



# Using Generative AI for Reconstructing Cultural Artifacts: Examples Using Roman Coins

RESEARCH ARTICLE

MARK ALTAWEEL

ADEL KHELIFI

MOHAMMAD HASHIR ZAFAR

\*Author affiliations can be found in the back matter of this article

ubiquity press

## ABSTRACT

Generative AI, propelled by innovations like ChatGPT, has gained widespread recognition. In the realm of archeology, there exists significant potential for generative AI, particularly in reconstructing the appearance of cultural artifacts through the introduction of 2D or 3D renderings derived from damaged or degraded objects. In this study, we showcase and evaluate the practical application of Generative Adversarial Networks (GANs), harnessing the power of deep learning, for 2D image reconstruction of ancient Roman coins, aimed at aiding their improved visualization. Roman coins are chosen as our focal point due to their relative abundance and accessibility through online repositories and datasets. Our results demonstrate improved ability to enhance damaged or degraded coins, rendering them more similar to their better-preserved counterparts. In some instances, generated coins are virtually indistinguishable from the originals. The contribution of this work showcases the potential of GANs in assisting cultural heritage specialists and archeologists in recreating the appearance of damaged objects, thereby facilitating improved visualization of coins that are not well preserved. However, we also discuss when there might be limitations to using GANs in reconstructions. Although this work is tailored for ancient coins, GANs hold promise in application for other artifacts provided sufficient training data are available. We discuss how GANs can be applied and improve the appearance of artifact reconstructions, where we also provide relevant data used in this research.

## CORRESPONDING AUTHOR:

**Mark Altaweel**

Institute of Archaeology,  
University College London, UK  
m.altaweel@ucl.ac.uk

## KEYWORDS:

generative; artificial  
intelligence; adversarial  
network; coins; Roman;  
artifacts; reconstruction

## TO CITE THIS ARTICLE:

Altaweel, M, Khelifi, A  
and Zafar, MH. 2024.  
Using Generative AI for  
Reconstructing Cultural  
Artifacts: Examples Using  
Roman Coins. *Journal of  
Computer Applications in  
Archaeology*, 7(1): 301–315.  
DOI: [https://doi.org/10.5334/  
jcaa.146](https://doi.org/10.5334/jcaa.146)

## 1. INTRODUCTION

The reconstruction and enhanced visualization of damaged or incomplete cultural artifacts has been a long-sought goal within cultural heritage and archeology. This has included varied techniques applied such as 3D object reconstruction, virtual reality, mathematical algorithms, including diffusion models, and other digital reconstruction techniques (Koutsoudis *et al.* 2013; Bruno *et al.* 2010; Kleber & Sablatnig 2009). One recent trend is using generative artificial intelligence (AI) to create new 2D or 3D images, using algorithms such as generative adversarial networks (GANs; Aggarwal *et al.* 2021). Generative AI has been associated with popular large language models such as ChatGPT or even so-called 'deepfakes' that create false images that often appear very similar to a likely or original image (Wach *et al.* 2023).

While the desire to reconstruct objects serves purposes ranging from aiding in object identification to facilitating public presentation, visualization and conservation efforts, generative AI approaches for the restoration of incomplete or deteriorated artifacts is only now beginning to be explored in archeology (Colmenero-Fernández & Feito 2021). We posit that there is substantial untapped potential in harnessing Generative Adversarial Networks (GAN)-based methodologies to improve the appearance and quality of damaged artifacts, enabling degraded objects to be reconstructed so that they might be more clear or improve their physical appearance. We seek to evaluate the potential of generative AI in reconstructing and enhancing visualization of cultural objects. The unique strength of GANs lies in their ability to autonomously generate lifelike images that closely emulate real-world objects, drawing from comprehensive training datasets (Choi *et al.* 2019). This benefits viewers through enhanced visualization and providing a more clear picture of artifacts that might be damaged or degraded. Although 3D imagery can be employed for object reconstruction, 2D images often represent the most prevalent and easily accessible form of data. As with any generative AI approach, extensive training data are essential; data augmentation techniques further enhance the quality of training datasets (Farahanipad *et al.* 2022).

We explore and present a GAN-based methodology tailored to reconstruct artifacts, imbuing them with a likeness to better-preserved objects or improved appearance. Our objective is to unveil and evaluate the potential of GANs for artifact and object reconstruction and ascertain their relevance for reconstructing objects so that overall appearance is improved for damaged objects. Our work includes demonstrating GANs' ability to improve images so that features are more clear for general visualization. We have chosen to illustrate our

approach using 2D images of Roman coins because of their prevalence and the wealth of available data on both fully preserved and partially intact specimens. Although we use coins in our example, we are not numismatics experts, but we incorporate archaeologists familiar with coins and use coins as an illustrative example for the potential of GANs in cultural object reconstructions and visualization.

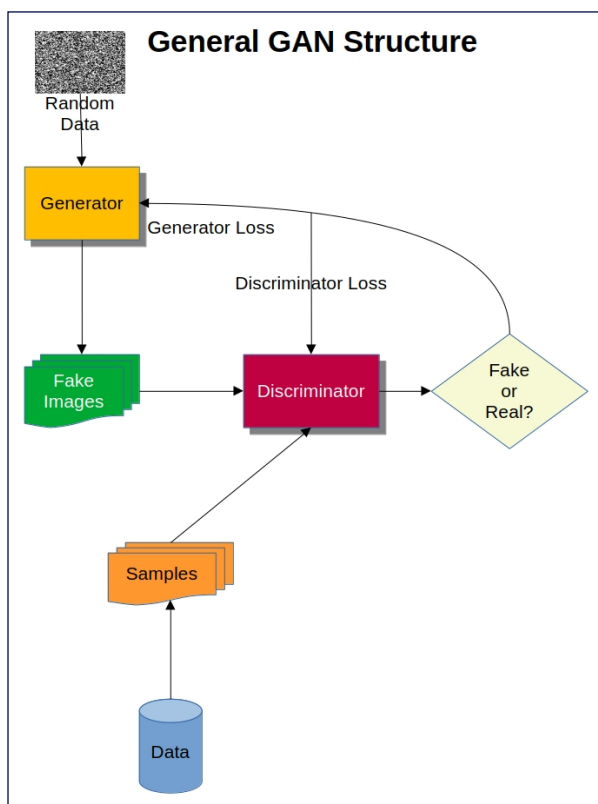
We did conducted a literature search and evaluated different GAN-based methods. Prior to delving into the specifics of our approach, we provide an overview of the broader landscape of generative AI, with a focus on its relevance for object reconstruction in archeology. Subsequently, we describe the methodology and the algorithms employed. Following this, we present our findings. We discuss the broader implications of our results for the integration of generative AI approaches for image-based object reconstruction and provide insights into the benefits and limitations of the approach for visualization. This includes how generative AI approaches could be improved for artifact reconstruction while also developing clear guidelines for their use so that the creation of fake images for deception is avoided.

## 2. BACKGROUND

### 2.1. CONCEPT BEHIND GANs

Given that GANs are relatively new in archeology, it is essential to understand the fundamental principles of GANs and what they can potentially provide the field. At their core, GANs represent a form of deep learning that harnesses the power of two artificial neural networks: the discriminator and the generator. These networks engage in a dynamic interplay resembling a zero-sum game. To comprehend this concept fully, artificial neural networks need to be understood, which are a subset of deep learning approaches. Artificial neural networks aim to replicate the intricate processes of human learning by creating interconnected nodes, analogous to neurons in the human brain. These nodes are structured in layers, facilitating data flow through the network, thereby enabling the identification and learning of specific patterns, such as those found in images or objects within images (LeCun *et al.* 2015). In the realm of GANs, the generator's primary function is to produce new, generated images. The generator's output is fed into the discriminator's input. The discriminator endeavors to differentiate between newly generated images and those from a training dataset. The ultimate goal is for the generator to create data that closely mimics the training data to the extent that the discriminator cannot distinguish between the two. When this equilibrium is achieved, this signifies a converged GAN model (Goodfellow *et al.* 2014).

