DeepHadad: Enhancing the Readability of Ancient Northwest Semitic Inscriptions

A. C. Aioanei†[1] and J. Klein[2] and K. M. Klein[3] and R. R. Hunziker-Rodewald[1] and D. L. Michels[2]

[1]Strasbourg University, Faculty of Theology and Religious Science, France
[2]KAUST Computational Sciences Group, KSA
[3]Amsterdam University, Faculty of Humanities, History, European Studies and Religious Studies, Netherlands

Abstract

We present DeepHadad, a novel deep learning approach to improve the readability of severely damaged ancient Northwest Semitic inscriptions. By leveraging concepts of displacement maps and image-to-image translation, DeepHadad effectively recovers text from barely recognizable inscriptions, such as the one on the Hadad statue. A main challenge is the lack of pairs of well-preserved and damaged glyphs as training data since each available glyph instance has a unique shape and is not available in different states of erosion. We overcome this issue by generating synthetic training data through a simulated erosion process, on which we then train a neural network that successfully generalizes to real data. We demonstrate significant improvements in readability and historical authenticity compared to existing methods, opening new avenues for AI-assisted epigraphic analysis.

CCS Concepts

• Computing methodologies** → **Mesh geometry models; Reconstruction; Neural networks; • Applied computing** → **Arts and humanities;

1. Introduction

Ancient inscriptions are invaluable sources of historical and cultural information, but their degradation over time poses significant challenges for epigraphic analysis. The Hadad statue inscription, a monumental artifact from the mid-8th century BCE discovered in Zincirli (ancient Sam'al), exemplifies these challenges [Gze15]. The Hadad statue, measuring 2.85 meters in height, contains an inscription by Panamuwa (I) —a crucial text for understanding Sam'alian culture and the development of the Aramaic script [Nie14].

However, centuries of weathering and intentional damage have made many glyphs barely discernible, hindering traditional epigraphic methods. Specific challenges include:

- Varying depth of the relief.
- Vesicular nature of the basalt material.
- Cracks and intentional destruction, interrupting the continuity of the text.
- Severe erosion, blurring the boundaries between glyphs and background.

Although computational techniques have shown promise in the restoration of damaged texts [ASP19, FLAG20], they have limits when dealing with severely eroded 3D inscriptions.

To address these limitations, we present *DeepHadad*, an innovative deep learning approach designed specifically to improve the readability of severely damaged ancient inscriptions. Our method combines:

- 3D geometrical data extraction via displacement maps, which encode depth information as 2D images.
- Synthetic Data Generation using real-world damage patterns.
- A custom train neural network for the restoration of inscriptions.

Using these techniques, *DeepHadad* recovers text from barely recognizable Aramaic inscriptions, improving existing methods in terms of readability and historical authenticity. This paper details our approach, presents quantitative and qualitative results, and discusses implications for future epigraphic studies and the preservation of cultural heritage.

2. Related Work

Recent advancements in computational epigraphy have opened new avenues for the restoration and analysis of ancient texts. These approaches can be broadly categorized into text-based restoration methods and image-based enhancement techniques.

† Arab-German Young Academy of Sciences and Humanities (AGYA), Vorderasiatisches Museum Berlin (VAM)

2.1. Text-based Restoration

Deep learning models have been developed to restore missing or damaged text in ancient inscriptions:

- Assael et al. [ASP19, ASS*22] introduced Pythia, a neural network trained on ancient Greek inscriptions to predict missing characters.
- Fetaya et al. [FLAG20] developed a restoration model for Babylonian cuneiform tablets using contextualized language models.

Although these methods show promise in filling textual lacunae, they primarily operate on transcribed text or 2D images, limiting their applicability to severely eroded 3D inscriptions like on the Hadad statue.

2.2. Image-based Enhancement

Image processing and computer vision techniques have been applied to enhance the legibility of deteriorated inscriptions:

- Nguyen et al. [NBLS21] applied Generative Adversarial Networks (GAN) to denoise ancient Cham inscription images.
- Pan et al. [PLY*20, PLY*22] applied GANs and monocular vision techniques for the restoration of the Borobudur temple.

