

FLESHING OUT THE BONES: STUDYING THE HUMAN REMAINS TRADE WITH TENSORFLOW AND INCEPTION

Damien Huffer, Shawn Graham



CASE STUDY

Fleshing Out the Bones: Studying the Human Remains Trade with Tensorflow and Inception

Damien Huffer* and Shawn Graham†

There is an active trade in human remains facilitated by social media sites. In this paper we ask: can machine learning detect visual signals in photographs indicating that the human remains depicted are for sale? Do such signals even exist? This paper describes an experiment in using Tensorflow and the Google Inception-v3 model against a corpus of publicly available photographs collected from Instagram. Previous examination of the associated metadata for these photos detected patterns in the connectivity and rhetoric surrounding this ‘bone trade’, including several instances where ‘for sale’ seemed to be implied, though not explicitly stated. The present study looks for signals in the visual rhetoric of the images as detected by the computer and how these signals may intersect with the other data present.

Keywords: neural networks; Tensorflow; human remains; machine learning; Instagram; social media

Introduction

The modern online trade in anatomical, ethnographic and archaeological human remains, especially using social media and e-commerce platforms such as Instagram, Facebook, Etsy, Marktplaats, Amazon.com, etc., is becoming documented and exposed (e.g. Huxley & Finnegan 2004; Huffer & Chappell 2014; Halling & Seidemann 2016; Huffer et al. in press). In our first paper on data mining the human remains trade on Instagram (Huffer & Graham 2017), we scraped several thousand posts to explore and map this well-connected network of collectors and dealers. We found individuals in this trade who are ‘specialist’ (exclusively specific categories) and ‘generalist’ (taxidermy, medical implements, bones of animals and humans, etc.) in their collecting focus. Furthermore, the display of their collections and the “true” stories of the acquisition of rare items seems to influence another group we identified, the ‘enthusiasts’, or those who do not seem themselves to collect, but rather who re-broadcast the photographs of those who do to help form tastes. Our previous work is limited in that it only mined the *text* of the posts, neglecting the rich visual data of the photographs.

The problem is one of identifying, at scale, the legal and ethical ramifications of a given photograph when the language of the accompanying post itself is occluded, in code or innuendo. Not everything for sale is marked as such. What were we missing? It takes a long time to manually study each of the thousands of images in our dataset to understand broader patterns. The manual nature of this type of content analysis has been a major obstacle to

analyzing images at a large scale in many fields. However, in recent years computational methods have become available that make it possible to train a typical desktop computer to recognize content in large collections. This technological breakthrough makes it possible for us to efficiently identify the contents of the thousands of images we have collected. In particular, the release of the Tensorflow library and the Google Inception-v3 convolutional neural network model (2017) trained on Imagenet (Stanford Vision Lab 2016) allows us to identify clusters of visually similar images, at scale, relatively quickly. These clusters can then be referenced against the data mining of the text in the accompanying posts.

Can we teach machines to automatically identify from photographs alone patterns in the ‘visual rhetoric’ which signal that an item is for sale? Our first step in analyzing visual rhetoric in our sample of Instagram images is to develop methods for the automatic clustering of images with similar patterns of features, not all of which are necessarily apparent to the human eye. We can then identify visual rhetoric through the choices made about how and what objects the photographer included in their image to convey meaning to their audience.

We first introduce neural networks and computer vision technology in general, and in archaeology in particular. Then we provide some brief background on the trade in human remains. We then walk the reader through our process using Tensorflow and Inception-v3 to turn our collection of images into image vectors. We describe how we reduced and visualized the multi-dimensional vector space in two dimensions using t-SNE, and clustered the results using affinity propagation. We close with remarks on the visual rhetoric of the human remains trade on Instagram.

* Stockholm University, SE

† Carleton University, CA

Corresponding author: Shawn Graham (shawn.graham@carleton.ca)

Neural Networks and Computer Vision

Computer vision technologies are typically based on a neural network (NN), a system that allows a computer to learn how to perform some task by analyzing training examples. The NN consists of layers of computational 'neurons', which are mathematical functions that mimic the functioning of a biological neuron. In our case, we want the computer to learn to recognize the contents of images in our sample, so it can identify large numbers of visually similar images. Our NN has been trained (not by us) by inputting a large set of images with known and labeled contents, so that the NN has developed a system for recognizing image contents by itself. During the training process, the labeled images trigger the neurons to fire in particular ways depending on the visual information contained in the image; images that are similar in some fashion cause similar cascades of firing. These cascades are represented as mathematical functions that the NN stores to use when we give it the unlabeled images in our sample. The NN uses these functions on our images to detect the content that we have trained it to recognize. This is a necessarily simplified description of what NN do. For an extended technical description see Bishop (1995). In our particular case, while it is the final layer of the NN that applies the label to the images, we can stop the process at the penultimate layer to retrieve the mathematical description of the images and use this data to identify visually similar materials.

