Ceramic Fabric Classification of Petrographic Thin Sections with Deep Learning

MIKE LYONS

RESEARCH ARTICLE

]u[ubiquity press

ABSTRACT

Classification of ceramic fabrics has long held a major role in archaeological pursuits. It helps answer research questions related to ceramic technology, provenance, and exchange and provides an overall deeper understanding of the ceramic material at hand. One of the most effective means of classification is through petrographic thin section analysis. However, ceramic petrography is a difficult and often tedious task that requires direct observation and sorting by domain experts. In this paper, a deep learning model is built to automatically recognize and classify ceramic fabrics, which expedites the process of classification and lessens the requirements on experts. The samples consist of images of petrographic thin sections under cross-polarized light originating from the Cocal-period (AD 1000–1525) archaeological site of Guadalupe on the northeast coast of Honduras. Two convolutional neural networks (CNNs), VGG19 and ResNet50, are compared against each other using two approaches to partitioning training, validation, and testing data. The technique employs a standard transfer learning process whereby the bottom layers of the CNNs are pre-trained on the ImageNet dataset and frozen, while a single pooling layer and three dense layers are added to ‘tune’ the model to the thin section dataset. After selecting fabric groups with at least three example sherds each, the technique can classify thin section images into one of five fabric groups with over 93% accuracy in each of four tests. The current results indicate that deep learning with CNNs is a highly accessible and effective method for classifying ceramic fabrics based on images of petrographic thin sections and that it can likely be applied on a larger scale.

CORRESPONDING AUTHOR:

Mike Lyons

German Archaeological
Institute, University of Bonn, DE
s5milyon@uni-bonn.de

KEYWORDS:

AI; Ceramic Fabrics;
Convolutional Neural Network;
Deep Learning; Petrography;
Pottery

TO CITE THIS ARTICLE:

Lyons, M. 2021. Ceramic Fabric Classification of Petrographic Thin Sections with Deep Learning. *Journal of Computer Applications in Archaeology*, 4(1), 188–201. DOI: <https://doi.org/10.5334/jcaa.75>

1 INTRODUCTION

General interest in deep learning has been growing considerably over the last several years and is currently being applied in virtually every industry and professional domain. This certainly includes archaeology, which has been experiencing an influx of new methodologies with the incorporation of deep learning into various projects and sub-disciplines. In this paper, this ‘high-tech’ application is applied to what some may call one of the more ‘low-tech’ methods of understanding our cultural heritage: ceramic petrography of pottery in thin section. Petrography is a well-established and effective means of characterizing and interpreting pottery assemblages. It is, however, a highly specialized field that requires significant time expenditure and expertise to engage. This paper aims to develop a model that can automate one of the primary aspects of petrographic work: the classification of ceramic sherds into specific fabric groups. This is accomplished via training a convolutional neural network (CNN) with images of thin sections labeled with their respective fabric and testing it on images of new sherds. Based on the results of testing this method with the material from an excavation at the site of Guadalupe in Honduras, it appears not only to be possible but highly effective.

It should be noted that this technique does not replace the need for experts in petrography. Instead, it allows their expertise to be embedded in a deep learning model and applied to a more extensive set of data than would otherwise be feasible. This technique’s strength lies in using an *established* classification of ceramic fabrics to build a model that can automatically recognize and presort new ceramic samples. This depends on the data being labeled with its proper classification before using it to build and test a model, a type of machine learning called supervised learning. However, with further development, it is entirely conceivable that an unsupervised learning approach in which data is not labeled before training and testing could be used to build or significantly assist with an original classification of ceramic fabrics.

1.1 THE ARCHAEOLOGICAL SITE OF GUADALUPE

The site of Guadalupe is located in the department of Colón on the northern coast of Honduras ([Figure 1](#)) in the Northeast Honduras Archaeological Zone. The entire region has received relatively little archaeological attention, and thus little is known about it as a whole. The Guadalupe Archaeological Project conducted recent excavations of a mound in the local schoolyard of the village of Guadalupe under the direction of Dr. Markus

Figure 1 Overview map of the northern coast of Honduras showing Guadalupe’s location. Elevation data: ASTGTM Version 3 (NASA et al. 2019).

Reindel of the German Archaeological Institute (DAI) and Dr. Peter Fux of the Museum Rietberg. The project is funded by the DAI and the Swiss-Liechtenstein Foundation for Archaeological Research Abroad and is in cooperation with the University of Zurich, Museum Rietberg Zurich, the Honduran Institute of Anthropology and History, and the National Autonomous University of Honduras.

The Guadalupe excavation unearthed a thick concentration of ceramic material, and other artifact remains. Evidence of occupation, such as an accumulation of *bajareque* (a dried-mud building material) and various postholes and pits, has been identified below this concentration. Towards the mound's eastern periphery, several burials were uncovered in association with prestige objects that indicate ritual activity. Carbon-14 dating and stylistic analysis of the pottery assemblage indicate that the site dates to the Cocal period (AD 1000–1525) and possibly the Transitional Selin period (AD 800–1000) (Reindel and Fecher 2017; Reindel, Fux, and Fecher 2018; Reindel, Fux, and Fecher 2019).

The ceramic material of Guadalupe consists of a pottery assemblage largely attributed to the Cocal period (Figure 2). This style comprises several vessel types ranging from deep storage jars to shallow bowls that often exhibit tripod appliqué supports.¹ As for decoration, they are adorned with a wide range of incisions and appliques that broadly conform to several specific motif types (Healy 1993; Fecher 2021). Several painted examples

were found as well. Production techniques likely involved coiling or slab-building, while the vessels almost certainly underwent open firing either in a pit or above ground. Vessel surfaces are often unevenly oxidized and reduced, with only a few specimens exhibiting characteristics of completely reduced firing conditions.

Of particular interest for this paper are the various clay materials and aplastic inclusions used in the production of Guadalupe's pottery. From 2018 to 2021, the author participated in the excavation of Guadalupe by analyzing the composition of the ceramic material. Given the site's relatively small size, a diverse variety of clay material would not necessarily be expected. However, the preliminary analysis shows that there are upwards of 15 clearly distinguishable ceramic fabrics at the site. It is currently unclear if any of these fabrics represent imported materials, but it is certainly possible given the large variety and the fact that long-distance exchange of various objects such as jade and obsidian has already been established (Fecher 2021; Reindel, Fux, and Fecher 2019; Stroth et al. 2019).

