



HAL
open science

Homer restored: Virtual reconstruction of Papyrus Bodmer 1

Léopold Cudilla, Simon Perrin, Yejing Xie, Harold Mouchère, Isabelle
Marthot-Santaniello

► **To cite this version:**

Léopold Cudilla, Simon Perrin, Yejing Xie, Harold Mouchère, Isabelle Marthot-Santaniello. Homer restored: Virtual reconstruction of Papyrus Bodmer 1. HIP '23: 7th International Workshop on Historical Document Imaging and Processing, Aug 2023, San Jose, France. pp.37-42, 10.1145/3604951.3605518 . hal-04264768

HAL Id: hal-04264768

<https://hal.science/hal-04264768>

Submitted on 30 Oct 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Homer restored: Virtual reconstruction of Papyrus Bodmer 1

Léopold Cudilla*
Simon Perrin*
leopold.cudilla@etu.univ-nantes.fr
simon.perrin@etu.univ-nantes.fr
Nantes Université
Nantes, France

Yejing Xie
Harold Mouchère
yejing.xie@ls2n.fr
harold.mouchere@ls2n.fr
Nantes Université, École Centrale
Nantes, CNRS, LS2N, UMR 6004
Nantes, France

Isabelle Marthot-Santaniello
i.marthot-santaniello@unibas.ch
University of Basel
Basel, Switzerland

ABSTRACT

In this paper, we propose a complete method to reconstruct a damaged piece of papyrus using its image annotated at the character level and the original ancient Greek text (known otherwise). Our reconstruction allows us to recreate the written surface, making it readable and consistent with the original one. Our method is in two stages. First, the text is reconstructed by pasting character patches in their possible locations. Second, we reconstruct the background of the papyrus by applying inpainting methods. Two different inpainting techniques are tested in this article, one traditional and one using a GAN.

This global reconstruction method is applied on a piece of Papyrus Bodmer 1. The results are evaluated visually by the authors of the paper and by researchers in papyrology. This reconstruction allows historians to investigate new paths on the topic of writing culture and materiality while it significantly improves the ability of non specialists to picture what this papyrus, and ancient books in general, could have looked like in Antiquity.

CCS CONCEPTS

• Applied computing → Digital libraries and archives; • Computing methodologies → Reconstruction.

KEYWORDS

Generative Adversarial Networks, papyrus, inpainting

ACM Reference Format:

Léopold Cudilla, Simon Perrin, Yejing Xie, Harold Mouchère, and Isabelle Marthot-Santaniello. 2018. Homer restored: Virtual reconstruction of Papyrus Bodmer 1. In *Proceedings of 7th International Workshop on Historical Document Imaging and Processing (HIP'23)*. ACM, New York, NY, USA, 6 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 INTRODUCTION

Documents from Antiquity rarely reach us in a complete state. Therefore, one of the tasks of archaeologists and historians is to

propose hypotheses for reconstructing the original state of the objects they study. The Papyrus Bodmer 1 is an exceptional piece: it preserves a part of a land register written in Middle Egypt (Panopolis) around 215 AD that is 31 cm high and 5.70m long while it could have been up to 11m long when it was complete [3]. At an unknown moment after 215 AD, a scribe cut this roll into two and copied on the back Book 5 and 6 of Homer's *Iliad*. For yet unexplained reasons, he chose a different number of verses per column between the two books (29 to 31 in Book 5, 38 to 40 in Book 6). For this work, we chose to focus on one column of Book 5 (see Figure 1). Images of the other parts and information on the papyrus are available on the online catalog of the Bodmer foundation¹.



Figure 1: Part of Papyrus Bodmer 1 (*Iliad* 5.645-678) from the 3rd century AD and the proposed reconstruction.