Figure 1 provides an example of a basic (i.e., so-called ‘vanilla’) GAN structure. In the initial stages of training, the generator produces data that significantly deviate from the discriminator’s data. Random input serves as the basis for generating these data. However, as training progresses and more adjustments are made based on previous training, the discriminator becomes more error-prone in distinguishing between generated and real images. Through a process known as backpropagation, the generator adjusts its weights based on feedback from the discriminator. This iterative process leads to the GAN’s refinement in subsequent iterations. The generator incorporates a loss function that penalizes it for failing to deceive the discriminator, while the discriminator acts as a classifier that distinguishes between generated and real input data. The discriminator’s loss function measures classification accuracy, penalizing it for classifying real instances as fakes, and vice versa. Weight adjustments based on outcomes are integral to this process. Both the generator and discriminator are typically trained separately and kept constant while the other is trained. As the generator improves, the discriminator’s performance deteriorates. A converged model is typically achieved when the discriminator’s accuracy oscillates around 50%; this signifies the inability to confidently distinguish real from generated images (Gonog & Zhou 2019; Langr & Bok 2015).



**Figure 1** General structure and workflow of a basic ‘vanilla’ GAN showing the role of the generator and discriminator in generating ‘fake’ images and using real image data to compare with generated images.

Various GAN types exist, primarily differing in their underlying structures and generator and discriminator configurations, including how image reconstruction is generated based on given input. Some GANs focus on certain aspects of reconstruction, such as adding or infilling missing areas in images, while others focus on enhancing provided features such as faces or clothing (Gagnaniello *et al.* 2021).

## 2.2. USE OF GANs IN HERITAGE

Generative AI, particularly using GANs, has experienced a surge in popularity across diverse domains, including advertising, filmmaking, and autonomous driving (Yan *et al.* 2023). GANs are often popular for reconstructing and enhancing imagery; they also can create fake or derived images which can be used to train other machine learning models (Zhang *et al.* 2024; Wang *et al.* 2024; Hermoza & Sipiran 2018a). While GANs have gained traction in numerous industries, their adoption in archeology, or heritage more broadly, has been relatively limited until recently. There is recognition that generative AI will become increasingly adopted for a variety of areas in heritage, including in artifact reconstructions (Münster *et al.* 2024). Material culture reconstruction, in fact, appears to be among the more likely areas of application. GANs are now increasingly being adopted and finding favor within some archeological or heritage-related areas due to their proficiency in augmenting or recreating missing data. Other approaches have also looked at improving visual experience for viewers. Garozzo *et al.* (2021) introduced a GAN-based approach to transform unrealistic representations into realistic classical architectural scene reconstructions. Their methodology enables the synthesis of objects and their spatial alignment within predefined ontological frameworks, contributing to the creation of more authentic visual narratives. The work was deployed to create realistic images so that they can be used for object classification. In another context, the Z-GAN architecture was employed to transfer 2D image features onto a 3D voxel model, facilitating the learning of intricate 3D shapes commonly encountered in architectural structures (Kniaz *et al.* 2019). This model generates voxel models using object silhouettes from an input image and information during a training stage. A U-Net generator is used to translate between 2D color images and 3D models.

For smaller cultural objects, perhaps one of the earliest applications is a GAN called ORGAN, which is a conditional GAN variant used in reconstructing cultural objects. This basically applies a shape completion network that represents 3D objects as a voxel grid; it uses a 3D encoder that compresses input voxels using a series of 3D convolutional layers. Notably, ORGAN demonstrates the ability to reconstruct cultural artifacts, specifically ceramics missing up to 50% of their surface area, with more damaged objects less clearly reconstructed. This approach facilitated 3D reconstructions, employing dual

loss functions to train on missing ceramic parts (Hermoza & Sipiran 2018b). GANs have also been instrumental in tasks such as ceramic profile reconstruction and 3D volumetric reconstruction, streamlining the recreation of ceramic profiles and aiding in volume determination for various vessels (Navarro *et al.* 2022; Colmenero-Fernández & Feito 2021). Recent work has also applied GAN-based architecture for profile drawings in small objects. This deploys a multi-branch feature cross fusion (U<sup>2</sup>FGAN) algorithm for generating line drawings for different cultural remains, which enables line extraction and edge detection (Zeng *et al.* 2024).

For numismatics, Cycle-Consistent GANs, or CycleGANs, have been deployed to reconstruct Roman coins from degraded 2D images. Although this approach predominantly enhanced surface details for eroded coins, such as facial features and legends, it did not address other forms of damage including missing structural components, cracking, and discoloration. In this approach, we attempt to improve upon this as outlined and discussed below by broadening the types of reconstructions applied. Nevertheless, previous CycleGAN application demonstrates how surface details that are deteriorated on coins could be enhanced and made clearer for visual presentation (Zachariou *et al.* 2020). The architecture of CycleGANs is discussed below.

In addition to physical artifacts, GANs have found utility in ancient written content reconstruction. For example, GANs have been used to recreate bone inscriptions and illuminate their evolutionary processes. By simulating different character development stages, GANs can trace the evolution of script forms, shedding light on what these ancient inscriptions might have originally looked like (Chang *et al.* 2022). This highlights also the potential for GANs in contributing to areas such as ancient languages or inscriptions' histories. Overall, the last few years have demonstrated an accelerated use of GANs in heritage and archeology, but work has still been generally limited. In particular, data availability limitations and the technical nature of applying GANs are likely preventing their wider application, although we anticipate the adoption of GANs will accelerate in coming years. For many GAN-based efforts, specific GAN architectures enable these works to focus on varied reconstruction, that is how given and different data are generated to produce a new image. For instance, in the case of voxels, generating from 2D images is possible by focusing on the translator and image generation capabilities. These all vary from the basic generator and discriminator networks in 'vanilla' GANs.

### 3. DATA AND METHODS

We procured data through web scraping from the Online Coins of the Roman Empire (OCRE) site, which

accesses Nomisma-based websites. Coins range from being well preserved to heavily degraded. This enabled other repositories to be accessed to provided data used for model training and qualitative validation and testing. Furthermore, we utilize lightly to heavily eroded or damaged coins in our testing phase to assess our model's capability to reconstruct coin images and determine the ability of the deployed GAN to improve coin reconstructions. We considered damaged coins as those that are clearly broken from their original appearance, eroded, or cracked. The degree of damage varied in our recovered coin samples; we utilize light to severely damage coins based on the surface area affected in the results discussed below. Such damage could have been deliberate in antiquity; however, we aim to reconstruct damaged areas to an appearance as close as possible to when the object was made. In total, our training data comprises 351 coin images, while 110 coin images are reserved for validation and testing purposes. Our intention is to use a GAN to both recreate damaged or broken areas, including cracks, as well as enhance coin features that might be degraded.

#### 3.1. DATA

Our dataset, utilized for both training and validation to assess the model's fidelity to real coin images, is provided in the supplementary data. We sourced the dataset from Nomisma-based websites containing Roman coin images. We deliberately confined the dataset to 2D images; these are the most prevalent and align with our primary goal: facilitating the reconstruction and visualization of eroded or degraded coins. Our dataset encompasses coin images from the Roman Empire (31 BCE – 476 CE), including different denominations (e.g., *denarius*, *dupondius*), and includes coins from mostly Roman emperors. We do not focus on any period or coin types. However, we searched different periods and emperors within the date range to obtain a diverse dataset. We selected coins based on good quality images, well preserved condition, coming from known collections, and those that are not well preserved (Figure 2). This is in order for training to have a variety of available data that aids reconstruction, based on less well preserved coins but also using better preserved coins to improve reconstruction. Data restrictions limit potential to reconstruct coin images. For instance, relatively shiny well preserved images of complete coins informs the trained GAN to then attempt to generate coins that appear more shiny. Based on this criteria, we were able to collect our images, including all training and validation data, which are made available in the supplementary data. While we restricted the coin variety for better control of image quality during testing and validation, it is entirely feasible to expand the dataset to encompass coins from different historical periods and types. The dataset comprises colored images, including obverse and reverse views, with annotated features for

enhanced analysis. In our examples provided, we focus on the obverse. We acknowledge that because we are limited to what is available online, both training and validation are affected. Very poorly preserved coins, where identification is nearly impossible, commonly found in excavations are not as prevalent in our data, although examples of coins not well preserved are used in this work.