1.2 CURRENT METHODOLOGIES OF CERAMIC FABRIC CLASSIFICATION

Ceramic petrography employs thin sections to better understand the composition of ceramic artifacts through characterization and interpretation. This can help identify provenance, exchange relationships between sites or regions, production technologies and preferences, and

Figure 2 Selection of pottery examples from Guadalupe. **a)** Tripod vessel with geometric incisions, punctates, and zoomorphic appliqué lug (PAG-15-683). **b)** Constricted vessel with wave patterns and zoomorphic appliqué lug (PAG-40-3). **c)** Rim sherd with geometric paint pattern (PAG-43-3). **d)** Shallow tripod vessel with zoomorphic appliqué lug and anthropomorphic appliqué supports (PAG-53-29). **e)** Turtle-shaped ocarina (PAG-43-1). **f)** Roller stamp with geometric pattern (PAG-120-1). Photo credits: a) and b) F. Fecher, c) K. Engel, d) T. Remsey, e) P. Bayer, and f) M. Lyons.

other relevant information regarding archaeological research questions (Quinn 2013). Thin sections themselves are 30 µm thick slices of ceramic material made for study under a polarized light microscope. The polarized light conditions, plane-polarized (PL) and cross-polarized (XP), allow the diagnostic features of mineral inclusions to be identified and various other characteristics, such as matrix birefringence, to be illuminated. In this process, three major characteristics are noted: inclusions, clay matrix, and voids. Together, these characteristics are the core means of defining ceramic fabrics and grouping them, a task generally carried out by manual investigation of thin section samples.

Various approaches have been pursued to automate or speed up the classification or grouping of ceramic fabrics (see Cau et al. 2004; Middleton, Freestone, and Leese 1985; Middleton, Leese, and Cowell 1991; Reedy 2006; Schubert 1986; Stoltman 1989). Two familiar approaches include textural analysis and modal analysis, which depend on statistical interpretation. Textural analysis, also known as grain-size analysis, attempts to quantify the size distribution of aplastic inclusions. This quantitative data can then be used as a basis for the interpretation of clay preparation, such as levigation or intentional addition of inclusions, which ultimately help to create an identifiable signature of a fabric. Modal analysis, also known as frequency analysis, characterizes ceramics based on their proportions of inclusions, voids, and matrix. Again, this quantitative technique aids in defining or grouping fabrics (Quinn 2013). These approaches are still viable techniques and can often lead to a deeper understanding of the ceramic material at hand. However, they are rather labor-intensive and depend on an adequate level of domain expertise.

Chemical analysis is another approach used since the 1970s in the pursuit of fabric classification. Generally, neutron activation analysis (NAA) or X-ray fluorescence (XRF) is employed along with several other methods, such as X-ray diffraction (XRD). Multivariate data analysis techniques, such as principal component analysis (PCA), are used to interpret this chemical data and cluster ceramic artifacts into chemically similar groupings (Baxter 2006; Papageorgiou 2020; Quinn 2013). These statistics-based groupings and characterizations have proven time and again to be effective methods. However, the information gleaned can differ from a strictly petrographic approach, and thus a 'mixed-mode' approach has been argued for, in which both petrographic data and chemical data are interpreted in concert (Baxter et al. 2008).

1.3 NEURAL NETWORKS: A BRIEF BACKGROUND

Although deep learning with artificial neural networks (ANNs) may initially appear to be a new technology, and it certainly is a relatively new technique being explored in archaeology, its history stretches back to the 1940s.² However, the first major example of a CNN being

employed for image recognition is more recent and was developed by LeCun et al. (1989) in which hand-written ZIP Code digits were able to be recognized with a high level of accuracy, a technology that still sees widespread use today. Specifically, CNNs are a class of ANN that specialize in analyzing structured arrays of data, such as images, and excel at identifying features and patterns within data. More recently, Geirhos et al. (2018) have shown CNNs to have a strong bias towards recognizing textures rather than shapes, which has direct implications to the recognition of ceramic fabrics, which are highly textural in nature. Generally, CNNs consist of several groups comprised of convolutional, pooling, and fully connected layers that recognize increasingly complex features as the layers are traversed. In convolutional layers, a kernel of a defined width and height 'convolves' through an array and returns an activation when a recognized pattern is matched. For an in-depth description that includes the mathematical definitions of CNNs and deep learning in general, see Goodfellow, Bengio, and Courville (2016). Currently, CNNs are one of the state-of-the-art methods in computer vision.

Aside from the general development of deep learning technology, techniques in the way CNNs can be employed have also developed. One of the critical concepts in this work is the use of transfer learning (Yosinski et al. 2014). CNNs learn to classify images by processing large sets of labeled images (for supervised learning) and essentially modifying the internal parameters to recognize images outside of this so-called 'training data' effectively (LeCun, Bengio, and Hinton 2015). Transfer learning uses the saved lower-level parameters of a network learned on a set of images (often ImageNet) and applies them as the foundation layers for a new recognition task. Several additional layers, usually fully-connected dense layers, are placed on top of this architecture to facilitate learning the new set of images. Because these lower-level layers tend to recognize low-level image features, such as edges and corners, this already-learned recognition can be transferred to novel tasks that learn higher-level features with the additional layers. By leveraging a network's understanding of basic features, recognition of more complex features of one's own data is accelerated. Importantly, this not only aids in feature recognition, but it has also been demonstrated to improve accuracy when dealing with small datasets (LeCun, Bengio, and Hinton 2015), such as that of this work.

1.4 PREVIOUS WORK: CNN APPLICATIONS IN ARCHAEOLOGY AND GEOLOGY

Examples of CNNs being applied in archaeology as a whole are increasing substantially. In the following, several examples are illustrated that lead from more general applications to those explicitly concerned with material in thin section. Bogacz and Mara (2020) use CNNs in 3D modeling classifications. They attempt to

classify cuneiform tablets into specific periods. They do this on a subset of the HeiCuBeDa³ collection of well-labeled 3D models of cuneiform tablets that have at least 30 examples from a specific period. They compare various CNN methods that include not only image-based recognition (ResNet50 in their case), but also CNNs designed to work with 3D point clouds, such as PointNet++ and SplineNet. They achieve an 84% accuracy with a unique approach that uses components of both PointNet++ and SplineNet. Caspari and Crespo (2019), on the other hand, have used CNNs in the pursuit of identifying archaeological features in satellite image data. Specifically, they attempt to identify Iron Age burial mounds. They compare their results with other methods, such as support vector machines, and find that the CNN performs the best. Similar examples of CNNs being used in combination with satellite imagery are abundant.

