Drawing the Line: Deep Segmentation for Extracting Art from Ancient Etruscan Mirrors

By Rafael Sterzinger, Simon Brenner, and Robert Sablatnig

INTRODUCTION

- **Overview**
  - 3,000+ specimens; ~60 in AUT
  - 6th century to 1st century BCE
  - Published in *Corpus Speculormu Etrscorum*
- **Features**
  - Front: Polished mirror
  - Back: Engraved Greek mythology/Etruscan inscriptions
- **Challenges:**
  - Labor-intensive tracing
  - Damage leads to subjectivity

METHODOLOGY

- **Architecture**
  - EffcientNet-B6 encoder, UNet decoder
- **Training:**
  - Patch-level training: 512x512, resized to 256x256 pixels
  - Generalized Dice Loss [1]
- **Inference:**
  - Custom weight map to recombine patches
  - SAM [2] to remove false positives from the background

RESULTS

- **Human Baseline**
  - Annotator traced mirror twice, four months apart
  - Achieves a pFM of 56.8, highlighting problem of subjectivity
- **Our Model:**
  - Performs comparably to human annotator
  - Outperforms traditional methods Otsu and Sauvola drastically
  - Performance gains from threshold adjustments are negligible

A typical Etruscan mirror: fine drawings of Greek mythology or Etruscan inscriptions adorn its backside

Data Acquisition

- 3D Reconstruction via **Photometric Stereo**
  - Captures local surface details
  - Uses multiple images with consistent camera settings but varying lighting
  - **Produces normal, depth, and albedo maps**
- Imaging Equipment
  - Camera: Phase One IQ260 Achromatic, with Sensor: 8,964x6,716 pixel medium-format
  - Lens: Schneider-Kreuznach 120mm LS Macro

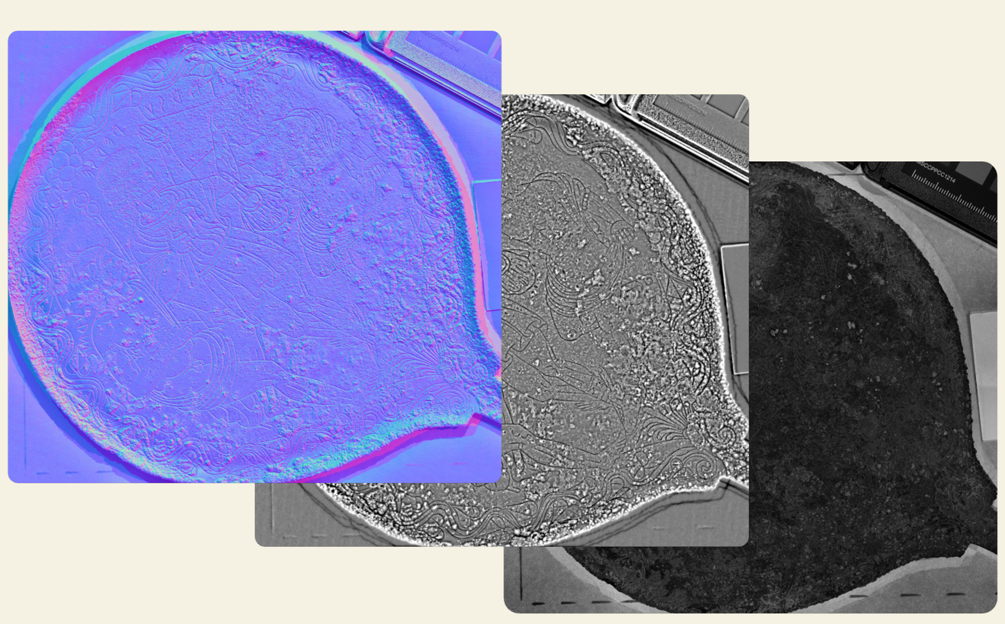

Exemplary normal, depth, and albedo map obtained via photometric stereo

Dataset

- Collection
  - **59 Etruscan mirrors from Austria:** 53 from Kunsthistorisches Museum Wien, 6 from other locations
  - Total: 29 examples annotated; 19 backsides & 10 fronts
- Datasplit
  - Difficult due to limited data, engraving density, varying conditions
  - **25 examples for training, 3 for validation/testing,** 1 outlier
  - Validation/Testing: mirrors of different conditions and engraving density; **create non-overlapping patches, shuffle, and split into validation and test sets**

