Enhancing the Readability of Palimpsests Using Generative Image Inpainting

Mahdi Jampour[a], Hussein Mohammed[b] and Jost Gippert[c]
Cluster of Excellence, Understanding Written Artefacts, Universität Hamburg, Hamburg, Germany

Keywords: Generative AI, Inpainting, Deep CNN, Historical Manuscripts, Image Enhancement, Palimpsests.

Abstract: Palimpsests are manuscripts that have been scraped or washed for reuse, usually as another document. Recovering the undertext of these manuscripts can be of significant interest to scholars in the humanities. Multispectral imaging is a technique often used to make the undertext visible in palimpsests. Nevertheless, this approach is not sufficient in many cases, due to the fact that the undertext in resulting images is still covered by the overtext or other artefacts. Therefore, we propose defining this issue as an inpainting problem and enhancing the readability of the undertext using generative image inpainting. To this end, we introduce a novel method for generating synthetic multispectral palimpsest images and make the generated dataset publicly available. Furthermore, we utilise this dataset in the fine-tuning of a generative inpainting approach to enhance the readability of palimpsest undertext. The evaluation of our approach is provided for both the synthetic dataset and palimpsests from actual research in the humanities. The evaluation results indicate the effectiveness of our method in terms of both quantitative and qualitative measures.

1 INTRODUCTION

Palimpsests are documents or writing substrates where new text has been written over previously erased content. Historical constraints, including the cost and complexity of production, led to repurposing manuscripts by removing old content through mechanical or chemical methods. The erased text, termed undertext, is replaced with new handwriting, known as overtext. Scholars in the humanities find significant interest in retrieving the undertext from palimpsests.

Revealing the undertext in palimpsest manuscripts poses challenges due to degradation, occlusion by overtext, and other visual artefacts. Multispectral imaging (MSI) is a common technique, capturing light within specific wavelength ranges to make the undertext visible. While MSI aids visibility, it may not render the text readable in many cases. Figure 1 illustrates the difference between MSI and enhanced readability for a palimpsest document. Parts of the undertext obscured by overtext and artifacts cannot be fully revealed by MSI alone. To address this, we

Figure 1: An example of a palimpsest image. (left) Original document, (middle) Multispectral image, and (right) Enhanced undertext.

propose employing a generative inpainting approach to generate these occluded pixels, guided by information about the shape of letters in a particular script. Efforts to separate undertext from overtext based on aging differences and visual contrast using statistical methods have been made. However, these methods are effective only when clear distinctions exist between undertext and overtext. In contrast, this study employs deep learning to reconstruct separated undertext using models trained on the same letterforms. An end-to-end approach is proposed to reconstruct the undertext in multispectral palimpsest images within its original visual context. Defining this task as an image inpainting problem, we treat overtext and vi-

[a] https://orcid.org/0000-0002-1559-1865
[b] https://orcid.org/0000-0001-5020-3592
[c] https://orcid.org/0000-0002-2954-340X

Jampour, M., Mohammed, H. and Gippert, J.
Enhancing the Readability of Palimpsests Using Generative Image Inpainting.
DOI: 10.5220/0012347100003654

In *Proceedings of the 13th International Conference on Pattern Recognition Applications and Methods (ICPRAM 2024)*, pages 687-694
ISBN: 978-989-758-684-2; ISSN: 2184-4313

sually similar artefacts as masked regions, aiming to reconstruct the underlying undertext portions.

Image inpainting involves reconstructing regions within an image by filling in areas that need modification. This process is considered ill-posed because there's no single definitive solution for the pixel values in the target regions. Image inpainting is essential in applications such as image/video restoration, removal, and etc.

We reframe the problem as an image restoration task, treating overtext and visually similar artefacts as concealed areas, with the objective of reconstructing the undertext beneath them. To achieve this, we employ the MAT model (Li et al., 2022), a state-of-the-art transformer-based approach, as an end-to-end solution to enhance palimpsest manuscript readability. Additionally, we introduce a method to generate a training dataset of multispectral palimpsest images for fine-tuning. Our third contribution is the provision of a dataset containing 1000 synthetic samples of multispectral Georgian palimpsest images, each paired with a corresponding mask and ground-truth, resulting in 3000 images. The Synthetic Georgian Palimpsest (SGP) images dataset is publicly available for research purposes in our research data repository (Jampour et al., 2023a).

