Interferometric SAR and Machine Learning: Using Open Source Data to Detect Archaeological Looting and Destruction

HASSAN EL-HAJJ

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

RESEARCH ARTICLE

]u[ubiquity press

ABSTRACT

Archaeological heritage in the Near East is under an ever increasing threat from multiple vectors such as looting and systematic destruction, militarization, and uncontrolled urban expansion in the absence of governmental control among others. Physically monitoring endangered sites proves to be infeasible due to the dangerous ground conditions on the one hand, and the vast area of land on which they are dispersed. In recent years, the abundant availability of Very High Resolution (VHR) imaging satellites with short revisit times meant that it was possible to monitor a large portion of these sites from space. However, such images are relatively expensive and beyond the means of many researchers and concerned local authorities. In this paper, I present an approach that uses open source data from two of the European Space Agency's (ESA) Copernicus Constellation, Sentinel-1 and Sentinel-2 in order to generate disturbance patches, from which looting and destruction areas are classified using Machine Learning. Such an approach opens the door towards sustainable monitoring over large swaths of land over long periods of time.

CORRESPONDING AUTHOR:

Hassan El-Hajj, PhD Candidate

Classical Archaeology
Department, Philipps-
Universität Marburg,
Biegenstrasse 10, 35037
Marburg, Germany

hassanhajj910@gmail.com

KEYWORDS:

Coherence Maps; Near East;
Archaeological Looting;
Archaeological Destruction;
Cultural Heritage; Open Source

TO CITE THIS ARTICLE:

El-Hajj, H. 2021. Interferometric SAR and Machine Learning: Using Open Source Data to Detect Archaeological Looting and Destruction. *Journal of Computer Applications in Archaeology*, 4(1), pp. 47–62. DOI: <https://doi.org/10.5334/jcaa.70>

1 INTRODUCTION

Cultural heritage is generally seen in a positive light, and its preservation seen as a step that benefits everyone (Silverman et al., 2007: 3). This heritage is increasingly under threat in the Middle East and North Africa (MENA), as well as other regions around the world, where a power vacuum means that the relevant authorities are not able to protect these sites. Militants are trying to alter the region's history by destroying major archaeological sites such as Palmyra in Syria, looters and extremists are gutting archaeological sites for their artefacts which are sold on the black market, as was the case during the Lebanese Civil war (Fisk, 1991), military bases and outposts are encroaching on ancient tells such as at Tell Jifar in Syria (Casana and Panahipour, 2014) and urban development is spreading within previously protected areas as in the vicinity of Cyrene in Libya (Rayne et al., 2017).

The inability to physically monitor a large number of sites spread over a vast area, coupled with the volatile situation in the region meant that it was not possible to comprehensively assess the damage in real time. In response to the difficult ground situation, archaeologists turned to optical remote sensing to get a better view of the situation in Syria (Casana and Panahipour, 2014; Cunliffe, 2014; Casana, 2015), Iraq (Stone, 2004), Egypt (Parcak, 2015) and Libya (Rayne et al., 2017) as well as several large scale projects such as the Endangered Archaeology in the Middle East and North Africa (EAMENA) (Bewley et al., 2016) and the American School of Oriental Research's Project (Casana and Lauiger, 2017) among others. The majority of looting monitoring research relies primarily on Very High Resolution (VHR) imagery provided by commercial satellites (such as WorldView, GeoEye, etc.) which features an ever increasing ground resolution coupled with a high temporal resolution. Visual inspection of the satellite images is one of the methods that yields the most accurate results at the moment as trained operators can directly distinguish the presence of damage at the observed sites, as well as

the damage type. However, this method comes with two major disadvantages; the high prices of up-to-date VHR imagery, and the large number of sites to be monitored by operators.

Obtaining up-to-date imagery to cover a large area on a regular basis yields an enormous cost (up to 750 EUR per 25 square kilometres), far beyond the means of many researchers and concerned local authorities. To bypass this issue many researchers opted to use freely accessible VHR imagery provided by Google Earth, which also provides historical imagery allowing researchers to assess the damage to certain sites over a long period of time. One of the downsides of this approach is the fact that Google Earth (and similar engines) do not provide up to date results, leaving researchers oblivious to the present ground situation for months or years depending on the frequency of imagery updates, which can vary from one region to another. An example of such a long gap is the destruction of a fortified water hole in Barramiya, Egypt, which was highlighted by the EAMENA project. The destruction took place at some point between September 2012 and July 2013, a time span of almost a year (see [Figure 1](#)). This lack of information is due to the fact that Google Earth only updated their images at this specific location after this long gap. Such long waits between image updates mean that it was impossible to neither record the date of the destruction nor notice the destruction in-progress and attempt to interfere in such cases.

