Australian Archaeology

ISSN: (Print) (Online) Journal homepage: https://www.tandfonline.com/loi/raaa20

Reconstructing rock art chronology with transfer learning: A case study from Arnhem Land, Australia

Jarrad Kowlessar, James Keal, Daryl Wesley, Ian Moffat, Dudley Lawrence, Abraham Weson, Alfred Nayinggul & Mimal Land Management Aboriginal Corporation

To cite this article: Jarrad Kowlessar, James Keal, Daryl Wesley, Ian Moffat, Dudley Lawrence, Abraham Weson, Alfred Nayinggul & Mimal Land Management Aboriginal Corporation (2021): Reconstructing rock art chronology with transfer learning: A case study from Arnhem Land, Australia, Australian Archaeology, DOI: 10.1080/03122417.2021.1895481

To link to this article: https://doi.org/10.1080/03122417.2021.1895481

Published online: 30 Mar 2021.

Submit your article to this journal

Article views: 105

View related articles

View Crossmark data

Full Terms & Conditions of access and use can be found at
https://www.tandfonline.com/action/journalInformation?journalCode=raaa20

AUSTRALIAN ARCHAEOLOGY
https://doi.org/10.1080/03122417.2021.1895481

Routledge
Taylor & Francis Group

Check for updates

ARTICLE

Reconstructing rock art chronology with transfer learning: A case study from Arnhem Land, Australia

Jarrad Kowlessar[a], James Keal[b], Daryl Wesley[a], Ian Moffat[a], Dudley Lawrence[c], Abraham Weson[c], and Alfred Nayinggul[d], Mimal Land Management Aboriginal Corporation[c]

[a]Archaeology, College of Humanities, Arts and Social Sciences, Flinders University, Bedford Park, South Australia, Australia; [b]School of Physical Science, The University of Adelaide, Adelaide, South Australia, Australia; [c]Mimal Land Management Aboriginal Corporation, Nightcliff, Northern Territory, Australia; [d]Njanjma Aboriginal Corporation, Gunbalanya, Northern Territory, Australia

ABSTRACT

In recent years, machine learning approaches have been used to classify and extract style from media and have been used to reinforce known chronologies from classical art history. In this work we employ the first ever machine learning analysis of Australian rock art using a data efficient transfer learning approach to identify features suitable for distinguishing styles of rock art. These features are evaluated in a one-shot learning setting. Results demonstrate that known Arnhem Land Rock art styles can be resolved without knowledge of prior groupings. We then analyse the activation space of learned features and report on the relationships between styles and arrange these classes into a stylistic chronology based on distance within the activation space. By generating a stylistic chronology, it is shown that the model is sensitive to both temporal and spatial patterns in the distribution of rock art in the Arnhem Land Plateau region. More broadly, this approach is ideally suited to evaluating style within any material culture assemblage and overcomes the common constraint of small training data sets in archaeological machine learning studies.

ARTICLE HISTORY
Received 8 December 2020
Accepted 23 February 2021

KEYWORDS
Rock art; machine learning; Arnhem Land; style

Introduction

The Arnhem Land Plateau region of northern Australia has a rich and detailed history of rock art inscription (Chaloupka 1993) (Figure 1). The earliest dated pictograph in this region was inscribed around 28,000 years ago (David et al. 2013) and human occupation has been dated to 65,000 years ago (Clarkson et al. 2017). With the practice of rock art production continuing through to the current time, the longevity of this artistic culture speaks to the complexity and diversity of artistic styles observed across this landscape. The frequency of paintings along with their stylistic diversity in the context of the deep history of human activity has created a great challenge for those wishing to unravel the antiquity of the region's art (Chippindale and Taçon 1998).

This challenge is further complicated by the limited availability of rock art for which direct dating methods can be applied (Jones et al. 2017). In cases of repeated site usage, the chronology of the art can be derived from relative sequences in the superimposition, with older motifs appearing beneath newer. This method of investigation can also be supplemented with inferences drawn from environmental and technological details depicted in the art (Chaloupka 1993; Chippindale and Taçon 1998; Lewis 1988). For the majority of rock art, however, evidence sufficient to pinpoint the time of inscription does not exist.

Style is an ever-present attribute of how different items of material culture relate to one another. However, classifying style and quantifying stylistic relationships are time-consuming challenges in the field of archaeology, particularly rock art studies (Schaafsma 1985). The rock art assemblage of the Arnhem Land Plateau region is stylistically diverse and has long been the focus of archaeological debate. By organising the assemblage into distinct styles, archaeologists seek to understand the chronological evolution of art in the region. These styles are difficult to define and even more difficult to organise into a relative sequence. The complexity of rich and nuanced information encoded in style is central to the study of archaeology (Conkey and Hastorf 1993). Style is the variation in material culture over time and space which occurs when human activity is conducted in a particular way and in the context of alternative ways (Clegg 1987; Hegmon 1992). Following this definition, 'stylistic behaviour is founded on the basic human cognitive process of

CONTACT Daryl Wesley daryl.wesley@flinders.edu.au Archaeology, College of Humanities, Arts and Social Sciences, Flinders University, GPO Box 2100, Bedford Park, SA 5001, Australia
© 2021 Australian Archaeological Association Inc.