2 RELATED WORK

Most prior methods to improve the legibility of underlying text in palimpsest multispectral images have concentrated on segregating undertext from overtext. In contrast, our approach aims to offer an end-to-end solution, focusing on the complete reconstruction of undertext within the original visual context of multispectral palimpsest images. This objective is cast as an image inpainting challenge, where overtext and visually analogous artefacts are regarded as masked regions, and the occluded undertext undergoes reconstruction.

2.1 Readability Enhancement of Palimpsests

Numerous studies have tackled the challenge of improving the readability of underlying text in palimpsests, given their significance in manuscript research. Early works, such as that of Salerno et al. (Salerno et al., 2007), introduced statistical methodologies, notably Principal Component Analysis (PCA) and Independent Component Analysis (ICA), to separate undertext from overwritten text in palimpsest multispectral images. These techniques aimed to enhance the legibility of concealed text in such images. It is crucial to note, however, that PCA, while effective for noise reduction and data simplification, operates under the assumption of linear relationships between variables, a limitation that may not align with the inherent complexities of palimpsest enhancement challenges.

In a different study, Easton et al. (Easton et al., 2011) applied classical optical imaging and well-established image processing techniques to enhance the readability of the Archimedes Palimpsest. Similar methods, grounded in traditional image processing, significantly contributed to the decipherment of obscured text.

More recently, Starynska et al. (Starynska et al., 2021) leveraged deep generative networks to separate undertext from overtext in palimpsests. This approach necessitated a significant volume of data similar to the typeface of the undertext, which they generated synthetically. The authors used the generative model to separate undertext and restore the original text. On the other hand, a recent study by Obukhova et al. (Obukhova et al., 2023) rightly underscored the data requirements associated with deep learning approaches. Recognizing the need for extensive data, they evaluated the legibility of undertext in palimpsest images using several classical statistical methods, including PCA and Linear Discriminant Analysis (LDA).

2.2 Image Inpainting

Traditional image inpainting methods, including exemplar-based techniques, typically rely on patch-based algorithms that primarily utilize local information around the target region. This approach has limitations when tasked with inpainting extensive image regions or when considering global visual features, such as a dominant background texture. Both traditional and deep-learning-based approaches face challenges posed by the visual complexity of images (Xiang et al., 2023). In the dataset used for this study, an additional significant challenge arises due to the necessity of achieving a very high level of fidelity. This is primarily attributed to the fact that even slight additions, removals, or alterations in the orientation of a few pixels can result in entirely incorrect characters within the underlying text. Moreover, the target regions constitute a substantial portion of the image, yet are distributed across a large number of small local sections.

In works, Bertalmio et al. (Bertalmio et al., 2000) introduced the concept of utilizing information from the surroundings of inpainting regions. Subsequently,

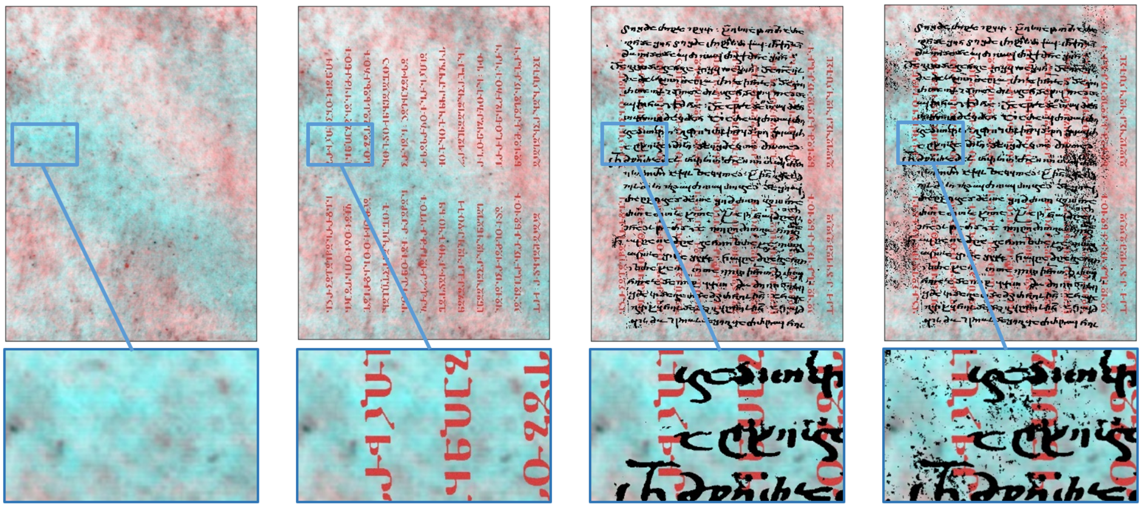

(a) (b) (c) (d)