Different emperors and dates were searched, although our search was somewhat random because quality, that is well preserved coins, and variety of samples differed based on availability. Coin collections are accessible via [Nomisma.org](http://Nomisma.org) SPARQL, where we accessed the [numismatics.org/ocre](http://numismatics.org/ocre) site for coin data (Nomisma 2024). To facilitate data collection, we developed our own custom scraper for this site (see supplementary data) that also enabled access to other coin repositories in museum collections that linked back via [numismatics.org](http://numismatics.org). This is provided as an additional output of this paper and information about its use is discussed in the given link. Our data search and acquisition process spanned and connected to the British Museum, the

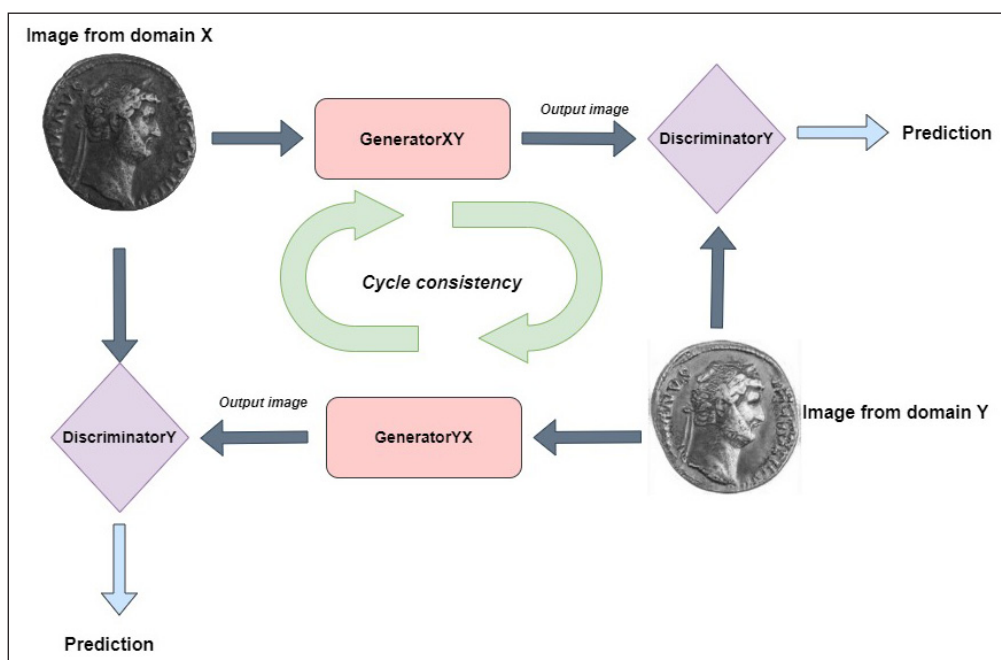
Portable Antiquities Scheme, the University of Vienna’s numismatic collection, the Kunsthistorisches Museum in Vienna, and the Münzkabinett in Berlin (American Numismatic Society 2023a, b; British Museum 2023; Portable Antiquities Scheme 2023; University of Vienna 2023; Kunsthistorisches Museum Wien 2023; Berlin 2023). We used the American Numismatics Society’s referenced coins for training and testing given the quality, diversity, and provided coin reference. All coin data utilized in this effort have metadata associated with their identifier numbers; these identifiers can be found in the image titles and link to their original files found in [numismatics.org](http://numismatics.org). The supplementary data contain training and generated output used for this effort. The link provided further explains our data.

### 3.2. GAN-BASED METHOD

We deployed a CycleGAN (i.e., short for Cycle-Consistent GAN; code and documentation link provided in the supplementary data), similar to prior deployments (Zachariou *et al.* 2020), as this model proved to be useful for our own goals in enhancing visualization



**Figure 2** Example coins divided into deteriorated and well preserved coin categories (or ‘bad’ and ‘good’ coins) used to train our GAN. Others can be found in the supplementary data.



**Figure 3** The CycleGAN-based approach applied in reconstructing coins.

(Figure 3). This is mainly because CycleGANs deal well with multifaceted reconstructions and one of the contributions of this paper is to demonstrate how multiple enhancements are possible. We had searched for the best underlying architecture to use and this architecture proved to be the most capable in coin reconstruction quality. Among various GANs evaluated, this GAN-based method seemed the best option given its ability to address different forms of degradation in reconstructing coins. We should note that a GAN used in this case is for enhancing given input data; other data may exist for the same artifact that could of course prove to be of a higher quality image and may show additional detail the GAN does not initially see in its enhancements. In addition to improving surface details, we wanted to be able to provide a more complete reconstruction of 2D coin images that enhanced broken or missing areas. This includes functionality such as infilling missing areas on coins. Previous work on CycleGANs focused on a limited range of enhancements, particularly eroded surface detail improvements. We wanted to determine if a broader set of feature enhancements is possible and CycleGANs proved to have this capability. The degree and extent of how damaged coins could be for reconstruction to be possible or the degree to which coins can be reconstructed based on how damaged they are is important to determine where advantages and limitations of CycleGAN are. Overall, this allows us to determine how effective CycleGAN is for damaged object reconstructions and address various types of coin damage.

The CycleGAN approach harnesses both generator and discriminator networks for training purposes; it is a generative adversarial tool for image-to-image translation tasks, where the goal is to learn the mapping between two image domains using paired training data. Obtaining such paired data can be difficult and expensive. CycleGAN can translate images between two domains without the need for paired data. CycleGAN is used in popular tools such as ArcGIS; between two image domains  $X$  and  $Y$ , CycleGAN learns a mapping  $G : X \rightarrow Y$  where the generated images from  $G : X$  are intended to be images comparable to domain  $Y$ . Another mapping  $F : Y \rightarrow X$  is also learned in order to translate an image from domain  $Y$  to domain  $X$  (ArcGIS 2023; Harms et al. 2019). This feedback between  $X$  and  $Y$  generated images becomes cycle consistent.

In using augmentation for training, several steps are conducted. This includes random crop positioning, that is, selecting random sub-regions of the image to ensure that the same image position in training does not always appear. The method also uses random horizontal flip which flips the given image and helps vary samples' orientation for training. Resizing is applied to deploy consistent image dimensions; normalizing pixel values helps to facilitate faster convergence in training. We also

use standard deviation for pixel values to scale pixels for training and create variety image data for training.

Table 1 lists key hyperparameters employed in CycleGAN and that are required to execute training. We utilized a variety of settings in the hyperparameters and the results given reflect our best output. The number of epochs refers to training cycles or passes for the GAN model. Image dimensions relates to the image resolution used for training data. The loss function used is a least squares function used in the discriminator for the GAN and used for evaluating the training process. This penalizes any misclassification between real and generated images during training. The patch size refers to the pixel dimensions used in training. The batch size refers to the samples in training that are propagated through the network. The initial learning rate defines the pace where the algorithm updates/learns and responds to the estimated error in training when updating the training model. Both the discriminator and generator have their own learning rates, where the discriminator responds to classification between real and generated images and the generator produces a randomized sampling to make a new image. This helps define step size during training; both the discriminator and generator apply the same learning rate in this case (Ghosh et al. 2020; Zachariou et al. 2020; Kurach et al. 2019). Data augmentation is used to avoid overfitting and poor results from limited data.

The following summarizes the key steps involved in training and deploying our model. For CycleGAN to accurately produce translations of the input image, what is known as cycle consistency, the idea is that if an image is translated from domain  $X$  to  $Y$ , it should be possible to translate the image back to domain  $X$  from  $Y$ . In our case, one generator ( $XY$ ) converts images of deteriorated coins to good coins, and vice versa for the other generator  $YX$ . Generator  $XY$  excels in recreating surface details and reconstructing worn or missing parts of the coin. If a deteriorated coin is fed into generator  $XY$  producing an image of a good coin, generator  $YX$  will use the good coin as input to produce an initial deteriorated coin similar to the original. Cycle consistency uses an additional loss to measure the difference between the initial deteriorated coin and the output produced by generator  $YX$ , thus

INPUT PARAMETER	VALUE
Epochs	155
Image Dimensions	256 × 256
Loss Function	Least Squares GAN
Patch Size	196 × 196
Batch Size	1
Initial Learning Rate	0.0002

**Table 1** Relevant hyperparameters for the CycleGAN deployed.

training the generators to produce accurate translations. The steps, as indicated in Figure 3, are summarized below.

*Adversarial Loss Process:*

1. GeneratorXY translates an image of a deteriorated coin to a good coin.
2. DiscriminatorY compares the translated image from generatorXY with original images of good coins and outputs whether the generated coin is fake or authentic.
3. GeneratorYX translates an image of a good to a deteriorated coin.
4. DiscriminatorX compares the image from generatorYX with original images of deteriorated coins and outputs whether the generated coin is fake or authentic.