Schmidhuber (2014) provides a comprehensive overview of NN from its origins in the 1960s to the present day. It highlights how NN used for 'deep learning' today differ from earlier 'shallow' NN (that is, the number and complexity of layers of neurons, and the 'credit assignment paths' or weights between neurons in these layers that make the NN perform the desired behavior). Whereas shallow NN typically require a great deal of supervision and training, the power of deep NN lies in the ability for them to learn without such supervision. Google's Inception-v3 model is a convolutional neural network (CNN), part of this class of deep NN (Szegedy et al. 2014). Because of their architecture explicitly engineered to mimic the connectivity in the visual cortex (which is referenced by the word 'convolutional'), convolutional networks have performed particularly well in image recognition tasks (cf. Olah 2014).

The 'Inception-v3 model' (Szegedy et al. 2014) was trained against the benchmark dataset, 'ImageNet' (Deng et al. 2009). Imagenet contains over 1000 classifications of nouns (i.e. items often found in images) where each noun is represented by hundreds or thousands of tagged images. While that *particular* dataset (ImageNet) has not been trained specifically on human remains, it is useful for us in that the *vector* representations of images used by the Inception-v3 model (from the penultimate layer before labeling) can be used to determine visual similarity. These vectors can be clustered. By finding clusters of similar images according to their contents, we can elucidate patterns in the visual rhetoric of these images which may then become the basis for filtering and tracking this trade.

Neural Networks in Archaeology

With the advent of powerful consumer-grade graphical processing units and open-source software, applications of NN to scientific problems have proliferated alongside bodies of code and other packages. Neural networks have appeared in the archaeological literature from at least the 1990s (cf. Baxter 2014 who provides an overview). Some studies concluded that the method did not have any more explanatory power than other more common models (Gibson 1996; Everitt & Dunn 2001), though work such as Bell and Croson (1998) concluded that early NN were particularly well suited for sparse datasets. Important work on NN in general in archaeology has been done by Barceló from the mid-1990s onwards (in particular, 1995, 2004, 2008; Barceló & Faura, 1999). Aprile, Castellano and Eramo (2014) found success using NN to classify mineral inclusions in potsherds. Ma et al. (2000) employed NN to aid in the classification of pottery. Kadar et al. (2004) used NN for archaeometric work with copper artefacts. Other work uses NN to enhance information retrieval on, for instance, pottery databases (Benhabiles & Tabia, 2016), or to determine whether or not degraded statuary belongs to a particular 'school' for the purposes of restoration (Wang et al. 2017). The above are just examples of a growing sub-field of research.

The Human Remains Trade

The questionable or illicit sale of cultural property is frequently estimated to be the third most profitable black-market industry following narcotics and weapons trafficking, "bring[ing] in \$2 billion to \$6 billion annually" (Choi 2011). While the exact dollar value of the antiquities trade is difficult to quantify, the amount of money involved in the human remains trade itself, at least on first reckoning (Huffer & Chappell 2014; Huffer & Graham 2017), appears low compared to other categories of antiquities trafficking. However, human remains trade research remains in its infancy. Financial considerations aside, indications are that its impact could be much larger. Recent research is bringing to light issues of damage to archaeological and ethnographic knowledge, vandalism of cemeteries historic and modern, theft from collections, (dealer) alleged complacency of professional scholars, loss of Indigenous cultural memory, and violation of the rights of descendent communities (Huffer et al. in press; Huffer & Charleton in review).

Redman (2016) provides a popular account of the various (often dubious) routes through which American and European museums and governments acquired human remains in the past. As some of these collections have been deaccessioned, individual items from them can sometimes find their way onto the commercial market. Even when circulating between collectors for years or decades, numerous questions remain surrounding how and when the skeletal remains appeared on the market, and (if archaeological in nature) what crucial contextual information could have been lost due to the looting act itself, regardless of when it took place (see for instance BABA0 2017).