Figure 2: Synthesizing MSI images of palimpsests. (a) Generated background, (b) rendered undertext on generated background, (c)Adding randomly generated overtext, (d) Adding randomly generated noise to mimic the typical texture in MSI images. The second row shows zoomed regions in the images.

enhanced variations of Bertalmio's approach were developed, known as exemplar-based image inpainting (Abdulla and Ahmed, 2021; Zhang et al., 2019), Markov Random Field (MRF) techniques (Ružić and Pižurica, 2015; Jampour et al., 2017), which rely on surrounding information for pixel value estimation. In certain scenarios, high-level attributes are harnessed to improve inpainting results. For example, (Jampour et al., 2017) proposed employing a guided face with MRF for inpainting facial features. However, the challenge of inpainting large regions presents a significant hurdle, where deep-learning-based approaches yield more favourable results. For instance, MAT (Li et al., 2022) stands out as one of the state-of-the-art method for large-hole image inpainting, harnessing transformers to handle high-resolution images and effectively integrate local and global information. Similarly, recent deep-learning approaches, including generative models (e.g., GANs), effectively address the challenge of missing region size, achieving notable success in image inpainting.

3 PROPOSED APPROACH

This section will commence with a presentation of the problem statement. Subsequently, we will outline the process of generating synthetic MSI images of palimpsests and introduce the Synthetic MSI Images of Georgian Palimpsests (SGP dataset). Following this, we will describe the process of selecting a pre-training dataset for the MAT model and elaborate on the fine-tuning of the pre-trained MAT model.

3.1 Problem Statement

In contrast to Starynska et al.'s approach (Starynska et al., 2021), which involves a separating undertext model, and statistical models like Obukhova et al.'s (Obukhova et al., 2023), we propose to reconstruct the undertext in multispectral palimpsest images within its original visual context. We frame this task as an image inpainting problem, where overtext and visually similar artefacts are considered masked regions, and the underlying undertext portions need reconstruction. Let P_{image} denote the set of pixel values in a palimpsest multispectral image, including background, undertext, overwritten text, and noise. A possible formulation of the problem can be expressed as:

$$P_{image} = \Psi(P_{target} \oplus P_{noise}) \tag{1}$$

Where P_{target} and P_{noise} are two sets of pixel values in the image representing the undertext and the overtext with the noise correspondingly, and both are subject to the $\Psi(.)$ function representing general background texture of MSI images. We define symbol $\oplus$ to show that these two sets (i.e., P_{target} and P_{noise}) are mixed non-linearly, and their intersection is non-empty, in another word $P_{target} \cap P_{noise} \notin \emptyset$. Therefore, there is a subset of pixel values J which belongs to both sets P_{target} and P_{noise} where:

$$\begin{aligned} &J \subset P_{target},\ and\ J \subset P_{noise} \\ &J = P_{target} \cap P_{noise} \end{aligned} \tag{2}$$

Our objective is to estimate J, and the primary challenge arises from the absence of information regarding the pixel values of the undertext (i.e., P_{target}).