*Cycle consistency process:*

1. GeneratorXY inputs deteriorated coinA and generates good coinB. CoinB is used as input to generatorYX which generates deteriorated coinC. Both coinA and coinC are compared and their loss is calculated.
2. GeneratorYX inputs good coinX and generates deteriorated coinY. CoinY is used as input to generatorXY which generates good coinZ. Both coinX and coinZ are compared and their loss is calculated.

In reconstruction training, instance normalization plays a pivotal role, ensuring that each example coin feature significantly contributes to the training process. This approach accommodates a diverse range of coins, with mean and variance calculated over spatial locations in the feature map (Ulyanov *et al.* 2017). We employ ResNet as the foundational network architecture, incorporating nine residual layers for given images. The discriminator network, following the PatchGAN paradigm, consists of five layers using pixel input (Ya-Liang *et al.* 2019). The pixel quality for input training images constrains the overall quality of the reconstruction outputs. Coins are also converted to grayscale prior to training to minimize light-based effects on generated coins; we noticed lighting can vary for coins in databases, which can alter detail and appearance of coins both in the generated and empirical data. Grayscale allowed images to be more equal and minimizes this effect.

We conducted 155 training epochs, which is somewhat different than the approach employed by Zachariou *et al.* (2020). Based on our observations, this number of epochs typically generates synthetic images closely resembling real coins. To optimize overall performance, the Adaptive Moment Estimation (ADAM) optimizer is used, with parameters set to 0.5 and 0.99, respectively (Jiang & Sweetser 2022). The learning rate is set at 0.0002 for the discriminator and generator. Consistent with prior

CycleGAN applications, we identified that discriminator learning can lead to overfitting, while variations in color and lighting conditions impact learning rates and accuracy. Further detail on the CycleGAN method are provided by Zhu *et al.* (2020) and Zachariou *et al.* (2020). We provide the final trained model in the supplementary data.

### 3.3. MODEL EVALUATION

After completing the results, reviewers are used to evaluate output generated images and their quality. Two distinct tests are applied in this work. Evaluators are asked to view coloration, edges and surface reconstruction, reconstruction of writing, reconstruction of cracks and chips, improvement on surface wear, and overall quality. These evaluators are familiar with ancient coins but are not specialists. The evaluation of training and output results is critical in demonstrating the utility of GANs to specific problems, where the generated output is evaluated against real images (Park *et al.* 2023; Parmar *et al.* 2022). Researchers have used and propose various methods for evaluation of GAN results using quantitative and qualitative measures. Traditional methods for accuracy and sensitivity of AI-based methods, such as precision-recall, are not feasible here because no coin is exactly the same in the Roman period. GAN-based methods are not always easy to evaluate using quantitative methods as results can be subjective or relative to user perceptions.

In our first test, qualitative accuracy is measured by presenting 20 generated and real coins and asking 5 independent judges to determine if a coin is real or fake using rating and preference judgment in a so-called rating and preference judgment method (Borji 2019). This allowed us to see if generated images are seen as similar or different from real coin images, where the independent judge rates the images. The first test shows how well generated coins appear realistic to viewers, that is if they are able to easily distinguish between real or fake coins. For all tests conducted, data and the evaluation forms used are provided in the supplementary data. A second, final test is used. In this case, the test is conducted to check the generated images' quality against real coin images that are deteriorated. This test asked 3 respondents if they notice any improvement on 22 coin images relative to the original, degraded coins used as the basis for generated images. Respondents rate reconstructions on a scale from 1 to 5; 1 reflects no visible improvement, 2 represents slight improvement, 3 is clear improvement, 4 is good quality improvement, and 5 is excellent improvement. This test specifically focuses on reconstruction quality. In summary, two tests are used. The first simply differentiated if the generated coins are realistic or not and the second focuses on the quality of reconstruction and if they improved deteriorated coins' features.

### 4. RESULTS

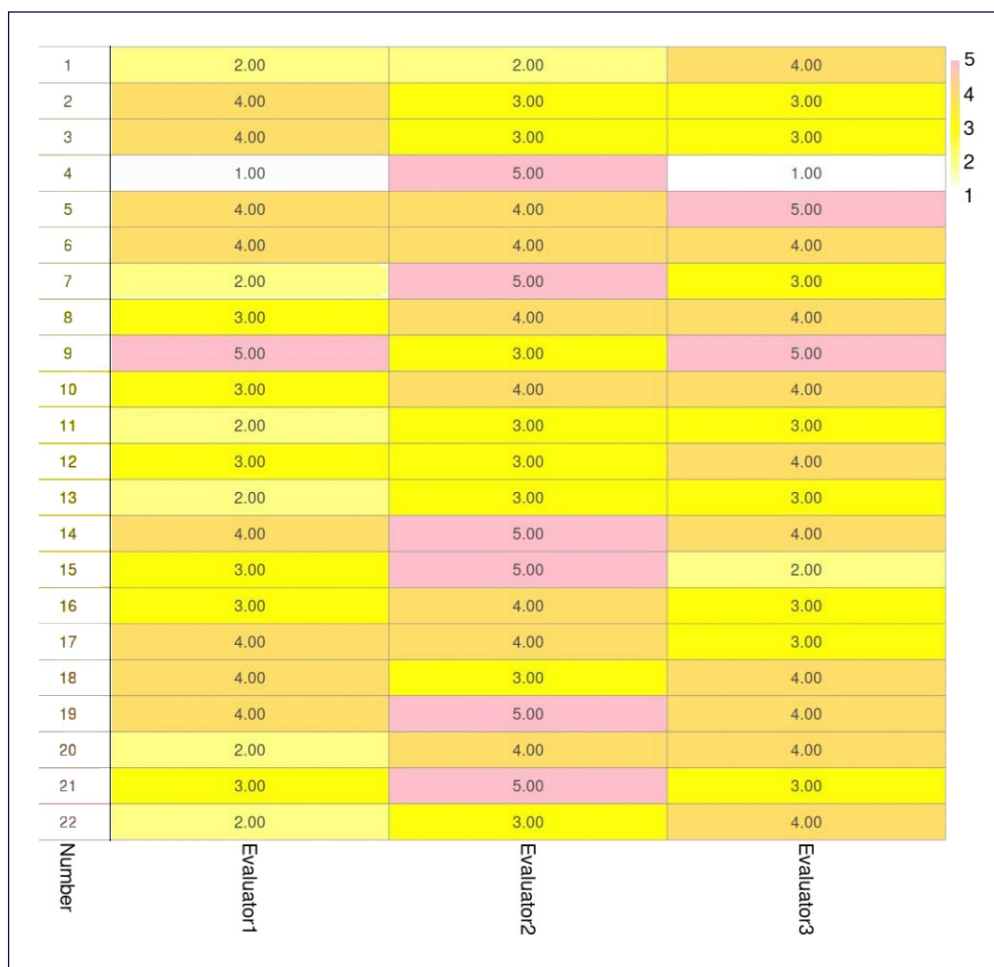
We present results that underscore the capabilities of our approach (Table 2). For the first test discussed in the methods, positive identifications, that is correctly distinguishing if a coin is real or fake, range between 55% and 35%, with the average being 46% accuracy in distinguishing real and fake coins. This shows that the ability to distinguish real coins from generated ones is not easy and demonstrates our GAN can produce realistic-looking coins not easily distinguishable from real coins. We acknowledge that our chosen testers are not numismatic experts, but all evaluators are

EVALUATOR	ACCURATE IDENTIFICATION
Evaluator 1	35%
Evaluator 2	50%
Evaluator 3	45%
Evaluator 4	45%
Evaluator 5	55%

**Table 2** Results (in percent) from the first test checking accuracy in distinguishing real and generated coins using 20 coins.

knowledgeable about heritage and archeology, including coins used in the Roman and other periods. The aim we had is to produce realistic looking coins that enhances appearance and reconstruction, which is how we focused our effort and evaluation. Coin experts may be better able to find differences between generated and real coins, but our goal is to generate realistic looking artifacts that improve visual appearance. The result achieved, we believe, demonstrates realistic-looking output; we want to also determine if our generated coins could be used to enhance or improve what damaged coins should look like if they were better preserved.