Examples of CNNs being applied in ceramic classification exist as well, but they generally tend toward classifications outside of petrography. Three notable contributions include ArchAide, the ARCADIA Project, and the recent work by Pawlowicz and Downum (2021). ArchAide (Gualandi et al. 2016; Wright and Gattiglia 2018) is a well-funded project with one of its primary goals being the use of image recognition technology, such as CNNs, to develop a mobile app that can recognize and classify various ceramic types. As of yet, their work has allowed for the automatic recognition of Terra Sigillata and Roman Amphorae based on shape and Majolica of Montelupo based on decoration. The ARCADIA Project (Chetouani et al. 2018; Chetouani et al. 2020) is pursuing a similar goal of using CNNs to classify ceramics, specifically the decorative patterns engraved on their surfaces. They first create 3D models of the sherds, extract the relevant features and use a combined transfer learning technique based on the VGG19 and ResNet50 models combined with compact bilinear pooling. They have thus far been highly successful with an accuracy of 95+%. Lastly, Pawlowicz and Downum (2021) successfully use CNNs to classify various types of Tusayan White Ware from Northeast Arizona based on decoration. They compare the CNN classification accuracy results to four experts in the field and note that they are comparable to and, in some cases, better than the experts.

Image recognition using deep learning with CNNs as an approach for the classification of ceramic fabrics in thin section has not yet been realized in the field of archaeology. Aprile, Castellano, and Eramo (2014) exemplify a related but ultimately different approach to using ANNs with thin section images. Their goal is the automatic identification of mineral inclusions, specifically quartz and calcareous aggregates and voids. They use a step-wise approach in which they first segment each mineral type using ImageJ⁴ and later use a modular ANN for feature identification. This is not a CNN but rather

a fully connected ANN designed to interpret a feature vector. They achieve an accuracy of 90%. Several other notable examples of image-based classification of ceramics that do not fall under the use of CNNs include Lopez, Lira, and Hein (2015), Hein et al. (2018), and Tyukin et al. (2018).

Geology, which can be said to have a more extensive working relationship with thin section petrography, has seen more applications of CNNs over the last several years. However, it is important to note that they are working with thin sections of rocks, not pottery. Some examples include Cheng and Guo (2017), Karimpouli and Tahmasebi (2019), Pires de Lima et al. (2019a), Pires de Lima et al. (2019b), and Su et al. (2020). These approaches are generally geared towards mineral identification and segmentation and identifying specific rock types, where CNNs are the primary means of classification. Pires de Lima et al. (2019a: 4) attempt to classify ‘microfossils, core images, petrographic photomicrographs, and rock and mineral hand sample images.’ They demonstrate a similar transfer learning approach to that used in this paper by comparing the MobileNetV2 architecture to that of the InceptionV3. Classification accuracies of 97–100% are achieved in each case except for petrographic thin sections, in which they achieved an accuracy of 81%. The maximum number of classes used for thin sections is just three, exemplifying the difficulty of the task.

The use of CNNs in archaeology and archaeologically relevant fields is gaining momentum. The above examples are an incomplete list of the ever-growing number of projects and researchers exploring this topic. In most cases, they deal with automated classification problems on the macroscopic scale, such as decoration-based or shape-based classifications to determine a type of pottery. This is an extremely useful field that should continue to be explored. However, it exists in contrast to the more microscopic-scale classification of ceramic fabrics presented here, which provides different but valuable information about ceramic material. Furthermore, the studies that do involve the analysis of ceramic fabrics use other methods that do not include CNNs and generally focus on characterization or feature extraction, whereas few focus on classification. The work here fills this gap as the first use of CNNs for the automated classification of ceramic fabrics types.

2 METHODOLOGY

2.1 SAMPLE PREPARATION

The material used in this study consists of ceramic sherds collected directly from the excavation in Guadalupe, with find contexts ranging from the trench surface down to the occupation layers. Sixty-six of these sherds were chosen for thin sectioning based on their macroscopic appearance. The goal was to acquire samples representing each possible fabric and to gather three or more examples

of each. Thin sections of samples were then prepared by the author with the aid and facilities of the University of Bonn's Department of Geology. After preparation, each thin section was observed under a polarized light microscope and classified into a fabric group. Fabric groups were defined based on Quinn's (2013) modified approach of Whitbread's (1989) proposal for a systematic description of thin sections.

2.2 SAMPLE DOCUMENTATION AND SELECTION

Once the thin sections had been grouped into a defined fabric, each was systematically photographed at 100x magnification under both cross-polarized light (XP) and plane-polarized light (PPL) to capture the entire surface. Each resulting image is a 2048 × 1536 pixel JPEG with a scale of 2.8 mm in length and 2.1 mm in height. The microscope used is a Keyence VHX-7000 with a built-in camera. However, a standard polarized light microscope equipped with any modern digital camera should be entirely sufficient. Only images photographed under XP lighting conditions and those with at least one-third of the image occupied by ceramic material in thin section are used in this analysis. It is important to note that the photos were not initially taken with the goal of using them in deep learning applications but rather for documentation purposes. They are, however, remarkably effective. This should be encouraging for employing similar techniques with already existing collections of thin section photos.

Ceramic artifacts recovered from Guadalupe are relatively abundant—over 55,000 were recovered from

just one of the four 2 × 3 m excavated units. However, the various fabrics are not equally represented among them; several are only accounted for by one or two examples and thus could not be included in this study. For this reason, the methodology is catered to a somewhat limited set of data, an unfortunate but very real feature of many archaeological endeavors. From the 66 selected sherds, only fabric groups represented by three or more individual sherds are included in this analysis. This ensures that a second approach to testing can occur where at least one example of each fabric is excluded from the training and validation datasets, thus providing a guaranteed 'unfamiliar' example of each fabric with which to test the model. This is important due to the nature of the data used. The data consists of multiple photographs of each sherd, not the sherd itself. The second approach removes the possibility that the model is only tested on sherds that were also used to train it.