Methods

Tensorflow is a software library developed by Google for machine learning using NN. It was open-sourced and released to the public in 2015. One of Google's tutorials for Tensorflow (2017) walks the user through the process of classifying a folder of images on the user's machine using the Inception-v3 CNN model. Enthusiasts worked out that the final step – classifying the images according to the vector representations in the model – could be skipped, leaving the user with the vector representations of similarity. Douglas Duhaime of Yale University's DH Laboratory wrote a blog post (2017) explaining this process in detail, including Python code, which we follow and employ here. We provide the computed image vectors and code (including copies of Duhaime's) used for the analysis we report here in our OSF research compendium (Graham 2018). We do not provide the images themselves in the compendium nor do we reproduce detailed versions of the images here, given the ethics of researching human remains (cf. BABA0 2010, item 6).

Our method can be summarized by the following steps:

- 1) **Turn our folder of images into vectors using Tensorflow and the Inception-v3 model.** The result of this process is that for each image we have a vector, the mathematical description of the image. We first installed Tensorflow using an Anaconda install of Python 3 on a Mac OS 10.10.5 laptop. Then, to classify a folder of images using Tensorflow we followed Duhaime (2017) and ran his modified 'classify-images.py' script to write the second-to-last layer of
- neural network weights from the model to a folder containing our 12010 image files. This took approximately 2 hours.
- 2) **Visualize those vectors by reducing the complexity down to two dimensions using t-SNE** (t-distributed stochastic neighbor embedding). Dimension reduction is a common step in handling NN output, regardless of the input type. It is necessary to visualize the data with currently available methods. This is a dimension reduction technique similar to Principal Component Analysis, but it uses non-linear methods that make it especially effective for high dimensional datasets such as images (Maaten & Hinton 2008). t-SNE optimizes for keeping points close to their neighbors, so it is an effective tool if we want to visualize which images are close together in our sample. We compute t-SNE by running Duhaime's 'get_tsne_vector_projections.py' script on our newly-created folder of image vector files. The .json formatted results of this operation are also in our repository.
- 3) **Identify clusters using affinity propagation.** Affinity propagation is a clustering algorithm that identifies exemplars among data points and forms clusters of data points around these exemplars. K-means is often used for clustering, but it is sensitive to the initial random selection of exemplars, and does not necessarily select the best representation of clusters; in Frey and Duecke's approach, all data points are considered as possible exemplars (2007). We imported the vector projections resulting from




Figure 1: t-SNE projection of image vectors of data collected for Huffer and Graham 2017.