For the second test, which checks reconstruction quality, damaged real coins are used for reconstructing coins (i.e., each test set had 1 degraded and 1 generated coins based of the degraded coin; Figure 4). Degraded coins vary from only minor degradation to a high degree of degradation (i.e., more than 30% surface damage or degradation as well as missing parts). We evaluate if generated coin images demonstrate improved facial features, more clear letters, clothing and other items appearing on coins. Overall, the average was 3.5, using the 1–5 scale for reconstruction quality, for the randomly selected (see supplementary data) coins generated, where almost all of the coins show at least some



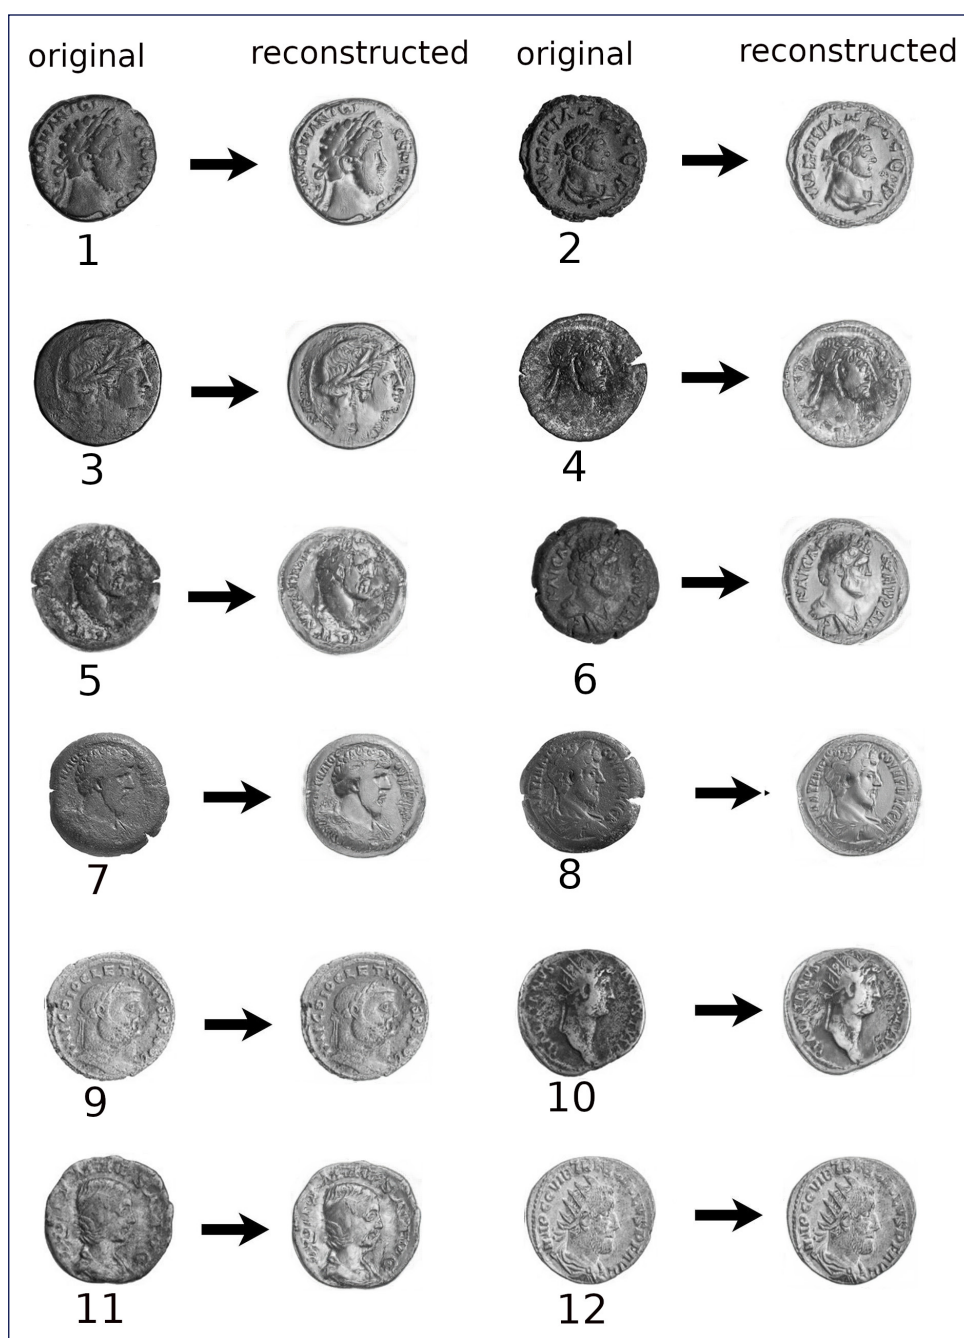
**Figure 4** Second test performed showing respondent results evaluating reconstructions and checking for visual improvement quality (1–5; 1 reflects no improvement and 5 reflects excellent improvement).



improvement. Coin 4 (CoinsEvaluation file), if we exclude Evaluator 2, does not show improvement, but this is mainly because the coin was already not very damaged. The majority of the results indicate that our approach is seen to improve coin reconstruction quality, where the samples had varying degrees of different damage or wear. However, ranges in improvements vary between 1–5, indicating there is variability in our random sample, reflecting that not all reconstructions improved clarity to a great extent. Generally, coins that are heavily damaged (e.g., 11) are likely to yield poorer results, but in most cases those varied between 2–4 in ratings, showing at least some improvement. Well preserved coins also only showed slight improvement in quality (e.g., 4). Evaluator 2 seemed potentially overly positive, but for many

examples the evaluators had similar opinions. Most of the testers are not coin experts, but the overall result helps to demonstrate visual improvement in the coins as identified by individuals with some knowledge on coins. These reviewers are comparable to informed, but non-expert viewers determining if the overall output is an enhancement.

We want to further demonstrate current capability by providing some reconstructed coin examples from our total output that contains original degraded or damaged ancient coins and new generated coins; the obverse side is used in the examples. Figure 5's numbers are linked to their numismatics.org identifiers; these identifiers can be found in the supplementary data. These original coin examples have light to moderate



**Figure 5** Example GAN reconstructions showing original (left) and reconstructed (right) Roman coins (obv.).

damage, where damage can be evaluated based on overall evident surface deterioration and damage. This includes damage such as cracking, pieces missing, discoloration, and wear. Metadata providing provenance details and the original source data for our examples can be found using the coin identifier record. The CycleGAN demonstrates coin reconstructions based on the final trained model. The examples demonstrate the various types of enhancements made using our CycleGAN approach. In particular, visual improvements are made in reconstructions that incorporate infilling (2, 4, 5, 6, 7, 8), that is filling in broken areas or cracks, feature enhancements where evidence of improved appearance are evident on clothing, hair, or faces (2, 3, 4, 5, 7, 11), as well as providing clarity in the text (2, 6, 8, 10). This improvement includes adjusting worn surfaces to improve visualization. Another area where we see improvement from the original coins is where discoloration, or toning, is evident. In particular, coins 1–8 and 10–11 from Figure 5 produce generated coins from darker appearances (grayscale) to lighter shades comparable to other, better preserved coins used in training. In coins 4 and 6, this can be seen as helping in feature identification as more detail becomes evident from the lighter shading.