These approaches have shown success in enhancing 2D representations of inscriptions but often struggle with the complex geometry and varied damage patterns found in 3D artifacts.

2.3. Synthetic Data

A significant challenge in applying deep learning techniques to the restoration of ancient inscriptions is the scarcity of training data. Large datasets of damaged and intact inscriptions are non-existent due to the nature of archaeological artifacts. Moreover, we never have pairs of intact and broken versions of the same glyph instance, which is crucial for training supervised models. This limitation necessitates the creation of synthetic data that can accurately simulate various damage patterns while preserving the characteristics of real inscriptions.

Our work addresses these challenges by combining 3D geometric information (via displacement maps) with a novel approach to generate synthetic data. This novel approach enables the training of restoration models for accurate and historically faithful restorations of damaged 3D inscriptions.

3. Methodology

DeepHadad employs a multi-stage process to restore damaged inscriptions:

3.1. Displacement Map Extraction

We extract displacement maps (DMs) from the 3D scan of the Hadad statue using high-resolution 3D scanning techniques and Blender. Displacement maps encode depth information from the 3D surface as grayscale images, where the intensity of the pixels corresponds to the height of the surface. This allows us to capture the subtle topographical variations of the inscription while working in a 2D image space, which is more amenable to deep learning techniques. We generate DMs for both well-preserved and damaged parts of the inscription, representing individual glyphs and small patches (see Fig. 1).

3.2. Damage Simulation

We employ a multi-stage approach to simulate various types of damage commonly observed in ancient inscriptions. The process can be mathematically represented as:

$$S(I) = C(E(D(I))) \quad (1)$$

Where $S(I)$ is the final synthetic damaged image, I is the original undamaged image, and D, E and C represent simulation functions of deformation, erosion, and cracks.

3.2.1. Erosion Simulation

We simulate erosion using morphological operations [SSP03, Zha22], primarily the erosion operation defined as:

$$E(I) = I \ominus K \quad (2)$$

where $\ominus$ denotes the erosion operation and K is a structuring element. We vary the kernel size and iteration count to simulate different degrees of erosion.

3.2.2. Elastic Deformation

Elastic deformation [SSP03, Zha22, RGG24] is applied to simulate warping and distortion:

$$D(I) = I(x + \delta x, y + \delta y) \quad (3)$$

where $\delta x = G_\sigma(R) * \alpha$ and $\delta y = G_\sigma(R) * \alpha$. G_σ is a Gaussian filter with standard deviation σ, R is a random displacement field, and α controls the intensity of the deformation.

3.2.3. Crack Simulation

Cracks are simulated through a process that adaptively integrates synthetically generated fracture patterns into the displacement map (see Fig. 2). This method modulates crack depth based on local surface features, ensures smooth transitions at crack edges, and maintains physical constraints, resulting in a realistic simulation of damage on the inscription surface.

This multi-stage process allows us to generate a diverse dataset of synthetically damaged inscriptions that closely resemble real-world deterioration, enabling robust training of our restoration model.

3.3. Synthetic Data Generation

To address the scarcity of training data, we develop a procedural approach to generate synthetically damaged glyphs. This process involves:

- Simulating various types of damage (erosion, cracks) on well-preserved glyph samples (see Fig. 3).
- Blending multiple glyphs to create larger, realistic inscription patches.