One of the interests of this long and rather well-preserved papyrus is to offer evidence of how a unique document (the land register) of major historical value was reused to bear a key piece of ancient literature (the *Iliad*). The semi-cursive handwriting of the *Iliad* part is more challenging than calligraphic scripts where

¹<https://bodmerlab.unige.ch/fr/constellations/papyri/barcode/1072205286>

*Both authors contributed equally to this research.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).
HIP'23, August 25-26, 2023, San José, California, USA
© 2018 Association for Computing Machinery.
ACM ISBN 978-1-4503-XXXX-X/18/06...\$15.00
<https://doi.org/XXXXXXX.XXXXXXX>

all letters are detached, but is also more comparable with dated specimens. It thus forms an ideal case study for computational restoration assisting current research in ancient History.

Image reconstruction methods have become increasingly powerful in recent years, thanks in particular to advances in the field of deep learning. As far as the restoration of old documents is concerned, work in the field of automatic image processing focuses mainly on improving the overall quality of documents, by removing blur, increasing the definition of the document or retrieving a deleted text [7]. Very recently, some work has also focused on assembling fragments for the reconstruction of old documents. In 2022, a PhD thesis [18] focused on the pairing of papyrus fragments using Siamese networks. Another work from 2020 [15] deals with the reconstruction of a papyrus using its fragments. This one proposes a solution using a graph neural network to create a graph representing the reconstruction of the papyrus by linking each fragment to its neighbors.

Despite these advances, the reconstruction of damaged ancient documents with known text but missing fragments has, to our knowledge, not been studied much. There is currently no machine learning solution for this task. Moreover, in the case of papyri, there are little data to train a model.

The main contributions of the paper are summarized as follows:

- (1) We propose a novel method for reconstructing heavily degraded papyri by restoring the original text with knowledge of it,
- (2) We show that it is possible to use GANs or more handcraft inpainting methods to restore heavily degraded papyri.

In the following section, we cover the related works about inpainting of image and documents. In section 3, we explain the proposed methods to complete the missing parts of one column of the Papyrus Bodmer 1. Then, in section 4, we present our experiments, and finally, we show and discuss different results in section 5.

2 RELATED WORKS

Inpainting is the process of filling in the missing parts of an image. It can be done physically as an artwork by a conservator-restorer, or digitally, as a computer vision task, using various techniques.

Among the first digital inpainting techniques proposed are Efros et al. [4] and Bertalmio et al. [1]. The first proposed a method which consists in extending an image made of a regular pattern, by synthesizing the texture. This method uses statistical information on the image patterns, modeled by a Markov random field model. A pixel is synthesized at a time, according to its neighboring pixels, by finding in the image a neighborhood of similar pixels. The second used diffusion process (Laplacian) to propagate pixels surrounding the area to inpaint.

Deep learning techniques have allowed for semantically meaningful inpainting [17, 25]. Many model architectures have been used to perform this task.

Liu et al. proposed in [12] an encoder-decoder based convolutional neural network (CNN) method to inpaint images with missing parts of various shapes and sizes.

A method of inpainting on historical documents to restore irregular holes using a U-Net [20] and Partial Convolution [11] is

proposed in [8]. The U-Net generates a mask of the irregular holes from the historical degraded document. Partial convolution is used to inpaint the holes using the generated mask and the document image.

Autoencoders are neural networks that learn two functions: encode the input and decode to recreate the input from the code generated just before. Some inpainting methods have been developed using an autoencoder [5, 29].

Generative Adversarial Networks (GAN) [6] are neural networks that generate data by making two networks confront each other. The first network, called the “generator”, attempts to generate the most realistic data possible, in order to make the second network, the “discriminator”, believe that the generated data are real data. The discriminator is trained with real data, and must detect whether a sample is real or generated by the generator. Several GAN-based methods were developed for inpainting [2, 13, 22, 28].

Kim et al. focused in [9] on using a Super-resolution GAN to inpaint the fold lines of old documents. The fold lines to be inpainted are manually marked. The method used does not aim to reconstruct heavily degraded documents, as inpainting is applied on very small areas of the documents.

