

Drawing the Line: Deep Segmentation for Extracting Art from Ancient Etruscan Mirrors^{*}

Rafael Sterzinger[✉], Simon Brenner[✉], and Robert Sablatnig[✉]

Computer Vision Lab, TU Wien, Vienna, AUT
{firstname.lastname}@tuwien.ac.at

Abstract. Etruscan mirrors constitute a significant category within Etruscan art and, therefore, undergo systematic examinations to obtain insights into ancient times. A crucial aspect of their analysis involves the labor-intensive task of manually tracing engravings from the backside. Additionally, this task is inherently challenging due to the damage these mirrors have sustained, introducing subjectivity into the process. We address these challenges by automating the process through photometric-stereo scanning in conjunction with deep segmentation networks which, however, requires effective usage of the limited data at hand. We accomplish this by incorporating predictions on a per-patch level, and various data augmentations, as well as exploring self-supervised learning. Compared to our baseline, we improve predictive performance w.r.t. the pseudo-F-Measure by around 16%. When assessing performance on complete mirrors against a human baseline, our approach yields quantitative similar performance to a human annotator and significantly outperforms existing binarization methods. With our proposed methodology, we streamline the annotation process, enhance its objectivity, and reduce overall workload, offering a valuable contribution to the examination of these historical artifacts and other non-traditional documents.

Keywords: Binarization · Photometric Stereo · Image Segmentation · Limited Data · Etruscan Art · Cultural Heritage

1 Introduction

Comprising more than 3,000 identified specimens, Etruscan hand mirrors represent one of the most extensive categories in Etruscan art. These ancient artworks, which date from the second half of the sixth century to the first century B.C.E., are predominantly made from bronze. On the front, they feature a highly polished surface acting as a mirror, whereas the back side typically depicts engraved and/or chased linear figurative illustrations of Greek mythology which may be supplemented by inscriptions in Etruscan [9]. An exemplary mirror showcasing such drawings is depicted in Fig. 1.

^{*} We gratefully acknowledge funding from the Austrian Science Fund/Österreichischer Wissenschaftsfonds (FWF) under grant agreement No. P 33721. Additionally, we thank Ms. Sindy Kluge for her time and expertise in annotating the mirrors.



Fig. 1: A typical Etruscan mirror: fine drawings adorn its backside.¹

These mirrors have been systematically examined since the 1980s to better understand the Etruscan culture, with findings published in the scholarly series *Corpus Speculorum Etruscorum* (e.g., [25,30]). One main component of their examination is the labor- and cost-intensive task of tracing the linear figurative artworks typically featured on their backside. Although performed by experts, i.e., archaeologists specialized in Etruscan culture, this task is prone to subjectivity: over time, these mirrors experienced erosion/corrosion, which increases the difficulty of distinguishing between artwork, damage, and background.

Based on these challenges, this work presents the first approach for automated segmentation of Etruscan hand mirrors to enhance annotation quality and objectivity while reducing the workload of manual tracing. For this, we collect high-resolution 3D surface models via photometric stereo and train a deep segmentation network to recognize intentional lines over scratches. Our approach is grounded in a thorough ablation study where we evaluate various surface modalities and explore different techniques to address the scarcity of data and annotations within this domain, including performing predictions on

¹ ©KHM-Museumsverband (Kunsthistorisches Museum Vienna)

a per-patch level, incorporating various data augmentations, and self-supervised learning, as well as examine different architectural components, such as encoders, decoders, and loss functions. Compared to an initial baseline, our modifications improve performance w.r.t. the pseudo-F-Measure by around 16%. Moreover, we juxtapose our method with a human baseline to underscore its efficacy as it achieves results on par with a human annotator.

Summarizing our contributions: our methodology streamlines the overall annotation process of Etruscan hand mirrors; It not only expedites the process of manual tracing but also offers a supplementary perspective, contributing to increased quality and objectivity. Finally, we provide public access to both the code and data utilized in this work (see github.com/RafaelSterzinger/etmira-segmentation) to promote transparency and reproducibility.

2 Related Work

Photometric Stereo (PS) is a powerful tool to capture unique insights into an object’s surface geometry. McGunnigle et al. [15] demonstrate its usefulness in segmentation by extracting handwriting on paper based on depth profiles. Similarly, Landström et al. [10] have applied this technique to detect cracks in steel surfaces, showing its versatile application. Moreover, the work by Zhang et al. [35], where they compute shape index and curvedness from PS for leaf venation extraction, or by Tao et al. [27], in which they detect air voids in concrete surfaces, further exemplify its broad applicability. Its ability to distinguish subtle surface variations makes it particularly suitable for analyzing complex patterns such as those found on Etruscan hand mirrors. However, traditional segmentation methods might fall short due to damage these mirrors experienced, increasing the difficulty of distinguishing between fore- and background.

