AI for Memory Preservation: Automated Restoration of Photographs Damaged by Floods

Jose Daniel Galvan Suazo, Hugo Albert Bonet, Carlos Monserrat-Aranda, and Cèsar Ferri

Valencian Research Institute for Artificial Intelligence (VRAIN), Universitat Politècnica de València, Spain

Abstract

The floods caused by the Isolated High-Level Depression (DANA) in the Valencian Community in October 2024 destroyed and damaged hundreds of thousands of personal photographs, erasing key pieces of collective and emotional memory. Within the project *Recuperar las Memorias*, we present an AI-based system for automated photo reconstruction, designed to support the recovery of more than 200,000 affected images. The system integrates YOLOv8 and SAM2 for automatic detection of damaged regions, followed by context-aware inpainting to restore visual coherence. Special modules are included for facial restoration, preserving identity in one of the most emotionally critical aspects of personal photographs. The tool is deployed as a web application that enables both single-image and batch restoration, making it accessible to non-expert users. Preliminary evaluation, combining human perceptual studies and automatic metrics (LPIPS), shows consistent alignment between subjective and objective assessments of quality. This demonstration highlights how advances in computer vision can be mobilised in real-world crisis contexts, placing AI at the service of cultural heritage, dignity, and memory preservation.

Code — https://github.com/jd-galvan/rebrot

Introduction

Extreme meteorological events, such as Isolated High-Level Depressions (DANAs), are recurrent in the Mediterranean region and often result in devastating floods. On October 29, 2024, one of the most severe DANAs in recent history struck the Valencian Community in Spain, with rainfall exceeding 600 litres per square meter in less than eight hours. The consequences were catastrophic: over 200 fatalities, tens of thousands of evacuations, and large-scale material destruction. Beyond the loss of infrastructure, a less visible but deeply human impact emerged—the destruction of personal and documentary photographs, many of which are irreplaceable carriers of cultural and emotional memory.

The project *Recuperar les Memòries*, coordinated by the Institut Universitari de Restauració del Patrimoni (IRP) of the Universitat Politècnica de València, has since collected more than 200,000 damaged photographs from affected families. A volunteer-based laboratory has been established to clean, disinfect, and digitise these photographs. However, the restored and digitised photographs kept damaged parts that were physically impossible to repair, given the scale of the task. Therefore, the need arose to develop an automatic method that would ensure consistent and timely recovery. This paper presents a demonstration of an AI-based system (named *Rebrot*[1]) for photo restoration that combines object detection, segmentation, and generative inpainting to reconstruct damaged areas of photographs in a manner consistent with the original content (see Figure 1).

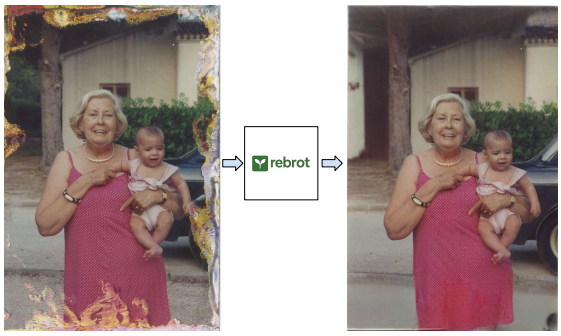

Figure 1: Damaged image and its reconstruction.

This work builds upon recent advances in automatic inpainting (Lugmayr et al. 2022; Avrahami, Fried, and Cohen-Or 2023), segmentation models (Kirillov et al. 2023), and detection architectures (Jocher, Chaurasia, and Qiu 2023), but differs in three key aspects. First, while prior work often benchmarks performance on synthetic or limited datasets, our system is deployed in a real-world humanitarian setting with a large volume of heterogeneous, flood-damaged images. Second, although the combination of the aforementioned models allows for a fairly consistent reconstruction of damaged images, it also affects more delicate areas (such as the faces of people in the images) in a more or less subtle way. This has necessitated the creation of specific pipelines to prevent this type of behaviour. Finally, our evaluation combines automatic perceptual similarity metrics (LPIPS) with expert-based human perceptual assessments, allowing us to analyse correlations between objective and subjective

[1]Rebrot (Catalan/Valencian: "to sprout again") means regrowth / new shoot—the first green that comes back after damage.

measures of restoration quality.