In total, the outlined selection criteria result in 24 applicable sherds represented by five well-defined fabric groups where n equals the number of sherds per individual fabric (*Figure 3*): Fabric a ($n = 5$), Fabric c ($n = 3$), Fabric d ($n = 8$), Fabric e ($n = 3$), and Fabric w ($n = 5$). Fabric d is the most abundant fabric found in the Guadalupe assemblage and comprises over half of the recovered sherds. The other four fabric types represent less than five percent each. Furthermore, the fabrics can largely be considered relatively distinct except for Fabrics c and d. As opposed to the other groups, c and d would likely require a more experienced petrographer if classifying them manually.

Figure 3 Examples of the five ceramic fabric types analyzed. Row 1) Fabric a: Inclusion-sparse birefringent, Row 2) Fabric c: Amphibole-rich, Row 3) Fabric d: Common, fine inclusions, Row 4) Fabric e: Microfossil-rich, and Row 5) Fabric w: Poorly sorted angular inclusions. Column a) Exterior surface; scale: in cm, Column b) fresh break; scale: in mm, Column c) close-up of thin section under cross-polarized light; scale: image width = 2.8 mm, and Column d) thin section under cross-polarized light; scale: see image.

2.3 DEEP LEARNING MODEL ARCHITECTURE

The model architectures used in this work are relatively standard mixes of blocks of convolutional layers, max pooling layers, and fully connected layers (Figure 4). As previously mentioned, the models take advantage of transfer learning. In this paper, two architectures are utilized: the VGG19 developed by Simonyan and Zisserman (2015) of the University of Oxford's Visual Geometry Group, and the ResNet50 developed by He et al. (2016) of Microsoft. Using two 'base models' provides a point of comparison between test results in order to verify that the use of CNNs in this fashion is generally applicable. The models are pre-trained with the ImageNet dataset, and the resulting weights are frozen. The top three and top one fully connected dense layers are removed from the standard VGG19 and ResNet50 architectures, respectively, and replaced with several layers for tuning to the thin section images. The additional layers include: 1) a 2D global average pooling layer, 2) a fully connected dense layer with an output of 1024, ReLU activation, 3) a fully connected dense layer with an output of 512, ReLU activation, and 4) a fully connected dense layer with an output of five, softmax activation. The models use stochastic gradient descent optimization in the form of root mean square propagation with a learning rate of 0.003 and a decay of 1×10^{-6} . The loss function is categorical cross-entropy.

The models are designed to be reproduced as easily as possible by interested parties. They are based on the open-source library TensorFlow using the highly abstracted API, Keras, and can be easily replicated in the programming languages R or Python. Thus, the entire workflow can be performed without the use of proprietary software. Furthermore, the VGG19 and ResNet50 are well-documented and can be implemented relatively easily into a variety of workflows using the mentioned tools. Specialized hardware is also not required to build and run such a model. Although a GPU can significantly accelerate training and testing speed, a standard CPU should suffice for all required processing. The models

presented here are trained with a consumer-grade CPU (i7-7700HQ) and GPU (RTX 2080 ti).

2.4 PREPROCESSING AND DATA AUGMENTATION

In total, 1069 images were used to train each model. In each case, the size of the image is reduced to 512 by 512 pixels. The training data images are then fed through an image data generator that creates multiple copies of the original images as they are introduced to the CNN model. Image augmentation whereby transformations including random rotation within 360 degrees, horizontal and vertical flipping, and random image brightness adjustments ranging from 0.5 to 1.5 times the original image brightness was tested but did not positively affect accuracy. Random zoom augmentation ranging from 0.5 to 1.5x was also tested but ultimately reduced the validation data's classification accuracy by around 5% and the training data by around 10%. In theory, augmentation allows a model to generalize better when classifying new images and avoids over-specialization towards the set of images used to train it (overfitting). The validation and test data in this study did not undergo any augmentation or data generation.

2.5 MODEL TRAINING: APPROACHES A AND B

In most CNN image classification applications, images are randomly assigned to 'train,' 'validation,' and 'test' sets at an 80:10:10 ratio, respectively. This is another method of ensuring that the model is able to generalize and is not overfitted to the training set. However, as noted above, the 1069 images used in this dataset are derived from 24 ceramics sherds belonging to just five ceramic fabrics. Due to the small number of accessible images here and the fact that several images represent the same sherd, two approaches are tested: the traditional approach (Approach A), in which images of each fabric are randomly split into training, validation, and testing sets at an 80:10:10 ratio, and a modified approach (Approach B) in which 1) a single sherd of each fabric is randomly excluded from the image set, 2) the remaining images are randomly split into training, validation, and testing sets at an 80:10:10 ratio,

Figure 4 Conceptual schematic of the model architecture in use. Transfer learning combines the 'frozen' base layers of the VGG19 model with the additional trainable layers used to 'tune' the model to the thin section dataset.

and 3) the previously excluded sherds are reintroduced into the test data. The second approach acts to further verify the model's effectiveness and ability to generalize, as all images of a single sherd from each fabric type are excluded from the training and validation datasets and only included in the test set. This ensures that a model can identify the fabrics of sherds that were not included in those used to train the model by simulating the introduction of 'foreign sherds.' This approach, however, leads to a more unconventional split between training, validation, and testing data, which is why Approach A is included as well.

3 RESULTS

3.1 APPROACH A

As mentioned above, Approach A randomly splits the images of each fabric into an 80:10:10 ratio of training, validation, and testing data. This results in 853 training, 107 validation, and 109 testing images. [Table 1](#) provides a detailed summary of how the images are split per fabric.

3.1.1 VGG19

Using Approach A with the VGG19 base model, training requires roughly 80 epochs until convergence, after which the model is essentially stable ([Figure 5a](#)). Training, validation, and testing images are classified with 100%,

100%, and 99.1% accuracy, respectively. [Figure 5b](#) shows the distribution of thin section images from the test dataset as the model classified them with respect to their actual classification.

3.1.2 ResNet50

Using Approach A with the ResNet50 base model, training requires roughly 90 epochs until convergence, after which the model is essentially stable ([Figure 6a](#)). Training, validation, and testing images are classified with 99.6%, 98.2%, and 100% accuracy, respectively.

[Figure 6b](#) shows the distribution of thin section images from the test dataset as the model classified them with respect to their actual classification.

FABRIC	TRAINING	VALIDATION	TEST	TOTAL
a	80	10	11	101
c	127	16	16	159
d	323	40	41	404
e	140	18	18	176
w	183	23	23	229
total	853	107	109	1,069
percent	79.79%	10.01%	10.20%	100%

Table 1 Approach A data partition of training, validation, and test image counts per fabric.