In the field of image segmentation, the state of the art is dominated by deep learning architectures such as the UNet [19], DeepLabV3++ [3], Pyramid Scene Parsing Network [36], Pyramid Attention Network [11], or the Feature Pyramid Network [13]. A driving force for such innovations is, for instance, the medical sector, which demands high-detailed segmentations as well. One example is the segmentation of vascular structures found in the retina (e.g., see [12]), which is key in determining retinal diseases. In this context, Kamran et al. [6] proposed a multiscale, GAN-based segmentation system, designed to reduce loss in fidelity. More recent works propose transformer-based [28] as well as diffusion-based architectures [32] for brain tumor, gland, or nuclei segmentation.

Saiz et al. [20,21] showcased the effectiveness of combining PS scanning with deep segmentation networks. With this paradigm, they perform defect detection in steel component manufacturing, paving the way for future applications where intricate and detailed segmentation is vital. Lastly, to give an example in the context of cultural heritage, Redon et al. [18] present a deep learning-based reconstruction pipeline that transforms RGB images into high-resolution normal maps to digitize Bayeux tapestry. Additionally, they couple the reconstruction with segmentation, enabling cost-effective and quick production of 3D-printed




Fig. 2: Exemplary albedo, depth, normal map, and corresponding ground truth.

bas-reliefs to enhance tactile exploration for visually impaired individuals. Such examples showcase the potential efficacy of PS when combined with deep learning to segment figurative drawings such as the ones found on Etruscan mirrors.

3 Methodology

In the following, we describe the setup for data acquisition, the captured dataset, and outline our pipeline in detail, focusing on preprocessing and the subsequent segmentation. For the latter, we summarize our explored augmentation/self-supervised learning techniques as well as point out training and inference specifics.

3.1 Data Acquisition

We employ PS, a 3D reconstruction technique introduced by Woodham [31] in 1980. Compared to other 3D acquisition methods, e.g., structured-light scanning or photogrammetry, it is especially adept at capturing local surface details [5]. PS accomplishes this by capturing a series of images with consistent camera settings but variable lighting conditions to then, for each pixel, determine surface orientation, distance, and illumination-independent brightness. These properties can be encoded across the entire image to *normal*, *depth* (computed from the normal map via integration [23]), and *albedo maps*, illustrated in Fig. 2.

For recording the images, we employ a PhaseOne IQ260 Achromatic camera with an $8,964 \times 6,716$ pixel medium-format sensor and a Schneider-Kreuznach 120 mm LS Macro lens. Using this setup, we obtained a surface resolution of 38.5 ppm (~ 978 dpi) for most of our mirrors. However, as imaging distance changed between locations, this value is not consistent throughout the images, ranging from 29.8 ppm to 58.3 ppm. For more information, we refer the reader to Brenner et al. [1] from where we derived our setup.

3.2 Dataset

Our dataset includes a wide variety of Etruscan mirrors preserved in public collections in Austria. Its core comprises 53 specimens located at the Kunsthistorischen Museum (KHM) Vienna and 6 additional ones scattered throughout Austria. We acquired annotations for 19 out of these 59 mirrors, whereby 19 backsides and 10 fronts were annotated, leading to 29 examples in total. Note that engravings mainly decorate the backside where they do not interfere with reflectance. Sometimes, however, they can also be found on the front near the handle or around the border, though with less dense drawings.

We split these labels into training, validation, and test sets, however, given their limited amount, their varying density of engravings per mirror, and their condition, this step is not straightforward: first, mirrors with dense engravings offer a stronger learning signal and are hence more preferable for the training, second, using mirrors of different condition for validation and testing, will cause discrepancies in performance, rendering model selection difficult. Therefore, we select the backsides of three mirrors of different conditions and densities in engravings for evaluation, one mirror from Wels and two from the KHM Vienna. We create non-overlapping patches of size 512×512 pixels, shuffle them, and split them in half to yield our validation and test set, to create splits with similar underlying distributions. Out of the remaining 26 examples, we remove one, representing an outlier, due to its different art style (points instead of lines), leaving us with 25 annotated examples for training.

3.3 Preprocessing

With three modalities at hand, albedo, depth, and normal, it is unclear which ones obtain the best predictive performance. In order to allow for a fair performance evaluation, preprocessing is crucial as it directly influences the result of




Fig. 3: Comparing the effect of preprocessing on the depth modality, i.e., employing a high-pass filter to remove low frequencies and to make lines more visible.

subsequent tasks. For instance, employing the depth modality and not performing any preprocessing renders the segmentation of fine artistic lines impossible; its effect is illustrated in Fig. 3. Here, we filter low frequencies by subtracting a Gaussian-filtered version of the depth map. Afterward, we cap the depth values to $\mu - 3\sigma$ as the minimum and $\mu + 3\sigma$ as the maximum intensity, where μ is the mean and σ is the standard deviation. Regarding the other two modalities, i.e., normal and albedo, we standardize the normal vectors to be of unit length, whereas we keep the latter modality unmodified.