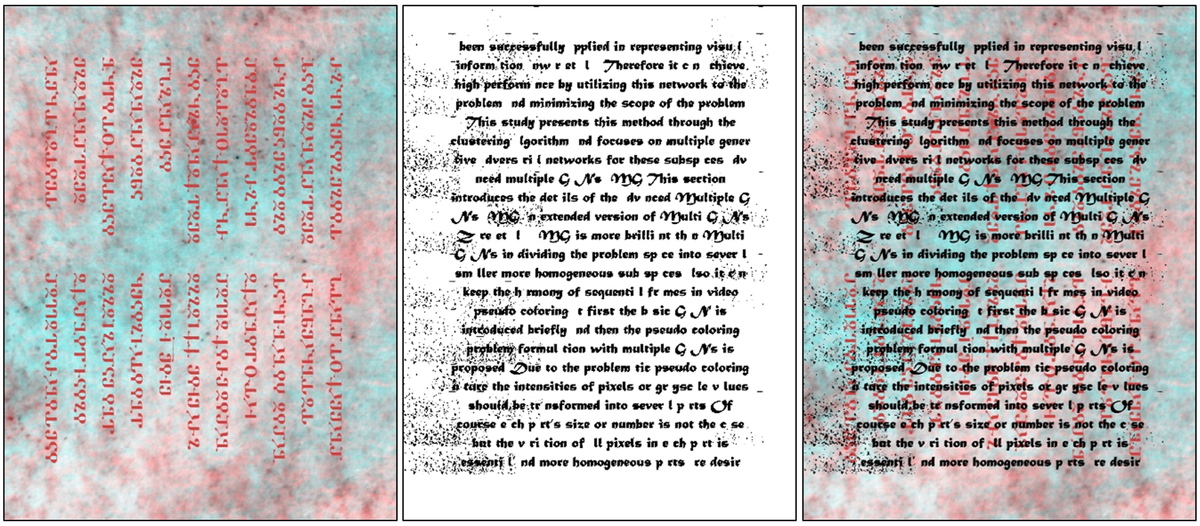

Figure 3: A generated sample from the proposed dataset. Left to right: (a) Ground-truth, (b) Mask, and (c) Target image.

As an inverse problem, we suggest generating a set of pixel values in P_{target} and P_{noise} that are very similar and emulating the $\Psi(.)$ function. This approach enables the model to learn the estimation of the overlapping regions (i.e., J) and subsequently perform inpainting on the palimpsests to reconstruct the missing portions of the undertext.

3.2 Generating Synthesised MSI Images of Palimpsests

Generative impainting approaches, such as GANs, yield state-of-the-art performance in various aspects (Jampour et al., 2023b; Zare et al., 2019). Nevertheless, they require a large amount of training data with ground-truth images and pixel-level annotated masks. Creating this kind of training data is not only time consuming, but also not possible due to the unavailability of such information. Therefore, we propose to generate synthetic MSI images of palimpsests. To this end, we create images in four steps as follows:

1. Generating synthesised background-texture of MSI images (i.e., $\Psi(.)$).
2. Rendering undertext using letters with similar typeface and mixing the rendered text with the generated background (i.e., $\Psi(P_{target})$).
3. Generating random overtext and mixing it with the resulting image from the previous step (i.e., $\Psi(P_{target} \oplus P_{noise})$).
4. Adding randomly generated noise to mimic the typical texture in MSI images.

For the generation of MSI-like background in the training images, regions without text were cropped from actual MSI images of palimpsests, and well-known post-processing steps, including stitching, scaling, and cleaning, were applied. Figure 2a illustrates an example of the background preparation. The undertext is rendered using a typeface that was created from these palimpsest images. Figure 2b showcases an example of the resulting images.

The MSI images used pertain to Georgian palimpsests from the 11^{th} century, with undertext in the Caucasian Albanian language from the 7^{th}-8^{th} centuries, stored in the library of St. Catherine's Monastery on Mount Sinai (Sin. georg. NF 13 and 55). The source of this MSI data is furnished by the Sinai Palimpsests Project via the Sinai Manuscripts Digital Library at UCLA. The initial draft of the typeface was crafted by Jost Gippert in 2005, and the final version was prepared by Andreas Stötzner in 2007 (Gippert and Dum-Tragut, 2023).

Subsequently, random overtext is generated and added to our synthetic images (i.e., $\Psi(P_{target} \oplus P_{noise})$). Figure 2c presents an example of the resulting images. In these images, the content of the overtext is randomly selected, but the direction, line spacing, and scale are made as similar as possible to the overtext found in the actual palimpsest images. Finally, additional noise is introduced to replicate the random artifacts found in the MSI images of palimpsests, as depicted in Figure 2d. We note that the amount, distribution, and location of noise are considered entirely random.

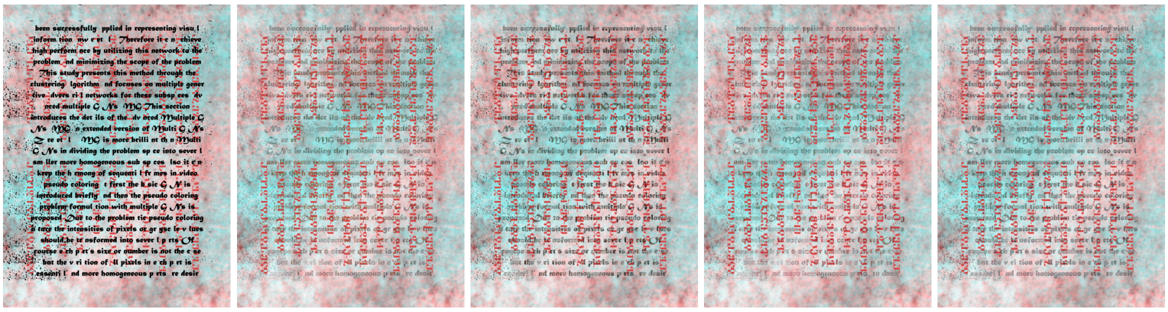

(a) (b) (c) (d) (e)

Figure 4: Visual comparison between the generated images using the MAT model pretrained on different datasets without further fine-tuning on the SGP dataset: (a) the synthetic input sample. The result of MAT pretrained on (b) CelebA-HQ, (c) FFHQ, (d) Places365, (e) Places365˙FullData.