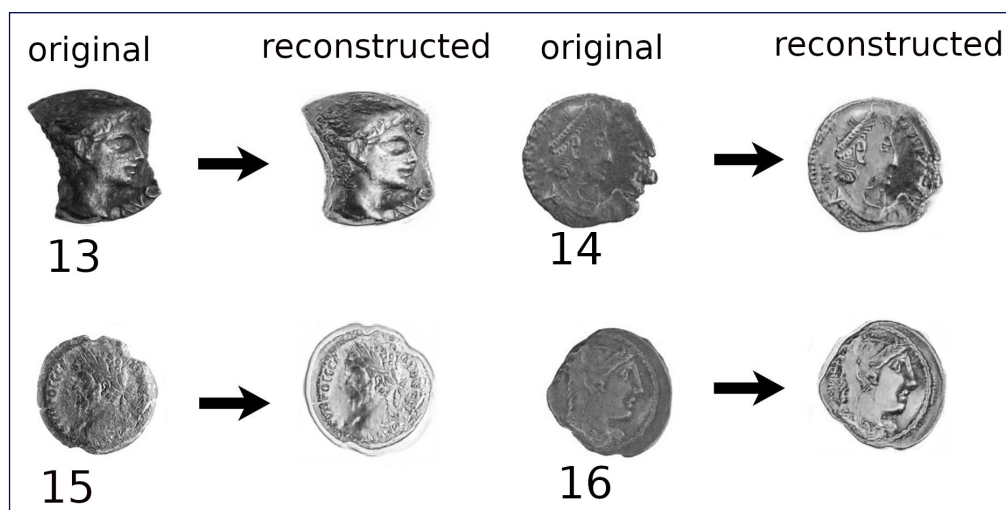
Generally, the results show coins are improved in appearance but there are areas where enhancements are not evident or are not very significant. In some cases, even if some parts of a coin are clearer other features are not improved. For instance, Figure 5:9 does not appear to greatly improve; Figure 5:5 & 6 enhances the coin image but the clarity still does not make it easier to identify features on the coin (e.g., the mouth or chin on the coins' faces). In this case, some distortion is noted and that can be due to insufficient training or poor reconstruction from what is evident on coins. Image quality is also likely to play a role in distorting output reconstruction. In cases where a coin did not need significant enhancement,

such as 12, the GAN did not distort the original image, indicating the model recognizes the coin is relatively undamaged and does not need significant enhancement. The examples demonstrate that CycleGAN works well in enhancing or clarifying many of the degraded coins, but there is a limit to this where more degraded coins are not greatly improved. For the most part, our assessment is qualitative rather than quantitatively measured (Zhang & Banovic 2021).

To better understand the limitations of the approach, a further analysis on more heavily degraded coins is assessed. Figure 6 provides some examples where large areas (i.e., more than 40%) are missing (13), facial features are very unclear (14&16), edges are broken (14, 15 & 16), and there is a lot of darkening or toning (16). In these cases, these qualities affect large areas or all of the coin surfaces. It is clear reconstructing a coin that is heavily damaged is not always successful (13), but where there is less damage or darkening the reconstruction is better and more clear (15). Furthermore, edges and features are improved (14) even with some moderate damage. Overall, what we see is if damage is severe (i.e., more than 40%), the reconstruction is incomplete, but with more minor damage then reconstruction is generally better with some feature enhancement and improved visualization. For feature enhancements, objects that are severely damaged (i.e., more than 40% surface damage) still show enhancements; however, deep surface damage that removes most of the surface details limits reconstruction.

## 5. DISCUSSION AND CONCLUSION

Our approach showcases the potential of GANs and generative AI for coin image enhancement. This application has utility in artifact visualization provided sufficient data are available from given repositories. We



**Figure 6** More heavily degraded real coins (left) and reconstructed (right) coins (obv.) using the CycleGAN. These examples highlight how the GAN addresses coins having 40% or more surface damage in cases.

intentionally confined our efforts to 2D coin images, as they constitute the most prevalent form of data and align with our primary objective: providing visual reconstructions to demonstrate visual enhancement for cultural objects.

Overall, our CycleGAN application demonstrates utility in creating good 2D reconstructions. We demonstrate multifaceted reconstruction that addresses infilling of broken parts and cracks, improvements on toning, and feature enhancements, where worn hair, face, clothing, writing, and other surface features are improved. Provided damage was less than 30–40% for coin surfaces, we see general improvement in coin appearance. Results show there is improvement in visualizing coins, including making degraded coins visually clear. In cases where no or very minimal damage (i.e., less than 1% surface damage) is evident, the GAN often did not overly reconstruct a given coin, indicating the model also recognises when there is minimal reconstruction needed or even no need to reconstruct an object. Currently, we do not think reconstructions are of the quality that enable easier identification of heavily eroded or poorly preserved coins to interpret what they might be, although enhancement and general visualization improvement is evident. This is among the main contributions of this work. The approach does show some weaknesses in cases where coins are heavily degraded, with reconstructions not showing much benefit in some cases. Coins missing more than 40% percent of their material, facial features heavily degraded by more than 40%, and if letters are similarly erased from coins then reconstructions proved difficult. At times, reconstructions are not accurate, where facial features appear distorted (e.g., Figure 5:6). This is mainly due to heavy surface damage that made a given feature difficult to reconstruct. Nevertheless, we see many cases where CycleGAN performs well in reconstructing coins and helping improve visualization.

Previously, CycleGAN was used to improve coins' general appearance, focusing mainly on worn surface features (Zachariou *et al.* 2020). We improve on earlier results by employing new training data that focus on a variety of different forms of damage, using CycleGAN's capabilities in global image transfer, and hyperparameter settings that enable the approach to do more than only limited feature enhancement. We demonstrate that our approach combined several types of enhancements in reconstructions as discussed above. This includes infilling of cracks, adding to broken parts, improvement of discoloration, and surface detail enhancements. This demonstrates CycleGAN's potential for artifact enhancement and visualization along multiple areas where coin degradation is evident. This we see as a new contribution from this work.

Weaknesses such as facial features not reconstructed completely or infilling more heavily damaged coins could be addressed. One way to improve results would be to use

better quality image data to enable more feature details to be captured in training and reconstructions. However, we noticed in attempts to improve reconstructions when addressing details on coins, including infilling and general appearance of features by introducing new training data or increasing training epochs, new problems arise, particularly overfitting becomes a concern. Improvements in some aspects, such as infilling, sacrifice accuracy and improvement in other areas, such as feature details in the coins' faces, as new coins are introduced. In particular, GANs are prone to overfitting, including in cases where reconstructions require multiple enhancement types (Tang 2020).

### 5.1. FURTHER RESEARCH AND ARCHITECTURE ENHANCEMENTS

Having conducted this work, we can now propose ways in which improvements can be made. One possibility is to use different reconstruction methods, including diffusion models (Lee & Yun 2024). These have potential in generating accurate reconstructions using underlying statistical sampling and have been shown to require less data and time in producing results. Another possibility is to develop a new GAN-based approach that emphasizes different steps or addresses limitations in given GAN architecture. For instance, infilling, various feature reconstruction, including improving writing visibility and facial or clothing features, might be better by enabling a step-by-step process in training that focuses on individual feature enhancements that then combines all features enhanced at a final stage. Rather than enhancing and training on an entire image, individual areas within an image can be used to train the GAN so that it learns well the types of evident damage, with the final stage combining the various enhancements so that a complete reconstruction is possible. CycleGAN is well suited for global image transfer, which enables learning from a given image used in training, but has weakness in object transfiguration. That is, attempts to segment only part of an image and use that to then enhance or improve on an output image without affecting other desired parts of a given image background proves difficult. Weaknesses in CycleGAN include the fact that it assumes intrinsic dimensions for a given output image equals the input image; variations in input size might, however, need to be accounted for in creating variations in output size (Zhou *et al.* 2017). For infilling and feature reconstruction, one key challenge is to develop an approach that can handle and better reconstruct objects with large areas missing, particularly where the input shape could significantly vary from the output shape (Hedjazi & Genc 2021). This would also be needed for obverse and reverse sides. Providing more enhanced optimization on training data using augmentation could improve reconstruction capabilities that are weak or have limited training data by increasing image variation for

reconstruction (Moreno-Barea *et al.* 2020). The challenge is to find similar coins to the desired reconstruction in training from a given input, which can also lead to overfitting if too many coins from a given type are used. For now, we see CycleGAN as a good approach that balances and accomplishes different improvements in coin reconstruction, even if it sacrifices some accuracy in specific features (e.g., infilling or enhanced facial details). Nevertheless, new GAN-based approaches might be a better solution in the future as various features may need simultaneous improvement on coins, where step-by-step enhancements on individual features in training could lead to an overall improved reconstruction.

## 5.2 ETHICAL CONSIDERATIONS

The topic of GANs also brings up an important point that archeologists and cultural experts should consider. There is potential that GANs can be used to create fake 2D/3D objects to deceive, such as generated objects being printed by 3D printers and being displayed as real (Xu *et al.* 2023). We believe guidelines need to be created for using GANs in cultural fields. This includes clear warnings that a given image/object is generated using a GAN and the reconstruction purpose should be defined. Although one can possibly identify a generated image, this is not always easy. One can create computational techniques deploying AI to find fake images; however, works should declare their purpose and intent if they utilize generative AI methods. Therefore, we suggest that works using GANs 1) disclose their intent in deploying generative AI, 2) clearly state specific architecture (e.g., type of GAN) used, and 3) generated images should be clearly labelled (e.g., as reconstructed, fake, etc.). Using such guidelines can help ensure GANs are used to enhance our knowledge and understanding rather than for nefarious purposes.

