

ARTIFICIAL INTELLIGENCE TO FIGHT ILLICIT TRAFFICKING OF CULTURAL PROPERTY

D. Abate^{1,6*}, A. Agapiou¹, K. Toumbas¹, A. Lampropoulos², K. Petrides², R. Pierdicca³, M. Paolanti⁴,
F. Di Stefano³, A. Felicetti⁵, E. S. Malinverni³, P. Zingaretti⁵

¹The Cyprus Institute, Athalassa Campus, 2121, Nicosia, Cyprus - (d.abate, an.agapiou, k.toumbas)@cyi.ac.cy

²HYPERTECH S.A., 32 Perikleous St, 15232 Athens, Greece - (a.lampropoulos, k.petrides)@hypertech.gr

³Università Politecnica delle Marche, Dipartimento di Ingegneria Civile, Edile e dell'Architettura, 60131 Ancona, Italy –
(r.pierdicca, f.distefano, e.s.malinverni)@staff.univpm.it

⁴University of Macerata, Department of Political Sciences, Communication and International Relations, VRAI Lab, 62100
Macerata, Italy – marina.paolanti@unimc.it

⁵Università Politecnica delle Marche, Dipartimento di Ingegneria dell'Informazione (DII), VRAI Lab, 60131 Ancona, Italy –
(a.felicetti, p.zingaretti)@staff.univpm.it

⁶Eratosthenes Centre of Excellence, Department of Resilience Society, 3102, Limassol, Cyprus -
dante.abate@eratosthenes.org.cy

KEY WORDS: Artificial Intelligence, Social Media Intelligence, Dark Web, Illicit Traffic, Cultural Heritage

ABSTRACT:

The use of artificial intelligence (AI) has the potential to be highly effective in detecting and monitoring illegal trafficking of cultural heritage (CH) goods through image classification techniques, particularly on online marketplaces where the trade of stolen CH objects has become a major global issue. Traditional investigation methods are no longer adequate, but with the assistance of AI, law enforcement agencies and CH organizations can now boost monitoring capabilities to detect, track, and possibly recover stolen objects more efficiently. AI algorithms can indeed analyze images to identify unique features and characteristics that can be used to determine their authenticity and provenance. Additionally, AI can detect patterns and networks of illicit trafficking, and link stolen objects to their places of origin, facilitating the recovery process. In this context, the SIGNIFICANCE project (Stop Illicit Heritage Trafficking with Artificial Intelligence) has been specifically designed to increase the response capabilities of public authorities and police corps against the illicit trafficking of cultural goods perpetrated through internet channels (i.e., social platforms, web, and dark web). By leveraging the power of Deep Learning (DL), AI can help prevent the loss of invaluable cultural artifacts and ensure that they are returned to their rightful owners and places of origin. This paper presents the results reached by the SIGNIFICANCE AI framework on image datasets collected over the web and social media through crawling algorithms.

1. INTRODUCTION

According to Hulkevych (2023), the trafficking of cultural artifacts is more rampant than ever before. While the consequences of cultural property crimes have been acknowledged since the "The Hague Convention for the Protection of Cultural Property in the Event of Armed Conflict" in 1954, the issue has become increasingly urgent in recent decades (Manacorda and Chappell, 2011).

Various factors have contributed to the expansion of this illicit phenomenon. The globalization of the art market has created fresh opportunities for criminals to profit from selling stolen cultural objects (Yates, 2015). Online marketplaces have made it easier to anonymously trade cultural goods without the need for physical storefronts or auction houses.

Political instability and conflicts in certain regions have resulted in a rise in looting and theft of cultural artifacts (Kar and Spanjers, 2017). During times of war or civil unrest, looters and traffickers view valuable cultural assets as a lucrative source of income or a means of money laundering (Davis and Mackenzie, 2014). Countries with compromised security conditions, particularly those that were once the territory of an-

cient civilizations, are particularly susceptible to these illegal activities.

Furthermore, the growth of illicit trafficking of cultural goods has been fueled by limited awareness and resources dedicated to prevention and combatting such crimes (Polk and Chappell, 2019; Delepierre and Schneider, 2015; Bokova, 2021). The inability to effectively prevent archaeological looting or trafficking is not confined to failed-state environments or specific regions; it is a challenge faced by many EU Member States (Casertano, 2020).

