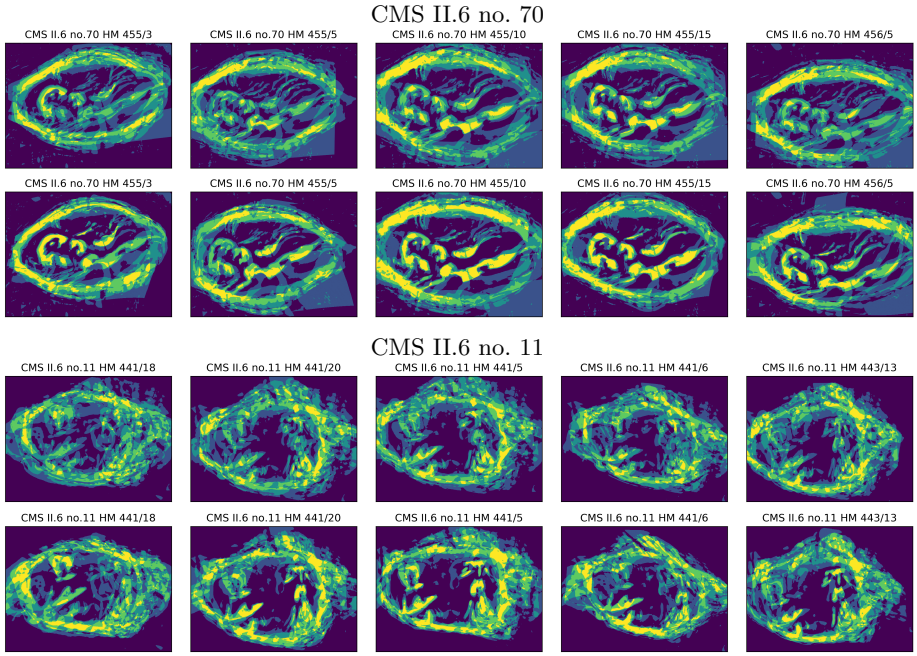

Fig. 4. Images shown are sums of the groups of impressions CMS II.6 no. 70 (top two rows) and CMS II.6 no. 11 (bottom two rows), aligned rigidly (first and third row) and warped (second and forth row), to match their respective target. All images were first binarized; bright yellow color depicts a large agreement between the curvatures of the impressions.

The seal impressions in Figure 5, drawn from the 4 groups of sealings impressed by gold sealing rings, were found at two different sites on Crete, namely Haghia Triadha for CMS II.6 no. 19/HMs 591 and 516 and Sklavokambos for CMS II.6 no. 260/HMs 632-635. While the seal impressions all depict the same

motif, there are visible small differences between each sealing, for example the presence or absence of the chariot reins or the position of the horse's head; the difficulty of explaining the cause of these very small differences had led, in the past, to uncertainty as to whether all the impressions in this set were stamped with the same seal or not.

We analyze how well the impressions can be aligned to each other, and expect a high degree of agreement in the binary overlay visualization if they were stamped by the same seal. Figure 5 shows that the impressions can be aligned well when using warping transformations. All impressions in the set match CMS II.6 no. 19 HM 591 and CMS II.6 no. 260 HM 634 particularly well, indicating that the practice of impressing these seals is similar, which would suggest that the same person was using this seal each time.

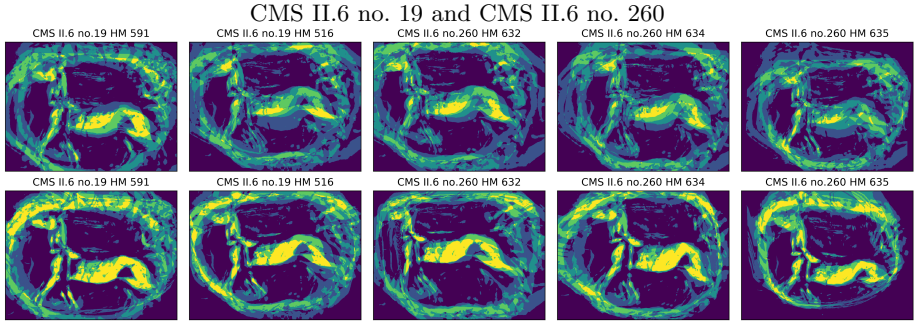

Fig. 5. Images shown are the sums of images in the groups CMS II.6 no. 19 and CMS II.6 no. 260, aligned rigidly at the top and warped at the bottom, to match the respective target image. All images were first binarized; bright yellow color depicts a large agreement between the curvatures of the impressions. HMs 591 and 516 (first two columns) and HMs 632-635 (last three columns) were found at two different sites. The juxtaposition shows that no amount of rigid rotation and translation aligns the motifs, while local deformations enable such an alignment.

7 Conclusion

In this work, we introduced the use of image registration techniques to decompose the differences in pairs of seal impression images into local visual features and global deformation. We employed a feature transform into an assignment space with outlier rejection based on visual feature kernel density and assignment manifold stress. Then, the final alignment functions are regressed with radial-basis function (RBF) support vector regressors (SVR). We tailored our visualizations to meet the needs of archaeological experts: (i) through binarization to highlight motifs, (ii) pairwise overlays for easy comparison, (iii) sum overlays to find representative seal impressions, (iv) and quiver plots to find

patterns in the sealing practice. The result of our process was concrete findings of examples, within sets of sealings impressed by the same seal, of seal impressions that were probably made by the same person, including a group of sealings from two different archaeological sites.