## 6. SUPPLEMENTARY DATA

Data shared here include the following: model training and validation data, test data used for reconstructions given to judges, real and reconstructed coins used to test quality of reconstructions, all of the paper's coin outputs (real and generated images), the final CycleGAN-derived model, and the code notebook. They can be found here: <https://figshare.com/s/187bc45ca8d3a04caddf>. All training and evaluation coins are named using the identifier on <http://numismatics.org>. The names of the files can be used to correspond to the metadata and data found there for individual coins (e.g., <http://numismatics.org/collection/1976.82.25> for 1976.82.25). The identifiers used for Figures 5–6 are also included. The GAN code, along with documentation, used for this effort is based on Zhu *et al.* (2020). The scraper created by this work and used for the coin collection effort for training

can be accessed here: [https://github.com/maltaweel/Coin\\_Scrape\\_CGAN](https://github.com/maltaweel/Coin_Scrape_CGAN).

## ACKNOWLEDGEMENTS


We thank Abu Dhabi University for generously funding this research.

## COMPETING INTERESTS

The authors have no competing interests to declare.

## AUTHOR AFFILIATIONS

**Mark Altaweel**  [orcid.org/0000-0003-1807-3737](https://orcid.org/0000-0003-1807-3737)  
Institute of Archaeology, University College London, UK

**Adel Khelifi**  [orcid.org/0000-0002-7844-9452](https://orcid.org/0000-0002-7844-9452)  
Faculty of Engineering, Abu Dhabi University, AE

**Mohammad Hashir Zafar**  
Faculty of Engineering, Abu Dhabi University, AE

## REFERENCES

- Aggarwal, A, Mittal, M and Battineni, G.** 2021. 'Generative adversarial network: An overview of theory and applications.' *International Journal of Information Management Data Insights*, 1(1): 100004. DOI: <https://doi.org/10.1016/j.jjime.2020.100004>
- American Numismatic Society.** 2023a. 'American numismatic society.' <https://numismatics.org/> [Accessed: (December 06, 2023)].
- American Numismatic Society.** 2023b. 'Online coins of the roman empire.' <https://numismatics.org/ocre/>, [Accessed: (September 06, 2023)].
- ArcGIS.** 2023. 'How cyclegan works?' <https://developers.arcgis.com/python/guide/how-cyclegan-works/> [Accessed: (December 18, 2023)].
- Berlin, M.** 2023. 'Münzkabinett online catalogue.' <https://ikmk.smb.museum/home>, [Accessed: (September 06, 2023)].
- Borji, A.** 2019. 'Pros and cons of gan evaluation measures.' *Computer Vision and Image Understanding*, 179: 41–65. DOI: <https://doi.org/10.1016/j.cviu.2018.10.009>
- British Museum.** 2023. 'British museum collections.' <https://www.britishmuseum.org/collection> [Accessed: (September 06, 2023)].
- Bruno, F, Bruno, S, De Seni, G, Luchi, ML, Mancuso, S and Muzzupappa, M.** 2010. 'From 3d reconstruction to virtual reality: A complete methodology for digital archaeological exhibition.' *Journal of Cultural Heritage*, 11(1): 42–49. DOI: <https://doi.org/10.1016/j.culher.2009.02.006>
- Chang, X, Chao, F, Shang, C and Shen, Q.** 2022. 'Sundial-gan: A cascade generative adversarial networks framework for

- deciphering oracle bone inscriptions.' In: *Proceedings of the 30th ACM International Conference on Multimedia, Lisboa Portugal*, ACM, 10 October 2022, pp. 1195–1203. DOI: <https://doi.org/10.1145/3503161.3547925>
- Choi, S-Y, Jeong, H-J, Park, K-S and Ha, Y-G.** 2019. 'Efficient driving scene image creation using deep neural network.' In: *2019 IEEE International Conference on Big Data and Smart Computing (BigComp)*, Kyoto, Japan, IEEE, February 2019, pp. 1–4. DOI: <https://doi.org/10.1109/BIGCOMP.2019.8679269>
- Colmenero-Fernández, A and Feito, F.** 2021. 'Image processing for graphic normalisation of the ceramic profile in archaeological sketches making use of deep neuronal net (dnn).' *Digital Applications in Archaeology and Cultural Heritage*, 22: e00196. DOI: <https://doi.org/10.1016/j.daach.2021.e00196>
- Farahanipad, F, Rezaei, M, Nasr, MS Kamangar, F and Athitsos, V.** 2022. 'A survey on gan-based data augmentation for hand pose estimation problem.' *Technologies*, 10(2): 43. DOI: <https://doi.org/10.3390/technologies10020043>
- Garozzo, R, Santagati, C, Spampinato, C and Vecchio, G.** 2021. 'Knowledge-based generative adversarial networks for scene understanding in cultural heritage.' *Journal of Archaeological Science: Reports*, 35: 102736. DOI: <https://doi.org/10.1016/j.jasrep.2020.102736>
- Ghosh, B, Dutta, IK, Carlson, A, Totaro, M and Bayoumi, M.** 2020. 'An empirical analysis of generative adversarial network training times with varying batch sizes.' In: *2020 11th IEEE Annual Ubiquitous Computing, Electronics and Mobile Communication Conference (UEMCON)*, IEEE, pp. 643–648. DOI: <https://doi.org/10.1109/UEMCON51285.2020.9298092>
- Gonog, L and Zhou, Y.** 2019. 'A review: Generative adversarial networks.' In: *2019 14th IEEE Conference on Industrial Electronics and Applications (ICIEA)*, Xi'an, China, IEEE, June 2019, pp. 505–510. DOI: <https://doi.org/10.1109/ICIEA.2019.8833686>
- Goodfellow, I, Pouget-Abadie, J, Mirza, M, Xu, B, Warde-Farley, D, Ozair, S, Courville, A and Bengio, Y.** 2014. 'Generative adversarial nets.' In: *Advances in neural information processing systems*, Mit Press, pp. 2672–2680.
- Graganiello, D, Cozzolino, D, Marra, F, Poggi, G and Verdoliva, L.** 2021. 'Are gan generated images easy to detect? a critical analysis of the state-of-the-art.' In: *2021 IEEE International Conference on Multimedia and Expo (ICME)*, pp. 1–6. DOI: <https://doi.org/10.1109/ICME51207.2021.9428429>
- Harms, J, Lei, Y, Wang, T, Zhang, R, Zhou, J, Tang, X, Curran, W, Liu, T and Yang, X.** 2019. 'Paired cycle-gan-based image correction for quantitative cone-beam computed tomography.' *Medical Physics*, 46: 3998–4009. DOI: <https://doi.org/10.1002/mp.13656>
- Hedjazi, MA and Genc, Y.** 2021. 'Efficient texture-aware multi-gan for image inpainting.' *Knowledge-Based Systems*, 217: 106789. DOI: <https://doi.org/10.1016/j.knosys.2021.106789>
- Hermoza, R and Sipiran, I.** 2018a. '3d reconstruction of incomplete archaeological objects using a generative adversarial network.' In: *Proceedings of Computer Graphics International 2018*, CGI 2018, New York, NY, USA, Association for Computing Machinery, p. 5–11. DOI: <https://doi.org/10.1145/3208159.3208173>
- Hermoza, R and Sipiran, I.** 2018b. '3d reconstruction of incomplete archaeological objects using a generative adversarial network.' In: *Proceedings of Computer Graphics International 2018, Bintan Island Indonesia*, ACM, 11 June 2018, pp. 5–11. DOI: <https://doi.org/10.1145/3208159.3208173>
- Jiang, Z and Sweetser, P.** 2022. 'Gan-assisted yuv pixel art generation.' In: Long, G, Yu, X and Wang, S (eds.) *AI 2021: Advances in Artificial Intelligence. Lecture Notes in Computer Science*. Springer International Publishing, pp. 595–606. DOI: [https://doi.org/10.1007/978-3-030-97546-3\\_48](https://doi.org/10.1007/978-3-030-97546-3_48)
- Kleber, F and Sablatnig, R.** 2009. 'A survey of techniques for document and archaeology artefact reconstruction.' In: *10th International Conference on Document Analysis and Recognition*, Barcelona, Spain, 2009, IEEE, pp. 1061–1065. DOI: <https://doi.org/10.1109/ICDAR.2009.154>
- Kniaz, V, Remondino, F and Knyaz, VA.** 2019. 'Generative adversarial networks for single photo 3d reconstruction.' *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 42(2): 403–408. DOI: <https://doi.org/10.5194/isprs-archives-XLII-2-W9-403-2019>
- Koutsoudis, A, Vidmar, B and Arnautoglou, F.** 2013. 'Performance evaluation of a multi-image 3d reconstruction software on a low-feature artefact.' *Journal of Archaeological Science*, 40(12): 4450–4456. DOI: <https://doi.org/10.1016/j.jas.2013.07.007>
- Kunsthistorisches Museum Wien.** 2023. 'Kunsthistorisches museum wien.' <https://www.ikmk.at/home?lang=en> [Accessed: (September 06, 2023)].
- Kurach, K, Lučić, M, Zhai, X, Michalski, M and Gelly, S.** 2019. 'A large-scale study on regularization and normalization in GANs.' In: Chaudhuri, K and Salakhutdinov, R (eds.) *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, PMLR, pp. 3581–3590.
- Langr, J and Bok, V.** 2015. *GANs in Action: Deep Learning with Generative Adversarial Networks*. Shelter Island, New York: Manning Publications.
- LeCun, Y, Bengio, Y and Hinton, G.** 2015. 'Deep learning.' *Nature*, 521(7553): 436–444. DOI: <https://doi.org/10.1038/nature14539>
- Lee, K-H and Yun, GJ.** 2024. 'Microstructure reconstruction using diffusion-based generative models.' *Mechanics of Advanced Materials and Structures*, 31(18): 4443–4461. DOI: <https://doi.org/10.1080/15376494.2023.2198528>