Recent photo-restoration research tackles mixed degradations on digitised prints, but none is tailored to flood-soaked, mud-smeared disaster imagery. Wan et al. (Wan et al. 2020) propose a triplet-domain translation network to restore old photos with both structured (scratches, blotches) and unstructured (noise, fading) defects, achieving strong generalisation from synthetic to real inputs, yet without modelling water-stain patterns or post-disaster artefacts. PikFix (Xu et al. 2023) combines restoration with reference-guided colourisation and introduces a paired real-old-photo dataset; again, the degradations are generic age-related issues rather than moisture-induced damage. Mendoza-Dávila et al. (Mendoza-Dávila, Porta-Montes, and Ugarte 2023) present a practical pipeline that classifies damage type (blur/cracks) and dispatches pretrained restorers for portraits, but do not address large, irregular stain masks or sediment streaks common after floods. In contrast, our tool targets prints affected by the 2024 DANA in Valencia: we segment flood-specific damage (water rings, mud trails, torn edges).

System Description

Our demonstration system consists of four main components: damage detection, region segmentation, context-aware inpainting, and web-based interaction. Figure 2 illustrates the overall workflow:

1. **Damage Detection**. The first stage applies a YOLOv8 (Varghese and Sambath 2024) model trained on a manually annotated dataset of damaged photographs. This model identifies candidate regions of degradation resulting from mud, moisture, and tearing. YOLOv8 was selected for its balance of detection accuracy and inference speed, which is essential for batch processing of large datasets.
2. **Segmentation Refinement**. Detected bounding boxes are refined using SAM2 (Ravi et al. 2025), which generates binary masks with high spatial precision. This step reduces the inclusion of undamaged regions and ensures that the inpainting stage operates on accurately localised damage.
3. **Inpainting and Reconstruction**. For reconstruction, we employ Stable Diffusion Inpainting XL 1.0 (Podell et al. 2024), a generative diffusion model specialised in filling masked areas while maintaining semantic and tonal consistency with surrounding content. Hyperparameters (e.g., guidance scale, number of inference steps) were tuned to balance fidelity and visual coherence.
4. **Special path for facial regions**. As Figure 2 shows, the treatment of faces follows a special path. In this case, once a face has been located, a mask is generated to remove its treatment from the overall image, and its reconstruction (the bottom path in Figure 2) is performed in a special way to ensure its identity preservation. Finally, once reconstructed, it is integrated into the final image.

The system is deployed as a web application built with Gradio (Abid et al. 2019), enabling users to upload images, preview automatic restorations, and process images in batch mode. The interface was designed with accessibility in mind, requiring minimal technical expertise from users, making it suitable for volunteers and affected families.

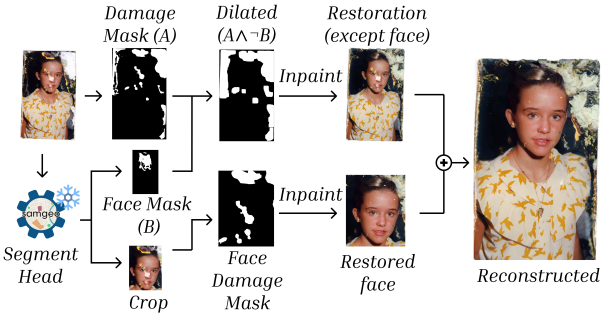

Figure 2: Image reconstruction pipeline. The upper path is the one followed to reconstruct the whole image. The lower path is the one followed to reconstruct faces.

To evaluate the system, we have used a dual evaluation pipeline: (i) human perceptual evaluation conducted with restoration experts using criteria such as identity preservation, stain removal, and structural coherence; and (ii) automatic evaluation using LPIPS, a learned perceptual similarity metric. Correlation analysis between human and automatic scores provides a robust validation framework. Details of this evaluation can be found in (Galván 2025; Albert Bonet 2025)

Conclusions

This demonstration presented an AI-based system for restoring photographs damaged during the 2024 DANA floods in the Valencian Community. By combining YOLOv8 for damage detection, SAM2 for precise segmentation, and Stable Diffusion Inpainting XL for context-aware reconstruction, we developed a complete pipeline capable of processing large collections of heterogeneous, real-world damaged images. The integration of automatic and human-centred evaluation provides a robust framework for assessing reconstruction quality, with preliminary results showing promising alignment between objective metrics and expert judgments.

Beyond its technical contribution, the system illustrates how advances in computer vision and generative modelling can be mobilised in socially meaningful contexts. In particular, it highlights the potential of AI for preserving cultural heritage and promoting emotional resilience in the aftermath of natural disasters.