DeepFillv2 is presented in [26]. *DeepFillv2* is a GAN-based architecture used to inpaint images with free-form mask and guidance. The generator is composed of two networks, the first generates a “coarse” inpainted version of the masked image. The second takes as input the “coarse” inpainted image and the masked image to produce a more detailed inpainted image. The generator uses Gated Convolution and a contextual attention block.

Transformers models [24] are deep learning models based on attention mechanisms. Inpainting models based on transformers are able to inpaint large missing parts of images and output semantically plausible results [10, 27].

Diffusion models are probabilistic machine learning models, using Markov chains. Very recently, different methods based on diffusion models for inpainting, denoising, image generation, and image resolution enhancement have appeared. The method proposed by Saharia et al. [21] is based on a conditional diffusion model, which allows coloring of images, increasing the quality of images, or inpainting. It is able to restore complex structures in missing areas of the images. Lugmayr et al. [14] proposed a denoising diffusion probabilistic model to inpaint images with a free-form mask.

While some methods presented above, especially neural network-based ones, can sometimes reconstruct the missing text of an image, none of them preserves the semantics of that text. They are therefore unsuitable for the text reconstruction of a heavily degraded document.

3 PROPOSED METHOD

In this section, we detail our two phases methodology to reconstruct the papyrus. The first phase aims to reconstruct the text of the papyrus, which is no longer present due to damage. In order to do this, since we otherwise know the text of the *Iliad*, we have used the preserved characters on the papyrus and pasted them at possible locations on the image to obtain a readable text. The second phase aims to reconstruct the papyrus background (the papyrus paper) in order to obtain a realistic reconstruction.