In future work, our method and visualization approach requires more experiments in differing domains to be validated, e.g. similarity and deformation analysis on digitalized Old Egyptian cursive handwriting [5]. In addition, we will investigate the usage of visual features common to all seal impressions in a set to aid in the detection of patterns in damaged seals and to automatically derive a prototypical impression that all seal impressions can be aligned to. Using such an approach will remove the need to manually inspect the resulting pairwise visualizations, growing quadratically with count of impressions in a set, to inspecting how each seal relates to the derived prototype, growing only linearly.

Acknowledgements

This work is partially supported by the *Federal Ministry of Education and Research* (BMBF) eHeritage II programme, grant no. `01UG1880X` for supporting the *3D forensic analysis and contextualisation of aegean seals and sealings* (ErKon3D) project. Furthermore we thank Dr. Maria Anastasiadou for practical help with the CMS collection.

References

1. Bogacz, B., Finlayson, S., Panagiotopoulos, D., Mara, H.: Quantifying deformation in aegean sealing practices. Computer Vision, Imaging and Computer Graphics Theory and Applications (2019)
2. Chui, H., Rangarajan, A.: A new point matching algorithm for non-rigid registration. Computer Vision and Image Understanding (2003)
3. Drucker, H., Burges, C.J.C., Kaufman, L., Smola, A., Vapnik, V.: Support vector regression machines. Neural Information Processing Systems (1996)
4. Fischler, M.A., Bolles, R.C.: Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography. SRI International (1980)
5. Gülden, S., Krause, C., Verhoeven, U.: Digital palaeography of hieratic. The Oxford Handbook of Egyptian Epigraphy and Palaeography (2020)
6. Li, X., Hu, Z.: Rejecting mismatches by correspondence function. International Journal of Computer Vision (2010)
7. Mara, H., Krömker, S.: Vectorization of 3D-Characters by Integral Invariant Filtering of High-Resolution Triangular Meshes. International Conference on Document Analysis and Recognition (2013)
8. Mead, A.: Review of the development of multidimensional scaling methods. Journal of the Royal Statistical Society (1992)
9. Papadaki, A.I., Hansch, R.: Match or No Match: Keypoint Filtering based on Matching Probability. Image Matching: Local Features & Beyond, CVPR Workshop (2020)

10. Tola, E., Lepetit, V., Fua, P.: DAISY: An Efficient Dense Descriptor Applied to Wide-Baseline Stereo. Pattern Analysis and Machine Intelligence (2010)
11. Tran, Q.H., Chin, T.J., adn M. S. Brown, G.C., Suter, D.: In Defence of RANSAC for Outlier Rejection in Deformable Registration. European Conference on Computer Vision (2012)
12. Verdie, Y., Yi, K.M., Fua, P., Lepetit, V.: TILDE: A Temporally Invariant Learned DEtector. Conference on Computer Vision and Pattern Recognition (2015)

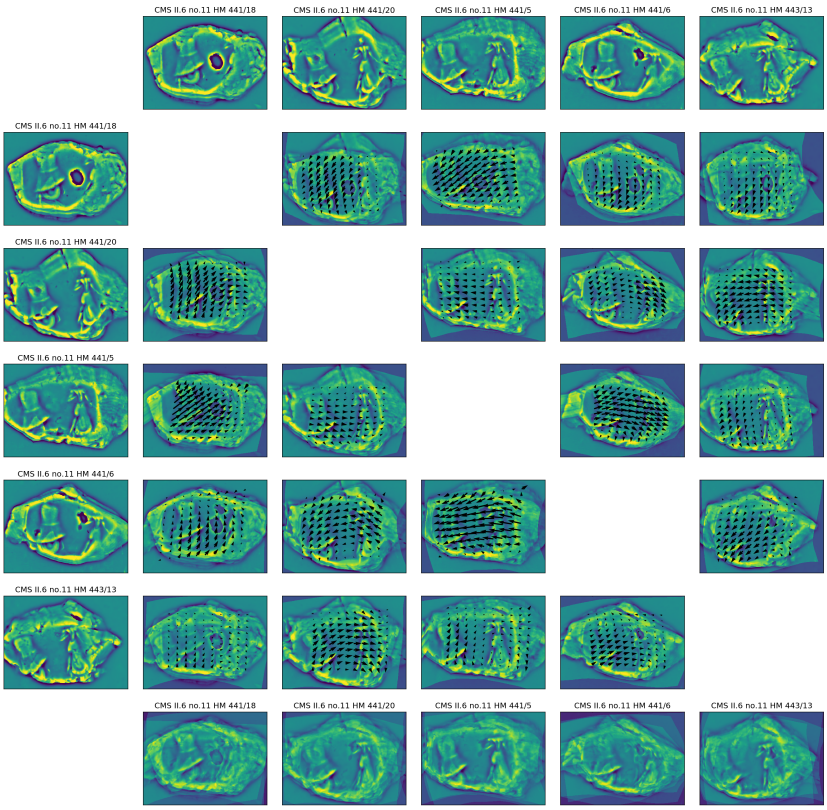

Fig. 6. Comparison matrix of a part of the CMS II.6 no 11 seal impression set. Top row and left column depict the impressions as rasterized MSII images being compared. The images within depict an additive overlay of the curvature images with a quiver overlaid to show the warping necessary to align the images. The bottom row depicts accumulated overlays of all warped source images on the respective target. CMS II.6 no. 11 HM 441/20 and HM 441/05 need only minimal warping to align visually very well, suggesting that these pairs of sealings could have been impressed by the same person.