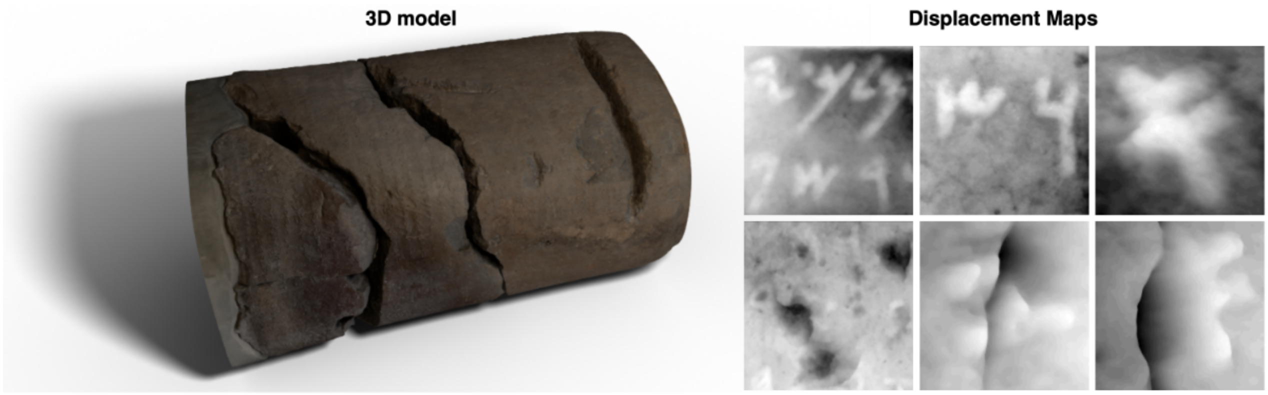

Figure 1: *On the left side, the 3D model of the Hadad statue is rendered. On the right are samples of displacement maps extracted from the model, with the well-preserved sections at the top and the damaged sections at the bottom.*

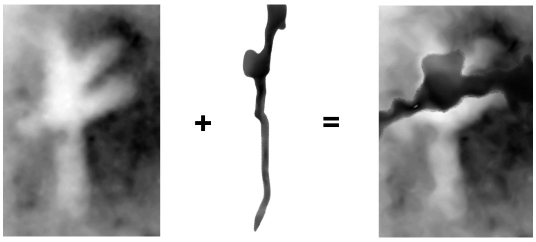

Figure 2: *Realistic crack integration in a displacement map (DM). Left: Original preserved DM of a glyph. Center: Synthetically generated crack pattern. Right: Result of adaptive integration of the crack pattern into the preserved DM, simulating topography-aware damage on the inscription surface.*

- Applying controlled augmentation techniques to increase dataset diversity.

This approach allows us to dynamically generate paired damaged and preserved DM datasets for training.

3.4. Network Architecture

DeepHadad introduces an image-to-image translation network designed specifically for the complex task of reconstruction of ancient inscription displacement maps. Our architecture combines and adapts several state-of-the-art deep learning techniques to address the unique challenges posed by damaged epigraphic artifacts. The core of our model consists of two main components: a generator and a discriminator, both tailored for the complexities of displacement map processing (see Fig. 4).

Key building blocks in our architecture include:

- **Spatially-Adaptive Normalization (SPADE)**: We incorporate SPADE blocks [PLWZ19] in our generator, allowing for spatial adaptivity crucial in preserving the fine geometric details of inscriptions. This technique enables our model to maintain local structural information often lost in traditional GANs.
- **Self-Attention Mechanisms**: Inspired by Vaswani et al. [VSP*17], we integrate self-attention layers in both the generator and discriminator. These capture long-range dependencies within the displacement maps, which is essential to maintaining global consistency in reconstructed inscriptions.
- **Skip Connections**: Adapted from U-Net architectures [RFB15], our skip connections ensure the preservation of fine-grained details throughout the reconstruction process, critical for accurately restoring damaged glyphs.
- **Residual Blocks**: Incorporated in both networks, residual connections [HZRS16] facilitate deeper architectures, allowing our model to learn more complex features of displacement maps.

Although our approach builds upon advances seen in models such as SPADE-GAN [PLWZ19] and U-Net [RFB15], *DeepHadad* is optimized for displacement map reconstruction. This specialization allows us to address challenges specific to epigraphic restoration, such as handling varying degrees of damage, preserving subtle surface textures, and maintaining the historical accuracy of reconstructed inscriptions. For a detailed description of the network architecture and implementation details, we refer readers to our publicly available source code repository at `https://github.com/aioaneia/deep-hadad`.