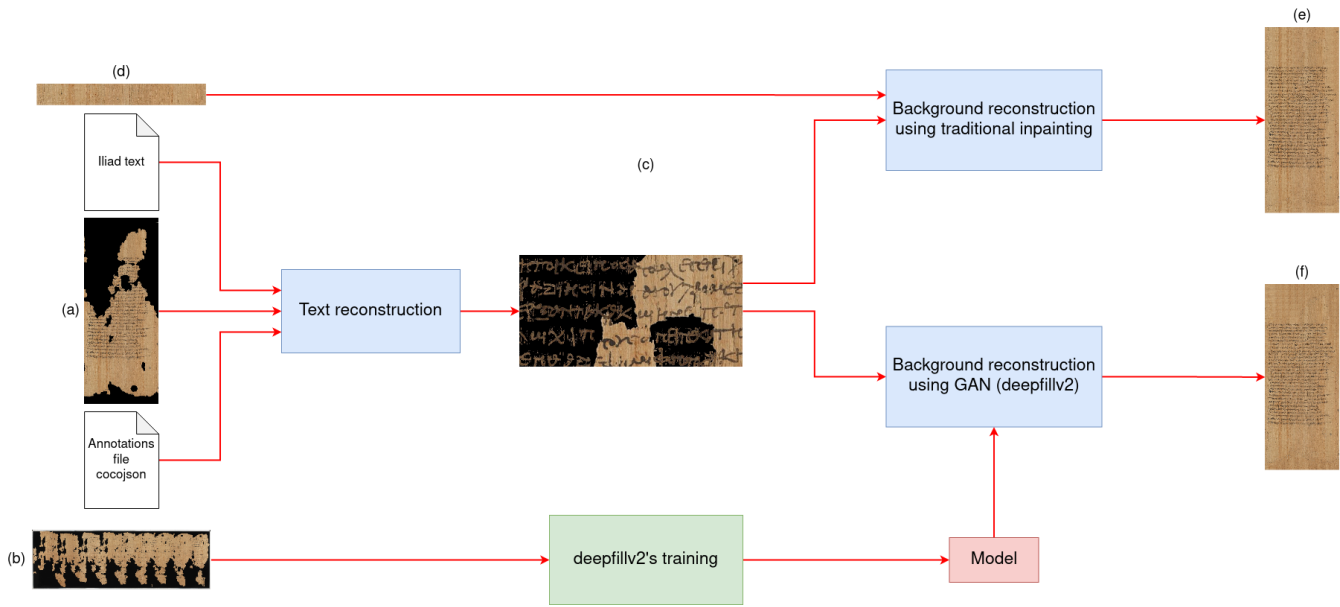


Figure 2: Global scheme of the proposed method. (a) Papyrus column to reconstruct. (b) Papyrus images to train *deepfillv2*. (c) Papyrus with reconstructed text. (d) Papyrus strip for traditional inpainting. (e) Papyrus column reconstructed with traditional inpainting. (f) Papyrus column reconstructed with a GAN (*deepfillv2*).

3.1 Text reconstruction

In the first phase, we aim at reconstructing the missing parts of the *Iliad* in order to recreate the full original text.

3.1.1 Data processing. The image of one column of the papyrus, a text file containing its transcription, and the annotations of the visible characters are used as input. The text file indicates which parts of the text are visible and which are missing on the image of the papyrus. The annotations are in COCO format (in a JSON file) exported from the Research Environment for Ancient Documents (READ)² interface and can be seen online³.

A character patch dictionary is created using the annotation file provided by the papyrologists (this file describes the position, the associated Greek letter and the state quality of the visible characters, from BT1 for complete to BT3 for not recognizable anymore). For each character present in the annotation file, a patch, associated with its Greek letter, is extracted. Only the characters with the best quality (tagged BT1) and not exceeding a width of 50 pixels are extracted. Thanks to these two conditions, we avoid characters that would hardly be readable or too wide to fit in the reconstruction. The dictionary allows us to have several patches for a letter, and thus to vary the patch used. Several reconstructions with randomly selected characters were tested.

The papyrus background is removed during extraction so that only the character ink remains (see Figure 3). This step is done with a basic method, only the darkest pixels (corresponding to the ink) are selected, defined as having a value lower than the average pixel of the patch.



Figure 3: Example of papyrus background removal.

3.1.2 Character placement. Once all the character patches have been extracted and the text file has been read, the characters can be placed in their probable locations.

For each area where characters are missing, the characters that should be present are determined, and corresponding patches for each of them are randomly drawn from the character patch dictionary.

To take advantage of all the available space in each zone, the characters are placed so that there is an equal amount of space between each character patch from the same row. If it is not possible to have any space between the patches (the total width of the patches is greater than the space available), then the patches will overlap slightly. Only pixels with minimum values are kept on the overlapping areas so that the characters are visible despite the overlap.

A part of the papyrus column at this phase with only the reconstructed text is shown in Figure 4. Some noise comes from the darkest papyrus fibers and illustrate the issue of ink fragmentation on papyri as seen during the DIBCO competition at ICDAR 2019 [19], where binarizing papyri images proved difficult.

²<https://github.com/readsoftware/read>

³<https://showcase.d-scribes.philhist.unibas.ch/ViewerBodmer>




Figure 4: Extract of the papyrus column at the text reconstruction phase.

3.2 Background reconstruction

In order to reconstruct the background of the papyrus, inpainting methods are used throughout the document. Two methods have been used. The first is a more traditional method, and the second involves deep learning with a GAN.

3.2.1 Traditional inpainting. The first method is extremely simple. A strip without writing from the column to be reconstructed is extracted and used to complete the document. An image of a blank papyrus leaf, of the same size as the input image, is created in which the extracted strip is pasted several times until it covers the entire image. Then, the image containing the reconstructed text is pasted over the new image, except the areas in black without papyrus.

3.2.2 GAN. The second method is the one using a GAN to reconstruct the papyrus background. For this purpose, we decided to use *deepfillv2* [26] for technical reasons (amount of data available and computational power for training) and its ability to do inpainting with free-form masks.