- Moreno-Barea, FJ, Jerez, JM and Franco, L.** 2020. 'Improving classification accuracy using data augmentation on small data sets.' *Expert Systems with Applications*, 161: 113696. DOI: <https://doi.org/10.1016/j.eswa.2020.113696>
- Münster, S, Maiwald, F, Di Lenardo, I, Henriksson, J, Isaac, A, Graf, MM, Beck, C and Oomen, J.** 2024. 'Artificial intelligence for digital heritage innovation: Setting up a r&d agenda for europe.' *Heritage*, 7(2): 794–816. DOI: <https://doi.org/10.3390/heritage7020038>
- Navarro, P, Cintas, C, Lucena, M, Fuertes, JM, Segura, R, Delrieux, C and González-José, R.** 2022. 'Reconstruction of iberian ceramic potteries using generative adversarial networks.' *Scientific Reports*, 12(1): 10644.
- Nomisma.** 2024. 'Nomisma.org.' <http://nomisma.org/>, [Accessed: (February 27, 2024)]. DOI: <https://doi.org/10.1038/s41598-022-14910-7>
- Park, S, Kim, J, Park, J, Jung, S-H and Sim, C.** 2023. 'How to train your pre-trained gan models.' *Applied Intelligence*, 53(22): 27001–26. DOI: <https://doi.org/10.1007/s10489-023-04807-x>
- Parmar, G, Zhang, R and Zhu, J-Y.** 2022. 'On aliased resizing and surprising subtleties in gan evaluation.' In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, IEEE, pp. 11410–20. DOI: <https://doi.org/10.1109/CVPR52688.2022.01112>
- Portable Antiquities Scheme.** 2023. 'Portable antiquities scheme.' <https://finds.org.uk/database/search/results/> [Accessed: (September 06, 2023)].
- Tang, S.** 2020. 'Lessons learned from the training of gans on artificial datasets.' In: *IEEE Access*, IEEE, pp. 165044–55. DOI: <https://doi.org/10.1109/ACCESS.2020.3022820>
- Ulyanov, D, Vedaldi, A and Lempitsky, V.** 2017. 'Improved texture networks: Maximizing quality and diversity in feed-forward stylization and texture synthesis.' In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Honolulu, HI, IEEE, July 2017, pp. 4105–4113. DOI: <https://doi.org/10.1109/CVPR.2017.437>
- Univerity of Vienna.** 2023. 'University of vienna digital coin cabinet.' <https://ikmk-ing.univie.ac.at/home?lang=en> [Accessed: (September 06, 2023)].
- Wach, K, Duong, CD, Ejdy, J, Kazlauskaitė, P, Korzynski, R, Mazurek, G, Paliszkiwicz, J and Ziemba, E.** 2023. 'The dark side of generative artificial intelligence: A critical analysis of controversies and risks of chatgpt.' *Entrepreneurial Business and Economics Review*, 11(2): 7–30. DOI: <https://doi.org/10.15678/EBER.2023.110201>
- Wang, Y, Luo, Y, Zu, C, Zhan, B, Jiao, Z, Wu, X, Zhou, J, Shen, D and Zhou, L.** 2024. '3d multi-modality transformer-gan for high-quality pet reconstruction.' *Medical Image Analysis*, 91(102983): 7–30. DOI: <https://doi.org/10.1016/j.media.2023.102983>
- Xu, D, Fan, S and Kankanhalli, M.** 2023. 'Combating misinformation in the era of generative ai models.' In: *Proceedings of the 31st ACM International Conference on Multimedia. Presented at the MM '23: The 31st ACM International Conference on Multimedia*, ACM, Ottawa, ON, Canada, ACM, pp. 9291–9298. DOI: <https://doi.org/10.1145/3581783.3612704>
- Ya-Liang, C, Zhe Yu, L, Kuan-Ying, L and Winston, H.** 2019. 'Free-form video inpainting with 3d gated convolution and temporal patchgan.' In: *IEEE/CVF International Conference on Computer Vision (ICCV)*, IEEE, pp. 9066–9075.
- Yan, N, Mei, Y, Xu, L, Yu, H, Sun, B, Wang, Z and Chen, Y.** 2023. 'Deep learning on image stitching with multi-viewpoint images: A survey.' *Neural Processing Letters*, 55(4): 3863–3898. DOI: <https://doi.org/10.1007/s11063-023-11226-z>
- Zachariou, M, Dimitriou, N and Arandjelović, O.** 2020. 'Visual reconstruction of ancient coins using cycle-consistent generative adversarial networks.' *Sci*, 2(3): 52. DOI: <https://doi.org/10.3390/sci2030052>
- Zeng, X, Cheng, L, Li, S and Liu, X.** 2024. 'Archaeology drawing generation algorithm based on multi-branch feature cross fusion.' This is a preprint but not yet published. DOI: <https://doi.org/10.21203/rs.3.rs-4409621/v1>
- Zhang, E and Banovic, N.** 2021. 'Method for exploring generative adversarial networks (gans) via automatically generated image galleries.' In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. DOI: <https://doi.org/10.1145/3411764.3445714>
- Zhang, Y, Seibert, P, Otto, A, Raßloff, A, Ambati, M and Kästner, M.** 2024. 'Da-vegan: Differentiably augmenting vae-gan for microstructure reconstruction from extremely small data sets.' *Computational Materials Science*, 232(112661). DOI: <https://doi.org/10.1016/j.commatsci.2023.112661>
- Zhou, S, Xiao, T, Yang, Y, Feng, D, He, Q and He, W.** 2017. 'Genegan: Learning object transfiguration and attribute subspace from unpaired data.' *CoRR abs/1705.04932*. DOI: <https://doi.org/10.5244/C.31.111>
- Zhu, J-Y, Park, T, Isola, P and Efros, AA.** 2020. 'Unpaired image-to-image translation using cycle-consistent adversarial networks.' In: *2017 IEEE International Conference on Computer Vision (ICCV)*, Venice, Italy, pp. 2242–2251. DOI: <https://doi.org/10.1109/ICCV.2017.244>

---

**TO CITE THIS ARTICLE:**

Altaweel, M, Khelifi, A and Zafar, MH. 2024. Using Generative AI for Reconstructing Cultural Artifacts: Examples Using Roman Coins. *Journal of Computer Applications in Archaeology*, 7(1): 301–315. DOI: <https://doi.org/10.5334/jcaa.146>

**Submitted:** 02 January 2024    **Accepted:** 31 July 2024    **Published:** 10 September 2024

**COPYRIGHT:**

© 2024 The Author(s). This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC-BY 4.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited. See <http://creativecommons.org/licenses/by/4.0/>.

*Journal of Computer Applications in Archaeology* is a peer-reviewed open access journal published by Ubiquity Press.