Figure 5 Training and testing results for VGG19 with Approach A. **a)** Training and validation accuracy (bottom) and loss (top) per epoch. **b)** Confusion matrix of test data showing the model's predicted fabric types (Prediction) vs. actual fabric types (Reference).

Figure 6 Training and testing results for ResNet50 with Approach A. **a)** Training and validation accuracy (bottom) and loss (top) per epoch. **b)** Confusion matrix of test data showing the model's predicted fabric types (Prediction) vs. actual fabric types (Reference).

3.2 APPROACH B

As mentioned above, Approach B randomly excludes a single sherd of each fabric from the image set, randomly splits the remaining images into training, validation, and testing sets at an 80:10:10 ratio, and then reintroduces the excluded sherds back into the test data. This results in 689 training, 85 validation, and 295 testing images. [Table 2](#) provides a detailed summary of how the images are split per fabric.

3.2.1 VGG19

Using Approach B with the VGG19 base model, training requires roughly 70 epochs until convergence, after which the model makes a minor improvement in accuracy at about 100 epochs ([Figure 7a](#)). Training, validation,

FABRIC	TRAINING	VALIDATION	TEST	TOTAL
a	58	7	36	101
c	102	13	44	159
d	284	35	85	404
e	92	11	73	176
w	153	19	57	229
total	689	85	295	1,069
percent	64.45%	7.95%	27.60%	100%

Table 2 Approach B data partition of training, validation, and test image counts per fabric.

and testing images are classified with 99.5%, 96.6%, and 96.3% accuracy, respectively. [Figure 7b](#) shows the distribution of thin section images from the test dataset as the model classified them with respect to their actual classification.

3.2.2 ResNet50

Using Approach B with the ResNet50 base model, training requires roughly 65 epochs until convergence, after which the model is essentially stable ([Figure 8a](#)). Training, validation, and testing images are classified with 99.5%, 97.7%, and 93.6% accuracy, respectively. [Figure 8b](#) shows the distribution of thin section images from the test dataset as the model classified them with respect to their actual classification.

3.3 SUMMARY OF RESULTS

In each test scenario, model accuracy in the classification of the test set of ceramic thin section images is above 93% (see [Table 3](#)). Approach A, in which no sherds were exclusively placed in the test set, achieved the highest accuracies at 99.1% for the VGG19 model and 100% for the ResNet50 model. Approach B, in which all images of a randomly selected single sherd from each fabric were excluded from training and validation but included in the test data, performed slightly worse with accuracies of

Figure 7 Training and testing results for VGG19 with Approach B. **a)** Training and validation accuracy (bottom) and loss (top) per epoch. **b)** Confusion matrix of test data showing the model's predicted fabric types (Prediction) vs. actual fabric types (Reference).

Figure 8 Training and testing results for ResNet50 with Approach B. **a)** Training and validation accuracy (bottom) and loss (top) per epoch. **b)** Confusion matrix of test data showing the model's predicted fabric types (Prediction) vs. actual fabric types (Reference).

APPROACH	BASE MODEL	TRAINING	VALIDATION	TEST
A	VGG19	100%	100%	99.1%
A	ResNet50	100%	99.1%	100%
B	VGG19	99.5%	96.6%	96.3%
B	Resnet50	99.5%	97.7%	93.6%

Table 3 Summary of results showing the accuracy of fabric predictions for training, validation, and testing images with respect to each combination of data partitioning approach and base model.

96.3% for the VGG19 model and 93.6% for the Resnet50 model. However, Approach B proves that the technique can effectively classify thin section images of sherds not previously exposed to the models. Between both approaches, the VGG19 base model performed slightly better than the ResNet50 base model.

Regardless of approach and model used, differentiating between fabrics c and d, the two most texturally similar fabrics, proved the most difficult, resulting in the highest misclassification rate of any fabric type. In Approach B, seven images of Fabric c were misclassified as Fabric d with the VGG19 model, while eleven were misclassified with the ResNet50 model. However, the majority of images from a single sherd, regardless of fabric always classify correctly (see Appendix A).

4 DISCUSSION

4.1 INTERPRETATION OF RESULTS AND ISSUES

The results indicate that the classification of ceramic fabrics in thin section with CNNs is generally effective. In each of the four tests, a high level of accuracy in classifying fabrics was achieved. Even when unfamiliar sherds were tested (Approach B), accuracy remained above 93%. These results have implications for continued archaeological investigation on the north coast of Honduras. As more petrographic samples of ceramics are made, they can first be evaluated using this model, verified, and then implemented into the model itself to improve its base of data. The model will continue to be expanded upon to integrate images of each identified fabric.

There are, of course, several issues. Firstly, the sample is small and narrowly defined. It comprises just over 1000 images of five fabrics comprising just 24 sherds originating from a single site. In virtually all deep learning tasks, a large sample size is paramount. Increasing the sample in each of the mentioned terms would greatly improve the significance of findings and increase the applicability to the broader region. The model also struggled slightly to differentiate between Fabrics c and d, two visually similar fabrics. Improving upon these results would likely depend on including additional samples of both fabrics—again, a problem relating to sample size.

Arguably, the most important aspect of this study considering the promising results and mentioned issues

is the proof of concept. It demonstrates the potential of applying deep learning to archaeological questions to which the composition of ceramic material can contribute. The encouraging results imply that this technique is effective and should be further explored with the ceramic assemblages of Guadalupe and northeast Honduras, but also in regions with more extensively studied and defined ceramic assemblages. Most importantly, it must be performed on a larger, more comprehensive scale if it is to be considered generally useful.

4.2 OUTLOOK

There are many directions that future work can take. Larger datasets that have been well-studied and well-labeled (i.e., classified) would be a good starting point. These would allow for the construction of models applicable to many different pottery assemblages. Combining these datasets could then potentially allow for the construction of more generalized models able to classify fabrics across multiple pottery assemblages, culture groups, or regions.

An important application of this work is that of provenance. Pottery has long been understood as a means of understanding intra- and intercultural exchange between past peoples. In some cases, stylistic analysis and chemical analysis do not suffice to differentiate between local and imported pottery. This can be due to a lack of diagnostic features or taphonomic processes like chemical transformation due to local deposition conditions. In these cases, it is reasonable to assume that petrographic image-based deep learning techniques can assist with identification due to their lack of dependence on chemical composition data and characteristic styles.