The outbreak of the COVID-19 pandemic has exponentially worsened the situation. While the antiquities market denies the widespread occurrence of illicit trade in cultural goods, law enforcement agencies overwhelmingly acknowledge the problem and identify organized crime involvement at every stage (Votey, 2021).

Despite unprecedented technological innovations benefiting numerous sectors with new tools, equipment, devices, and specialized knowledge, the heritage sector continues to struggle in fully harnessing these advancements to control the loot-

* Corresponding author

ing and trafficking of cultural assets (Fincham, 2019; Muheidat and Tawalbeh, 2021; Spalazzi et al., 2021).

This paper is organized as follows. Section 2 provides a description of the approaches that were adopted for the prevention of illicit traffic of CH goods. Section 3 describes the proposed DL-based pipeline. In Section 4, an evaluation of the proposed approach is offered, as well as a detailed analysis of each component of our DL-based framework. Finally, in Section 5, conclusions and discussion about future directions for this field of research are drawn.

2. RELATED WORKS

Public authorities and scientific communities involved in the protection and preservation of CH are failing to conserve and persuade others to do the same, resulting in ongoing destruction due to natural disasters and human-made actions. A significant proportion of the loss is attributed to looters who act for commercial reasons and are indirectly financed by private and sometimes public collectors of antiquities. To address this issue, public authorities, law enforcement agencies, and the research community are joining forces to develop technological solutions. Thanks to the very high priority of the topic, the European Commission is currently developing an action plan against trafficking in cultural goods for 2022-25 as part of the EU strategy on organized crime which reflects in several funding opportunities.

2.1 Research Projects

In recent times, various funded projects have been implemented. One such project is RITHMS (Research Intelligence and Technology for Heritage and Market Security)¹, which aims to enhance the operational capacity of Police and Customs/Border Authorities in combating organized and multi-criminal trafficking in cultural goods through research, technological innovation, outreach, and training. Another project, ENIGMA (Endorsing Safeguarding, Protection, and Provenance Management of Cultural Heritage), focuses on protecting cultural goods and artifacts from man-made threats by contributing to research on identification, traceability, and provenance (Patiás and Georgiadis 2023).

2.2 Stolen Object Databases

Several databases of stolen objects are available today, but not all of them are fully accessible to the public and academia. One of the largest databases is LEONARDO, maintained by the Command for the Protection of Cultural Heritage of the Italian Carabinieri, which contains over 1.2 million stolen objects, images, and theft cases. The LEONARDO database is the reference point for Italian and foreign law enforcement agencies and allows for a careful analysis of criminal phenomena related to the illicit trafficking of cultural property. The SWOADS Project and the mobile iTPC App, developed by Carabinieri TPC, are improving and expanding the software components of the LEONARDO database in terms of big data, machine learning, and blockchain technology.

The INTERPOL Stolen Works of Art database is the only international database with certified police information on stolen and missing objects of art. It combines descriptions and pictures of more than 52,000 items from 134 countries and allows users to complement their search by uploading a picture of any object of art and checking it with image-matching software. Studies in machine learning have been recently conducted to explore its applications in various domains related to CH, such as historical document processing and analysis of dance techniques.

The National Stolen Art File (NSAF) functions as a centralized database containing information about stolen art and cultural property². In May 1979, the FBI's Laboratory Division initiated the NSAF during its research phase. This computerized index acts as a repository for data and images of artwork that has been reported as stolen by local law enforcement agencies and FBI field offices, as well as items that have been recovered.

The Lost Art Database³ is maintained by the German Lost Art Foundation. It documents cultural property that was either demonstrably seized from their owners between 1933 and 1945 as a result of Nazi persecution, or for which such a seizure cannot be ruled out. In addition, the Lost Art Database lists cultural property seized, relocated, or removed during the Second World War.

2.3 Networks

EU CULTNET is an informal network of law enforcement authorities and experts competent in the field of cultural goods. It aims to strengthen coordination between law enforcement and cultural authorities and private organisations, i.e., antique shops, auction houses, online auctions, by identifying and sharing information on criminal networks suspected of being involved in illicit trafficking of stolen cultural goods.

"NETCHER – NETWORK and digital platform for Cultural Heritage Enhancing and Rebuilding" implemented a strong trans sectoral network as well as Recommendations on the fight against looting and trafficking of cultural goods. The participants of the project set up a European network of relevant operators and a Europe-wide charter of good practices to efficiently fight the illicit trafficking of antiquities.

