Semantic Motif Segmentation of Archaeological Fresco Fragments

Università Ca'Foscari Venezia

RePAIR

Aref Enayati*, Luca Palmieri*, Sebastiano Vascon, Marcello Pelillo, Sinem Aslan

*Equal contribution

Motivation:

Locating motifs and predicting their class provide high-level representation of fragments which is a valuable resource for:

- Fragment recognition, style classification and clustering
- Targeted inpainting and restoration of damaged or missing portions of frescoes
- Fresco reconstruction

Grouped fragments according to the motifs on their surfaces

Reconstructed fresco by matching motif types

We explored semantic segmentation of ancient fresco fragments, which were initially restored by eliminating manual annotations on their surfaces through blind inpainting, through two scenarios: 1) *Fragment segmentation to distinguish* image and fragment backgrounds and motif regions, 2) *Semantic motif segmentation* for more extensive exploration of the diverse artistic motifs adorning fresco surfaces.

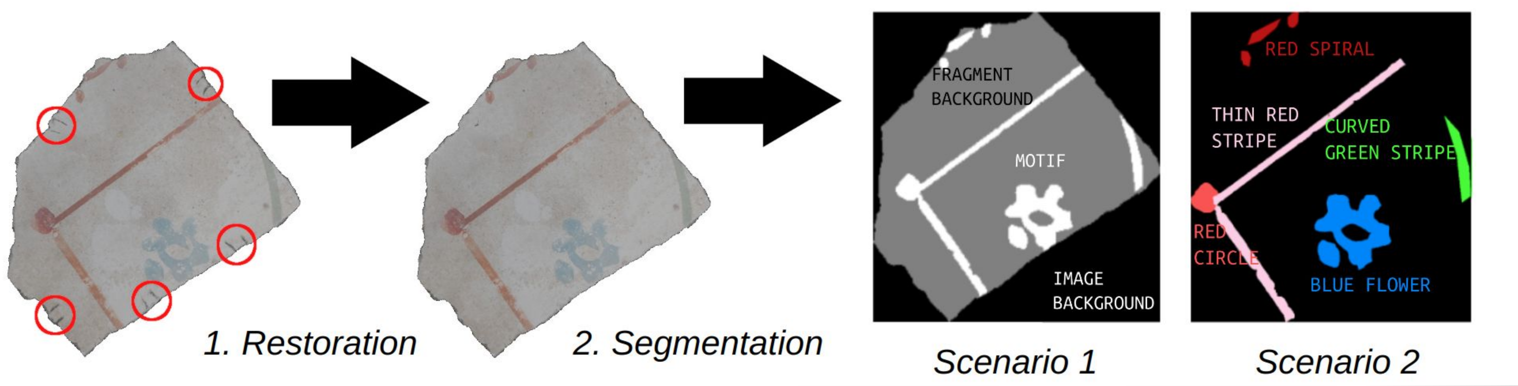

RESTORATION

SEGMENTATION

DATASET

BoFF

Black Marks Annotation on Fresco Fragments

Archaeologists make temporary markings on the fragments to indicate neighborship and aid in manual reconstruction.

The **BoFF dataset** includes 115 fragment images with 405 bounding boxes annotations and is designed for the automatic detection of manual markings to facilitate their removal through inpainting.

Example fragment images from the BoFF dataset, manual annotations indicated within red boxes

MoFF

Motifs Annotation on Fresco Fragments

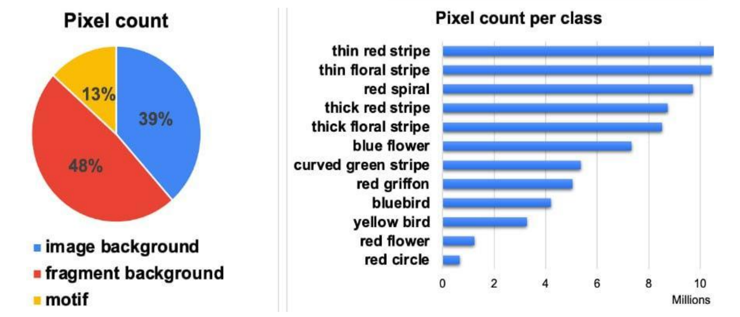

MoFF dataset is curated for motif extraction and categorization from fragmented frescoes. It consists of 405 high-quality images obtained from two separate ceiling frescoes located within the archaeological site of Pompeii. Accurate pixel-level segmentation masks are included, indicating motifs categorized into 12 distinct classes as identified by archaeologists.