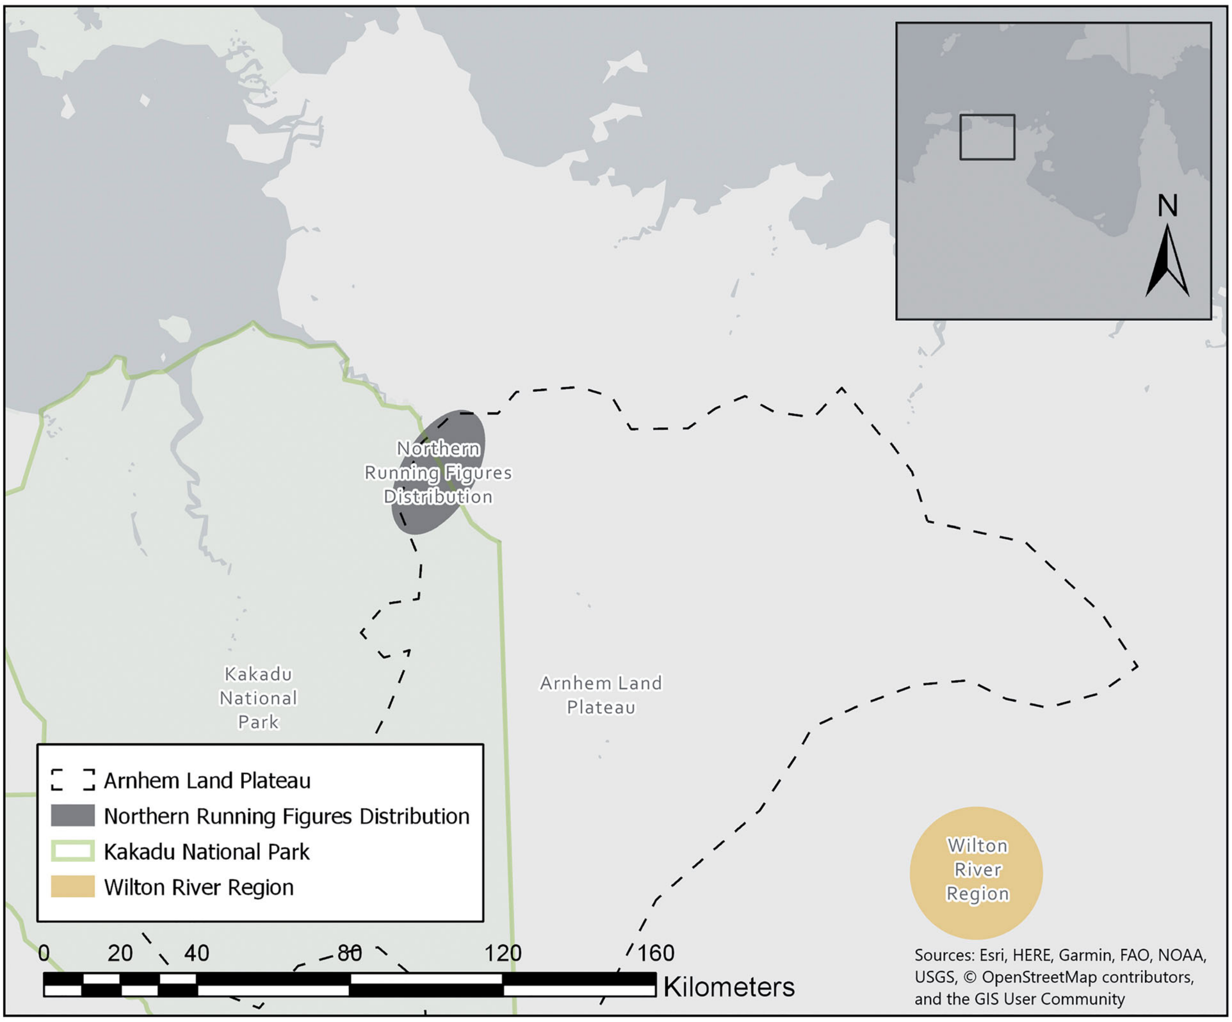

Figure 1. Map of the Arnhem land plateau rock art region highlighting the extents of the pan-Arnhem data sources as well as the geographically isolated distributions of the Northern Running Figures and the Wilton River motifs.

identification via comparison' (Wiessner 1984:190) and therefore any study of style is a study of the differences between styles. Style must also have both geographic and temporal controls, as does the social identity of the group of people for which a given style is characteristic (Sackett 1982).

In the context of art, when style is considered in this way, there are many disparate factors that could each describe style or could be combined to form different definitions of style. These factors include the subject matter being depicted (e.g. human figures, flora and fauna, geometric designs), different use of colours and design elements, placement and display of the subject (e.g. arrangement of the art in the depiction, the type of material on which the art is inscribed) and the method used to inscribe the art (brush work, line thickness, use of stencilling) (Domingo Sanz and Fiore 2014). Variations in any one or more of these elements could be used to describe stylistic classes (Domingo Sanz and Fiore 2014:7105). The application of stylistic analysis to rock art has been, and continues to be, critical in unravelling the overlapping history of graphic illustration in northern Australia (Layton 1992).

Observational approaches

Connoisseurship has proven to be one of the most effective means of detecting stylistic categories in rock art and forming those categories into larger movements or periods (Gunn 2018; Gunn et al. 2018). Styles are first defined by typological factors inferred by visual inspection of repeated motif patterns (Chaloupka 1993; Lewis 1988), after which they may be further organised into an inferred chronological order. A major concern for archaeological research in Arnhem Land has been the divisive chronological schema for styles of rock art. There have been competing rock art chronologies proposed for Arnhem Land each consisting of some variation on stylistic definitions (Brandl 1973; Chaloupka 1993; Chippindale and Taçon 1998; Haskovec 1992; Lewis 1988). These chronologies are strongly dependent on the identification of stylistic classes, technology and subject matter depicted in the rock art and subsequent superimposition. These competing relative chronological systems are therefore dependent on reliably identifying these characteristics within an assemblage of rock art.

Brandl (1973) produced one of the first detailed attempts at developing a chronology for Arnhem Land rock art with the introduction of eight stylistic periods. Brandl's (1973) styles became the basis for the chronology and styles later developed by Lewis (1988:107). The chronology developed by Lewis (1988:ix) consists of four major periods delineated by changes in material culture. These are known as the Boomerang, Hooked Stick, Broad Spearthrower, and Long Spearthrower Periods which have been named after the most distinctive material culture associated with each era (Lewis 1988:105). Lewis' (1988:13) methodology was based on depictions of human figures and their associated material culture items as he argued artefacts create identifiable, and therefore more reliable, temporal boundaries and criticises superimposition as a method for defining chronologies between styles (Lewis 1988:107).

In contrast Chaloupka (1993:89) anchored his proposed chronology to generalised timings coinciding with major environmental and climatic changes in Arnhem Land with periods divided into Pre-Estuarine, Estuarine, and Freshwater (and Contact) which contained at least 11 major styles or stylistic complexes. Lewis (1988:107) suggested that Chaloupka's Pre-Estuarine and Estuarine division does not exist and should be conflated with later periods. Chippindale and Taçon (1998) used Chaloupka's styles to build a chronology based on Old, Intermediate and New, identifying 10 major styles or stylistic complexes. Haskovec (1992:148) stated that material culture depicted in rock art changes less frequently than stylistic change. Therefore, he emphasises style should be the main approach towards identification and definition of a rock art assemblage (Haskovec 1992). To add to the chronological complexity of Arnhem Land rock art, Haskovec (1992:149) proposes six phases containing eight major styles or stylistic complexes. While there is a general acceptance for broad categories of rock art styles such as Dynamic Figures, it has been demonstrated that there is still significant variation within an overarching schema (Johnston 2017). Furthermore, the early Large Naturalistic Style has been noted as recursive in nature as large naturalistic depictions of subject matter occur throughout the Arnhem Land rock art sequence (Taçon and Chippindale 1994:214; Jones et al. 2020). The iterative and recursive nature of Indigenous rock art can be problematic for the identification of stylistic chronologies in Arnhem Land (Morphy 2012).