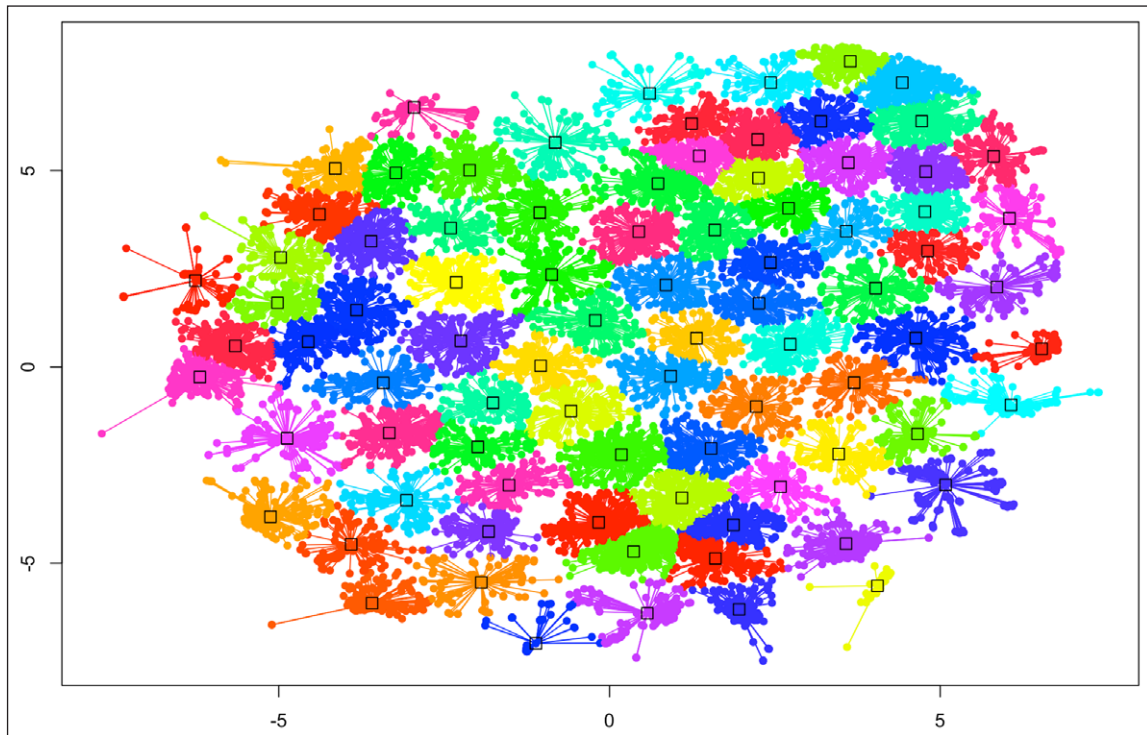


Figure 2: t-SNE projection of image vectors of data collected for Huffer and Graham 2017, colored to show the 84 clusters determined by affinity propagation. Squares indicate the location of the exemplar datapoint for a given cluster.




Figure 3: Html plot of the 84 exemplar images, zoomed out to obscure their details. Consideration of the associated post data, and similar plots of individual clusters within the broader dataset allows us to explore the meanings of the clusters.

the previous step into R to determine clusters with the affinity propagation algorithm using the R package `apcluster` (Bodenhofer et al. 2017), as described in Frey and Duecke (2007).

- 4) We **manually sampled** images in those clusters with the captions in the original posts to determine what the machine is seeing and why.

Results

Figure 1 shows the t-SNE projection of 12,010 image vectors. **Figure 2** is the same data, but colored according to the clustering determined by the affinity propa-

gation technique which also identifies ‘exemplar’ data points within clusters. We plotted these 84 exemplar images using the projection from the t-SNE and using the actual images as the data points as a locally-served website. **Figure 3** gives a zoomed-out view of this plot (thus obscuring the details of any one particular image; the html framework for that visualization is that used by Duhaime 2017). Looking at the html visualization of the data, we see that pencil-sketches and other ‘artistic’ interpretations of human bones are for the most part located in the bottom right quadrant, while photos of people with bones are in the top left. Taxidermy and animal skeletons