3.3 The Synthetic MSI Images of Georgian Palimpsests (SGP)

A dataset comprising 1000 synthetic MSI images of Georgian palimpsests has been created through the previously outlined steps. Each sample in the dataset is accompanied by a mask and a ground-truth image. The ground-truth image, P_{target}, serves as the target image for evaluating the generated image using the inpainting model. The mask image, P_{noise}, delineates the regions that require inpainting, encompassing the overtext and other noise artifacts. Figure 3 (a)-(c) displays a synthetic sample, including its corresponding mask and ground-truth images. All images are of uniform dimensions, measuring 2800×2100 pixels, and are stored in colour PNG format. Consequently, this dataset consists of a total of 3000 images across the 1000 samples and is made publicly available for research purposes at (Jampour et al., 2023a).

3.4 Selection of Pre-Training Datasets for the MAT Model

The Mask-Aware Transformer (MAT) model undergoes pre-training using various standard and publicly available datasets, including CelebA-HQ, FFHQ, Places365, and an extensive version referred to as Places365˙FullData. Initiating training from scratch on these large datasets can be time-consuming and computationally intensive. Therefore, the common practice is to use pre-trained models and fine-tune them for specific tasks. Our qualitative and quantitative assessments reveal that a MAT model pre-trained on the Places365 dataset produces the most promising results, as illustrated in Figure 4 and Table 1, respectively. In this quantitative evaluation, we employ the well-established Fréchet Inception Distance (FID) (Heusel et al., 2017) and the Structural Similarity Index Measure (SSIM) as our chosen metrics.

The FID for multivariate normal distribution is calculated as follows:

$$\begin{aligned} FID = & \left\|\mu_{I_{org}} - \mu_{I_{inp}}\right\|^2 \\ & - Tr(\sum I_{org} + \sum I_{inp} \\ & - 2\sqrt{\sum I_{org} + \sum I_{inp}}) \end{aligned} \tag{3}$$

Where I_{org} and I_{inp} are feature vectors of ground truth and inpainted images. $\mu_{I_{org}}$ and $\mu_{I_{inp}}$ are magnitude of the vectors I_{org} and I_{inp} and Tr(.), $\sum I_{org}$, and $\sum I_{inp}$ are respectively the trace and covariance matrix of vectors.

The SSIM on the other hand can be described as follows:

$$SSIM = \frac{(2\mu_{I_{org}}\mu_{I_{inp}} + c_1)(2\mu_{I_{org}}\mu_{I_{inp}} + c_2)}{(\mu^2_{I_{org}}\mu^2_{I_{inp}} + c_1)(\mu^2_{I_{org}}\mu^2_{I_{inp}} + c_2)} \tag{4}$$

Where $\mu_{I_{org}}$, $\mu_{I_{inp}}$ and $\sigma_{I_{org}}$, $\sigma_{I_{inp}}$ are mean and variance of feature vectors I_{org} and I_{inp} and c_1, c_2 are variables to stabilize the division.

The choice to employ the Fréchet Inception Distance (FID), in addition to it being a widely accepted evaluation metric, is motivated by the use of generative adversarial networks (GANs) in the MAT model. Research, as outlined by Lucic et al. (Lucic et al., 2018), has demonstrated that FID serves as an appropriate metric for quantifying the realism and diversity of images generated by GANs. FID measures differences in the representations of features, encompassing elements such as edges, lines, and higher-order phenomena.

In contrast, the Structural Similarity Index Measure (SSIM) is a widely-used metric approximating human visual perception of similarity. Three random SGP dataset samples underwent inpainting using the MAT model pre-trained on four distinct datasets. Figure 4 visually compares the results, while Table 1 provides a quantitative analysis. The MAT model

Table 1: Quantitative comparison on various pre-trained models executed on three synthetic documents before and after inpainting with FID and SSIM metrics.