In this work we examine styles associated with the depiction of human figures. Chaloupka (1984, 1993) described several styles of human figure depictions found in the Arnhem Land region. They are, in interpreted chronological order: Large Naturalistic Figures, Dynamic Figures, Post-Dynamic Figures, Northern Running Figures, and Simple Figures with Boomerangs. These styles are characterised using a variety of parameters including the complexity of lines used, patterns that fill the voids of the image, and the posture or activities typically depicted. The chronology put forth by Chaloupka has found broad agreement despite some challenges regarding superimposition and technological material culture associations (Lewis 1988). The geographic scale of this chronology, as well as its sensitivity to more nuanced styles, are poorly defined within these broad categories (Johnston 2017). For example, Chaloupka suggested that four distinct stylistic phases occur within the Dynamic Figures class, each represented as a minor variation in the style of depiction. Other researchers have argued that minor changes within Dynamic Figures are not evidence of more nuanced stylistic phases, however can instead be explained by individual artistic expression (Johnston 2017).

This debate raises the important question of how a distinct style should be defined. How major must a typological change be, and at what frequency must it occur, for it to be considered a new style? This also raises geographic questions about how regionalisation might affect motif style, even when these motifs still fit within existing broad styles. This research aims to provide insight into these questions through the use of machine learning methods to quantify style and locate it on a relative stylistic continuum.

Statistical approaches

Observational approaches to style identification have been extended to incorporate various statistical analyses (Travers and Ross 2016, 2017). These analyses are based on the measurement of artistic attributes identified by the researcher as important to classifying style. These approaches have proven effective in identifying patterns within the data that relate to style, as well as smaller patterns which can occur within an identified style, such as the changing use of technology and the social role of artefacts within the societies producing the art (Travers and Ross 2016, 2017). A recent technique called Geometric Morphometrics (GM) has been used to identify animal species in rock art motifs (Cobden et al. 2017). The approach locates the coordinates of features that are key to identifying species from known images and compares the relative geometric features in known and unknown depictions. This approach attempts to reduce researcher bias and make a mathematical comparison of the geometry of individual figures. Like the traditional statistical methodology, this approach relies on the identification of geometrically and biologically meaningful features in the data (Cobden et al. 2017). The approach is also supervised by classes defined by the researcher. This requires that the same features being identified are purposefully depicted by the original painter in the art, as opposed to a stylistic means of encoding species information that is unknown to the researcher. In this way statistical approaches to rock art style analysis are effective at identifying and quantifying known attributes, however these approaches remain unreceptive to the detection of unknown stylistic attributes.

Machine learning approaches

Machine learning approaches have been applied in archaeological context for object detection and more recently for stylistic analysis (Cintas et al. 2020; Horr et al. 2014; Tsigkas et al. 2020; Wang et al. 2017). Recent research efforts into computer vision and machine learning have shown the ability of learning algorithms to discriminate between stylistic categories of painted art, with reasonable accuracy (Karayev et al. 2014; Saleh et al. 2016; Shamir et al. 2010). However, applying machine learning approaches to the classification of Australian rock art is particularly challenging, due in part to the limited amount of data available to individual researchers. As such, any machine learning approach suitable for analysing rock art must be capable of learning meaningful features from a small dataset, consisting of only a few examples, or without training using rock art at all.

In practice, machine learning with limited domain data can be achieved using transfer learning. Transfer learning involves taking a model developed for one task and reusing it for a second task (Caruana 1997). The tasks may differ in problem domain, datasets, or both. The success of the approach is dependent on similarities between the two tasks, and on how related the datasets are. An initial training phase is performed using background data, with any further training on the evaluation dataset referred to as ‘fine-tuning’. For image analysis, transfer learning relies on the reuse of learned features that are common in image datasets, such as edges, corners, and other geometric structures comprised of the spatial arrangement of pixel values.

In this work we experiment with three models trained to classify images. We explore the viability of different background data by training each of the models on three well-established image classification datasets. We then assess each trained model’s applicability for analysing a dataset of rock art from Arnhem Land, northern Australia. We do this without necessitating retraining, which is largely impossible due to the limited availability of data, by evaluating the trained models in a one-shot learning setting. In one-shot learning, a model must correctly predict an image’s class given only one example of each possible class with which to compare it (Bromley et al. 1993; Fei-Fei et al. 2006; Koch et al. 2015; Lake et al. 2011). In this way, we demonstrate each model’s impressive ability to classify examples of rock art into predefined stylistic groupings without any fine-tuning.

Successfully classifying style does not conclude the analysis of the archaeologist. The more important question is what machine learning may tell us about how the identified styles relate to one another. If an algorithm can classify style, it implies that the machine has learned an internal representation that encodes relevant discriminative features through its visual analysis of the paintings. It follows that an analysis of the learned feature space of such an algorithm is, in essence, an analysis of the space of artistic style in the data. This has been demonstrated in the domain of classical art through an analysis of the activation space preceding a classifier’s output layer (Elgammal et al. 2018). This activation space describes a vector that represents the strength of activation across all the neurons of that layer. A visualisation of this activation space was able to portray the relative distances between stylistic groupings and even a correlation with the chronological order in which they were created.

Considering this, we go on to analyse our trained image encoding networks by visualising the stylistic clustering of Arnhem Plateau rock art in the activation space of encoded images. We quantify the stylistic distance between clusters, ultimately organising the styles into a single dimensional space which can be interpreted as a stylistic spectrum. If similarity in style is to be attributed to spatio-temporal proximity, then this stylistic spectrum may also implicate a relative inscription chronology. This approach to analysing style using learned representations has the potential to enhance our understanding in a wide range of archaeological applications.

Materials and method

Ethics

All work was conducted with approval from the Social and Behavioural Research Ethics Committee (SBREC) at Flinders University (project number 7704). Approval for this research was also provided by Marrku Traditional Owners who were present during all data collection and recording conducted on Marrku land. All necessary permits were obtained for the described study from the Northern

Land Council, which complied with all relevant regulations and agreements with Traditional Owner communities for the field recordings conducted in Arnhem Land.

Models

Network architecture has a significant role to play in the discrimination of image features. Transfer learning has never been applied to the domain of rock art so it is not initially clear which network architecture will have the best results at image classification with 'style' as a focussed internal metric. For this reason we considered a range of recent convolutional architectures as image encoders: AlexNet (Krizhevsky 2014), VGG (Simonyan and Zisserman 2014), and ResNet (He et al. 2016). All three networks were reconstructed as described in their respective papers, trained, then had their final layers removed.

Background data

The data used for training has a significant impact on the learning achieved by any of the given networks. Whilst each may be able to classify the data sets they have been trained on, the learned features used for this classification will differ. For this reason different datasets may learn features that are more or less applicable to the rock art domain. To better understand what type of image data might produce the most transferable learned features useful to distinguishing rock art styles, a number of different datasets were selected and used for background training data. The three background datasets used for training were MNIST, Quick Draw and ImageNet. The MNIST dataset is comprised of 70,000 handwritten digits. These digits have been drawn by hundreds of different people and so many variations exist within the basic form of each of the 10 digit characters.