The *deepfillv2* generator can be separated into two neural networks. The first network consists of gated convolution blocks that produce a “coarse” inpainted version of the masked image. The second network contains two branches of gated convolution blocks and a contextual attention block, to refine the “coarse” version resulting from the first network, and thus produce a detailed inpainting consistent with the rest of the image. The activation function between the layers is *Leaky ReLu*, and the output function of the first and second networks is *tanh*. To train the generator, we compute the hinge loss function on the output of the generator, the mean absolute error (MAE) on the output of the two networks, as well as a “perceptual” loss function, obtained by using a pre-trained CNN neural network to compute the MAE between the feature maps of the base image and the inpainted image. The discriminator is composed of six convolution blocks, to divide the image into 32×32 patches. For each patch, the discriminator determines if the patch is real or generated by the generator. The hinge loss function is then calculated to train the discriminator.

We scanned the image vertically and horizontally to apply inpainting to patches of the image. To do so, the patch must contain a piece of papyrus and an area to be filled in. For this reason, the inpainting application path starts at the center left of the image, on a perfectly completed area and not at the top left of the image (black

area without papyrus). Thus, we gradually extend the papyrus over the areas to be restored.

4 EXPERIMENTS

4.1 Dataset

The dataset used to train *deepfillv2* is composed of patches extracted from online images of Papyrus Bodmer 1. The patches are extracted from areas without writing.

The masks are created either by randomly selecting patches from the base papyrus where torn edges can be seen and using the part without papyrus of these patches as masks, or by generating black rectangles a quarter or an eighth of the size of a patch. The masks are then applied to a corner of a patch without writing.

For each mask, a patch is extracted, and a character is randomly pasted onto the patch. A total of 633 patches of size 256×256 on three channels (RGB) are extracted from Papyrus Bodmer 1 for training (see Figure 5 for an example) and 633 masks are generated.



Figure 5: Example of an image and mask from the dataset. From left to right: a blank papyrus patch; an image with the black background to be extracted to create the mask; the extracted mask; the first image with the mask.

4.1.1 Data augmentation. To augment the size of the dataset, 75% of the masks are re-added as flipped, top to bottom, left to right or both. This way, 475 new masks are created, and for each mask, a new papyrus patch is extracted.

4.2 Implementation Details

The *deepfillv2* discriminator and generator were trained with the Adam optimizer and a learning rate of 0.0001 for both. No learning rate decay was used, and the networks were trained during 430 epochs. The mini-batch size was set to 1.

The method was implemented using PyTorch 1.3.11[16], Python 3.9.10 and CUDA 11.6. The implementation of *deepfillv2* with PyTorch is the one from *csqiangwen* on GitHub⁴.

Training was carried out using an *NVIDIA RTX 3060 Laptop 6GB VRAM* GPU, 8GB of RAM and an *AMD Ryzen 5 5600H 3.3Ghz* CPU.

The removal of the character background (presented in this section 3.1.1) is deactivated for the reconstruction of the papyrus with the GAN, as the results were less relevant following different tests.

5 RESULTS AND DISCUSSIONS

The evaluation of the results was done in a subjective visual way and with the help of the experts’ opinion, as we do not have as Ground Truth the original papyrus to compare our result to.

⁴https://github.com/csqiangwen/DeepFillv2_Pytorch

The two complete reconstructions, a reconstruction of the text and the original image of the Papyrus Bodmer 1 column, can be found at the following [Link](#).

5.1 Subjective visual analysis

Inpainting results with the reconstructed text are shown on Figure 6 and Figure 1 (the full column).

The reconstruction results obtained on real data (Figure 6) show that the colors of the papyrus are decently synthesized, but the texture aspect of the papyrus is missing. The fibers of the papyrus are not reconstructed. Furthermore, the general color of the papyrus is reconstructed, but it is easy to see on the bottom image at epoch 430 that the color tone doesn't quite match the rest of the patch. A slight dark edge at the tear is also visible, due to compression artifacts of the papyrus column image, causing the papyrus to not be clearly separated from the black background.




Figure 6: Examples of inpainting using GAN on patches from the image containing the reconstructed text at different epochs.