		FID ↓		SSIM ↑	
Test Image	Pre-trainedModel	Before	After	Before	After
image 0920	CelebA-HQ	298	225	0.718	0.805
	FFHQ	298	262	0.718	0.781
	Places365	298	160	0.718	0.816
	Places365˙FullData	298	297	0.718	0.792
image 0930	CelebA-HQ	292	194	0.591	0.747
	FFHQ	292	290	0.591	0.709
	Places365	292	131	0.591	0.761
	Places365˙FullData	292	362	0.591	0.723
image 0940	CelebA-HQ	310	239	0.713	0.792
	FFHQ	310	362	0.713	0.772
	Places365	310	206	0.713	0.807
	Places365˙FullData	310	400	0.713	0.784

Table 2: The inpainting results of the SSIM metric on three randomly selected samples using the fine-tuned (FT) model on the SGP Dataset.

SSIM ↑	image 0920		image 0930		image 0940	
Model	Before	After	Before	After	Before	After
No FT	0.718	0.816	0.591	0.761	0.713	0.807
1^{st} FT	0.718	0.818	0.591	0.765	0.713	0.809
2^{nd} FT	0.718	0.821	0.591	0.767	0.713	0.812
3^{rd} FT	0.718	0.823	0.591	0.766	0.713	0.813
4^{th} FT	0.718	0.823	0.591	0.764	0.713	0.814
5^{th} FT	0.718	0.823	0.591	0.761	0.713	0.813
6^{th} FT	0.718	0.822	0.591	0.758	0.713	0.813

Table 3: The inpainting results using the FID metric on three randomly selected samples using the fine-tuned (FT) model on the SGP Dataset.

FID ↓	image 0920		image 0930		image 0940	
Model	Before	After	Before	After	Before	After
No FT	298	160	292	131	310	206
1^{st} FT	298	164	292	135	310	211
2^{nd} FT	298	171	292	142	310	228
3^{rd} FT	298	179	292	145	310	244
4^{th} FT	298	185	292	149	310	255
5^{th} FT	298	191	292	151	310	259
6^{th} FT	298	192	292	152	310	258

trained on Places365 outperforms others in both FID and SSIM measurements, as evident in Figure 4. Consequently, the MAT model pre-trained on Places365 will be fine-tuned on the SGP Dataset.

3.5 Fine-Tuning the Pre-Trained Model

The MAT model, initially pre-trained on Places365, underwent fine-tuning using the SGP Dataset to enhance undertext reconstruction quality. The dataset's 1000 samples were divided into subsets: 800 for training, 100 for validation, and 100 for testing. Given the MAT model's 512×512 input size, patches of this size were cropped from all subset images. Training utilized 50,400 patches, while each validation and test subset employed 6,300 patches.

Testing experiments involved dividing each input image into 512×512 patches, later merged after inpainting for image reconstruction. Metrics, including SSIM and FID, were assessed based on the reconstructed image. Fine-tuning spanned ten epochs until validation loss plateaued, with no substantial performance gain. An initial learning rate of $1e^{-4}$ prevented drastic changes in pre-trained model weights.

4 EXPERIMENTAL RESULTS

This section encompasses a quantitative evaluation of the MAT model, pre-trained on the Places365 dataset and fine-tuned on the SGP dataset. Furthermore, it includes a qualitative evaluation on images from the synthetic data within the SGP dataset as well as images from real palimpsests.

4.1 Quantitative Evaluation

The SSIM and FID results appear in Tables 2 and 3, respectively, using three randomly selected samples. FID computation (Seitzer, 2020) necessitates multi-

Table 4: Evaluation of the test subset containing 100 samples before and after inpainting with 4^{th} FT.

	FID ↓		SSIM↑	
	Before	After	Before	After
Mean	293.43	194.74	0.6952	0.8114
SD	26.58	28.81	0.049	0.028