3.5. Training Process

Our training process is designed to ensure robust and generalizable results:

- We generate a new synthetic dataset of 45,000 paired DMs at the beginning of each epoch.
- The training set comprises 40,000 pairs, with 5,000 pairs reserved for validation.
- We employ adaptive learning rate scheduling, halving the learning rate every 5 epochs.
- Progressive weight adjustments are applied to various loss components throughout training.
- The model is trained for 200 epochs with a batch size of 32.

This dynamic data generation and adaptive training strategy help

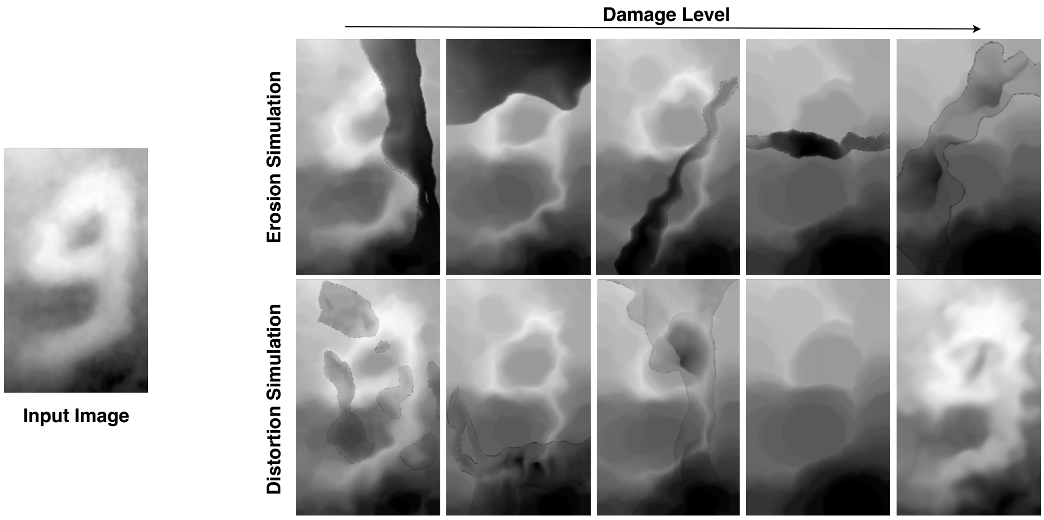

Figure 3: *Synthetic damage simulation on displacement maps of Aramaic glyphs. The* ***Input Image*** *shows the original undamaged displacement map of a glyph. In all images, lighter areas represent higher elevations, whereas darker areas indicate lower elevations or deeper engravings. Two types of damage are simulated:* ***Erosion Simulation*** *(top row) and* ***Distortion Simulation*** *(bottom row). The arrow labeled "Damage Level" indicates increasing severity or complexity of damage from left to right:* ***Erosion Simulation****: Progresses from light surface wear to extreme degradation featuring synthetic cracks.* ***Distortion Simulation****: Advances from minor to significant deformation and synthetic cracks. These diverse damage patterns, generated by our pipeline, enable our model to learn robust restoration techniques applicable to various real-world deterioration scenarios encountered in ancient inscriptions. The use of displacement maps allows us to simulate changes in the 3D structure of the glyphs, providing a more realistic representation of damage compared to 2D image processing alone.*