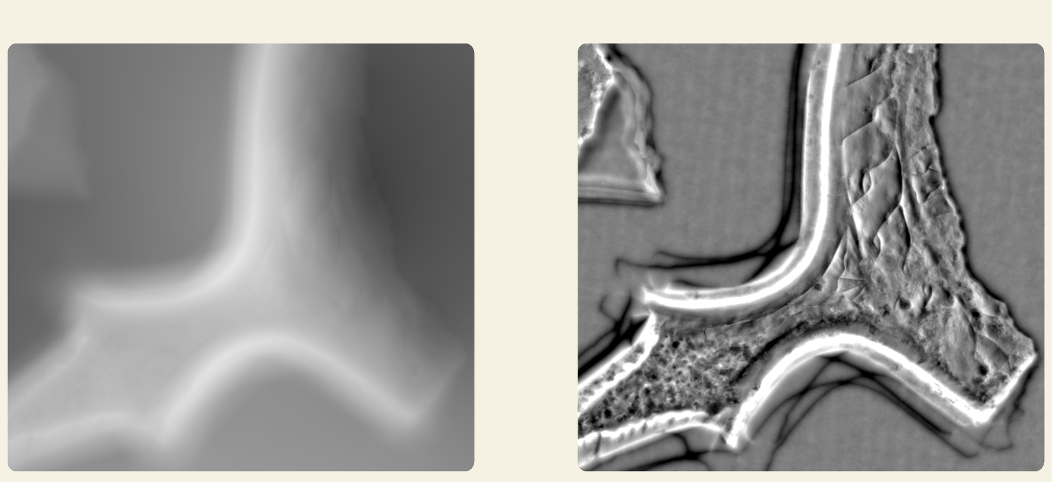

Comparing the effect of preprocessing: employing a high-pass filter to make lines more visible

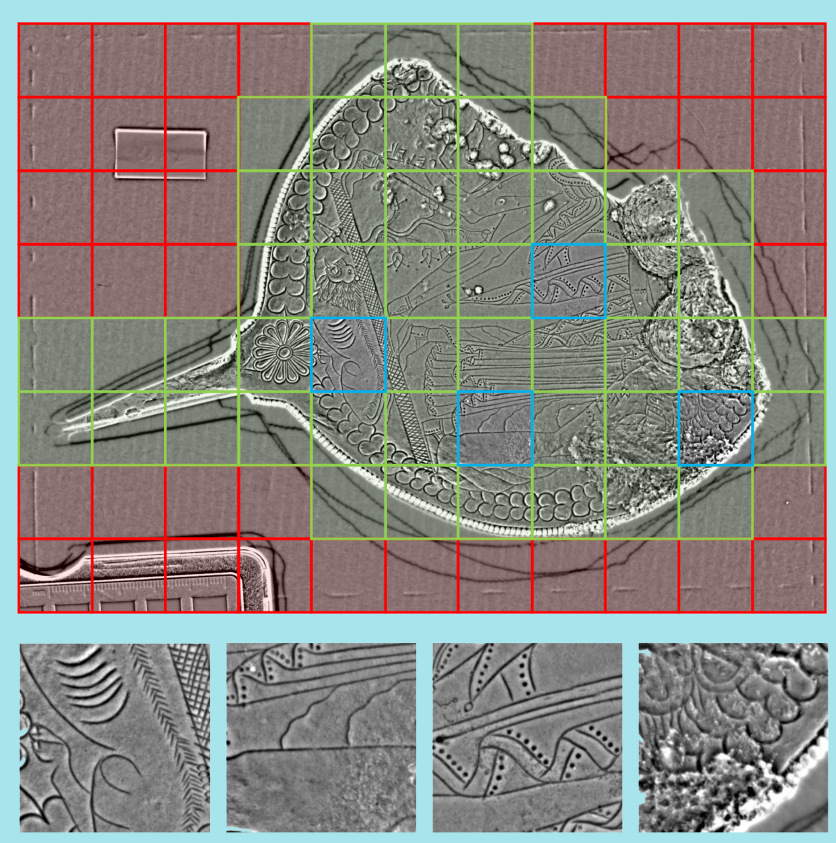

To circumvent data scarcity, we perform training on patches, where only mirror parts are used

Ablation Study

We measure performance using the **pseudo-F-Measure**

$$\text{pFM} = \frac{2 \times \text{p-Recall} \times \text{Precision}}{\text{p-Recall} + \text{Precision}}$$

a metric, less affected by small deviations on pixel-level. It uses pseudo-Recall (p-Recall), which is based on the skeletonized version of the ground truth.