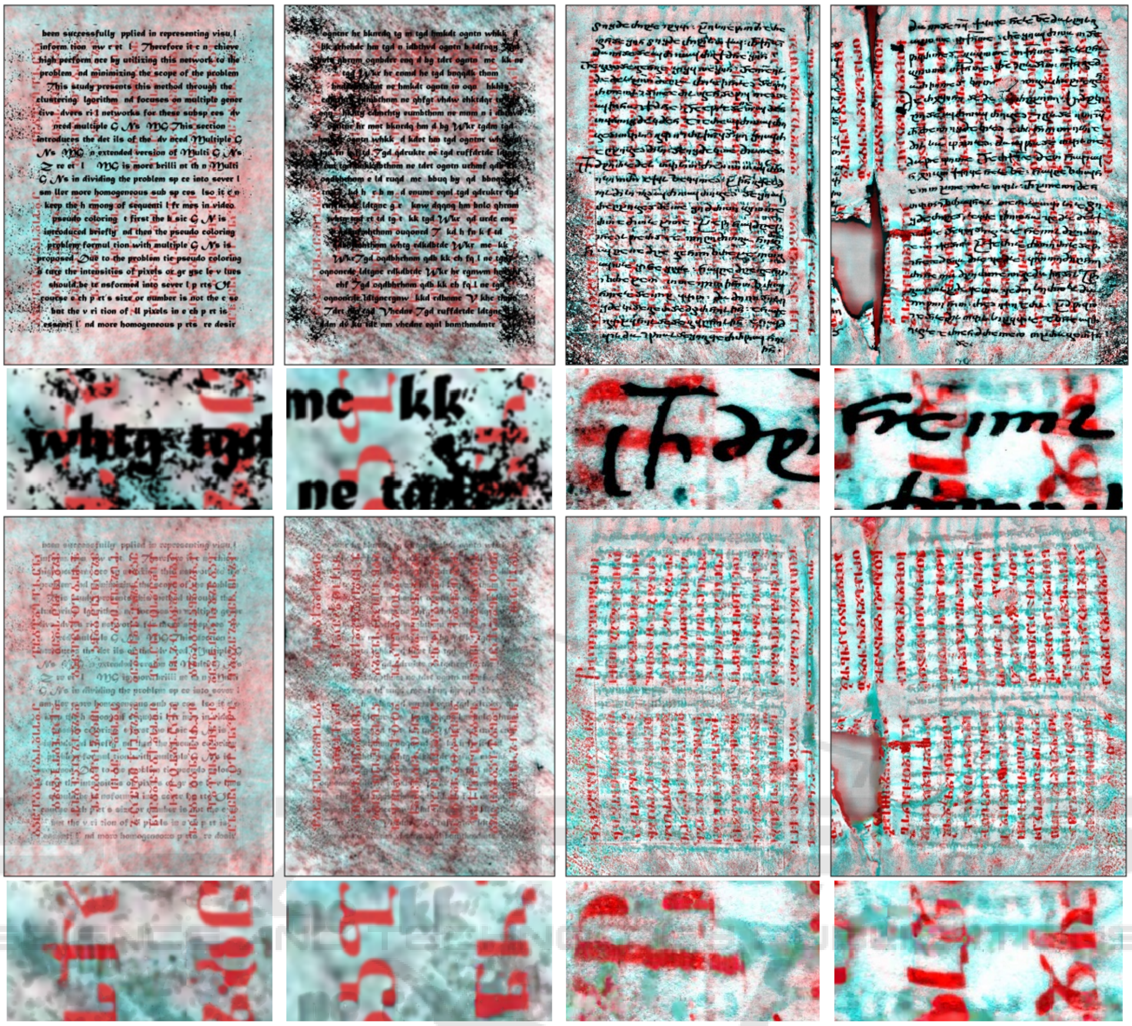

Figure 5: Real and synthetic palimpsest documents before and after inpainting. It can be seen that the readability of undertext has been clearly enhanced. The first and third (Zoomed) rows are the input images before inpainting, whilst the second and fourth rows are the output images after inpainting. The two right-most samples are real palimpsests documents along with their inpainting results.

ple input images. We extracted 368 patches per input image, cropped conventionally with a stride of 100 pixels and a consistent size of 512×512. This resulted in two folders for each ground truth and target image pair, each containing 368 patches, to compute FID. The FID results lack a discernible improvement, possibly due to added edges and lines affecting values. Conversely, Table 2 distinctly shows SSIM enhancements for all images using the fine-tuned model on the proposed SGP Dataset. The fine-tuned model is also assessed on the test dataset for future comparisons. Table 4 presents average values and standard deviations for FID and SSIM measures across all test samples in the SGP test set.

4.2 Qualitative Evaluation

The primary aim is to improve the readability of undertext in palimpsest images. Visual evaluation is crucial for validating this approach. The inpainting results, shown in the second row of Figure 5, demonstrate a substantial enhancement in undertext readability for both synthetic and real palimpsests. Notably, the reconstruction accurately restores large portions of occluded undertext letters, as evident in the zoomed regions in the two lower rows of Figure 5.