2.4 Recent Technological Advancement

Winterbottom et al. (Winterbottom et al., 2022) used machine learning to track and analyze the sales of antiquities on auction websites, such as eBay. The authors used a large CNN that was already trained on more generic image tasks as a starting point, and then further train it on a specific dataset composed by the digital collection of the Durham Oriental Museum. They achieved a 73% instance classification accuracy on a diverse dataset of images of artefacts with 24,502 images and 4332 unique object instances

¹ <https://rithms.eu>

² <https://artcrimes.fbi.gov>

³ <https://www.lostart.de/en/start>

3. MATERIAL AND METHODS

The growing significance of AI and DL techniques in Cultural Property Protection is partly attributed to their successful application in areas such as image classification. Convolutional Neural Networks (CNNs), for example, have the potential to surpass human capabilities in object recognition and image analysis, especially in terms of processing large volumes of data, which makes them particularly valuable in this field (Walter, 2023). Similar to other neural networks, a CNN consists of an input layer, an output layer, and several hidden layers in between. These layers perform operations that modify the data to learn specific data features. Once the features have been learned across different layers, the CNN architecture proceeds to classification.

In the context of identifying illicit traffic of Cultural Heritage (CH) goods from images, CNNs utilize the extracted features from these images to learn to distinguish different object classes and subsequently classify new images, enabling them to identify suspicious items based on detected patterns. Through training on a diverse and extensive dataset, a CNN can acquire the ability to recognize subtle differences between images that might be challenging for humans to discern.

The SIGNIFICANCE framework (Abate et al., 2022) harnesses AI to analyze and categorize images with the aim of extracting valuable insights about the artwork they depict. It comprises three distinct phases: Data extraction and storage, Deep Learning for CH Image Classification, and Performance Evaluation, each described in detail.

To conduct a comprehensive evaluation of the framework, the "SIGNIFICANCE dataset," which is a collection of CH-related images specifically compiled for this study using web and social media crawling algorithms (as outlined in Subsection 3.3), was employed.

3.1 Data Extraction and Storage

This initial stage is the most complex and critical phase of the entire process. The primary objective is to identify cultural heritage goods using reliable and up-to-date data. To achieve this goal, a web and Instagram crawling algorithm was developed and executed twice daily. The social media crawler leverages the Instaloader library to connect with Instagram and download posts. To collect relevant data, the algorithm searches for specific hashtags associated with different types of artworks as defined by domain experts (e.g., coins, frescoes, etc.). Only information pertinent to the research objective, such as the post image, description, hashtags, user information, and geolocation (if available), is downloaded.

The downloaded images are then stored in an AWS (Amazon Web Services) S3 (Simple Storage Service) bucket, while the accompanying text data is saved in a JSON file (which contains a reference to the image) and stored in MongoDB, a document-oriented database.

The purpose of collecting supplementary data is to analyze CH objects from a temporal and geographical perspective, as it is essential to consider not only the image itself but also the accompanying information. Once this stage is complete, a filtering process is initiated to refine the data further.

3.2 Deep Learning for Cultural Heritage Image Classification

VGG16, developed by Simonyan and Zisserman in 2014, is a remarkably powerful pre-trained convolutional neural network (CNN) widely recognized for its effectiveness in numerous computer vision tasks, with a specific emphasis on image classification. This architecture is composed of a total of 16 layers, featuring 13 convolutional layers, 5 max-pooling layers, and 3 fully connected layers.

When utilizing VGG16, the input image is expected to have a size of $224 \times 224 \times 3$, representing its width, height, and the three color channels (red, green, and blue). The network processes this input image through its layers, performing convolutions and pooling operations to extract hierarchical features at different levels of abstraction. The final output of VGG16 is a probability distribution across a predefined set of classes, indicating the network's confidence scores for each class.

Due to its architecture and the pre-training process, where it is initially trained on a large dataset like ImageNet, VGG16 has demonstrated impressive performance in image classification tasks. It has achieved notable success in accurately recognizing and categorizing objects within images, making it a popular choice among researchers and practitioners in the field of computer vision.

VGG16 can be a powerful tool for CH image classification that can effectively capture the underlying features of the images and achieve high accuracy (Paolanti et al., 2019).