This robust data set has been used for successful learning based classification since 1998 (LeCun et al. 1998). The images in the MINIST dataset were originally 28 by 28 pixels and were scaled to the input size of 224 by 224 pixels required by the chosen network architectures. This dataset was selected as a transfer learning candidate due to its parallels to the hand drawn figures in the primary dataset represented in monochromatic drawn lines. The Quick Draw dataset is a collection of 50 million drawings across 345 conceptual categories such as 'face', 'pizza', and 'fire hydrant' (Ha and Eck 2017). In contrast to MNIST, drawings in the Quick Draw dataset are pictographic, resulting in considerably more variation across any given category. Drawings were originally captured as timestamped vectors but for our purposes timing information was removed. Vectors were positioned and scaled to fill a 224×224 region, simplified using the Ramer–Douglas–Peucker algorithm with an epsilon value of 2.0, then rasterised with a stroke thickness of 16 pixels. Background and evaluation datasets were formed by randomly sampling 1000 images from each class.

Finally, we considered ImageNet, an image database organised according to concepts in the WordNet hierarchy (Deng et al. 2009). Currently, ImageNet provides an average of over 500 images per concept, each of which are quality controlled and human-annotated. In this work we did not handle ImageNet data directly. Rather, we made use of model parameters resulting from training a classifier on 1000 ImageNet concepts. Parameters were provided by the software package PyTorch (Paszke et al. 2017). ImageNet differs from the other two datasets in that it comprises photographs with three colour channels rather than binary masks or hand drawn representations.

Training

Training batches were formed by randomly sampling 64 images from a given training dataset. Irrespective of the dataset, an epoch was defined to comprise 100 batches. To determine loss, we imposed a cross-entropy objective to each classifier. This objective was combined with the standard back-propagation algorithm. The Adam optimisation algorithm (Kingma and Ba 2015) was applied with a learning rate of 1×10^{-4}. Weight decay regularisation was not used. The learning rate was reduced by a factor of 10 every time 10 epochs passed without observing a reduction in training loss. Training was terminated when the learning rate dropped below 1×10^{-8}. No training was performed using the primary rock art dataset.

Application to rock art

A primary data set was formed of Arnhem Plateau rock art from a variety of sources. A total of 98 motifs were sourced, each of which depicted a single human figure from one of six known stylistic classes: Dynamic Figures (DF), Post-Dynamic Figures (PDF), Northern Running Figures (NRF), Simple Figures with Boomerangs (SFWB), Simple Figures with Round Headdresses (RH), and Wilton River Region Simple Figures (SFWR). These classes were populated with examples both from published literature sources and from photographs taken of the rock art directly. The classes chosen represent

Table 1. Summary of the primary dataset formed from various sources.

Acr.	Style	Sources	Image count
DF	Dynamic Figures	(Chaloupka 1984, 1993; Chaloupka pers. comm.; May et al. 2018)	16
PDF	Post-Dynamic Figures	Marrku Traditional Owners	16
NRF	Northern Running Figures	Chaloupka 1993	16
SFWB	Simple Figures with Boomerangs	Chaloupka 1993	16
RH	Simple Figures with Round Headdress	Marrku Traditional Owners	18
SFWR	Wilton River Region Simple Figures	Marrku Traditional Owners	16

motif styles observed throughout Arnhem Land with the exception of the Northern Running figures which have an observed distribution limited to a small area on the northwestern corner of the Arnhem Plateau (Jones and May 2017; May et al. 2018).

The motifs gathered from photographs for this study are from the Wilton River Region of central Arnhem Land as part of an ongoing research project. Table 1 shows the classes chosen and the data source for each class. The classes chosen from the Wilton River Region were included to provide a geographical study area that falls outside of the areas represented by published sources. Figure 1 shows the Wilton River region and the location of the Northern Running Figure distribution relative to the Arnhem Plateau. Since the images in the primary dataset were gathered from multiple sources, they were first standardised. All images were traced to produce masked representations with no associated colour information. This also removed any rock surface information from consideration. All tracings were oriented so that the depicted figure was approximately upright and then zero-padded until square. Finally, the images were scaled to be 224×224 pixels in size. It should be noted that processing the images in this way discarded relative scale, rotation and colour information from the data that may have been pertinent to classifying style. We do not propose a solution to this limitation in the current work.

No fine tuning was performed using the rock art dataset. Rather, each convolutional encoder was used, as trained on background data, to encode rock art images into a vector space (activation space). For consistency, we chose to analyse the activation vector of the layer that follows the last convolutional layer of each network. In each case this activation vector had 4096 elements with the exception of ResNet, which had 2048 elements.

Evaluation was then performed by constructing a 6-way one-shot classification task (Bromley et al. 1993). For this, one test image and six additional images (one representing each of the known rock art styles) were chosen at random from the rock art dataset. The test image was then compared with each representative image in a pairwise manner and assigned the class of the representative image to which the network indicated the test image was most similar. Similarity was ascertained by taking the Euclidean distance between the pairs of vectors, with the vectors that were the least distant taken to be the most similar. The assigned class was compared with the test image's true class and this process was repeated to find the overall accuracy of each network when trained on each background dataset. Evaluation was performed on 25% of the data set. As a benchmark, the validation process was also conducted using raw pixel data from each image and the same after dimensionality reduction using PCA.

Analysis

Our analyses focussed on understanding and visualising the activation space of encoded rock art images. To help understand this space we visualise rock art in the activation space of the most accurate model by mapping it onto a single dimension. For this we used t-distributed stochastic neighbour embedding (t-SNE) (Van Der Maaten and Hinton 2008) for its ability to reduce dimensionality while maintaining the spatial relationships present in the original high dimensional space. This restructures the artwork representations into an inferred stylistic spectrum. The t-SNE data reduction was performed 100 times to ensure the resulting ordered data was repeatably reliable.

Results

Table 2 shows the accuracy of each encoder network model after training on each of the background datasets and then using the trained model to discern rock art styles in a one-shot setting. This accuracy score is based on the existence of known stylistic classes in the rock art data set. For comparison, a random guess achieves an accuracy of 16.67%, which all models were able to out-perform by a significant margin. Using the same nearest neighbour method but using the pixel data directly achieved an accuracy of 51.22%, which many model/dataset combinations were also able to out-perform. As a final comparison, the same method was applied to the pixel data after reducing their dimensionality using PCA to 98 dimensions, as dictated by the number of images in the dataset, achieving an accuracy of 43.00%.

Table 2. The accuracy of each encoder model after training on each background dataset when discerning rock art style in a one-shot setting.