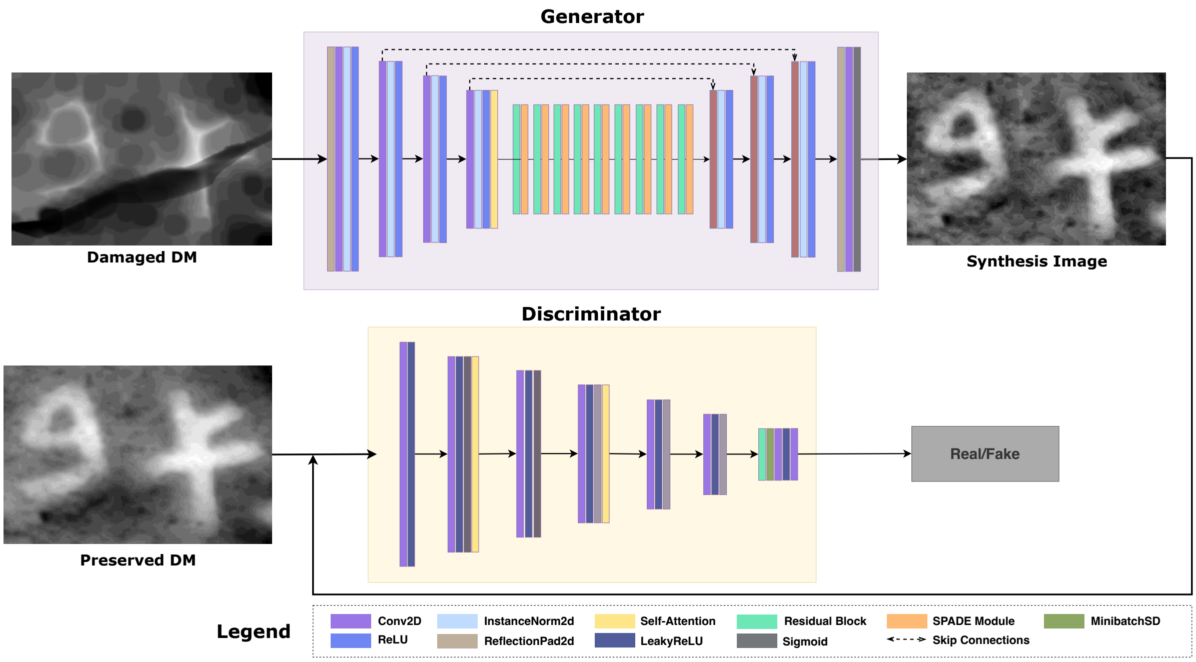

Figure 4: *DeepHadad network architecture. The generator (top) employs initial convolution, downsampling blocks, self-attention, SPADE blocks, and upsampling blocks with skip connections to reconstruct damaged displacement maps. The discriminator (bottom) utilizes convolutional blocks, self-attention layers, and a residual block to assess reconstruction quality. The color-coded legend details the various layer types.*

prevent overfitting and ensure the model's ability to handle diverse damage patterns.

4. Experiments and Results

We evaluate *DeepHadad* against state-of-the-art baselines using both quantitative metrics and expert qualitative assessment.

4.1. Quantitative Evaluation

Table 1 presents a comparison of *DeepHadad* with Pix2Pix and Pix2PixHD in key metrics:

Table 1: *Quantitative Evaluation Results*

Method	PSNR (dB)	SSIM	L1 Distance
Pix2Pix	27.89	0.845	0.021
Pix2PixHD	29.34	0.852	0.020
DeepHadad	**30.45**	**0.882**	**0.017**

DeepHadad outperforms existing methods across all metrics, demonstrating superior restoration quality. To provide insight into the training process, Figure 5 illustrates the progression of these metrics and the progression of the loss functions of the generator and discriminator during training (see below).

4.2. Qualitative Assessment

Figure 6 presents a sample restoration, demonstrating *DeepHadad*'s ability to enhance readability while maintaining historical authenticity.

5. Discussion and Future Work

DeepHadad represents a significant advancement in computational epigraphy, offering new possibilities for analyzing severely damaged inscriptions. The integration of 3D-aware processing with synthetic data and deep learning techniques enables more accurate and historically accurate restorations.

5.1. Limitations and Challenges

While *DeepHadad* shows promising results, it's important to acknowledge its current limitations:

- The model's performance may vary across different glyph styles not represented in the training data.
- Severe damage patterns might still pose challenges for accurate restoration.
- The synthetic data generation process, while highly effective, still produces detectable differences from real damaged inscriptions.

5.2. Future Directions

Building on the foundation of *DeepHadad*, future work will focus on:

- Expanding the synthetic dataset to cover a wider range of damage patterns, improving the model's generalization capabilities.
- Exploring graph neural networks to model spatial and semantic relationships between glyphs. This approach could improve restoration accuracy by leveraging contextual information and long-range dependencies within inscriptions.
- Developing advanced weathering simulations using physics-informed neural networks. This could lead to more realistic synthetic data generation.