A typical but valid complaint of CNN applications is that they appear to be a ‘black box’ technique. Understanding why a model chooses a specific class is difficult. For example, the classification of fabrics in ‘traditional’ ceramic petrography often depends on the identification of single elements, such as a specific mineral inclusion. At first glance, the method presented here may appear to classify based only on the general appearance of each thin section image and not on specific elements. It is more likely that both aspects play a role in the model’s classification, but it is difficult to say for certain. However, there are means of understanding the basic cues as to why images are classified as they are. Regarding images of thin sections, this would not only help us to understand why they are classified in a certain way but could also be used to identify the diagnostic features of each fabric. It is possible to specify and highlight the regions of each image that a model identifies as in favor of or against a specific class (Zintgraf, Cohen, and Welling 2016). Pawlowicz and Downum (2021) provide an excellent example of this application using gradient-based class activation maps on the decoration-based

classification of ceramics types from Northeast Arizona. With petrographic samples, this could be used as a new means of understanding the diagnostic features that differentiate one fabric from another that are not readily apparent to petrographers.

Lastly, the automated clustering of ceramic fabrics is one of the more exciting avenues of future research. Using unsupervised learning techniques, such as k-mean clustering or autoencoding, in combination with CNNs, could potentially allow fabrics to be classified without the explicit need for expert intervention other than verification of results. Caron et al. (2018) have recently demonstrated a highly effective approach to unsupervised learning with their 'DeepCluster' project that uses k-mean clustering with the ImageNet and YFCC100M (see Kalkowski et al. 2015) datasets. Hsu and Lin (2017), among others, have also explored this topic. These clustering techniques would work to group thin-section images based on texture similarity and could potentially classify fabrics of assemblages independently, meaning an already-established classification would not be needed. This classification could be further bolstered—or refuted—by comparison with geochemical analyses.

5 CONCLUSION

Ceramics are one of the most ubiquitous components of cultural material. Where they were once used or made, they almost certainly present themselves in the archaeological record today. The plethora of information gleaned from their form, context, and composition has made them one of the cornerstone materials for better understanding our collective past. This paper introduces a new method aimed at this pursuit by using machine learning and ceramic petrography to recognize and classify ceramic fabrics automatically. It demonstrates this on a ceramic assemblage originating in the archaeological site of Guadalupe on the north coast of Honduras. The presented method utilizes convolutional neural networks to classify ceramic fabrics based on images of petrographic thin sections. The results of 93+% classification accuracy in four tests indicate a highly effective means of classification, and although this use case is narrowly defined, it serves as a proof of concept to illustrate that this method should be widely applicable to any ceramic assemblage with definable ceramic fabrics. Continued investigation into this topic is sure to yield more exciting methods to better understand material culture and, more importantly, the people that produced it.

ACKNOWLEDGEMENTS

I would like to thank the entire team of the Guadalupe Archaeological Project, especially Franziska Fecher and

Markus Reindel, for their excavation work and welcoming me as a project member, Oscar Neil Cruz of the Honduran Institute of Anthropology and History for supporting our work in Honduras, the University of Bonn's Geology Department with particular consideration for Herald Euler and Nils Jung for allowing me to make thin sections and teaching me how in the process, and the Commission for Archaeology on Non-European Cultures of the German Archaeological Institute for providing me with an office and the infrastructure needed to conduct this research. A special thank you goes out to those who helped to discuss and critique this work: Bartosz Bogacz, Hubert Mara, Marcel Müller, Biagio Paparella, Markus Reindel, Matthew Reitzle, Jasmin Roth, and Ulrich Wölfel.

NOTES

- 1 For a catalog of the ceramic forms and styles found in Guadalupe, see Fecher et al. (2020) and Fecher (2021).
- 2 See Schmidhuber (2015) for a more comprehensive overview of the history of pioneering development in deep learning.
- 3 An open access collection of 1977 3D models of cuneiform tablets that include metadata, transcriptions, and transliterations for circa 700 examples (Mara 2019).
- 4 ImageJ is a freeware image processing application developed by the National Institute of Health (<https://imagej.nih.gov/ij/>).

ADDITIONAL FILE

The additional file for this article can be found as follows:

- **Appendix A.** Results Data. DOI: <https://doi.org/10.5334/jcaa.75.s1>

COMPETING INTERESTS

The author has no competing interests to declare.

AUTHOR AFFILIATIONS

Mike Lyons orcid.org/0000-0001-8299-2127
German Archaeological Institute, University of Bonn, DE

REFERENCES

- Aprile, A, Castellano, G and Eramo, G.** 2014. Combining image analysis and modular neural networks for classification of mineral inclusions and pores in archaeological potsherds. *Journal of Archaeological Science*, 50: 262–272. DOI: <https://doi.org/10.1016/j.jas.2014.07.017>
- Baxter, MJ.** 2006. A Review of Supervised and Unsupervised Pattern Recognition in Archaeometry. *Archaeometry*, 48(4): 671–694. DOI: <https://doi.org/10.1111/j.1475-4754.2006.00280.x>