	MNIST (%)	Quick draw (%)	ImageNet (%)
KochNet	38.60	49.21	–
AlexNet	43.22	49.96	72.92
VGG	42.00	54.01	69.86
ResNet	51.64	63.22	68.51

Figure 2 shows the results of t-SNE dimensionality reduction of the rock art data embedded in the activation space of the AlexNet model trained on ImageNet data. The dimensionality is reduced to a single dimension on the x-axis. Classes are artificially separated on the y-axis for ease of interpretation. Each class's mean is represented by a diamond while circular data points represent the class members. The order of the class means was consistent 100 out of 100 times when this activation space was reduced using t-SNE. Only minor variations in the placement of class outliers were observed in these 100 t-SNE iterations. Across the dataset of embedded images, the variance in each vector dimension was calculated. 25% of the total variance was accounted for by the 442 most dominant dimensions, 50% by 1212, and 75% by 2310. This indicates that the embedded vectors are not dominated by only a few dimensions.

Discussion

The accuracy attained by the networks tested demonstrate that the method presented can be used to separate previously identified stylistic classes of rock art without any prior training in the rock art domain. This demonstrates the capacity of transfer learning as a tool for further rock art analysis in the future. Of all the network architectures explored, AlexNet achieved the greatest classification accuracy on the rock art dataset in the one-shot setting. The ImageNet dataset served as the best background training data for all networks on which it was tested. This is a surprising result considering that the images in this dataset appear to be the least visually similar to the rock art data on which the networks were tested. This may be because the other datasets had fewer classes than ImageNet and may have caused the models to learn features that were too specific to be effectively transferred.

Most importantly, the method presented has the benefit of being less etic than existing observational or statistical approaches. While such an approach is by no means emic, the features of interest are extracted exclusively from data rather than being chosen by the researcher. For rock art analysis this means that any discriminating features that contributed to stylistic separation must be present and relevant within the rock art motifs and not independently identified by a researcher. This removes the bias inherent in manual feature selection. Analysis of the approach proves that learned features are sensitive to complex patterns that are not simply geometric but also stylistic. This is a key advantage over established statistical approaches such as Geometric Morphometrics which rely exclusively on geometric features distinguishing art forms and the correct identification of such features. The only disadvantage of the transfer learning approach over those that rely on hand-picked features is that this approach does not indicate precisely what it is that has been learned throughout the process. This contrasts with established statistical approaches which explicitly identify the nature of any patterns recognised (Travers and Ross 2016, 2017).

Researcher bias is further removed by the transfer learning approach whereby concepts of style and distance are learned from background data. This avoids the bias that would be introduced by defining rock art classes prior to training a machine learning approach on rock art data directly and disproves the contention of Dobrez and Dobrez (2019:15) that 'No appeal to Information Technology can override problems at the level of premises'. Without any fine tuning, the transfer learning approach was able to accurately make predictions about rock art style classes as inferred by Chaloupka (1984, 1993). By capturing clusters which distinctly correlate to established stylistic classes without being trained on this information directly strongly validates the choice of classes as distinct and separable styles. The ability to accurately separate these classes does not provide any further insight to the relationships between the styles each class represents. The analysis of the activation space does however clearly order these styles into a sequential gradient of similarity as shown in Figure 2. The dimensional reduction of the activation space has shown the order of stylistic similarity between classes matches the chronology originally developed by Chaloupka (1993) and further developed by Chippindale and Taçon (1998).

It should be noted that the styles identified through this approach are entangled with the class of images that are sampled within the data set. This is, in part, a limitation of the data set size, but more significantly a result of the observable trend in Arnhem Land art, which has only minor variation within a given class of art (such as human figures) for long time periods. This is particularly clear for the human class of figures where an identified stylistic phase is typically dominated by depictions of very specific technologies, postures, and activities. This is clearly demonstrated by the headdress depicted in motifs, which are distinct for each style

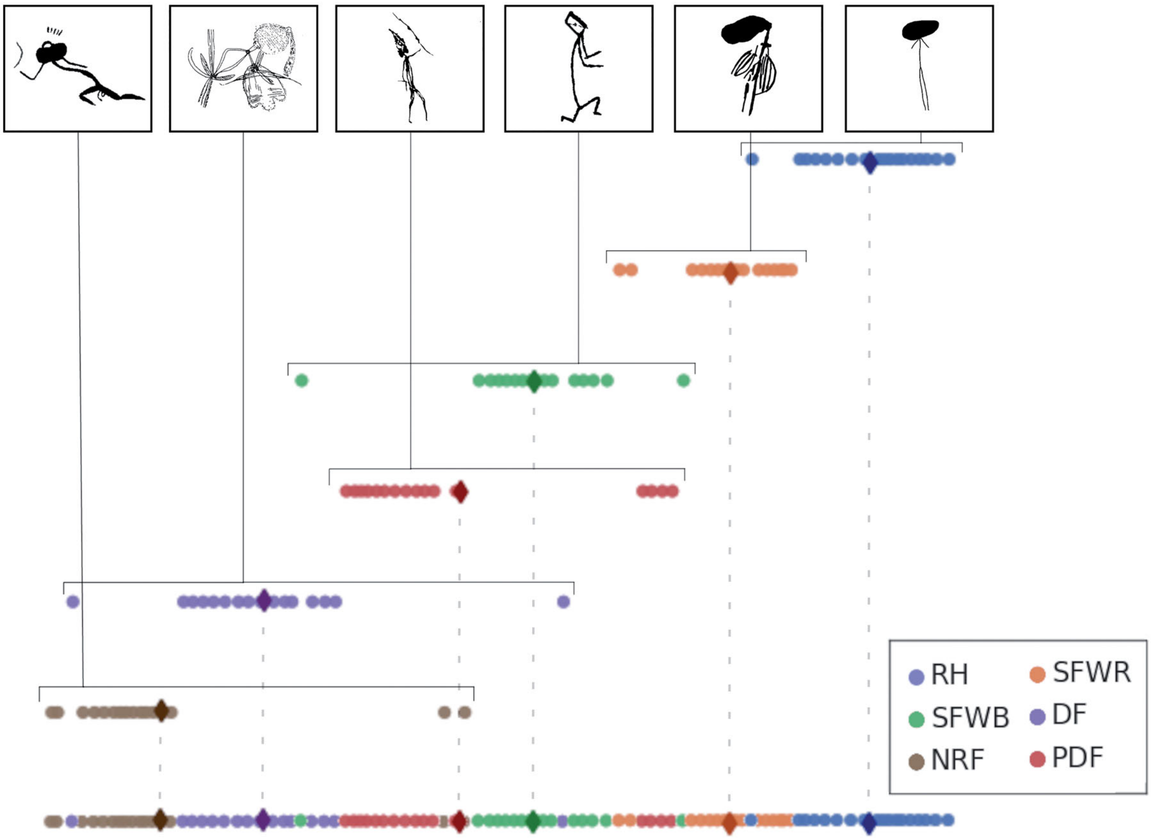

Figure 2. Rock art data embedded in the AlexNet/ImageNet activation space reduced to single dimension on the x-axis using the t-SNE method. Classes are separated on the y-axis for readability. Included human figures show examples of the artistic style of each class: Northern Running Figures (NRF), Dynamic Figures (DF), Post-Dynamic Figures (PDF), Simple Figures with Boomerangs (SFWB), Simple Figures from Wilton River (SFWR), Round Headdress figures from Wilton River (RH).