For CH image classification, the VGG16 network was fine-tuned on the SIGNIFICANCE dataset by replacing the last fully connected layer with a new one that has the appropriate number of output classes. To prevent overfitting and enable faster convergence, the pre-trained convolutional layers were frozen during training, while the new fully connected layer was trained using the SIGNIFICANCE dataset until convergence was achieved.

artefact	#samples	field	class	# samples for class			# all samples		
				total	train	test	total	train	test
icons	1678	location	Cyprus	839	671	168	1678	1342	336
			Other	839	671	168			
frescoes	718	location	Cyprus	248	198	50	718	574	144
			Italy	248	198	50			
		period	Byzantine	326	261	65	718	575	143
			Other	326	261	65			

Table 1. Artefact feature.

The classification task involved two phases: Multilabel classification and Artworks features. The goal was to classify artifacts and fields for each artwork. The dataset was split into a training set (80%) and a test set (20%), with class balance maintained in both datasets. To achieve this, a deep learning (DL) model was trained using transfer learning from a pre-trained model on ImageNet. The model was fine-tuned using an adaptive optimizer (Adam) with a learning rate of 10e-5, a mini-batch size of 32, and trained for 50 epochs.

3.3 SIGNIFICANCE Dataset

The classes that will be identified as the most important to recognize are the following: coins, frescoes, icons, manuscripts, and others. To illustrate the diversity of these classes, Figure 1 showcases a selection of images from the SIGNIFICANCE dataset. Each artwork in the dataset has been labeled according to various features such as location, period, material, and more. Each feature encompasses different classes, thereby providing a comprehensive overview of the dataset.

Table 1 and Table 2 report the data cardinality (number of images), for each class used for the training and testing of the models. The results are for both the phases defined.



Figure 1. Artefact Images, coin (top left), fresco (top right), icon (bottom left), manuscript (bottom right).

3.4 Performance Evaluation

Following, the standard metrics used to measure the performance of classification methods are described as follow.

Confusion matrix (Table 3) is a specific table layout that allows visualization of the performance of a deep learning algorithm. Each row of the matrix represents the instances in an actual class while each column represents the instances in a predicted class, or vice versa. The name stems from the fact that it makes it easy to see whether the system is confusing two classes (i.e. commonly mislabeling one as another). Support counts the samples belonging to a ground truth class.

Class	# samples for class			# all samples		
	total	train	test	total	train	test
coins	718	574	144			
frescoes	718	574	144			
icons	718	574	144	3590	3590	3590
manuscripts	718	574	144			
others	718	574	144			

Table 2. Artefact classes.

	Target class	Other Classes	Total
Target class	TP	FN	P
Other Classes	FP	TN	N
Total	P [^]	N [^]	

Table 3. Confusion Matrix.

True Positive counts the correctly matched samples. For example, it is predicted that an image shows a coin, and it truly is.

False Positive counts the false predicted samples belonging to other ground truth classes. For example, it is predicted that an image shows a coin, but it is not true.

False Negative counts the false predicted samples belonging to the same ground truth class. For example, it is predicted that an image shows another class, but it shows a coin.

Precision explains how many of the correctly predicted cases actually turned out to be positive. Precision is useful in cases where False Positives are a higher concern than False Negatives:

$$Precision = TP / (TP + FP) \quad (1)$$

Recall explains how many of the actual positive cases were predicted correctly with the developed model. It is a useful metric in cases where False Negatives are a higher concern than False Positives:

$$Recall = TP / (TP + FN) \quad (2)$$

F1 Score gives the harmonic mean between Precision and Recall metrics. It is maximum when Precision is equal to Recall:

$$F1Score = 2 * Precision * Recall / (Precision + Recall) \quad (3)$$

Accuracy simply measures how often the classifier correctly predicts. Accuracy can be defined as the ratio of the number of correct predictions and the total number of predictions:

$$Accuracy = (TP + TN) / (TP + FP + FN + TN) \quad (4)$$

4. RESULTS

The training phase of the experiments was performed using two GPUs in parallel to speed up the process. However, only one GPU was used in the testing phase to ensure fair and accurate evaluation.

Table 4 reports the results of the classification of Icons Location which examples are reported in Figure 2. Table 5 is the results of Frescos Location classification which examples are reported in Figure 3. Table 6 reports the results of Frescos Period which examples are reported in Figure 4. The results of the Multilabel classification are reported in Table 7.