6. Conclusion

We presented *DeepHadad*, a novel deep learning approach to improve the readability of ancient damaged inscriptions. Using advanced synthetic data generation and 3D-aware image processing, our method significantly improves existing restoration techniques. *DeepHadad* opens new avenues for AI-assisted epigraphic analysis, improving our ability to decipher and interpret ancient texts.

References

[ASP19] ASSAEL Y., SOMMERSCHIELD T., PRAG J.: Restoring ancient text using deep learning: a case study on Greek epigraphy. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)* (Hong Kong, China, nov 2019), Association for Computational Linguistics, pp. 6368–6375. URL: https://aclanthology.org/D19-1668, doi:10.18653/v1/D19-1668. 1, 2

[ASS*22] ASSAEL Y., SOMMERSCHIELD T., SHILLINGFORD B., BORDBAR M., PAVLOPOULOS J., CHATZIPANAGIOTOU M., ANDROUTSOPOULOS I., PRAG J., FREITAS N.: Restoring and attributing ancient texts using deep neural networks. *Nature 603* (03 2022), 280–283. doi:10.1038/s41586-022-04448-z. 2

[FLAG20] FETAYA E., LIFSHITZ Y., AARON E., GORDIN S.: Restoration of fragmentary babylonian texts using recurrent neural networks. *Proceedings of the National Academy of Sciences 117* (2020), 22743 – 22751. URL: https://api.semanticscholar.org/CorpusID:211990752. 1, 2

[Gze15] GZELLA H.: *A Cultural History of Aramaic: From the Beginnings to the Advent of Islam.* Brill, Leiden, Netherlands, 2015. 1

[HZRS16] HE K., ZHANG X., REN S., SUN J.: Deep residual learning for image recognition. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2016), 770–778. URL: https://api.semanticscholar.org/CorpusID:206594692. 3

[NBLS21] NGUYEN T.-N., BURIE J.-C., LE T.-L., SCHWEYER A.-V.: On the use of attention in deep learning based denoising method for ancient cham inscription images. In *Document Analysis and Recognition – ICDAR 2021* (Cham, 2021), Lladós J., Lopresti D., Uchida S., (Eds.), Springer International Publishing, pp. 400–415. 2

[Nie14] NIEHR H.: The katumuwa stele in the context of royal mortuary cult at samal. In *In Remembrance of Me: Feasting with the Dead in the Ancient Middle East*, Herrmann V. R., Schloen J. D., (Eds.), no. 37 in Oriental Institute Museum Publications. The Oriental Institute of the University of Chicago, 2014, pp. 57–62. 1

[PLWZ19] PARK T., LIU M.-Y., WANG T.-C., ZHU J.-Y.: Semantic image synthesis with spatially-adaptive normalization. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2019), pp. 2332–2341. doi:10.1109/CVPR.2019.00244. 3

[PLY*20] PAN J., LI L., YAMAGUCHI H., HASEGAWA K., THUFAIL F. I., TANAKA S.: Fused 3d transparent visualization for large-scale cultural heritage using deep learning-based monocular reconstruction.