5 CONCLUSIONS

Palimpsests, historical manuscripts repurposed through scraping or washing, hold significant value for humanities scholars. Scholars commonly use multispectral imaging to reveal obscured undertext in these manuscripts. However, this method often falls short in making the undertext readable due to occlusion by overtext and various artefacts. This paper proposes a novel approach to enhance undertext readability in palimpsests through generative image inpainting. We introduced a method to generate the Synthetic MSI Images of Georgian Palimpsests (SGP) dataset, publicly available. Additionally, we fine-tuned a pre-trained Mask-Aware Transformer (MAT) model using this dataset to improve its performance. The resulting model was quantitatively evaluated using FID and SSIM metrics on both synthetic SGP dataset images and real palimpsest images. Qualitative evaluation illustrated the practicality of our approach in enhancing undertext readability for manuscript research. The results clearly demonstrate the effectiveness of our approach in both quantitative and qualitative measures.

ACKNOWLEDGEMENTS

The research for this work was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2176 'Understanding Written Artefacts: Material, Interaction and Transmission in Manuscript Cultures', project no. 390893796. The research was conducted within the scope of the Centre for the Study of Manuscript Cultures (CSMC) at Universität Hamburg.

In addition, we thank St Catherine's Monastery on Mt Sinai, the Early Manuscripts Electronic Library and the members of the Sinai Palimpsest Project (http://sinaipalimpsests.org/), esp. Michael Phelps, Claudia Rapp and Keith Knox, for providing the MSI images.

REFERENCES

Abdulla, A. and Ahmed, M. (2021). An improved image quality algorithm for exemplar-based image inpainting. *Multimed Tools Appl*, 80:13143–13156.

Bertalmio, M., Sapiro, G., Caselles, V., and Ballester, C. (2000). Image inpainting. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '00, page 417–424. ACM Press/Addison-Wesley Publishing Co.

Easton, R. L., Christens-Barry, W. A., and Knox, K. T. (2011). Spectral image processing and analysis of the archimedes palimpsest. In *2011 19th European Signal Processing Conference*, pages 1440–1444.

Gippert, J. and Dum-Tragut, J. (2023). *Caucasian Albania: An International Handbook*. De Gruyter.

Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., and Hochreiter, S. (2017). Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS'17, page 6629–6640. Curran Associates Inc.

Jampour, M., Li, C., Yu, L.-F., Zhou, K., Lin, S., and Bischof, H. (2017). Face inpainting based on high-level facial attributes. *Computer Vision and Image Understanding*, 161:29–41.

Jampour, M., Mohammed, H., and Gippert, J. (2023a). Synthetic MSI Images of Georgian Palimpsests (SGP Dataset).

Jampour, M., Zare, M., and Javidi, M. (2023b). Advanced multi-gans towards near to real image and video colorization. *J Ambient Intell Human Comput*, 14:12857–12874.

Li, W., Lin, Z., Zhou, K., Qi, L., Wang, Y., and Jia, J. (2022). Mat: Mask-aware transformer for large hole image inpainting. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10748–10758. IEEE Computer Society.

Lucic, M., Kurach, K., Michalski, M., Bousquet, O., and Gelly, S. (2018). Are gans created equal? a large-scale study. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, NIPS'18, page 698–707. Curran Associates Inc.

Obukhova, N. A., Baranov, P. S., Motyko, A. A., Chirkunova, A. A., and Pozdeev, A. A. (2023). Palimpsest research based on hyperspectral information. In *2023 25th International Conference on Digital Signal Processing and its Applications (DSPA)*, pages 1–4.

Ružić, T. and Pižurica, A. (2015). Context-aware patch-based image inpainting using markov random field modeling. *IEEE Transactions on Image Processing*, 24(1):444–456.

Salerno, E., Tonazzini, A., and Bedini, L. (2007). Digital image analysis to enhance underwritten text in the archimedes palimpsest. *IJDAR*, 9:79–87.

Seitzer, M. (2020). pytorch-fid: FID Score for PyTorch. https://github.com/mseitzer/pytorch-fid. Version 0.3.0.

Starynska, A., Messinger, D., and Kong, Y. (2021). Revealing a history: palimpsest text separation with generative networks. *IJDAR*, 24:181–195.

Xiang, H., Zou, Q., Nawaz, M. A., Huang, X., Zhang, F., and Yu, H. (2023). Deep learning for image inpainting: A survey. *Pattern Recognition*, 134:109046.

Zare, M., Lari, K. B., Jampour, M., and Shamsinejad, P. (2019). Multi-gans and its application for pseudo-coloring. In *2019 4th International Conference on Pattern Recognition and Image Analysis (IPRIA)*, pages 1–6.

Zhang, N., Ji, H., and et al., L. L. (2019). Exemplar-based image inpainting using angle-aware patch matching. *J Image Video Proc*, 70:1.