	Precision	Recall	F1-score	Accuracy
Cyprus	0.98	0.99	0.99	
Others	0.99	0.98	0.99	
Overall	0.99	0.99	0.99	0.99

Table 4. Icons-Location Classification Results.



Figure 2. Images of icons location. Cyprus (left), Other (right).

	Precision	Recall	F1-score	Accuracy
Cyprus	0.96	0.99	0.97	
Italy	0.89	0.98	0.93	
Others	0.95	0.82	0.88	
Overall	0.93	0.93	0.93	0.93

Table 5. Frescoes-Location Classification Results.



Figure 3. Images of frescoes location. Cyprus (left), Italy (centre), Other (right).

	Precision	Recall	F1-score	Accuracy
Byzantine	0.94	0.92	0.93	
Others	0.94	0.95	0.94	
Overall	0.94	0.94	0.94	0.94

Table 6. Frescoes-period Classification Results.



Figure 4. Images frescoes' period. Cyprus (left), Other (right).

	Precision	Recall	F1-score	Accuracy
Coins	0.99	1.00	1.00	
frescoes	0.94	0.95	0.94	
icons	0.96	0.92	0.94	
manuscripts	0.97	0.98	0.98	
others	0.97	0.99	0.98	
Overall	0.97	0.97	0.97	0.97

Table 7. Artefacts Classification Results.

Table 8 shows the confusion matrix of the classification of Icons Location. Table 10 is the confusion matrix of Frescos Location classification. Table 9 shows the confusion matrix of Frescos Period. The confusion matrix of the Multilabel classification is reported in Table 11.

	Cyprus	Other	Total
Cyprus	166	2	168
Other	3	165	168
Total	169	167	

Table 8. Icons-Location Confusion Matrix.

	Byzantine	Other	Total
Byzantine	60	5	65
Other	4	74	78
Total	64	79	

Table 9. Frescos-Period Confusion Matrix.

	Target class	Italy	Other	Total
Cyprus	49	0	1	50
Italy	0	49	1	50
Other	2	6	36	44
Total	51	55	38	

Table 10. Frescoes-Location Confusion Matrix.

The results of this study demonstrate that the deep learning model achieved high accuracy in classifying the artifacts' typology and fields of each artwork. Specifically, the model accurately identified and classified each of the five classes of artwork (icons, frescoes, coins, manuscripts, and other) with an over- all accuracy of 93%. Additionally, the model extracted useful information about the depicted artwork, such as its location, period, and style.

The findings indicate that Deep Learning (DL) models hold great promise in automatically classifying images of artwork and extracting valuable information related to cultural heritage objects. These results have far-reaching implications across various domains, including art history, cultural heritage preservation, and Cultural Property Protection.

Given the escalating issue of illicit trade in cultural heritage goods, the need to develop novel tools and methodologies for identifying and tracking these objects has become crucial. In this context, the utilization of DL models presents a promising approach to address these objectives. By leveraging the capabilities of DL, it becomes possible to enhance efforts in identifying, cataloging, and monitoring cultural heritage items, ultimately contributing to the protection and preservation of cultural property.

The potential impact of employing DL models in this context extends beyond the realms of academia and research. It holds the potential to inform decision-making processes, aid law enforcement agencies, support cultural institutions, and raise public awareness about the significance of cultural heritage and the need for its safeguarding.

	Coins	Frescoes	Icons	Manu- scripts	Others	Total
Coins	144	0	0	0	0	144
Frescoes	3	137	4	0	3	144
Icons	0	7	132	4	1	144
Manuscripts	0	1	2	141	0	144
Others	1	1	0	0	142	144
Total	148	146	138	145	146	

Table 11. Artefact Confusion Matrix.

5. CONCLUSIONS AND FUTURE WORKS

In conclusion, this study demonstrates the potential of deep learning algorithms, particularly Convolutional Neural Networks (CNNs), in recognizing and identifying cultural heritage goods from images. By training a CNN on a large and diverse data- set of legit CH objects, it achieved a high level of accuracy in classifying new images as objects of interest for further investigations.