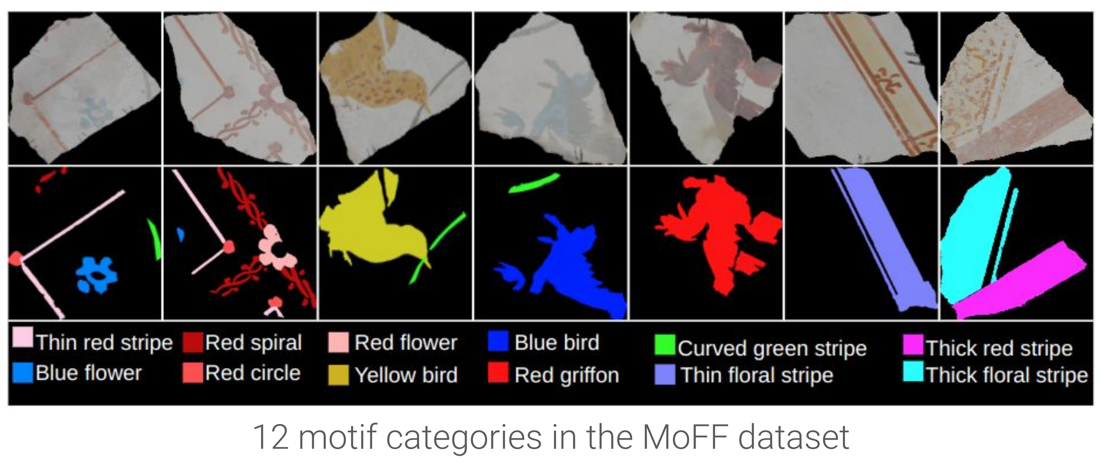

12 motif categories in the MoFF dataset

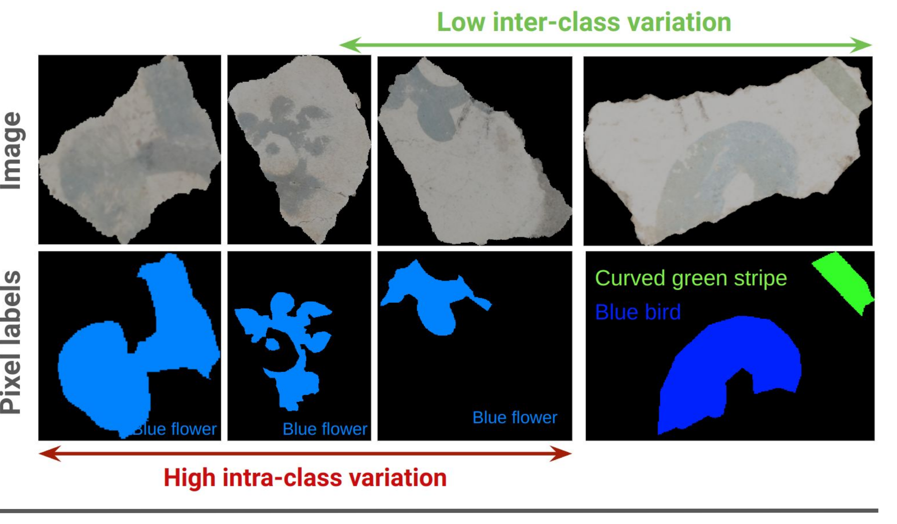

BENCHMARK

Detection of manual annotations

- YOLOv5 model [1] is used as a baseline for **detecting manual annotations** made by the archeologists on the fragments. Train, val, and test sets were split in an 80/10/10 ratio (91, 12, and 12 images in each set, along with 324, 42, and 39 annotated boxes.)
  - The images were resized to 416 × 416 pixels. A pre-trained model from Yolov5 is used for model initialization. Rotation-based data augmentation is used.

Model	Precision	mAp0.5	TP	FP	FN
YOLOv5	0.741	0.596	28	3	11

Restoration through Inpainting

Detected manual annotations in bounding boxes were inpainted to clean the fragment surface from them.

As a baseline, the exemplar-based inpainting method of Criminisi [2] is used in two iterations.