Future work will address three directions: (i) expanding the evaluation with larger and more diverse human studies, (ii) refining facial identity preservation through specialised generative models, and (iii) integrating the system into the operational workflow of the *Recuperar les Memòries* initiative to scale its impact. By making the tool accessible through a web interface and exploring open-sourcing options, we aim to ensure that the benefits of this research extend to both the scientific community and affected citizens.

Acknowledgments

We wish to express our deepest gratitude to the *Salvem les Fotos UPV / Recuperar les Memòries* project, an extraordinary initiative led by Pedro Vicente-Mullor, Pilar Soriano and Esther Nebot and supported by numerous volunteers, who dedicated countless hours to cleaning and physically restoring over 200,000 family photographs damaged by the DANA of October 29, 2024. Their commitment, patience, and technical skill not only preserved invaluable personal memories but also made this research on digital restoration possible. This work stands on their collective effort and generosity. This work has been partially supported by the grant PID2024-162030OB-100 funded by MCIN/AEI/10.13039/501100011033 and ERDF A way of making Europe and by the grant CIPROM/2022/6 funded by the Generalitat Valenciana. We acknowledge support from "Càtedra de Inteligencia Artificial aplicada a la Administración Pública" funded by Generalitat Valenciana.

References

Abid, A.; Abdalla, A.; Abid, A.; Khan, D.; Alfozan, A.; and Zou, J. 2019. Gradio: Hassle-free sharing and testing of ml models in the wild. *arXiv preprint arXiv:1906.02569*.

Albert Bonet, H. 2025. *Restauración de recuerdos tras la DANA con IA: Segmentación y reconstrucción preservando la identidad*. Master's thesis.

Avrahami, O.; Fried, O.; and Cohen-Or, D. 2023. Blended Latent Diffusion. In *ACM SIGGRAPH 2023 Conference Proceedings*, SIGGRAPH '23. ACM.

Galván, J. D. 2025. *Reconstrucción de imágenes afectadas por la DANA combinando detección, segmentación e inpainting con Stable Diffusion.* Master's thesis.

Jocher, G.; Chaurasia, A.; and Qiu, J. 2023. YOLO by Ultralytics. GitHub repository.

Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A.; Lo, W.-Y.; Dollár, P.; and Girshick, R. 2023. Segment Anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 3992–4003.

Lugmayr, A.; Danelljan, M.; Romero, A.; Yu, F.; Timofte, R.; and Gool, L. V. 2022. RePaint: Inpainting using Denoising Diffusion Probabilistic Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 11461–11471.

Mendoza-Dávila, C.; Porta-Montes, D.; and Ugarte, W. 2023. PhotoRestorer: Restoration of Old or Damaged Portraits with Deep Learning. In *Proceedings of the 19th International Conference on Web Information Systems and Technologies (WEBIST)*, 104–112. ISBN 978-989-758-672-9.

Podell, D.; English, Z.; Lacey, K.; Blattmann, A.; Dockhorn, T.; Müller, J.; Penna, J.; and Rombach, R. 2024. SDXL: Improving Latent Diffusion Models for High-Resolution Image Synthesis. In *The Twelfth International Conference on Learning Representations*.

Ravi, N.; Gabeur, V.; Hu, Y.-T.; Hu, R.; Ryali, C.; Ma, T.; Khedr, H.; Rädle, R.; Rolland, C.; Gustafson, L.; Mintun, E.; Pan, J.; Alwala, K. V.; Carion, N.; Wu, C.-Y.; Girshick, R.; Dollar, P.; and Feichtenhofer, C. 2025. SAM 2: Segment Anything in Images and Videos. In *The Thirteenth International Conference on Learning Representations*.

Varghese, R.; and Sambath, M. 2024. Yolov8: A novel object detection algorithm with enhanced performance and robustness. In *2024 International conference on advances in data engineering and intelligent computing systems (ADICS)*, 1–6. IEEE.

Wan, Z.; Zhang, B.; Chen, D.; Zhang, P.; Chen, D.; Liao, J.; and Wen, F. 2020. Bringing Old Photos Back to Life. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

Xu, R.; Tu, Z.; Du, Y.; Dong, X.; Li, J.; Meng, Z.; Ma, J.; Bovik, A. C.; and Yu, H. 2023. Pik-Fix: Restoring and Colorizing Old Photos. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*.