- Baxter, MJ, Beardah, CC, Papageorgiou, I, Cau, MA, Day, PM and Kilikoglou, V.** 2008. On Statistical Approaches to the Study of Ceramic Artefacts using Geochemical and Petrographic Data. *Archaeometry*, 50(1): 142–157. DOI: <https://doi.org/10.1111/j.1475-4754.2007.00359.x>
- Bogacz, B and Mara, H.** 2020. Period Classification of 3D Cuneiform Tablets with Geometric Neural Networks. In: *17th International Conference on Frontiers in Handwriting Recognition (ICFHR)*, 2020. IEEE, pp. 246–251. DOI: <https://doi.org/10.1109/ICFHR2020.2020.00053>
- Caron, M, Bojanowski, P, Joulin, A and Douze, M.** 2018. Deep Clustering for Unsupervised Learning of Visual Features. In: Ferrari, V, Hebert, M, Sminchisescu, C and Weiss, Y (eds.). *Computer Vision – ECCV 2018*. Cham: Springer, pp. 139–156. DOI: https://doi.org/10.1007/978-3-030-01264-9_9
- Caspari, G and Crespo, P.** 2019. Convolutional Neural Networks for Archaeological Site Detection – Finding “Princely” Tombs. *Journal of Archaeological Science*, 110: 104998. DOI: <https://doi.org/10.1016/j.jas.2019.104998>
- Cau, M-A, Day, PM, Baxter, MJ, Papageorgiou, I, Iliopoulos, I and Montana, G.** 2004. Exploring automatic grouping procedures in ceramic petrology. *Journal of Archaeological Science*, 31(9): 1325–1338. DOI: <https://doi.org/10.1016/j.jas.2004.03.006>
- Cheng, G and Guo, W.** 2017. Rock images classification by using deep convolution neural network. *Journal of Physics: Conference Series*, 887: 012089. DOI: <https://doi.org/10.1088/1742-6596/887/1/012089>
- Chetouani, A, Debroutelle, T, Treuillet, S, Exbrayat, M and Jesset, S.** 2018. Classification of Ceramic Shards Based on Convolutional Neural Network. In: *2018 25th IEEE International Conference on Image Processing (ICIP)*. IEEE, pp. 1038–1042. DOI: <https://doi.org/10.1109/ICIP.2018.8451728>
- Chetouani, A, Treuillet, S, Exbrayat, M and Jesset, S.** 2020. Classification of Engraved Pottery Sherds Mixing Deep-Learning Features by Compact Bilinear Pooling. *Pattern Recognition Letters*, 131: 1–7. DOI: <https://doi.org/10.1016/j.patrec.2019.12.009>
- Fecher, F.** 2021. Links and Nodes: Networks in Northeast Honduras during the Late Pre-Hispanic Period (AD 900–1525). PhD dissertation. University of Zurich. DOI: <https://doi.org/10.5167/uzh-203124>
- Fecher, F, Reindel, M, Fux, P, Gubler, B, Mara, H, Bayer, P and Lyons, M.** 2020. The Ceramic Finds from Guadalupe, Honduras: Optimizing Archaeological Documentation with a Combination of Digital and Analog Techniques. *Journal of Global Archaeology*, 2020: 1–53. DOI: <https://doi.org/10.34780/joga.v2020i0.1009>
- Geirhos, R, Rubisch, P, Michaelis, C, Bethge, M, Wichmann, FA and Brendel, W.** 2018. ImageNet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness. *CoRR, abs/1811.12231*. arXiv: 1811.12231.
- Goodfellow, I, Bengio, Y and Courville, A.** 2016. *Deep Learning*. Cambridge, MA: MIT Press.
- Gualandi, ML, Scopigno, R, Wolf, L, Richards, J, Garrigos, JBI, Heinzelmann, M, Hervas, MA, Vila, L and Zallocco, M.** 2016. ArchAIDE – Archaeological Automatic Interpretation and Documentation of cERamics. In: Catalano, CE and Luca, LD (eds.). *Eurographics Workshop on Graphics and Cultural Heritage*. The Eurographics Association. DOI: <https://doi.org/10.2312/gch.20161408>
- He, K, Zhang, X, Ren, S and Sun, J.** 2016. Deep Residual Learning for Image Recognition. In: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778. DOI: <https://doi.org/10.1109/CVPR.2016.90>
- Healy, PF.** 1993. Northeastern Honduras. In: Henderson, JS and Beaudry-Corbett, M (eds.). *Pottery of Prehistoric Honduras*. Los Angeles, CA: Institute of Archaeology, University of California, pp. 194–213. DOI: <https://doi.org/10.2307/j.ctvhhhg50.16>
- Hein, I, Rojas-Domínguez, A, Ornelas, M, D’Ercole, G and Peloscchek, L.** 2018. Automated classification of archaeological ceramic materials by means of texture measures. *Journal of Archaeological Science: Reports*, 21: 921–928. DOI: <https://doi.org/10.1016/j.jasrep.2017.12.032>
- Hsu, C-C and Lin, C-W.** 2017. Unsupervised Convolutional Neural Networks for Large-scale Image Clustering. In: *2017 IEEE International Conference on Image Processing (ICIP)*. IEEE, pp. 390–394. DOI: <https://doi.org/10.1109/ICIP.2017.8296309>
- Kalkowski, S, Schulze, C, Dengel, A and Borth, D.** 2015. Real-time Analysis and Visualization of the YFCC100m Dataset. In: Friedland, G, Ngo, C-W and Shamma, DA (eds.). *Proceedings of the 2015 Workshop on Community-Organized Multimodal Mining: Opportunities for Novel Solutions – MMCommons’15*. New York, NY: ACM Press, pp. 25–30. DOI: <https://doi.org/10.1145/2814815.2814820>
- Karimpouli, S and Tahmasebi, P.** 2019. Segmentation of digital rock images using deep convolutional autoencoder networks. *Computers & Geosciences*, 126: 142–150. DOI: <https://doi.org/10.1016/j.cageo.2019.02.003>
- LeCun, Y, Boser, B, Denker, JS, Henderson, D, Howard, RE, Hubbard, W and Jackel, LD.** 1989. Backpropagation Applied to Handwritten Zip Code Recognition. *Neural Computation*, 1(4): 541–551. DOI: <https://doi.org/10.1162/neco.1989.1.4.541>
- LeCun, Y, Bengio, Y and Hinton, G.** 2015. Deep learning. *Nature*, 521(7553): 436–444. DOI: <https://doi.org/10.1038/nature14539>
- Lopez, P, Lira, J and Hein, I.** 2015. Discrimination of Ceramic Types Using Digital Image Processing by Means of Morphological Filters. *Archaeometry*, 57(1): 146–162. DOI: <https://doi.org/10.1111/arcm.12083>
- Mara, H.** 2019. HeiCuBeDa Hilprecht – Heidelberg Cuneiform Benchmark Dataset for the Hilprecht Collection. Version V2. heiDATA. DOI: <https://doi.org/10.11588/data/IE8CCN>
- Middleton, AP, Freestone, IC and Leese, MN.** 1985. Textural Analysis of Ceramic Thin Sections: Evaluation of Grain Sampling Procedures. *Archaeometry*, 27(1): 64–74. DOI: <https://doi.org/10.1111/j.1475-4754.1985.tb00348.x>