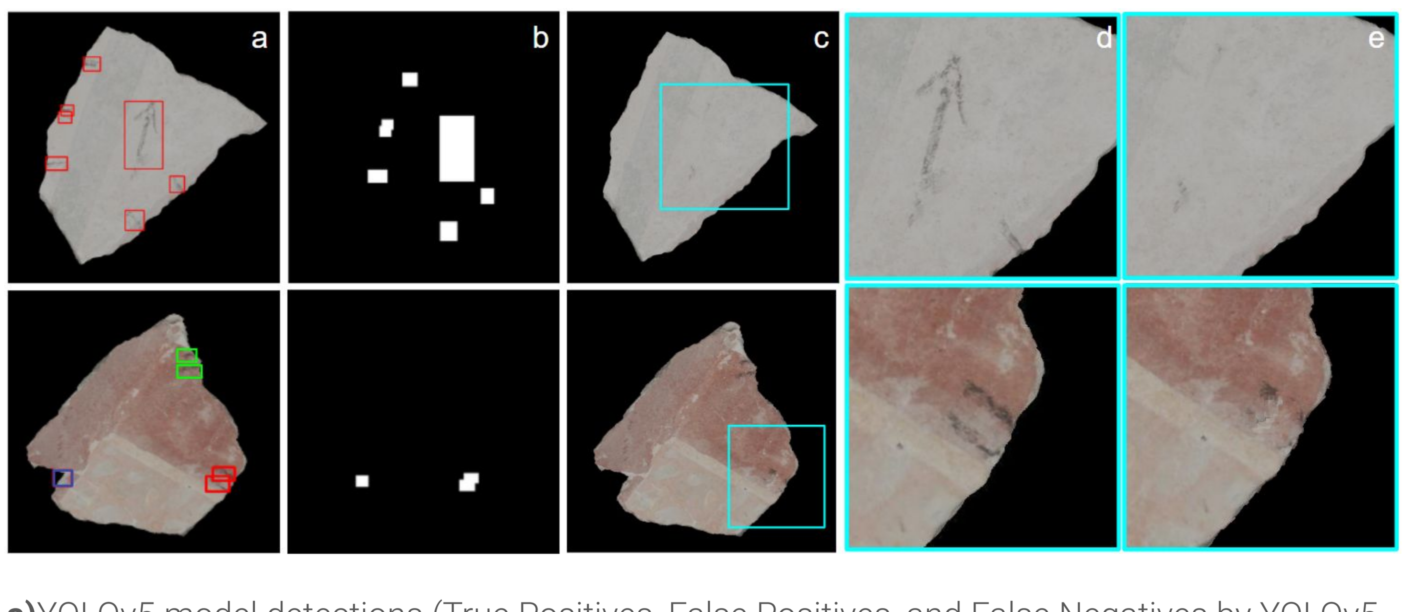

a)YOLOv5 model detections (True Positives, False Positives, and False Negatives by YOLOv5 are highlighted in red, blue and green boxes, respectively); **b)** generated inpainting masks; **c)** inpainting results; **(d)** and **(e)** are detailed views of **(a)** and **(c)**.

Scenario 1: Fragment Segmentation

We used two UNET models (original and modified architecture) for semantically segmenting the image background, fragment background, and motifs into a single unified motif class.

- Original U-Net architecture modified by removing two layers both in the contracting and expanding paths.
- Restored and cropped images are resized to 256x256 for reducing overfitting and shorter training time. Input images in various color spaces and image enhancement schemes were examined.

Table 2. Segmentation performance of different color spaces and image enhancement techniques on MoFF dataset. The first, second, and third best-performing configurations are shown using **red**, blue, and *green* fonts, respectively.

Configuration	$IoU_{average}$	IoU_{motif}	$PA_{average}$	PA_{motif}
RGB	0.82	0.32	*0.89*	0.39
HSV	**0.87**	*0.65*	**0.92**	0.81
YcrCb	0.80	0.45	0.88	0.72
RGB&CLAHE	0.47	0.14	0.63	0.14
RGB&HistEq	0.82	0.43	*0.89*	0.48
RGB&Gamma	0.86	0.48	0.91	0.58
HSV&CLAHE	0.56	**0.69**	0.71	0.83
HSV&HistEq	0.80	0.57	*0.89*	**0.98**
HSV&Gamma	0.86	0.66	**0.92**	0.90
YCrCb&CLAHE	**0.87**	0.59	**0.92**	0.70
YCrCb&HistEq	*0.85*	0.60	**0.92**	*0.88*
YCrCb&Gamma	*0.85*	0.51	0.91	0.67