yet have very limited variation within that style. In this way we can extrapolate that this method is only useful for distinguishing style within a specific class of artistic depiction (in this case human figures). This extends to suggest that if depictions of another class (such as macropods for example) from the same stylistic phases as represented by the human figures, it would be unlikely that they would be correctly associated with human figures of the same style. It is important to note that with traditional approaches to stylistic analysis the association of different classes of depictions belonging to the same style are heavily dependent on superimposition and only partly based on factors of artistic commonality such as brushwork (Chaloupka 1993; Chippindale and Taçon 1998; Lewis 1988). Where traditional approaches have identified such a commonality it is unclear if this observation is accurate or influenced by a knowledge of the motifs association to one another though the superimposition chronology. A detailed investigation of the activation space allows us to demonstrate that this method is still sensitive to subtle stylistic variations within a class and is not entirely focussed on macro features within each class (such as headdress).

The t-SNE dimensionality reduction technique proved effective for visualising the encoder's activation space so that learned style information could be viewed as a stylistic spectrum. Figure 2 shows that the AlexNet model trained on ImageNet data, which resulted in the highest accuracy of all tested model/data combinations, was able to independently reproduce Chaloupka's chronology from learned features. This implies that the progression of artistic style over time occurred on a stylistic gradient with similar human depictions occurring at similar times. It follows that a gradual development of style rather than sudden and significant innovation may best explain the broad changes in art style in the Arnhem Plateau region. This stylistic gradient represents the temporal transition between styles apart from the Northern Running Figures. The position of the Northern Running Figures in the stylistic spectrum does not agree with their temporal position in the broad Arnhem chronology. This stylistic uniqueness may be indicative of the geographic isolation of this style, with its limited distribution to the north-west corner of the plateau. The development of this style could be interpreted to have occurred in a way that was less influenced by art of previous times and is similarly less influential on future stylistic evolution of the region. An artistic isolation that matches the limited geographic distribution of this style would be difficult or impossible to infer via other methods.

A similar uniqueness can be observed in the rock art data from the Wilton River Region. Two classes were selected from this region that represented art that matched Chaloupka's Simple Figures class and those that distinguished themselves with distinct round headdress. The Simple Figures from the Wilton River Region closely relate to Chaloupka's complex of Simple Figures, however they were still distinctly separable. This is also true for those with the distinct round headdress which occur on the outer most extreme of the stylistic gradient observed in the network's activation space. This may show regional uniqueness developing later in the chronology. The Wilton River can be interpreted to be more closely related to the wider Arnhem Plateau styles during the Post-Dynamic Figures phase as the Post-Dynamic Figures class was only populated with art from Wilton River and fits appropriately in the chronology. During later Simple Figure phase of rock art production in Arnhem Land, the Wilton River motifs may have become more regionally distinct. These distinctions still fall into the larger Simple Figures stylistic category but can also be further distinguished through classification. This provides a means by which nuanced differences in stylistic changes can be detected, quantified and given a relativistic account of their presence in a broader spectrum. The observation of nuanced stylistic differences and their attribution to a geographic trend, a temporal stylistic locality or the individualistic expression of distinct artists has been raised in the literature without a well-established means by which to determine these differences (Johnston 2017).

The placement of style on a continuous spectrum organised by comparative stylistic distances, allows class definitions to be made with a mathematical basis. This not only removes the subjectivity of class definitions but also allows individual data points to be viewed in a local neighbourhood of stylistic similarity, its placement therein possibly indicating more nuanced relationships to its localised peers. It may be possible to gain insight into what distinguishing features the neural network has learned through inspection of the individual figures represented by each point in Figure 2. Similar features can be seen within individual classes as well as between adjacent classes. The Dynamic Figures can be seen to be largely clustered together with a gradation of features being present. The class cluster shows figures with bolder outstretched running legs with a greater degree of infill being plotted closer to or overlapping Northern Running Figures. Figures with detailed grass skirts and other adornments are also plotted closely together. One outlier plots closely to the Simple Figures with Boomerangs class. This outlier has a more upwards posture and simple line work than other Dynamic Figures and these features may indicate their alignment with the Simple Figures with Boomerangs class.

Post-Dynamic Figures are mostly plotted in a cluster that is close to Dynamic Figures with the exception of four outliers which plot closer to the Simple Figures recorded in the Wilton River Region. This may be indicating that the Post-Dynamic Figure images used were all recorded in the Wilton River Region. These four outliers may demonstrate some geographic rationalisation to the Wilton River area within the Post-Dynamic Figures class.

One notable outlier to the Simple Figures class is a figure with a grass skirt. This figure was plotted among the Dynamic Figures class closely fitting with the only other examples within this data set depicting grass skirts. This makes the characterisation by the neural network obvious for this figure.

Interpretations of the images that make up the stylistic gradient are limited by the data size for each class and inferences from visual inspection must be considered loosely. The outliers and separation of clusters within the limited data size may indicate individual artistic expression (Johnston 2017). However, the broad pattern and transitions between the styles are made clear and some significant indication of what features the machine may have considered can be clearly observed. This adds value to the approach as it allows some amount of interrogation of the results which could be further strengthened with the inclusion of more data. Finally, the success of this case study suggests a much wider use for machine learning approaches within archaeological research. The transfer learning methodology means that the training data set is generic and so can be applied to any archaeological assemblage, including those in 3 D and colour. It removes the need to use large, research question specific, training data and so opens this method for use on other material culture items such as pottery and stone artefacts.

Conclusion

This research demonstrate that machine learning provides a means by which minor stylistic patterns in rock art can be detected, quantified and analysed. Knowledge of other associated variables such as geographic distribution allows some elements of the unsupervised stylistic distinctions to be better understood. This suggests a considerable potential role for this approach in resolving chronological and stylistic questions in rock art research more generally. The application of transfer learning has been able to identify a new Arnhem Land rock art style. This method has demonstrated that Simple Figures from Wilton River are categorised separately to other early anthropomorphic figurative art styles in Arnhem Land. This way of organising style demonstrates the philosophical concept of styles existence as identification via comparison (Wiessner 1984). Organisation of style in such a continuum may remove the subjectivity of stylistic class definition and provide universal means by which style can be organised in a cultural context, for almost any archaeological materials.