Figure 7 compares the same section of the papyrus in the two reconstructions. The top one corresponds to the reconstruction using the traditional inpainting method (copying the strip). With this technique, the texture of the document is perfectly recovered, but a repetitive pattern can be observed and there is no continuity in the fibers between reconstructed and original parts. The bottom reconstruction is the one using GAN. The textures of the papyrus are largely missing in the reconstructed parts and patterns are repeated very strongly on the upper right part. As in the other reconstruction, the continuity of the fibers is not kept.

In both reconstructions, the text is correctly reconstructed and readable.

5.2 Analysis by papyrology experts

Both proposed methods to reconstruct a full column of *Iliad* allow specialists and lay audience picturing how the original papyrus could have looked like. In a detailed view, it cannot be mistaken for a genuine piece due to the still visible transitions between the reconstructed parts and the original ones as well as the erratic aspect of the reconstructed text that has not fully managed to mimic a cursive script (see Figure 8). At the column level, however, it provides a convincing general impression and ground for future

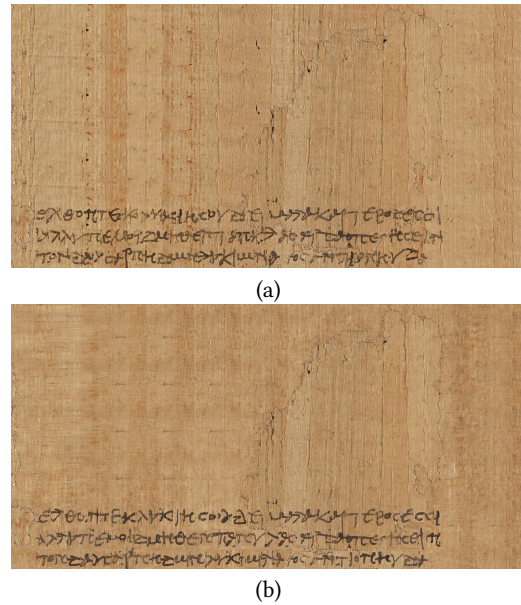


Figure 7: The same section of the papyrus in the two reconstruction methods: (a) traditional inpainting with the strip, (b) GAN. The reconstructed text is different in the two images due to randomness in character patches selection.

discussions among specialists, for instance on the motivation to chose a given number of verses per column. The most convincing restoration of the background is the traditional inpainting because it preserves the vertical pattern of the fibers. This is possible because a) it is written on the side where the fibers are vertical (on the other side of the roll, the fibers are horizontal, thus along and not across the writing) and b) it has a large, well-preserved lower margin to take a coherent strip from. We can thus speculate that this solution would be less successful and more time consuming if extended to the almost 6-meter long roll, and maybe not applicable on a text written along the fibers.

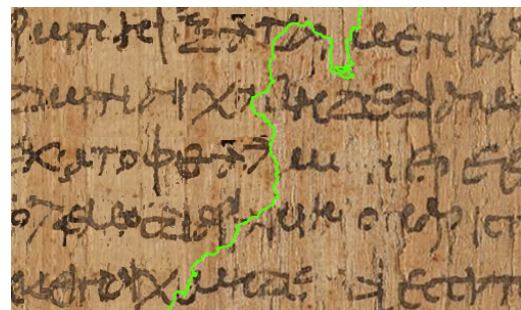


Figure 8: Part of the reconstruction of Papyrus Bodmer 1 column with the delimitation between the original (right) and the reconstructed part, using traditional inpainting (left).

6 CONCLUSION AND FUTURE WORK

In this paper, we proposed a complete method to reconstruct a column of papyrus with both its written text and background starting from the preserved parts of the original. The choices made to restore the text (random selection of individual letters to avoid repetition, ink segmentation method) and the two approaches tested for the backgrounds have yielded convincing results at the page (column) level that are already useful to foster discussions among specialists and provide educational information to a lay audience.