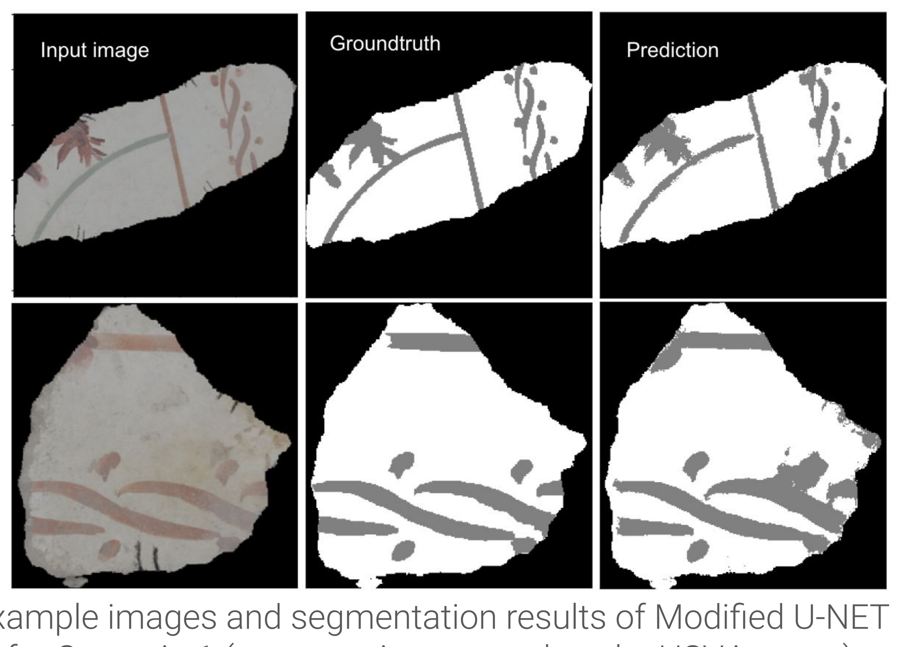

Example images and segmentation results of Modified U-NET for Scenario 1 (computations were done by HSV images)

Scenario 2: Semantic Motif Segmentation

In this scenario, we focused on the more challenging task of semantically segmenting motifs into 12 distinct classes.

- We used fragment images in HSV color space, resized to 512×512 pixels.
- Two U-NET architectures and YOLOv8 were used as baseline models..
- YOLOv8 [3] achieves the best performance in three metrics, except for IoU_{avg}.
- Qualitative results show that YOLOv8 localizes motifs better and predicts motif class with higher precision than two U-NET architectures.

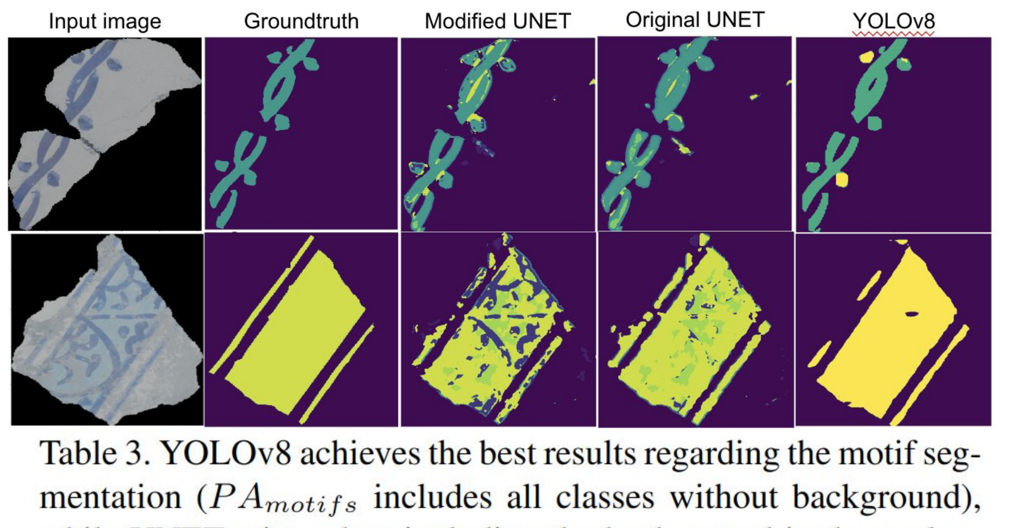

Table 3. YOLOv8 achieves the best results regarding the motif segmentation (PA_{motifs} includes all classes without background), while UNET wins when including the background in the evaluation (PA_{avg} refer to all classes including background, same for IoU).

Architecture	IoU_{motifs}	IoU_{avg}	PA_{motifs}	PA_{avg}
YOLOv8	**0.582**	0.538	**0.634**	**0.797**
Original U-NET	0.416	**0.606**	0.452	0.630
Modified U-NET	0.345	0.569	0.392	0.600

Conclusions

This work focuses on the unexplored task of semantic segmentation of ancient fresco fragments. Briefly,

- We introduced two new image datasets of curated archaeological data
- We defined and provided a baseline for two archaeology-related tasks, i.e., fragment restoration and semantic segmentation.
- We performed a comprehensive analysis to explore the diversity of pictorial contents on the fragments

Acknowledgements

This work is part of a project that has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 964854.

References

[1] Glenn Jocher, et al., ultralytics/yolov5: v7.0 - YOLOv5 SOTA Realtime Instance Segmentation, Nov. 2022

[2] Antonio Criminisi, Patrick P'erez, and Kentaro Toyama. Region filling and object removal by exemplar-based image inpainting. In IEEE Transactions on image processing. IEEE, 2004

Check out our project page!