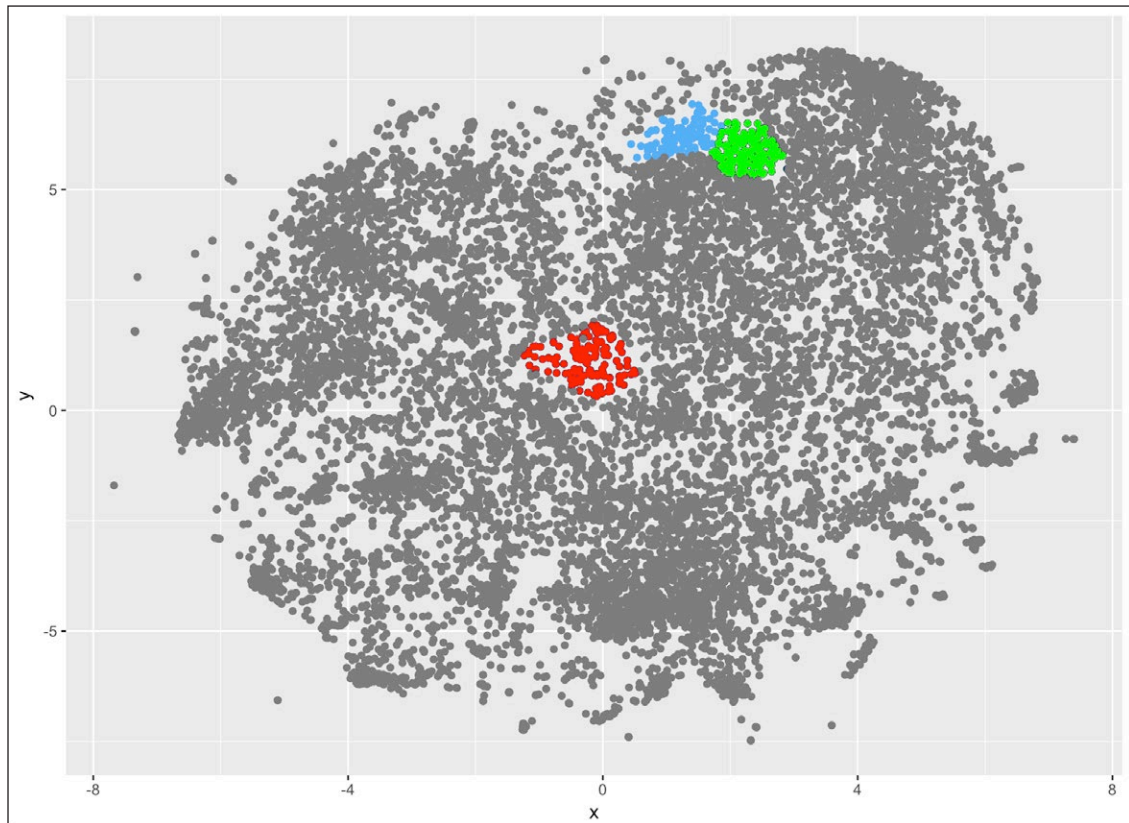


Figure 4: t-SNE projection of image vectors of data collected for Huffer and Graham 2017, where the red points indicate the location of cluster 35, the blue points indicate the location of cluster 80, and the green points indicate the location of cluster 82.




Figure 5: HTML plot of cluster 35, zoomed out to obscure details.

seem to be in the lower left, while human bones are in the upper right.

Discussion

In the very center of the complete projection is an image of a skull upon a shelf, with a price tag. This exemplar is from cluster 35. **Figure 4** plots clusters 35, 80 and 82, which we will discuss in more detail, in the context of the complete corpus as depicted in **Figure 1**. We plotted the 143 images assigned to cluster 35 in our html framework in order to explore it (**Figure 5**). It seems as if the criti-

cal feature that unifies these images is that they are of bones/skulls that are positioned on things – often, but not always, a shelf, (see **Figure 6**, a detail from the exemplar image for this cluster). The image is composed so that the foreground is in sharp focus and the other items on the shelf are blurred. They are reminiscent of mid-20th century museology, of items ranged in ordered rows, heightening the sense of ‘other’, distancing their humanity (cf. Redman 2016).

If we return to the html plot of the cluster exemplars, and consider the photos of human remains according to



Figure 6: Detail of the exemplar image for cluster 35, showing the placard indicating the human remains for sale. In the un-cropped image, a human skull is positioned to the right of the placard, facing the viewer's left.

their original associated text, two more clusters attract our immediate interest – cluster 80 (168 images), and cluster 82 (110 images). These clusters are in fact adjacent in the t-SNE plot, and quite distinct from the location of cluster 35. Cluster 80 is skulls that often have been photographed square to the face, and largely fill the frame, while cluster 82 seems to be skulls that are turned slightly to the left or right or upside down. In the associated posts for the images in cluster 80, the language used is of the ‘look what just arrived in my collection’ or ‘look what I just gave away’ discourse. Some photos are indicated as having been taken in a museum, and there is at least one photograph from a well-known business in this trade where the associated post advertises that the store is seeking skulls to buy. Items marked for sale are discussed obliquely, e.g., “new skull arrived... come pay your respects at [the] most amazing curiosity shop in Texas”.

In the associated posts for the images in cluster 82 are many explicit notices of materials for sale. A number of active business (predominantly in Canada, but also in the UK) account for several of the posts and clearly state the item is for sale, often naming a price. As was reported in Huffer and Chappell (2014), active businesses (with brick-and-mortar storefronts and/or online websites with catalogs, PayPal account details, contact information, etc.) are also known to exist in the US, the Netherlands, Belgium, Australia, and elsewhere. Perhaps there are national trends in visual rhetoric?

Visual Rhetoric in Images of Human Remains on Instagram

While the Inception-v3 model was never trained on human remains, the Tensorflow framework allows us to unpack its identification of visual similarities to the level where salient features of the images are identified. We can then use those vectors to explore patterns in the visual composition of these images that can then be

cross-referenced with the original language in the posts. This initial experiment does seem to support the idea that items for sale are displayed in ways that are discernible to the machine, and so, the machine can be taught to trawl other bodies of data for more evidence of the trade in human remains. The machine directs our attention to the framing of photographs, and the relationship of the human remains to other elements within the photograph. Exhibition design – rows of objects in cases on display – are recreated here. The interplay of foreground and background also seems to be important. Photos composed to show off a collection might also be subtly signaling that the item might also be for sale. These signals could be isolated, and used to train further iterations of a CNN, allowing a researcher to scale up their investigation. We intend to cross-reference this data with the network of followers and followed, to see how these visual clusters play out across networks of influence and on other platforms aside from Instagram.