This study shows that the proposed image classification solution is effective, especially when using the most recent AI architectures. The SIGNIFICANCE AI framework might have several potential contributions to Cultural Property Protection:

- *Automated processing:* the framework can automate the identification of artwork images, making it faster and more efficient than manual methods;
- *Artwork classification:* The framework can classify artwork images into different categories (e.g., icons, frescoes, coins, manuscripts, and others), making it easier for investigators to search and organize artwork.
- *Specific feature extraction:* The framework can extract specific features (e.g., location, period, etc.) for each type of artwork, providing useful information for investigators and curators to better understand the artwork and its history in case of theft.

Overall, the framework could provide valuable insights and data for law enforcement agencies, art historians, and cultural heritage organizations working to combat the illicit trade of heritage as- sets.

ACKNOWLEDGEMENTS

This Project is funded by the European Union’s Internal Security Fund — Police 101003625 — SIGNIFICANCE — ISFP-2019-AG-OPC.

REFERENCES

- Abate, D., Paolanti, M., Pierdicca, R., Lampropoulos, A., Toumbas, K., Agapiou, A., Vergis, S., Malinverni, E., Petrides, K., Felicetti, A. et al., 2022. SIGNIFICANCE. STOP ILLICIT HERITAGE TRAFFICKING WITH ARTIFICIAL INTELLIGENCE. The International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences, 43, 729–736.
- Bokova, I., 2021. UNESCO’S response to the rise of violent extremism: A decade of building international momentum in the struggle to protect cultural heritage. Getty Publications.
- Casertano, L., 2020. Combating the Illicit Trafficking of Cultural Property: The Multifaceted Response to a Complex Challenge. *Global Jurist*, 20(1).
- Davis, T., Mackenzie, S., 2014. 15 crime and conflict: Temple looting in Cambodia. *Cultural property crime*, Brill, 292–306.
- Delepierre, S., Schneider, M., 2015. Ratification and implementation of international conventions to fight illicit trafficking in cultural property. *Countering Illicit Traffic in Cultural Goods*, 129.
- Fincham, D., 2019. Assessing the viability of blockchain to impact the antiquities trade. *Cardozo Arts & Ent. LJ*, 37, 605.
- Hulkevych, V., 2023. Countering illicit trafficking of cultural property in Ukraine. *Global Perspectives on Cultural Property Crime*, Routledge, 204–219.

Kar, D., Spanjers, J., 2017. Transnational crime and the developing world. *Global Financial Integrity*, 53–59.

Manacorda, S., Chappell, D., 2011. *Crime in the art and antiquities world: Illegal trafficking in cultural property*. Springer Science & Business Media.

Muheidat, F., Tawalbeh, L., 2021. Artificial intelligence and blockchain for cybersecurity applications. *Artificial Intelligence and Blockchain for Future Cybersecurity Applications*, Springer, 3–29.

Paolanti, M., Pierdicca, R., Martini, M., Felicetti, A., Malinverni, E., Frontoni, E., Zingaretti, P., 2019. Deep convolutional neural networks for sentiment analysis of cultural heritage. *The International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences*, 42, 871–878.

Patias P, Georgiadis C. Fighting Illicit Trafficking of Cultural Goods—The ENIGMA Project. *Remote Sensing*. 2023; 15(10):2579. <https://doi.org/10.3390/rs15102579>

Polk, K., Chappell, D., 2019. Policing and prosecution of heritage crime: revisiting the Cordata just how organised is the international traffic in cultural heritage? *Research Handbook on Transnational Crime*, Edward Elgar Publishing.

Simonyan, K., Zisserman, A., 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.

Spalazzi, L., Paolanti, M., Frontoni, E., 2021. An offline parallel architecture for forensic multimedia classification. *Multimedia Tools and Applications*, 1–16.

Votey, M., 2021. *Illicit Antiquities and the Internet: The Trafficking of Heritage on Digital Platforms*. NYUJ Int'l L. & Pol., 54, 659.

Walter, B., 2023. Analysis of convolutional neural network image classifiers in a hierarchical max-pooling model with additional local pooling. *Journal of Statistical Planning and Inference*, 224, 109–126.

Winterbottom, T., Leone, A., Al Moubayed, N., 2022. A deep learning approach to fight illicit trafficking of antiquities using artefact instance classification. *Scientific Reports*, 12(1), 1–12.

Yates, D., 2015. Illicit cultural property from Latin America: Looting, trafficking, and sale. *Countering illicit traffic in cultural goods: the global challenge of protecting the world's heritage*. International Council of Museums, 33–45.