Acknowledgements

The language, images and information contained in this publication includes reference to Indigenous Knowledge including traditional knowledge, traditional cultural expression and references to biological resources (plants and animals) of the Marrku and Dakkal people. The source Indigenous Knowledge is Confidential Information, traditional law and custom applies to it and the Marrku and Dakkal people assert copyright over it in addition to any copyright in the complete work. The language, images and information are published with the consent of Mimal Aboriginal Corporation as the representative of the Marrku and Dakkal people pursuant to a limited licence for publication which is for the purposes of general education. No further use and absolutely no commercial use is authorized without the prior consent and agreement of the Marrku and Dakkal people. Such use would be beyond the scope of the licence granted. Please contact Mimal Aboriginal Corporation to request permission to refer to any Indigenous Knowledge in this publication and in particular for permission to use or reproduce the imagery of the rock art and its surrounding environment.

The authors would like to acknowledge Mimal Land Management Aboriginal Corporation, Marrku Traditional Owners, Elders Dudley Lawrence, Jack Nawilil, Robert Redford, Abraham Weson, and consultant Mr Peter Cooke for their contributions to this research. The authors thank Manilakarr Traditional Owner Alfred Nayinggul. All necessary permits were obtained for the described study from the Northern Land Council, which complied with all relevant regulations.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

This research was supported by Australian Research Council DECRA grants to DW [DE170101447] and to IM [DE160100703], and a Flinders University Early Career Researcher Impact Seed Grant and a Research Investment Fund Grant to IM.

ORCID

Jarrad Kowlessar http://orcid.org/0000-0002-1390-0265
James Keal http://orcid.org/0000-0002-2428-2239
Daryl Wesley http://orcid.org/0000-0002-3537-5668
Ian Moffat http://orcid.org/0000-0002-2171-7145

References

Brandl, E.J. 1973 *Australian Aboriginal Paintings in Western and Central Arnhem Land: Temporal Sequences and Elements of Style in Cadell River and Deaf Adder Creek Art*. Canberra: Australian Institute of Aboriginal Studies.

Bromley, J.A.N.E., J.W. Bentz, L.É.O.N. Bottou, I. Guyon, Y.A.N.N. Lecun, C. Moore, E. Säckinger and R. Shah 1993 Signature verification using a Siamese time delay neural network. *International Journal of Pattern Recognition and Artificial Intelligence* 07(04):669–688.

Caruana, R. 1997 Multitask learning. *Machine Learning* 28(1):41–75.

Chaloupka, G. 1984 *From Palaeoart to Casual Paintings: The Chronological Sequence of Arnhem Land Plateau Rock Art*. Monograph Series 1. Darwin: Northern Territory Museum of Arts and Sciences.

Chaloupka, G. 1993 *Journey in Time: The World's Longest Continuing Art Tradition: The 50,000 Year Story of the Australian Aboriginal Rock Art of Arnhem Land*. Chatswood, NSW: Reed.

Chippindale, C. and P.S. Taçon 1998 The many ways of dating Arnhem Land rock-art, north Australia. In C. Chippindale and P.S. Taçon (eds), *The Archaeology of Rock-Art*, pp.90–111. Cambridge: Cambridge University Press.

Cintas, C., M. Lucena, J.M. Fuertes, C. Delrieux, P. Navarro, R. González-José and M. Molinos 2020 Automatic feature extraction and classification of Iberian ceramics based on deep convolutional networks. *Journal of Cultural Heritage* 41(1):106–112.

Clarkson, C., Z. Jacobs, B. Marwick, R. Fullagar, L. Wallis, M. Smith, R.G. Roberts, E. Hayes, K. Lowe, X. Carah, S.A. Florin, J. McNeil, D. Cox, L.J. Arnold, Q. Hua, J. Huntley, H.E.A. Brand, T. Manne, A. Fairbairn, J. Shulmeister, L. Lyle, M. Salinas, M. Page, K. Connell, G. Park, K. Norman, T. Murphy and C. Pardoe 2017 Human occupation of northern Australia by 65,000 years ago. *Nature* 547(7663):306–310.

Clegg, J. 1987 Style and tradition at Sturt's Meadows. *World Archaeology* 19(2):236–255.

Cobden, R., C. Clarkson, G.J. Price, B. David, J.M. Geneste, J.J. Delannoy, B. Barker, L. Lamb and R.G. Gunn 2017 The identification of extinct megafauna in rock art using geometric morphometrics: A *Genyornis newtoni* painting in Arnhem Land, northern Australia? *Journal of Archaeological Science* 87:95–107.

Conkey, M.W. and C.A. Hastorf 1993 *The Uses of Style in Archaeology*. Cambridge: Cambridge University Press.

David, B., B. Barker, F. Petchey, J.-J. Delannoy, J.-M. Geneste, C. Rowe, M. Eccleston, L. Lamb and R. Whear 2013 A 28,000 year old excavated painted rock from Nawarla Gabarnmang, northern Australia. *Journal of Archaeological Science* 40(5):2493–2501.

Deng, J., W. Dong, R. Socher, L.J. Li, K. Li and L. Fei-Fei 2009 ImageNet: A large-scale hierarchical image database. IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops (CVPR Workshops), pp. 248–255. Miami, FL: IEEE. doi: 10.1109/CVPR.2009.5206848

Dobrez, L. and P. Dobrez 2019 Vasari, Sschapiro, Schaafsma: Three points of departure for a discussion of style in rock art. *Arts* 8(3):80.

Domingo Sanz, I. and D. Fiore 2014 Style: Its role in the archaeology of art. In C. Smith (ed), *Encyclopedia of Global Archaeology*, pp.7104–7111. New York: Springer.

Elgammal, A., M. Mazzone, B. Liu, D. Kim and M. Elhoseiny 2018 The Shape of Art History in the Eyes of the Machine. Paper presented to Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18). Retrieved 1 December 2020 from <http://arxiv.org/abs/1801.07729>.

Fei-Fei, L., R. Fergus and P. Perona 2006 One-shot learning of object categories. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 28(4):594–611.

Gunn, R.G. 2018 *Art of the Ancestors: Spatial and Temporal Patterning in the Ceiling Rock Art of Nawarla Gabarnmang, Arnhem Land, Australia*. Oxford: Archaeopress.

Gunn, R.G., L. Douglas and R. Whear 2018 The complexity of Arnhem Land rock art complexes. *Rock Art Research* 35(1):3–24.

Ha, D. and D. Eck 2017 A Neural Representation of Sketch Drawings. Paper presented to ICLR 2018. Retrieved 1 December 2020 from <http://arxiv.org/abs/1704.03477>.