- Middleton, AP, Leese, MN and Cowell, MR.** 1991. Computer-Assisted Approaches to the Grouping of Ceramic Fabrics. In: Middleton, A (ed.). *Recent developments in ceramic petrology*. London: British Museum, pp. 265–276.
- NASA, METI, AIST, Spacesystems, J., U.S. and Team, J. A. S.** 2019. ASTER Global Digital Elevation Model V003. NASA EOSDIS Land Processes DAAC. DOI: <https://doi.org/10.5067/ASTER/ASTGTM.003>
- Papageorgiou, I.** 2020. Ceramic investigation: how to perform statistical analyses. *Archaeological and Anthropological Sciences*, 12(9). DOI: <https://doi.org/10.1007/s12520-020-01142-x>
- Pawlowicz, LM and Downum, CE.** 2021. Applications of Deep Learning to Decorated Ceramic Typology and Classification: A Case Study using Tusayan White Ware from Northeast Arizona. *Journal of Archaeological Science*, 130: 105375. DOI: <https://doi.org/10.1016/j.jas.2021.105375>
- Pires de Lima, R, Bonar, A, Coronado, DD, Marfurt, K and Nicholson, C.** 2019a. Deep convolutional neural networks as a geological image classification tool. *The Sedimentary Record*, 17(2): 4–9. DOI: <https://doi.org/10.2110/sedred.2019.2.4>
- Pires de Lima, R, Suriamin, F, Marfurt, KJ and Pranter, MJ.** 2019b. Convolutional neural networks as aid in core lithofacies classification. *Interpretation*, 7(3): 27–40. DOI: <https://doi.org/10.1190/INT-2018-0245.1>
- Quinn, PS.** 2013. Ceramic petrography: The interpretation of archaeological pottery & related artefacts in thin section. Oxford: Archaeopress. DOI: <https://doi.org/10.2307/j.ctv1jk0jf4>
- Reedy, CL.** 2006. Review of Digital Image Analysis of Petrographic Thin Sections in Conservation Research. *Journal of the American Institute for Conservation*, 45(2): 127–146. DOI: <https://doi.org/10.1179/019713606806112531>
- Reindel, M and Fecher, F.** 2017. Archäologisches Projekt Guadalupe: Kulturelle Interaktion und vorspanische Siedlungsgeschichte im Nordosten von Honduras. *Zeitschrift für Archäologie Ausereuropäischer Kulturen*, 7: 349–356. DOI: <https://doi.org/10.5167/uzh-144745>
- Reindel, M, Fux, P and Fecher, F.** 2018. Archäologisches Projekt Guadalupe: Bericht über die Feldkampagne 2017. *Tech. rep.* Zürich: SLSA. Schweizerisch-Liechtensteinische Stiftung für archäologische Forschungen im Ausland. DOI: <https://doi.org/10.5167/uzh-158145>
- Reindel, M, Fux, P and Fecher, F.** 2019. Archäologisches Projekt Guadalupe: Bericht über die Feldkampagne 2018. *Tech. rep.* Zürich: SLSA. Schweizerisch-Liechtensteinische Stiftung für archäologische Forschungen im Ausland. DOI: <https://doi.org/10.5167/uzh-179198>
- Schmidhuber, J.** 2015. Deep Learning in Neural Networks: An Overview. *Neural Networks*, 61: 85–117. DOI: <https://doi.org/10.1016/j.neunet.2014.09.003>
- Schubert, P.** 1986. Petrographic Modal Analysis – a Necessary Compliment to Chemical Analysis of Ceramic Coarse Ware. *Archaeometry*, 28(2): 163–178. DOI: <https://doi.org/10.1111/j.1475-4754.1986.tb00384.x>
- Simonyan, K and Zisserman, A.** 2015. Very Deep Convolutional Networks for Large-Scale Image Recognition. In: Bengio, Y and LeCun, Y (eds.). *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7–9, 2015, Conference Track Proceedings*.
- Stoltman, JB.** 1989. A Quantitative Approach to the Petrographic Analysis of Ceramic Thin Sections. *American Antiquity*, 54(1): 147–160. DOI: <https://doi.org/10.2307/281336>
- Stroth, L, Otto, R, Daniels, JT and Braswell, GE.** 2019. Statistical artifacts: Critical approaches to the analysis of obsidian artifacts by portable X-ray fluorescence. *Journal of Archaeological Science: Reports*, 24: 738–747. DOI: <https://doi.org/10.1016/j.jasrep.2019.02.023>
- Su, C, Xu, S-J, Zhu, K-Y and Zhang, X-C.** 2020. Rock Classification in Petrographic Thin Section Images Based on Concatenated Convolutional Neural Networks. *Earth Science Informatics*, 13: 1477–1484. DOI: <https://doi.org/10.1007/s12145-020-00505-1>
- Tyukin, I, Sofeikov, K, Levesley, J, Gorban, AN, Allison, P and Cooper, NJ.** 2018. Exploring Automated Pottery Identification [Arch-I-Scan]. *Internet Archaeology*, 50. DOI: <https://doi.org/10.11141/ia.50.11>
- Whitbread, I.** 1989. A proposal for the systematic description of thin sections towards the study of ancient ceramic technology. In: Maniatis, Y (ed.). *Archaeometry, Proceedings of the 25th International Symposium*. Amsterdam: Elsevier, pp. 127–138.
- Wright, H and Gattiglia, G.** 2018. ArchAIDE: Archaeological Automatic Interpretation and Documentation of cERamics. In: *Workshop on Cultural Informatics co-located with the EUROMED International Conference on Digital Heritage 2018 (EUROMED 2018)*. Nicosia, Cyprus: Zenodo. DOI: <https://doi.org/10.5281/zenodo.3603090>
- Yosinski, J, Clune, J, Bengio, Y and Lipson, H.** 2014. How Transferable Are Features in Deep Neural Networks? In: *Proceedings of the 27th International Conference on Neural Information Processing Systems – Volume 2. NIPS'14*. Montreal, Canada: MIT Press, pp. 3320–3328.
- Zintgraf, LM, Cohen, TS and Welling, M.** 2016. A New Method to Visualize Deep Neural Networks. *CoRR*, abs/1603.02518. arXiv: 1603.02518.

TO CITE THIS ARTICLE:

Lyons, M. 2021. Ceramic Fabric Classification of Petrographic Thin Sections with Deep Learning. *Journal of Computer Applications in Archaeology*, 4(1), 188–201. DOI: <https://doi.org/10.5334/jcaa.75>

Submitted: 03 April 2021 Accepted: 25 July 2021 Published: 28 September 2021

COPYRIGHT:

© 2021 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.