Conclusions

Our research was motivated by the question of whether machine learning can detect visual signals in Instagram photographs indicating that the human remains depicted are for sale. We found that meaningful clusters (in terms of items for sale, or items for display) of similar images containing human remains can be identified by a neural network model, dimensionality reduction with t-SNE and affinity propagation clustering. We have demonstrated an approach to getting insights from large collections of images that may be useful in a variety of research contexts relating to cultural heritage and archaeology.

Can machine learning detect visual signals in Instagram photographs indicating that the human remains are for sale? The results of this initial experiment would seem to indicate ‘yes’. The positioning of a skull, for instance, relative to the plane of the camera; the arrangement of materials on shelving (or other objects), mimicking a museum display case; and foregrounds in sharp detail with backgrounds blurred all seem to be relevant signals. These all seem to be invitations to the viewer to consider the availability of an object for purchase or trade. No doubt we will discover more.

The results and discussion presented above are only the beginning of what we expect to be able to do using a neural network approach. Outside of human remains trade research itself, the use of this approach might be helpful in identifying many other kinds of materials bought and sold online, whether licit and with no public objections, or related to possible or confirmed illicit activity (such as drug or wildlife trafficking). The automated research into illicit markets of any kind, let alone cultural property, on social media, is still in its infancy (e.g. Yang & Luo 2017; Hernandez-Castro & Roberts 2015). Although our application of NN to the myriad questions raised by the human remains trade is still in its preliminary stages it adds to the growing body of research that is expanding the use of NN in general to studies of human history. Moving forward, we intend to replicate our studies on a growing data set obtained from monitoring various platforms.

Investigating the differences between social-media platforms could reveal how each platform is used by this particular community and give insights into why.

In addition to aiding other researchers investigating diverse forms of trafficking, the focus on using computer vision to study human remains places our work within the larger corpus of research that makes use of image classification, machine learning and NN in archaeology, bioarchaeology and forensic science in general. This corpus has included, for example, ceramic classification (e.g. Aprile et al. 2014), or sex, age and stature estimation of skeletons recovered from both archaeological and forensic contexts (e.g. Ionescu et al. 2016), and even examining whether or not the quick disposal of the dead during a conflict can leave a unique signature in the mortuary record of a place beyond the presence of a mass grave (Spars 2014). Other archaeological uses are possible, for example improving upon human abilities to detect new site looting from satellite images. The application of automated methods could speed up 'citizen science' methods to crowd-source data analysis efforts (Hersher 2017).

Acknowledgements

We would like to express our gratitude to Douglas Duhaime of the Yale Digital Humanities Lab, whose succinct and clear blog post will no doubt serve as the launch pad for several Tensorflow projects amongst digital humanists. We are grateful for the work of the anonymous reviewers whose comments make this work much better than it otherwise would have been. We are also indebted to Ben Marwick for reading various iterations of this piece, for his patient discussions with us regarding statistics, and for his general support and enthusiasm for open science methods in archaeology full stop. Thank you Ben! We hope to pay it forward someday in like manner. Errors are ours alone.

This research was supported by the Social Sciences and Humanities Research Council of Canada.

Competing Interests

The authors have no competing interests to declare.