A future advance for this method would be to propose a better generation of the background by including the texture of the papyrus. Since no relation is attested so far between the fibers and the appearance of the text, if the final purpose is to provide a hypothetical reconstruction of a roll in its complete state, one could choose in the future to segment all the written text and paste it on a homogeneous papyrus background that would not be the original. Another direction would be to generate characters as proposed in [23] to avoid reusing those already present in the document and better take into account the cursive aspects of the handwriting (each letter shape is influenced if not modified by its direct neighbors). Last, a metric could be developed to evaluate the result of a papyrus reconstruction without knowing the original document. The study of Papyrus Bodmer 1 shall continue in the future towards a complete reconstruction of this fascinating manuscript to better understand its production and shed new light on ancient books in general.

ACKNOWLEDGMENTS

This work was partially supported by the Swiss National Science Foundation as part of the project no. PZ00P1-174149 “Reuniting fragments, identifying scribes and characterizing scripts: the Digital paleography of Greek and Coptic papyri (d-scribes)”.

REFERENCES

- [1] Marcelo Bertalmio, Guillermo Sapiro, Vincent Caselles, and Coloma Ballester. 2000. Image inpainting. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*. 417–424.
- [2] Ugur Demir and Gozde Unal. 2018. Patch-based image inpainting with generative adversarial networks. *arXiv preprint arXiv:1803.07422* (2018).
- [3] Tomasz Derda. 2010. *P. Bodmer I Recto, A Land List from the Upper Panopolite Nome in Upper Egypt (after AD 213/4)*. Warsaw.
- [4] Alexei A Efros and Thomas K Leung. 1999. Texture synthesis by non-parametric sampling. In *Proceedings of the seventh IEEE international conference on computer vision*, Vol. 2. IEEE, 1033–1038.
- [5] Mohammad Hossein Givkashi, Mahshid Hadipour, Arezoo Parizanganeh, Zahra Nabizadeh, Nader Karimi, and Shadrokh Samavi. 2022. Image Inpainting Using AutoEncoder and Guided Selection of Predicted Pixels. In *2022 30th International Conference on Electrical Engineering (ICEE)*. IEEE, 700–704.
- [6] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2020. Generative adversarial networks. *Commun. ACM* 63, 11 (2020), 139–144.
- [7] Fabian Hollaus, Simon Brenner, and Robert Sablatnig. 2019. CNN based binarization of multispectral document images. In *2019 International Conference on Document Analysis and Recognition (ICDAR)*. IEEE, 533–538.
- [8] Amreen Kaur, Ankit Raj, N Jayanthi, and S Indu. 2020. Inpainting of irregular holes in a manuscript using unet and partial convolution. In *2020 Second International Conference on Inventive Research in Computing Applications (ICIRCA)*. IEEE, 778–784.
- [9] Yong Jun Kim, Debapriya Hazra, Yungcheol Byun, and Khi-Jung Ahn. 2019. Old document restoration using super resolution GAN and semantic image inpainting. In *Proceedings of the International Workshop on Artificial Intelligence and Education*. 34–38.
- [10] Wenbo Li, Zhe Lin, Kun Zhou, Lu Qi, Yi Wang, and Jiaya Jia. 2022. MAT: Mask-Aware Transformer for Large Hole Image Inpainting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 10758–10768.
- [11] Guilin Liu, Fitsum A Reda, Kevin J Shih, Ting-Chun Wang, Andrew Tao, and Bryan Catanzaro. 2018. Image inpainting for irregular holes using partial convolutions. In *Proceedings of the European conference on computer vision (ECCV)*. 85–100.
- [12] Hongyu Liu, Bin Jiang, Yibing Song, Wei Huang, and Chao Yang. 2020. Rethinking image inpainting via a mutual encoder-decoder with feature equalizations. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow*. 725–741.