In addition to that, we generate segmentation masks to capture the location of a mirror within a shot using the Segment Anything Model (SAM) [8]. Although these masks are not necessarily required, we employ them for Cross-Pseudo Supervision (CPS) [4] as well as during inference to mask out regions that impossibly contain engravings. Lastly, using the acquired masks, we ignore non-mirror parts for the calculation of per-channel means and standard deviations which are later on used to normalize the input.

As mentioned, to alleviate our lack of annotations, we perform inference on a per-patch level: for validation and testing, we extract non-overlapping, quadratic patches of size 512×512 pixels. Regarding our training data, we pad four pixels, since $6720 \equiv 0 \pmod{2240}$, to the original resolution ($8,964 \times 6,716$ pixels) to extract 25 overlapping tiles of size $2,988 \times 2,240$ pixels using a stride of half the size; tiles, containing no annotation, are discarded. Later, during training, to create varied training data for one epoch, we extract ten patches per tile, all of the same size as our validation/test patches. Finally, all patches are resized to a height and width of 256×256 pixels to reduce model complexity.

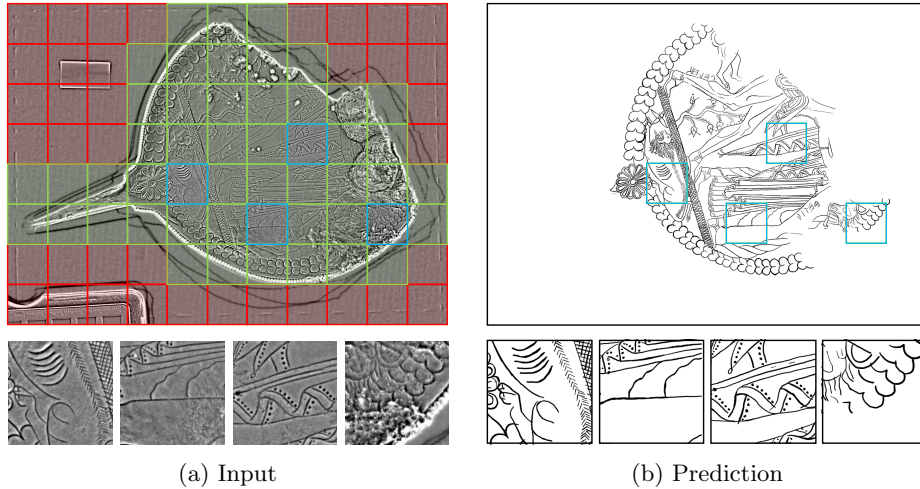


Fig. 4: In order to extract artistic lines from damaged Etruscan hand mirrors, we employ a deep segmentation network on a per-patch level: green denotes patches used during training, red, patches not used during training and blue, ones that are included within the current mini-batch.

3.4 Segmentation

Our proposed architecture consists of an EfficientNet-B6 [26] encoder combined with a UNet [19] decoder, employing augmentations such as rotations, flips, and shifts. A high-level overview of our methodology is depicted in Fig. 4, which illustrates the per-patch level training (512×512 , stride 256; resized to 256×256) that we employ to alleviate our lack of annotations. At inference, predictions are performed for all patches and are recombined to form the complete segmentation mask of a mirror.

Training All our models are trained on an NVIDIA RTX A5000 until convergence, i.e., no improvement $\geq 1e - 3$ w.r.t. the pseudo-F-Measure (pFM; see Equation 1) for ten consecutive epochs, using a batch size of 32 and a learning rate of $3e - 4$. As a loss function, we employ a generalized Dice overlap (Dice Loss; see Sudre et al. [24]), which is well suited for highly unbalanced data, and optimize it using Adam [7]. Additionally, we incorporate a learning rate scheduler that also monitors the pFM on our validation set. If there is no improvement for five consecutive training epochs, the learning rate is halved.

Augmentation In order to enhance the robustness and generalization capability of our deep learning model, we explore augmentation techniques during the training phase. In detail, we experimented with two techniques: first, the Cut-Mix augmentation strategy, proposed by Yun et al. [33]. This technique involves

cutting patches from one training image and pasting them onto another. Second, MixUp, as proposed by Zhang et al. [34], involves training our model on convex combinations of pairs of examples. For this, we randomly sampled interpolation levels between 0.4 and 0.6. Besides these two techniques, we also use *standard augmentations*, including flipping, shifting, and rotating.

Self-Supervised Learning Given the limited amount of annotations and the large amount of unlabelled data, we opt to explore a semi-supervised learning technique. In the works by Chen et al. [4], the authors introduce CPS, a new consistency regularization. CPS enforces consistency by perturbing two networks with different initializations for the same input image, i.e., the pseudo-one-hot segmentation generated by one network is employed to supervise the other network and vice versa. For this, we employ the masks generated by SAM to select only image patches that include mirror parts, and hence potentially include engravings. Lastly, an important hyperparameter is the weighing λ of this regularization term which needs to be tuned.