Haskovec, I. 1992 Northern Running Figures of Kakadu National Park: A study of a regional style. In J. McDonald and I. Haskovec (eds), *State of the Art: Regional Rock Art Studies in Australia and Melanesia*, pp.148–158. Melbourne: Australian Rock Art Research Association.

He, K., X. Zhang, S. Ren and J. Sun 2016 Deep residual learning for image recognition. IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops (CVPR Workshops), pp. 770–778. Las Vegas, NV: IEEE. doi: 10.1109/CVPR.2016.90.

Hegmon, M. 1992 Archaeological research on style. *Annual Review of Anthropology* 21(1):517–536.

Horr, C., E. Lindinger and G. Brunnett 2014 Machine learning based typology development in archaeology. *Journal on Computing and Cultural Heritage* 7(1):1–23.

Johnston, I.G. 2017 Body ornamentation in dynamic scenes: Ritual and headdresses in the dynamic figure rock art of Western Arnhem Land, Australia. *The Artefact: The Journal of the Archaeological and Anthropological Society of Victoria* 38:67–80.

Jones, T. and S.K. May 2017 Rock art and ritual function: The Northern Running Figures of Western Arnhem Land. *The Artefact: The Journal of the Archaeological and Anthropological Society of Victoria* 38:53–66.

Jones, T., V.A. Levchenko, P.L. King, U. Troitzsch, D. Wesley, A.A. Williams and A. Nayingull 2017 Radiocarbon age constraints for a Pleistocene–Holocene transition rock art style: The Northern Running Figures of the East Alligator River region, Western Arnhem Land, Australia. *Journal of Archaeological Science: Reports* 11:80–89.

Jones, T., D. Wesley, S.K. May, I.G. Johnston, C. McFadden and P.S.C. Taçon 2020 Rethinking the age and unity of large naturalistic animal forms in early Western Arnhem Land rock art, Australia. *Australian Archaeology* 86(3):238–252.

Karayev, S., M. Trentacoste, H. Han, A. Agarwala, T. Darrell, A. Hertzmann and H. Winnemoeller 2014 Recognizing image style. In M. Valstar, A. French and T. Pridmore (eds), *Proceedings of the British Machine Vision Conference*. Nottingham: BMVA Press.

Kingma, D.P. and J.L. Ba 2015 Adam: A Method for Stochastic Optimisation. Retrieved 1 December 2020 from <https://arxiv.org/abs/1412.6980>.

Koch, G., R. Zemel and R. Salakhutdinov 2015 Siamese neural networks for one-shot image recognition. In F. Bach and D. Blei (eds), *Proceedings of the 32nd International Conference on Machine Learning*. Lille: JMLR.

Krizhevsky, A. 2014 One Weird Trick for Parallelizing Convolutional Neural Networks. Retrieved 1 December 2020 from <http://arxiv.org/abs/1404.5997>.

Lake, B.M., R. Salakhutdinov, J. Gross and J.B. Tenenbaum 2011 One shot learning of simple visual concepts. *Proceedings of the Annual Meeting of the Cognitive Science Society* 33(33):2568–2573.

Layton, R. 1992 *Australian Rock Art: A New Synthesis*. New York: Cambridge University Press.

LeCun, Y.A., L. Bottou, Y. Bengio and P. Haffner 1998 Gradient-based learning applied to document recognition. *Proceedings of the IEEE* 86(11):2278–2246.

Lewis, D. 1988 *The Rock Paintings of Arnhem Land, Australia: Social, Ecological and Material Culture Change in the Post-Glacial Period. BAR International Series 415*. Oxford: British Archaeological Reports.

May, S.K., I.G. Johnston, P.S. Taçon, I.D. Sanz and J. Goldhahn 2018 Early Australian anthropomorphs: Jabiluka's Dynamic Figure rock paintings. *Cambridge Archaeological Journal* 28(1):67–83.

Morphy, H. 2012 Recursive and iterative processes in Australian rock art: An anthropological perspective. In J. McDonald and P. Veth (eds), *A Companion to Rock Art*, pp.294–305. West Sussex: Wiley-Blackwell.

Paszke, A., S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga and A. Lerer 2017 Automatic differentiation in PyTorch. In U.V. Luxburg, I. Guyon, S. Bengio, H. Wallach and R. Fergus (eds), *31st Conference on Neural Information Processing Systems*. Long Beach: Curran Associates Inc.

Sackett, J.R. 1982 Approaches to style in lithic archaeology. *Journal of Anthropological Archaeology* 1(1):59–112.

Saleh, B., K. Abe, R.S. Arora and A. Elgammal 2016 Toward automated discovery of artistic influence. *Multimedia Tools and Applications* 75(7):3565–3591.

Schaafsma, P. 1985 Form, content, and function: Theory and method in North American rock art studies. *Advances in Archaeological Method and Theory* 8: 237–277.

Shamir, L., T. Macura, N. Orlov, D.M. Eckley and I.G. Goldberg 2010 Impressionism, expressionism, surrealism: Automated recognition of painters and schools of art. *ACM Transactions on Applied Perception* 7(2):1–18.

Simonyan, K. and A. Zisserman 2014 Very Deep Convolutional Networks for Large-Scale Image Recognition. Retrieved 1 December from <http://arxiv.org/abs/1409.1556>.

Taçon, P. and C. Chippindale 1994 Australia's ancient warriors: Changing depictions of fighting in the rock art of Arnhem Land, NT. *Cambridge Archaeological Journal* 4(2):211–248.

Travers, M. and J. Ross 2016 Continuity and change in the anthropomorphic figures of Australia's Northwest Kimberley. *Australian Archaeology* 82(2):148–167.

Travers, M. and J. Ross 2017 The shifting function of artefacts in Australia's Northwest Kimberley rock art assemblage. *The Artefact: The Journal of the Archaeological and Anthropological Society of Victoria* 38:8–26.

Tsigkas, G., G. Sfikas, A. Pasialis, A. Vlachopoulos and C. Nikou 2020 Markerless detection of ancient rock carvings in the wild: Rock art in Vathy. *Pattern Recognition Letters* 135(1):337–345.

Van Der Maaten, L. and G. Hinton 2008 Visualizing data using t-SNE. *Journal of Machine Learning Research* 9: 2579–2605.

Wang, H., Z. He, Y. Huang, D. Chen and Z. Zhou 2017 Bodhisattva head images modeling style recognition of Dazu Rock Carvings based on deep convolutional network. *Journal of Cultural Heritage* 27(1): 60–71.

Wiessner, P. 1984 Reconsidering the behavioral basis for style: A case study among the Kalahari San. *Journal of Anthropological Archaeology* 3(3): 190–234.