- [13] Hongyu Liu, Ziyu Wan, Wei Huang, Yibing Song, Xintong Han, and Jing Liao. 2021. Pd-gan: Probabilistic diverse gan for image inpainting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 9371–9381.
- [14] Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. 2022. Repaint: Inpainting using denoising diffusion probabilistic models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 11461–11471.
- [15] Cecilia Ostertag and Marie Beurton-Aimar. 2021. Using Graph Neural Networks to Reconstruct Ancient Documents. In *Pattern Recognition. ICPRI International Workshops and Challenges: Virtual Event, January 10-15, 2021, Proceedings, Part VII*. Springer-Verlag, Berlin, Heidelberg, 39–53. https://doi.org/10.1007/978-3-030-68787-8_3
- [16] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems* 32 (2019).
- [17] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. 2016. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2536–2544.
- [18] Antoine Pirrone. 2022. *Apprentissage profond auto-supervisé de métriques : application à la prédiction d'assemblage de fragments de papyrus*. Theses. Université de Bordeaux. <https://theses.hal.science/tel-03685839>
- [19] Ioannis Pratikakis, Konstantinos Zagoris, Xenofon Karagiannis, Lazaros Tsochatzidis, Tammoy Mondal, and Isabelle Marthot-Santaniello. 2019. ICDAR 2019 Competition on Document Image Binarization (DIBCO 2019). In *2019 International Conference on Document Analysis and Recognition (ICDAR)*. 1547–1556. <https://doi.org/10.1109/ICDAR.2019.00249>
- [20] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. 2015. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III* 18. Springer, 234–241.
- [21] Chitwan Saharia, William Chan, Huiwen Chang, Chris Lee, Jonathan Ho, Tim Salimans, David Fleet, and Mohammad Norouzi. 2022. Palette: Image-to-image diffusion models. In *ACM SIGGRAPH 2022 Conference Proceedings*. 1–10.
- [22] Yong-Goo Shin, Min-Cheol Sagong, Yoon-Jae Yeo, Seung-Wook Kim, and Sung-Jea Ko. 2020. Peps++: Fast and lightweight network for image inpainting. *IEEE transactions on neural networks and learning systems* 32, 1 (2020), 252–265.
- [23] Matthew I Swindall, Timothy Player, Ben Keener, Alex C Williams, James H Brusuelas, Federica Nicolardi, Marzia D’Angelo, Claudio Vergara, Michael McOsker, and John F Wallin. 2022. Dataset Augmentation in Papyrology with Generative Models: A Study of Synthetic Ancient Greek Character Images. In *The 31st International Joint Conference on Artificial Intelligence. IJCAIECAI*.
- [24] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems*, I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.), Vol. 30. Curran Associates, Inc.
- [25] Raymond A Yeh, Chen Chen, Teck Yian Lim, Alexander G Schwing, Mark Hasegawa-Johnson, and Minh N Do. 2017. Semantic image inpainting with deep generative models. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 5485–5493.
- [26] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. 2019. Free-form image inpainting with gated convolution. In *Proceedings of the IEEE/CVF international conference on computer vision*. 4471–4480.
- [27] Yingchen Yu, Fangneng Zhan, Rongliang WU, Jianxiang Pan, Kaiwen Cui, Shijian Lu, Feiyang Ma, Xuansong Xie, and Chunyan Miao. 2021. Diverse Image Inpainting with Bidirectional and Autoregressive Transformers (MM ’21). Association for Computing Machinery, New York, NY, USA, 69–78. <https://doi.org/10.1145/3474085.3475436>
- [28] Shengyu Zhao, Jonathan Cui, Yilun Sheng, Yue Dong, Xiao Liang, Eric I Chang, and Yan Xu. 2021. Large scale image completion via co-modulated generative adversarial networks. *arXiv preprint arXiv:2103.10428* (2021).
- [29] Nina Žizakić, Izumi Ito, Laurens Meeus, and Aleksandra Pizurica. 2019. Autoencoder-learned local image descriptor for image inpainting. In *BNAIC/BENELEARN 2019*, Vol. 2491.

Received 20 February 2007; revised 12 March 2009; accepted 5 June 2009