Inference After training, we need to adjust the methodology for inference to obtain a segmentation not only on a patch level but for whole mirrors. After having acquired masks for individual patches, we recombine them using a custom weight map to avoid visual artifacts such as sharp edges. Finally, we use the masks from SAM to remove False Positives (FP) from non-mirror parts.

4 Evaluation

In this section, we evaluate our design choices. In detail, we consider the effectiveness of employing different modalities, augmenting the input via CutMix, MixUp, and applying standard augmentations, training the model in a semi-supervised manner, and making global decisions such as exchanging the encoder/loss function. Our initial *Baseline*, comprising an EfficientNet-B4 [26] as encoder and a UNet [19] as decoder, trains on all modalities initially optimizing the Binary Cross-Entropy (BCE) loss and does not incorporate augmentations nor semi-supervised learning.

During this evaluation process, we report standard segmentation metrics such as the Intersection-over-Union (IoU), the Sørensen-Dice coefficient (Dice), the F-Measure (FM), as well as the pseudo-F-Measure (pFM). The pFM is a metric commonly used for evaluating the binarization quality of handwritten documents and, hence, well-suited for our binarization task, i.e., a task where shifting the mask by a single pixel will have a major impact on per-pixel metrics such as the IoU. Compared to the traditional FM the pFM relies on an alternative metric for recall, known as the pseudo-Recall (p-Recall), which is calculated based on the skeletonized version of the ground truth mask [17]:

$$\text{pFM} = \frac{2 \times \text{p-Recall} \times \text{Precision}}{\text{p-Recall} + \text{Precision}} \quad (1)$$

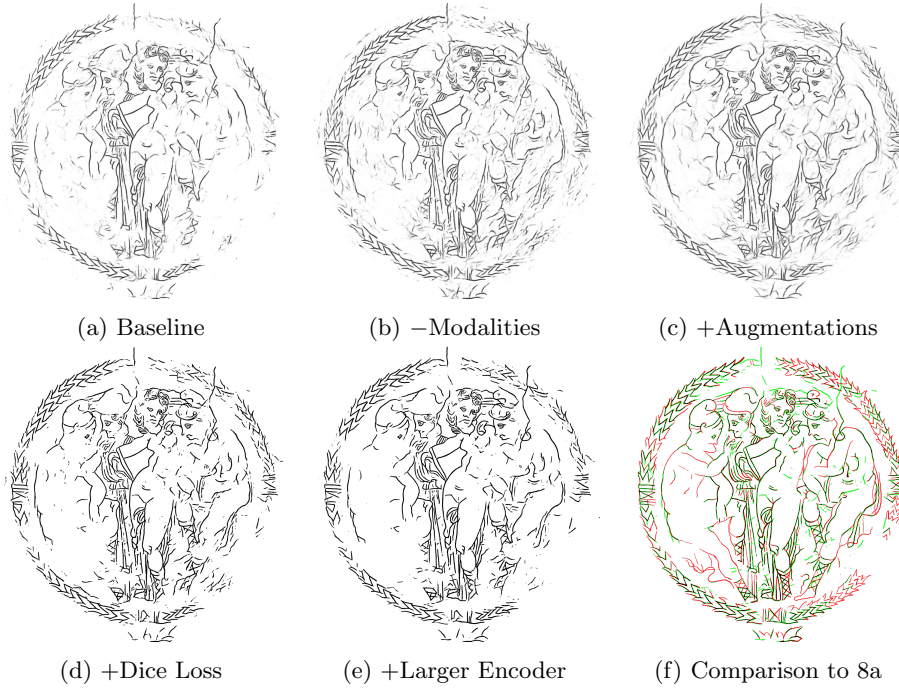


Fig. 5: Showcasing our ablation study (Table 1) on the exemplary mirror “ANSA-VI-1700”, which yielded an IoU of 38.3 and a pFM of 57.6; the union of green and black denotes the prediction, red and black Annotation A (Fig. 8a).

4.1 Ablation Study

Generally, for this section, tables are sorted in ascending order according to the pFM, excluding the first entry which denotes the previous best model, i.e., the model where the currently evaluated modification is not applied. An overview of kept design choices and their impact are denoted quantitatively in Table 1 and qualitatively in Fig. 5. First, we remove the albedo and normal modality to yield a pFM of 43.8. Next, we incorporate augmentations such as flipping, rotating, and shifting, improving pFM performance to 46.2. The most significant change visually is achieved by incorporating the Dice Loss, pushing performance to a pFM of 47.7. Finally, by employing a larger encoder (EfficientNet-B6), compared to our baseline, we improve predictive performance w.r.t. the pFM by around +16%, from 41.7 to 48.3.

Modalities Starting from our baseline, we first considered the usefulness of employing all modalities at hand versus different subsets. For this evaluation, we stack subsets of modalities and evaluate their corresponding predictive performance. The results of this can be viewed in Table 2. Inspecting the model’s performance on these different subsets, we found that training solely on the

Table 1: Overview of our ablation study; higher is better.