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>
- BABAO – British Association for Biological Anthropology and Osteoarchaeology.** 2010. *Code of Ethics*. Available at URL: <http://www.babao.org.uk/assets/Uploads/code-of-ethics.pdf> [Last accessed 1 March 2018].
- BABAO – British Association for Biological Anthropology and Osteoarchaeology.** 2017. *Statement on the Sale of Human Remains*. Available at URL: <http://www.babao.org.uk/assets/Uploads/Sale-of-Human-Remains-Final-V2.docx> [last accessed 1 March 2018].
- Barceló, JA.** 1995. Back-propagation algorithms to computer similarity relationships among archaeological artefacts. In: Wilcock, J and Lockyear, C (eds.), *CAA 93. Computer Applications and Quantitative Methods in Archaeology 1993*, 165–176. Oxford: Tempus Reparatum.
- Barceló, JA.** 2004. Beyond classification: automatic learning in archaeology and cultural studies. *Intelligenza Artificiale*, 1: 29–33.
- Barceló, JA.** 2008. *Computational Intelligence in Archaeology*. Hershey, NY: Information Science Reference.
- Barceló, JA and Faura, JM.** 1999. Time series and neural networks in archaeological seriations. An example on early pottery from the Near East. In: Dingwall, L, Exon, S, Gaffney, V, Laflin, S and van Leusen, M (eds.), *Archaeology in the Age of the Internet. CAA97. Computer Applications and Quantitative Methods in Archaeology. Proceedings of the 25th Anniversary Conference, University of Birmingham, April 1997*, 91–102. Oxford: Archaeopress.
- Baxter, M.** 2014. Neural networks in archaeology. Available at: https://www.academia.edu/8434624/Neural_networks_in_archaeology [Last accessed 1 November 2017].
- Bell, S and Croson, C.** 1998. Artificial neural networks as a tool for archaeological data analysis. *Archaeometry*, 40(1): 139–151. DOI: <https://doi.org/10.1111/j.1475-4754.1998.tb00829.x>
- Benhabiles, H and Tabia, H.** 2016. Convolutional neural network for pottery retrieval. *Journal of Electronic Imaging*, 26(1): 011005-1:011005-6. DOI: <https://doi.org/10.1117/1.JEI.26.1.011005>
- Bishop, CM.** 1995. *Neural networks for pattern recognition*. Oxford: Oxford University Press.
- Bodenhofer, U, Palme, J, Melkonian, C and Kothmeier, A.** 2017. Apcluster version 1.4.4. *R Project*. Available at: <https://cran.r-project.org/web/packages/apcluster/apcluster.pdf> [Last accessed 5 November 2017].
- Choi, CQ.** 2011. NY mummy smugglers reveal vast antiquities black market. *Live Science*, 26 July 2011. Available at: <https://www.livescience.com/15234-ny-mummy-smugglers-reveal-vast-antiquities-black-market.html> [Last accessed 3 November 2017].
- Deng, J, Dong, W, Socher, R, Li, LJ, Li, K and Fei-Fei, L.** 2009. Imagenet: A large-scale hierarchical image database. In: *IEEE Conference on Computer Vision and Pattern Recognition*, 248–255. IEEE. DOI: <https://doi.org/10.1109/CVPR.2009.5206848>
- Duhaime, D.** 2017. Identifying similar images with Tensorflow, 28 May 2017. Available at: <http://douglasduhaime.com/posts/identifying-similar-images-with-tensorflow.html> [Last accessed 1 November 2017].
- Everitt and Dunn.** 2001.
- Frey, B and Dueck, D.** 2007. Clustering by Passing Messages Between Data Points. *Science*, 315(5814): 972–976. DOI: <https://doi.org/10.1126/science.1136800>
- Gibson.** 1996.