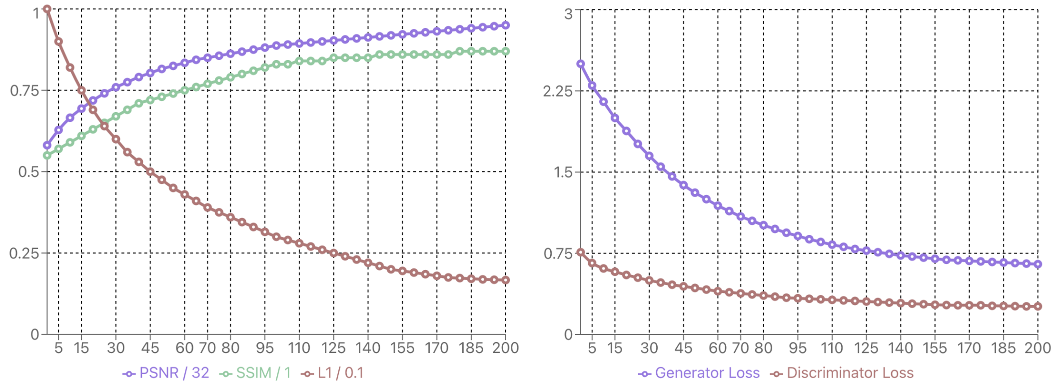

Figure 5: *Training progression of DeepHadad over 200 epochs. Left: Normalized performance metrics. PSNR is scaled by 1/32, SSIM by 1 (unchanged), and L1 distance by 1/0.1. This normalization allows for direct comparison of metrics' improvements. PSNR and SSIM show steady increases, while L1 distance decreases, indicating improving restoration quality. Right: Generator and Discriminator loss functions. Both losses decrease and converge, suggesting a balanced adversarial training process. The generator loss shows a more pronounced decrease, reflecting the model's improving ability to produce realistic restorations.*

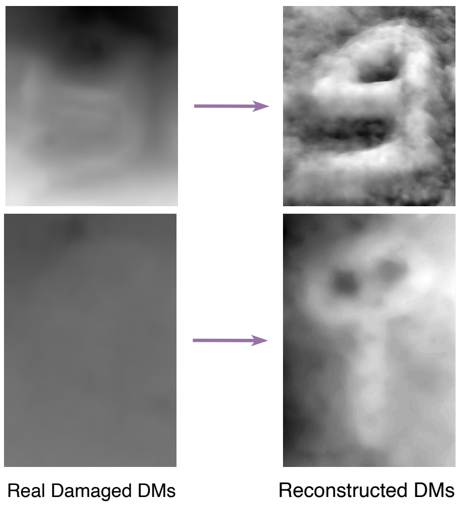

Figure 6: *Sample reconstructions of damaged Aramaic glyph displacement maps (DMs) using DeepHadad. Left column: Real damaged DMs showing severely eroded glyphs. Right column: Reconstructed DMs demonstrating the model's ability to restore glyph shapes. Top row shows a more legible initial state, while the bottom row illustrates reconstruction from a heavily degraded input.*

ISPRS Annals of Photogrammetry, Remote Sensing and Spatial Information Sciences V-2-2020 (08 2020), 989–996. doi:10.5194/isprs-annals-V-2-2020-989-2020. 2

[PLY*22] Pan J., Li L., Yamaguchi H., Hasegawa K., Thufail F. I., Brahmantara, Tanaka S.: 3d reconstruction of borobudur reliefs from 2d monocular photographs based on soft-edge enhanced deep learning. *ISPRS Journal of Photogrammetry and Remote Sensing 183* (01 2022), 439–450. doi:10.1016/j.isprsjprs.2021.11.007. 2

[RFB15] Ronneberger O., Fischer P., Brox T.: *U-Net: Convolutional Networks for Biomedical Image Segmentation*. Springer International Publishing, Cham, 2015, pp. 234–241. URL: https://doi.org/10.1007/978-3-319-24574-4_28, doi:10.1007/978-3-319-24574-4_28. 3

[RGG24] Risager K. V. E., Gholamalizadeh T., Ghazi M. M.: Non-reference quality assessment for medical imaging: Application to synthetic brain mris. URL: https://api.semanticscholar.org/CorpusID:271329036. 2

[SSP03] Simard P., Steinkraus D., Platt J.: Best practices for convolutional neural networks applied to visual document analysis. *7th International Conference on Document Analysis and Recognition (ICDAR 2003)* (01 2003), 958–962. doi:10.1109/ICDAR.2003.1227801. 2

[VSP*17] Vaswani A., Shazeer N., Parmar N., Uszkoreit J., Jones L., Gomez A. N., Kaiser Ł., Polosukhin I.: Attention is all you need. In *Advances in Neural Information Processing Systems* (2017), pp. 5998–6008. URL: https://api.semanticscholar.org/CorpusID:13756489. 3

[Zha22] Zhang Y.-J.: *2D Computer Vision: Principles, Algorithms and Applications*. World Scientific Publishing, Singapore, 2022. URL: https://www.worldscientific.com/worldscibooks/10.1142/12497. 2