Modification	IoU	Dice	FM	pFM
Baseline	25.5	35.5	40.6	41.7
–Albedo/Normal Modality	27.3	38.6	42.8	43.8
+Standard Augmentations	28.8	37.9	44.7	46.2
+Dice Loss	31.0	47.2	47.3	47.7
Final Model (+Larger Encoder)	31.7	48.1	48.1	48.3

depth modality performed significantly better (pFM of 43.8) than when trained on other combinations. For instance, training solely on surface normals achieves a lower pFM of 42.3. However, compared to other variations (all modalities and depth/normal plus albedo), the difference in performance is more noticeable. Hence, we opt for training solely on the depth modality due to its significantly better performance (two-tailed P-value of 0.032 over three runs) and simplicity. An additional benefit is that depth maps are readily suited for standard augmentations as opposed to normal maps.

Table 2: Results on the impact of employing different subsets of modalities during training on the predictive performance; higher is better.

Albedo	Depth	Normal	IoU	Dice	FM	pFM
✓	✓	✓	25.5	35.5	40.6	41.7
✓	-	✓	23.5	32.9	38.0	39.1
✓	✓	-	25.4	34.8	40.5	41.8
-	-	✓	25.4	35.1	40.5	42.3
-	✓	-	27.3	38.6	42.8	43.8

Augmentation Next, we consider the effectiveness of CutMix, MixUp, and standard augmentations applied combined or individually. We observe that augmentation is important in our setting as applying flips, rotations, and/or shifts is beneficial, resulting in an improved pFM of 46.2. This is in contrast to CutMix and MixUp, as their application did not achieve better results, at least w.r.t. the pFM. Considering the Dice score, however, we find that applying MixUp alone yields an improvement of ca. +5%. Hence, we evaluate employing MixUp in combination with standard augmentations which results in similar performance as applying only standard augmentations (two-tailed P-value of 0.6415 over three runs). For more detailed results, we refer to Table 3.

Table 3: Results on the impact of CutMix, MixUp, and standard augmentations on the predictive performance; higher is better.

CutMix	MixUp	Augment	IoU	Dice	FM	pFM
-	-	-	27.3	38.6	42.8	43.8
✓	-	-	26.6	36.8	42.0	42.7
✓	✓	-	27.2	37.6	42.8	43.6
-	✓	-	27.8	40.6	43.5	43.6
-	✓	✓	28.2	39.2	44.0	44.8
✓	✓	✓	28.5	39.9	44.4	45.0
-	-	✓	28.8	37.9	44.7	46.2

Architecture We continue with revising our choice in the loss function, testing replacing the BCE loss with a Focal loss [14], Dice loss [24], as well as a the multi-loss, consisting of all three (see Wazir and Fraz [29]). Here, we found the Dice loss to improve performance by around +3%, achieving a pFM of 47.7, qualitatively significantly improving appearance. Results of this experiment are denoted in Table 4. Furthermore, we consider different encoders/decoders for our architecture: for the encoder, we upgraded to an EfficientNet-B6 (pFM of 48.3; considered variants from B5 up to B8) but kept the decoder the same, as no performance improvements have been observed. Information on decoders tested and their corresponding performance are detailed in Table 5.

Table 4: Overview of different loss functions tested and their corresponding impact on the predictive performance; higher is better.

BCE-Loss	Focal Loss [14]	Dice-Loss [24]	IoU	Dice	FM	pFM
✓	-	-	28.8	37.9	44.7	46.2
-	✓	-	13.6	1.7	23.9	26.5
✓	✓	✓	30.6	44.7	47.0	47.5
-	-	✓	31.0	47.2	47.3	47.7

Cross-Pseudo Supervision Finally, to potentially make use of the large portion of unlabelled data, we explore semi-supervised learning via CPS. For this, we evaluate different weightings $\lambda \in \{0, 0.3, 0.5, 1\}$ for the CPS loss term: employing a weight of $\lambda = 0.5$ achieves a pFM of 48.1. Compared to this, using a weight of $\lambda = 0.3$ is slightly worse, achieving a pFM of 47.7, and $\lambda = 1$ yields a pFM of 44.8. However, the best performance is achieved by not employing CPS, as this obtains a pFM of 48.3. For this reason, we discard this option and

Table 5: Overview of different decoders evaluated and their corresponding impact on the predictive performance; higher is better.

Decoders	IoU	Dice	FM	pFM
UNet [19]	28.8	37.9	44.7	46.2
Pyramid Scene Parsing Network [36]	19.9	21.5	37.3	32.5
Pyramid Attention Network [11]	22.9	27.5	37.8	38.4
DeepLabV3++ [3]	24.9	33.3	39.9	40.2
Feature Pyramid Network [13]	24.1	29.6	38.8	40.5
LinkNet [2]	26.8	35.4	42.3	42.8
UNet++ [37]	27.9	36.6	43.7	45.4

mention these results only because we consider them valuable to others faced with similar tasks, e.g., binarization of intricate handwriting or calligraphy.