- Google.** 2017.
- Graham, S.** 2018. Fleshing Out the Bones (Project), 1 March 2018 [Last accessed 5 March 2018]. DOI: <http://doi.org/10.17605/OSF.IO/9CFJA>
- Halling, CL and Seidemann, RM.** 2016. They sell skulls online?! A review of internet sales of human skulls on eBay and the laws in place to restrict sales. *Journal of Forensic Sciences*, 61(5): 1322–1326. DOI: <https://doi.org/10.1111/1556-4029.13147>
- Hernandez-Castro, J and Roberts, DL.** 2015. Automatic detection of potentially illegal online sales of elephant ivory via data mining. *PeerJ Computer Science*, 1(e10). DOI: <https://doi.org/10.7717/peerj-cs.10>
- Hersher, R.** 2017. Space archaeologist wants citizen scientists to identify archaeological looting. *The Two-Way: Breaking News from NPR*, 31 January 2017. Available at: <http://www.npr.org/sections/thetwo-way/2017/01/31/512661370/space-archaeologist-wants-citizen-scientists-to-identify-archaeological-looting> [Last accessed 5 November 2017].
- Huffer, D and Chappell, D.** 2014. The mainly nameless and faceless dead: an exploratory study of the illicit traffic in archaeological and ethnographic human remains. *Crime, Law, and Social Change*, 62(2): 131–153. DOI: <https://doi.org/10.1007/s10611-014-9528-4>
- Huffer, D, Chappell, D, Charlton, N and Spatola, B.** In press. Bones of contention: The online trade in archaeological, ethnographic and anatomical human remains on Instagram. In: Chappell, D and Hufnagel, S (eds.), *Art Crime Handbook*. London: Palgrave Macmillan Press.
- Huffer, D and Charlton, N.** In review. Serious enquiries only, please: Ethical issues in the online human remains trade. In: Marquez-Grant, N and Squires, K (eds.), *Ethical Challenges in the Analysis of Human Remains*. London: Springer Press.
- Huffer, D and Graham, S.** 2017. The Insta-Dead: The rhetoric of the human remains trade on Instagram. *Internet Archaeology*, 45(5). DOI: <https://doi.org/10.11141/ia.45.5>
- Huxley, AK and Finnegan, M.** 2004. Human remains sold to the highest bidder! A snapshot of the buying and selling of human skeletal remains on eBay®, an internet auction site. *Journal of Forensic Science*, 49(1): 1–4.
- Ionescu, V-S, Teletin, M and Voiculescu, EM.** 2016. Machine learning techniques for age at death estimation from long bone lengths. In: *Proceedings of the 2016 IEEE 11th International Symposium on Applied Computational Intelligence and Informatics (SACI)*, 457–462. Timișoara: IEEE. DOI: <https://doi.org/10.1109/SACI.2016.7507421>
- Kadar, M, Ileana, I and Joldes, R.** 2004. Artificial neural networks used in forms recognition of the properties of ancient copper based alloys. In: Niccolucci, F and Hermon, S (eds.), *Beyond the artefact. Digital Interpretation of the Past. Proceedings of CAA2004. Prato 13–17 April 2004*, 448–450. Budapest: ArcheoLingua, 2010.
- Ma, Q, Yan, A, Hu, Z, Li, Z and Fan, B.** 2000. Principal component analysis and artificial neural networks applied to the classification of Chinese pottery of Neolithic age. *Analytica chimica acta*, 406(2): 247–256. DOI: [https://doi.org/10.1016/S0003-2670\(99\)00764-3](https://doi.org/10.1016/S0003-2670(99)00764-3)
- Maaten, LVD and Hinton, G.** 2008. Visualizing data using t-SNE. *Journal of machine learning research*, 9: 2579–2605.
- Olah, C.** 2014. Conv Nets: A Modular Perspective, 8 July. Available at: <https://colah.github.io/posts/2014-07-Conv-Nets-Modular/> [Last accessed 1 November 2017].
- Redman, SJ.** 2016. *Bone rooms: From scientific racism to human prehistory in museums*. Cambridge, MA: Harvard University Press. DOI: <https://doi.org/10.4159/9780674969711>
- Spars, S.** 2014. Exposing those who bury the dead: A new perspective in modern conflict archaeology. In: Cornish, P and Saunders, NJ (eds.), *Bodies in Conflict: Corporeality, Materiality and Transformation*, 103–122. London: Routledge Press.
- Stanford Vision Lab.** 2016. *ImageNet*. Available at URL: <http://image-net.org/> [Last accessed 5 November 2017].
- Szegedy, C, Liu, W, Jia, Y, Sermanet, P, Reed, S, Anguelov, D, Erhan, D, Vanhoucke, V and Rabinovich, A.** 2014. Going Deeper with Convolutions. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, arXiv:1409.4842. Available at: <https://arxiv.org/abs/1409.4842>.
- Wang, H, He, Z, Huang, Y, Chen, D and Zhou, Z.** 2017. Bodhisattva head images modeling style recognition of Dazu Rock Carvings based on deep convolutional network. *Journal of Cultural Heritage*, 27(3): 60–71. DOI: <https://doi.org/10.1016/j.culher.2017.03.006>
- Yang, X and Luo, J.** 2017. Tracking illicit drug dealing and abuse on Instagram using multimodal analysis. *ACM Transactions on Intelligent Systems and Technology (TIST) – Special Issue: Cyber Security and Regular Papers*, 8(4): 58:1–58:15.

How to cite this article: Huffer, D and Graham, S. 2018. Fleshing Out the Bones: Studying the Human Remains Trade with Tensorflow and Inception. *Journal of Computer Applications in Archaeology*, 1(1), pp. 55–63, DOI: <https://doi.org/10.5334/jcaa.8>

Submitted: 08 November 2017

Accepted: 06 May 2018

Published: 30 May 2018

Copyright: © 2018 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/>.

]u[

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

OPEN ACCESS 