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**

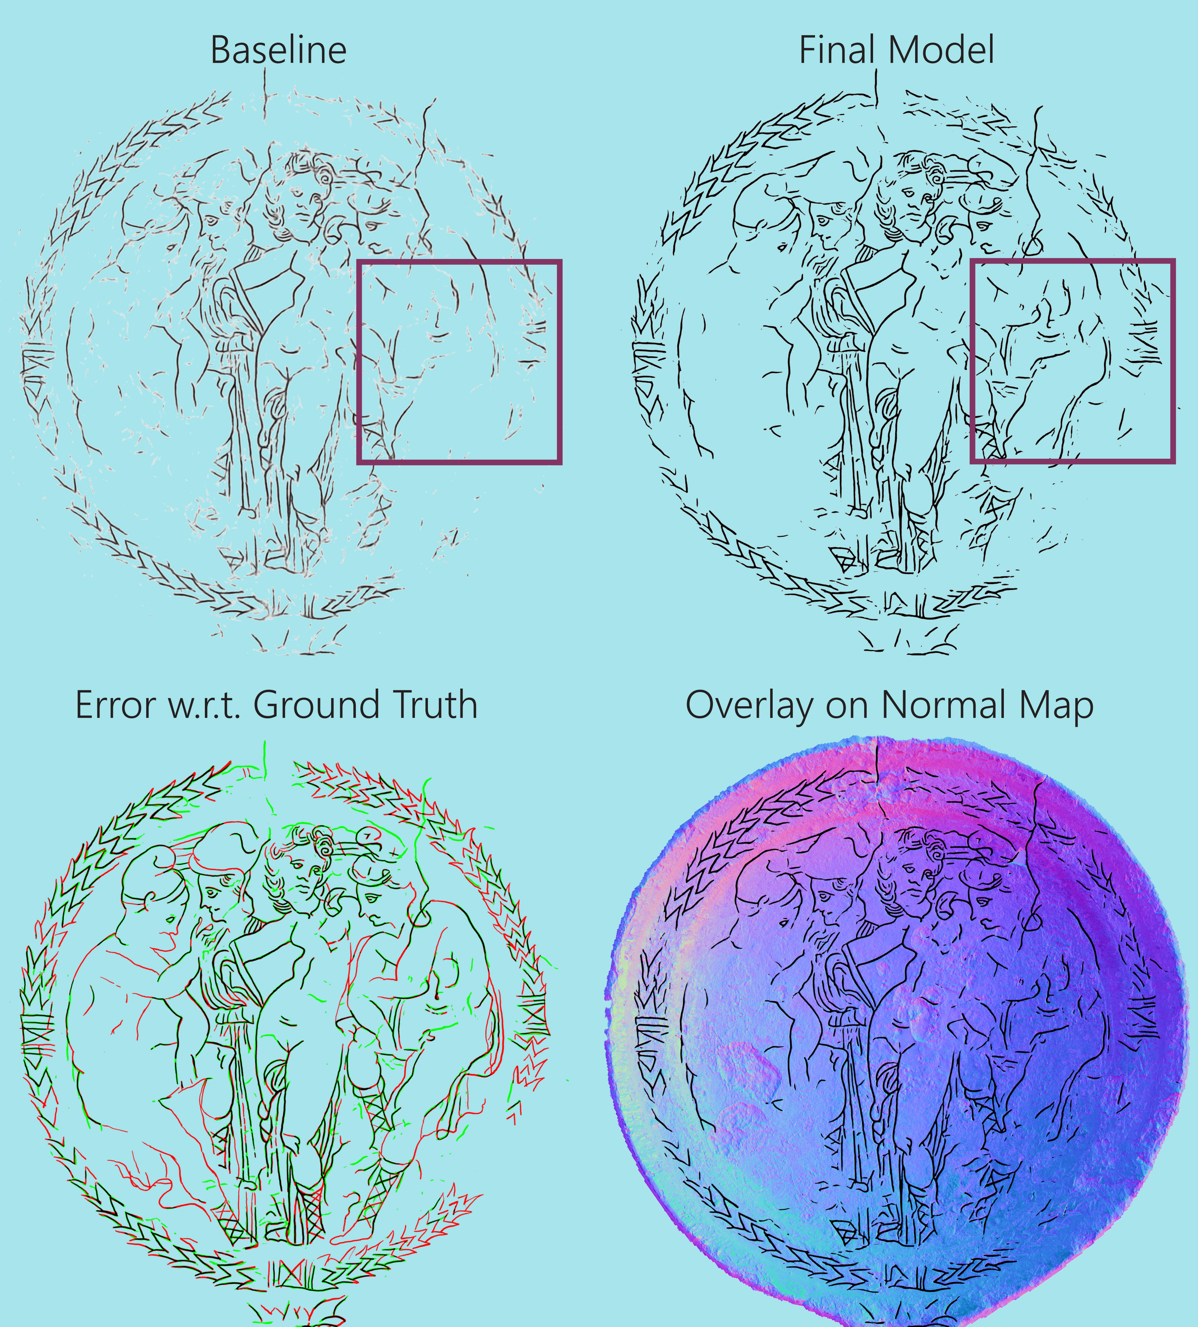

Predictive Performance

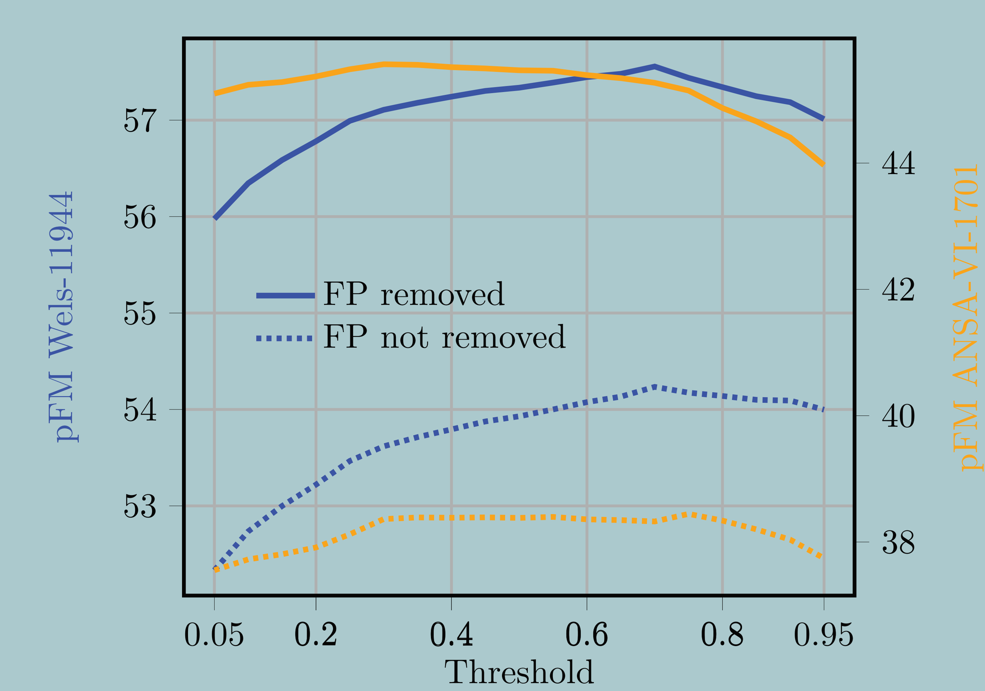

Performance varies with mirror conditions, evidenced by differences in pFM and optimal threshold values. Tuning the threshold shows minimal impact.