On a final note, we want to point out that, although our explored augmentation techniques, i.e., CutMix and MixUp, as well as CPS, did not improve performance, we still recommend them in settings where the normal modality, a modality where standard augmentations are not readily available, outperforms the depth modality. In experiments different to these, when we trained solely on surface normals, employing CutMix in conjunction with MixUp improved the IoU by ca. +4% and adding CPS on top gained ca. +2%.




Fig. 6: Comparing predictive performance on two whole mirrors: performance is heavily influenced by a mirror’s condition, noticeable by the difference in pFM and deviating optimal threshold values. Additionally, the mask acquired from SAM [8] aids in removing unwanted FP from the background.



Fig. 7: Showcasing segmentation quality on mirror “ANSA-VI-1700”.

5 Results

After verifying our design choices, we move away from evaluating the performance of our model on a per-patch level, towards testing it on whole mirrors: we select two mirrors of different conditions from our validation/test set and evaluate our model’s performance on them as well as consider the influence of the thresholding value. On top of this, we compare the impact of having a segmentation mask of the mirror object, in other words, a mask to remove FP from the background (non-mirror parts), onto the pFM.

Inspecting Fig. 6, we observe that the pFM is heavily dependent on a mirror’s condition: considering mirror “Wels-11944”, the pFM is significantly higher than for the more damaged mirror “ANSA-VI-1701”. As discussed in the beginning, this is expected as extracting drawings becomes ever more susceptible to subjectivity, depending on the level of damage a mirror has experienced. We also note the significant differences in optimal threshold values, showcasing the difficulty of finding one to fit them all; however, the difference in performance is negligible, ranging only between a difference in pFM of 1 and 1.5. Moreover, the use of a constant, naive threshold of 0.5 further narrows this range, suggesting skipping tuning altogether.

Next, under inspection of mirror “ANSA-VI-1700”, we assess the predictive performance qualitatively. Fig. 7 illustrates the mirror with the prediction overlaid. Here, we note three areas to discuss in more detail: first, at the left bottom




Fig. 8: An illustration of potential subjectivity when tracing mirror “ANSA-VI-1700”: the union of green and black denotes Annotation A, red and black Annotation B. With an IoU of ca. 37.0 and a pFM of ca. 56.8, creating ground truth annotations is clearly prone to subjectivity.

lies an area where the interpretation of lines is subject to debate. In spite of this ambiguity, our model maintains objectivity, extracting only the lines that are clearly discernible. Moving to the top, we encounter an area where the model struggles to distinguish between damage and deliberate lines, mistaking surface cracks for intentional engravings. Lastly, in the bottom right corner, the model overlooks unmistakable decorations of the border, failing to identify them as engravings, presumably attributable to the border’s sloping surface, which contrasts the overall flatness of the mirror disk.

5.1 Human Baseline

Finally, we evaluate our model against a human baseline. For this, we selected a mirror to be drawn twice to better understand the extent of subjectivity in annotations when an artwork is traced. Given that some of our mirrors are in pristine condition, we select a heavier damaged one, i.e., mirror “ANSA-VI-1700” from the KHM Vienna, such that the segmentation task is more prone to misinterpretation. Although the same annotator traced this mirror twice, we kept four months in between to reduce the influence of the first tracing on the second. With an IoU of around 37.0 and a pFM of 56.8, Fig. 8 illustrates the large room for interpretation when annotating engravings from Etruscan mirrors.

For comparison to these results, we evaluate the predictive performance of our model and contrast it with one local (Sauvola [22]) and one global (Otsu [16]) existing binarization method; results are denoted in Table 6. With a pFM of around 57.5/53.6, depending on the considered annotation, our proposed methodology scores better/slightly worse, achieving performance comparable to a human annotator. Furthermore, it clearly sets itself apart from existing binarization methods, performing ca. 76%/74% better.

Table 6: Comparing our model against a human baseline: our model is better/slightly worse depending on which mask is considered (Annotation A or B; see Fig. 8), achieving performance comparable to a human annotator. Furthermore, it performs significantly better than existing binarization methods.

Annotation	IoU		pFM	
	A	B	A	B
Human Baseline	37.03		56.78	
Otsu [16]	12.84	14.37	22.79	25.14
Sauvola [22]	19.08	17.46	32.72	30.76
Our Model	37.95	33.60	57.53	53.58

6 Conclusion

In conclusion, our work presents a pioneering effort in automating the segmentation of Etruscan hand mirrors, aiming to enhance objectivity, improve annotation quality, and reduce the associated workload. By leveraging photometric stereo alongside deep segmentation networks, we have tackled the labor-intensive task of manually tracing engravings from the backside of these historical artifacts. After conducting thorough experimentation and refining our methodology, including data augmentation, semi-supervised learning, as well as general architectural changes such as different loss strategies, we have surpassed our initial baseline by approximately +16% in terms of the pseudo-F-Measure. Moreover, compared to a human baseline, our approach demonstrates performance similar to a human annotator and also outperforms existing binarization methods significantly. By mitigating the inherent subjectivity introduced by damage on these mirrors, our methodology not only streamlines the annotation process by drastically reducing workload but also improves the quality and objectivity of annotations by offering a supplementary perspective, overall contributing to the examination process of Etruscan hand mirrors.

In summary, our work offers a practical solution to the challenges associated with analyzing these historic artifacts, representing advancements in the application of advanced imaging techniques and deep segmentation in the context of non-traditional documents and cultural heritage in general. For future work, we aim to further improve objectivity by addressing the subjectivity of underlying annotations by creating a golden standard, i.e., a set of annotations that experts generally agree on. Another avenue we want to explore is the idea of collaborative annotation, where our model works in tandem with human annotators to iteratively refine annotations, i.e., suggestions provided by humans serve as valuable priors to update the posterior of our model, to further accelerate and qualitatively enhance the annotation process.

References

1. Brenner, S., Frühwirth, T., Mayer, S.: Revealing ‘invisible’ poetry by w. h. auden through computer vision: Using photometric stereo to visualize indented impressions. *Digital Scholarship in the Humanities* **38**(4), 1404–1417 (Jun 2023)
2. Chaurasia, A., Culurciello, E.: Linknet: Exploiting encoder representations for efficient semantic segmentation. In: 2017 IEEE Visual Communications and Image Processing (VCIP). IEEE (Dec 2017)
3. Chen, L.C., Zhu, Y., Papandreou, G., Schroff, F., Adam, H.: Encoder-Decoder with Atrous Separable Convolution for Semantic Image Segmentation, p. 833–851. Springer International Publishing (2018)
4. Chen, X., Yuan, Y., Zeng, G., Wang, J.: Semi-supervised semantic segmentation with cross pseudo supervision. In: 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). IEEE (Jun 2021)
5. Herbort, S., Wöhler, C.: An introduction to image-based 3d surface reconstruction and a survey of photometric stereo methods. *3D Research* **2**(3) (Sep 2011)
6. Kamran, S.A., Hossain, K.F., Tavakkoli, A., Zuckerbrod, S.L., Sanders, K.M., Baker, S.A.: RV-GAN: Segmenting Retinal Vascular Structure in Fundus Photographs Using a Novel Multi-scale Generative Adversarial Network, p. 34–44. Springer International Publishing (2021)
7. Kingma, D., Ba, J.: Adam: A method for stochastic optimization. In: International Conference on Learning Representations (ICLR). San Diego, CA, USA (2015)
8. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., Dollar, P., Girshick, R.: Segment anything. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 4015–4026 (October 2023)
9. Kluge, S.: Through the Looking-Glass: der etruskische Spiegel VI 2627 aus dem Bestand des Kunsthistorischen Museums als Fallbeispiel für die rituelle Manipulation und Defunktionalisierung im Kontext der Grablege in der etruskischen Kultur. In: Weidinger, A., Leskovar, J. (eds.) *Interpretierte Eisenzeiten. Fallstudien, Methoden, Theorie. Tagungsbeiträge der 10. Linzer Gespräche zur interpretativen Eisenzeitarchäologie*, pp. 241–254. Studien zur Kulturgeschichte von Oberösterreich, Folge 55, Linz (2024)
10. Landstrom, A., Thurley, M.J., Jonsson, H.: Sub-millimeter crack detection in casted steel using color photometric stereo. In: 2013 International Conference on Digital Image Computing: Techniques and Applications (DICTA). IEEE (Nov 2013)
11. Li, H., Xiong, P., An, J., Wang, L.: Pyramid attention network for semantic segmentation. In: British Machine Vision Conference 2018, BMVC 2018, Newcastle, UK, September 3-6, 2018. p. 285. BMVA Press (2018)
12. Li, L., Verma, M., Nakashima, Y., Nagahara, H., Kawasaki, R.: Iternet: Retinal image segmentation utilizing structural redundancy in vessel networks. In: 2020 IEEE Winter Conference on Applications of Computer Vision (WACV). IEEE (Mar 2020)
13. Lin, T.Y., Dollar, P., Girshick, R., He, K., Hariharan, B., Belongie, S.: Feature pyramid networks for object detection. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). IEEE (Jul 2017)
14. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollar, P.: Focal loss for dense object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **42**(2), 318–327 (Feb 2020)