Human Baseline

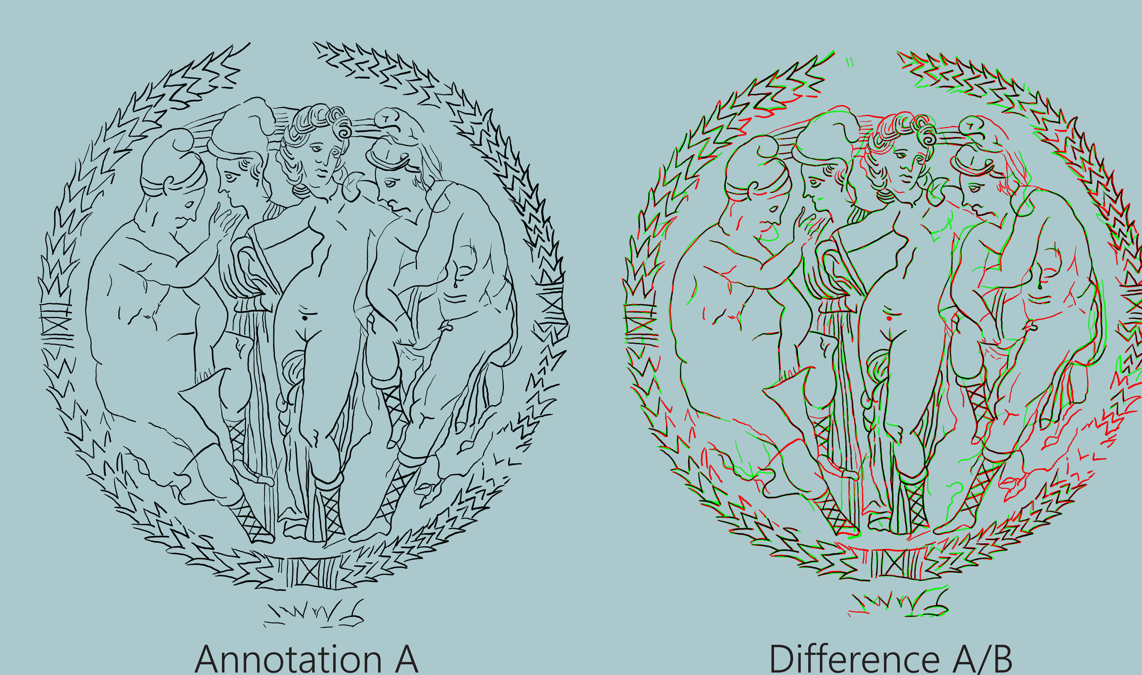

Tracing Etruscan mirrors is **prone to subjectivity,** as evidenced here. Note that Annotation B is not shown.

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

Our model's performance is **comparable to a human annotator,** with pFM scores of 57.5/53.6 depending on the annotation. The model significantly **outperforms existing binarization methods,** with improvements of approximately 76% / 74%.

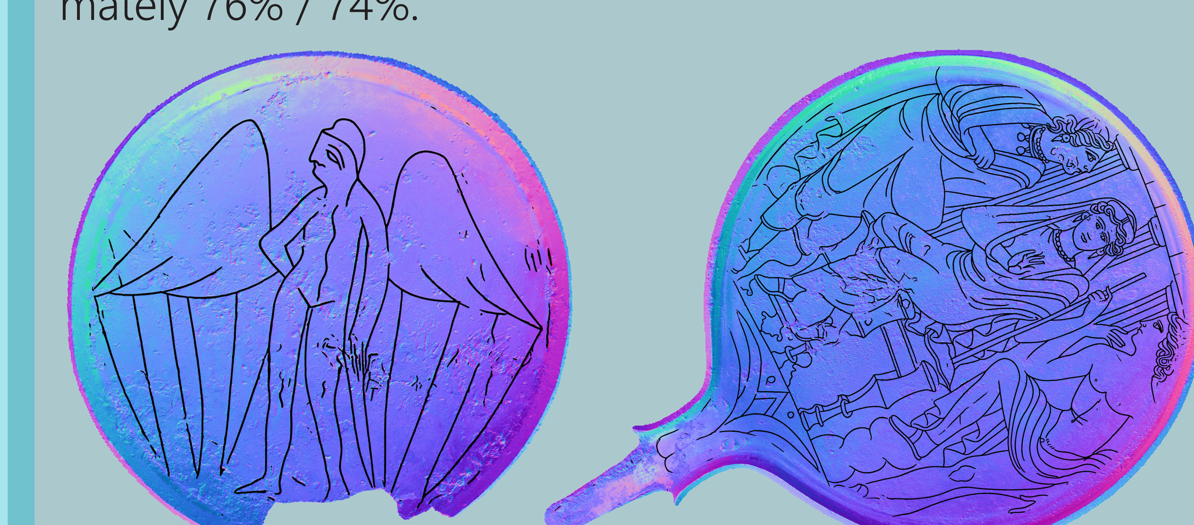

[1] Sudre, C.H., Li, W., Vercauteren, T., Ourselin, S., Jorge Cardoso, M.: Generalised Dice Overlap as a Deep Learning Loss Function for Highly Unbalanced Segmentations. (2017)
[2] Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., Dollar, P., Girshick, R.: Segment Anything. (2023)

Conclusion

Our approach automates the segmentation of Etruscan mirrors, improving annotation quality and reducing workload. With a +16% improvement in pFM over our baseline, our method rivals human performance and surpasses existing binarization techniques.

Paper Code Contact

This poster is presented at the International Conference on Document Analysis and Recognition 2024, ICDAR, Athens 2024.

This project has received funding from the Austrian Science Fund/Österreichischer Wissenschaftsfonds (FWF) under grant agreement No. P 33721.

CVL Computer Vision Lab

TU WIEN TECHNISCHE UNIVERSITÄT WIEN