15. McGunnigle, G., Chantler, M.: Resolving handwriting from background printing using photometric stereo. *Pattern Recognition* **36**(8), 1869–1879 (Aug 2003)
16. Otsu, N.: A threshold selection method from gray-level histograms. *IEEE Transactions on Systems, Man, and Cybernetics* **9**(1), 62–66 (Jan 1979)
17. Pratikakis, I., Gatos, B., Ntirogiannis, K.: Icfhr 2012 competition on handwritten document image binarization (h-dibco 2012). In: 2012 International Conference on Frontiers in Handwriting Recognition. IEEE (Sep 2012)
18. Redon, M., Pizenberg, M., Quéau, Y., Elmoataz, A.: 3d surface approximation of the entire bayeux tapestry for improved pedagogical access. In: 2023 IEEE/CVF International Conference on Computer Vision Workshops (ICCVW). IEEE (Oct 2023)
19. Ronneberger, O., Fischer, P., Brox, T.: U-Net: Convolutional Networks for Biomedical Image Segmentation, p. 234–241. Springer International Publishing (2015)
20. Saiz, F.A., Barandiaran, I., Arbelaz, A., Graña, M.: Photometric stereo-based defect detection system for steel components manufacturing using a deep segmentation network. *Sensors* **22**(3), 882 (Jan 2022)
21. Saiz, F.A., Alfaro, G., Barandiaran, I., Graña, M.: Generative adversarial networks to improve the robustness of visual defect segmentation by semantic networks in manufacturing components. *Applied Sciences* **11**(14), 6368 (Jul 2021)
22. Sauvola, J., Seppanen, T., Haapakoski, S., Pietikainen, M.: Adaptive document binarization. In: Proceedings of the Fourth International Conference on Document Analysis and Recognition. ICDAR-97, vol. 1, pp. 147–152 vol.1. IEEE Comput. Soc (1997)
23. Simchony, T., Chellappa, R., Shao, M.: Direct analytical methods for solving poisson equations in computer vision problems. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **12**(5), 435–446 (1990)
24. Sudre, C.H., Li, W., Vercauteren, T., Ourselin, S., Jorge Cardoso, M.: Generalised Dice Overlap as a Deep Learning Loss Function for Highly Unbalanced Segmentations, p. 240–248. Springer International Publishing (2017)
25. Swaddling, J., Nicholls, R., Rasmussen, T., Museum, F.: *Corpus Speculorum Etruscorum: Great Britain. No. Band 1, Teil 1 in Corpus Speculorum Etruscorum: Great Britain*, Fitzwilliam Museum (1993)
26. Tan, M., Le, Q.V.: Efficientnet: Rethinking model scaling for convolutional neural networks. In: Chaudhuri, K., Salakhutdinov, R. (eds.) Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9–15 June 2019, Long Beach, California, USA. Proceedings of Machine Learning Research, vol. 97, pp. 6105–6114. PMLR (2019)
27. Tao, J., Gong, H., Wang, F., Luo, X., Qiu, X., Huang, Y.: Automated image segmentation of air voids in hardened concrete surface using photometric stereo method. *International Journal of Pavement Engineering* **23**(14), 5168–5185 (Nov 2021)
28. Valanarasu, J.M.J., Oza, P., Hacihaliloglu, I., Patel, V.M.: Medical Transformer: Gated Axial-Attention for Medical Image Segmentation, p. 36–46. Springer International Publishing (2021)
29. Wazir, S., Fraz, M.M.: Histoseg: Quick attention with multi-loss function for multi-structure segmentation in digital histology images. 2022 12th International Conference on Pattern Recognition Systems (ICPRS) pp. 1–7 (2022)
30. Wiman, I.: *Corpus Speculorum Etruscorum: Norway and Sweden. Corpus Speculorum Etruscorum. Norway-Sweden*, L’Erma di Bretschneider (2018)
31. Woodham, R.J.: Photometric method for determining surface orientation from multiple images. *Optical Engineering* **19**(1) (Feb 1980)

32. Wu, J., Fu, R., Fang, H., Zhang, Y., Yang, Y., Xiong, H., Liu, H., Xu, Y.: Medsegdiff: Medical image segmentation with diffusion probabilistic model. In: Oguz, I., Noble, J.H., Li, X., Styner, M., Baumgartner, C., Rusu, M., Heimann, T., Kontos, D., Landman, B.A., Dawant, B.M. (eds.) *Medical Imaging with Deep Learning, MIDL 2023*, 10-12 July 2023, Nashville, TN, USA. *Proceedings of Machine Learning Research*, vol. 227, pp. 1623–1639. PMLR (2023)
33. Yun, S., Han, D., Chun, S., Oh, S.J., Yoo, Y., Choe, J.: Cutmix: Regularization strategy to train strong classifiers with localizable features. In: *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*. IEEE (Oct 2019)
34. Zhang, H., Cissé, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: Beyond empirical risk minimization. In: *6th International Conference on Learning Representations, ICLR 2018*, Vancouver, BC, Canada, April 30 - May 3, 2018, *Conference Track Proceedings* (2018)
35. Zhang, W., Hansen, M.F., Smith, M., Smith, L., Grieve, B.: Photometric stereo for three-dimensional leaf venation extraction. *Computers in Industry* **98**, 56–67 (Jun 2018)
36. Zhao, H., Shi, J., Qi, X., Wang, X., Jia, J.: Pyramid scene parsing network. In: *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE (Jul 2017)
37. Zhou, Z., Siddiquee, M.M.R., Tajbakhsh, N., Liang, J.: Unet++: Redesigning skip connections to exploit multiscale features in image segmentation. *IEEE Transactions on Medical Imaging* **39**(6), 1856–1867 (Jun 2020)