The very large number of endangered sites in the MENA region means that many operators are needed in order to consistently keep an eye on the ground situation. This is often unfeasible due to the high financial cost of such an approach, which means that researchers often focus on a handful of large sites while neglecting smaller sites that could be looted. Several projects reverted to crowd sourcing, such as TerraWatchers (Savage et al., 2017), using volunteers online to help detect looting in VHR imagery derived from Google Earth or similar. This has the advantages of covering a large area and numerous sites but falls again to the same disadvantages that

Figure 1 Destruction of the fortified water hole in Barramiya, Egypt as seen by consecutive Google Earth images (acquisition dates 11-09-2012 and 16-07-2013) – Original state shown in Blue box, destruction results in Red box.

come with using Google Earth images, as well as the lack of domain specific knowledge by general public members who participate in such projects.

More sophisticated approaches using VHR imagery were developed in order to automatically detect ground changes in and around cultural heritage sites. Looting was detected around Nicosia, Cyprus by using a combination of images acquired from World View-2 and Google Earth. These images were then enhanced by manipulating their histograms, as well as pan sharpening the World View-2 derived images and deriving other products (e.g. PCA, vegetation indices, vegetation suppression) in order to better detect soil disturbances, which are considered to be associated with archaeological looting (Agapiou et al., 2017). Other VHR approaches relied on texture analysis based on Gabor Filters (Manjunath et al., 1996) in order to detect the destruction of archaeological structures in Syria and Iraq (Cerra, Plank, et al., 2016; Cerra and Plank, 2020). Finally, VHR Google Earth images were also used to quantify the damage to archaeological sites, using an Automatic Looting Feature Extraction Approach (ALFEA), which aims at accurately map looting holes in and around archaeological sites (Masini et al., 2020). However, the fact that these approaches rely on VHR imagery, either acquired from a commercial provider (World View-2) or through Google Earth leads to the same problem of cost and contemporaneity of the data in question.

To circumnavigate the problems posed above, some researchers are reverting to open source imagery with a constant stream of images, albeit with lower resolution, such as ESA's Sentinel-2 or NASA's Landsat program. These platforms were used by the EAMENA team to map the urban expansion around Cyrene, Libya as well as the destruction around Homs, Syria (Rayne et al., 2017). A similar approach that uses color composites from open source Landsat-7 ETM+ was recently proposed to detect changes in and around the archaeological site of Apamea, Syria (Agapiou, 2020). Rayne et al. (2020) used the reflectance change between different Sentinel-2 passes in order to automatically highlight areas of disturbances around archaeological sites with good results. On the other hand, Synthetic Aperture Radars (SAR) sensors are less frequently used in archaeological contexts, however, the number of publications utilizing SAR in archaeological contexts is on the rise (Tapete and Cigna, 2019). SAR data was used to detect the ground disturbances on the Nasca Lines UNESCO World Heritage Site in Peru (Cigna et al., 2018), as well as to map the disturbances in Apamea (Tapete, Cigna & Donoghue, 2016). The archaeological applications of SAR also extend to the survey of large desert regions such as the case in Sinai (Stewart et al., 2016), as well as highlighting potential areas of buried archaeological remains (Stewart, 2017). One is encouraged to check the article by Tapete and Cigna (2019) which offers a review of the ever growing body of published work in this field.

While most of the above approaches rely on Very High Resolution data, especially when it comes to SAR, in this article, I propose an approach that relies on High Resolution open source data from ESA's Copernicus Sentinel-1 and Sentinel-2, as well as auxiliary geospatial data in order to filter the probable archaeological looting disturbances from vast monitored areas. It is not the aim of this approach to directly identify archaeological disturbances and their types; instead, this approach is to act as a filter, helping to narrow down the area to be monitored by highlighting potential sites that could be damaged between two successive satellite passes. It is hoped that by narrowing the search to a small number of sites per pass, it is then possible to acquire a small number of VHR images of the selected location which can then be assessed either by human operators or computer vision algorithms. Section 2 discusses the Methodology, while Section 3 explains the training data before presenting the results in Section 4. A general discussion of the approach, along with its strong points and weaknesses, is presented in Section 5 and concluding remarks in Section 6.

2 METHODOLOGY AND DATA

The proposed approach relies on the products derived from the Sentinel-1 constellation comprising of two satellites, Sentinel-1A and Sentinel-1B, each carrying a C-Band Synthetic Aperture Radar at 5.405 GHz, corresponding to 0.055 m wavelength (λ), with a temporal resolution of 12 days per satellite (ESA, 2020b). The first of the two Sentinel-1 satellites was launched in April 2014, while the second was launched two years later in April 2016, meaning that looting activities prior to 2014 are not investigated in this paper. All the used Sentinel-1 data is acquired as a Single Look Complex (SLC) in Interferometric Wide Swath (IW) mode which results in a 250 km swatch width and a 5 m by 20 m spatial resolution and an incidence range that varies between 29.10 degrees and 46.00 degrees at the near and far ranges (ESA, 2020b). While the Sentinel-1 data constitutes the core of this approach, supplementary data was collected from multispectral imaging satellites such as Landsat-8 (for periods prior to 2015) and Sentinel-2 (for observations after 2015). Both satellites carry optical instruments that sample multiple bands, in case of Landsat-8, 11 bands at a resolution that varies between 15 to 100 meters while in the case of Sentinel-2, 13 bands that vary between 10 to 60 meters. Besides the above mentioned satellite derived data, detailed plans (such as roads, urban areas, etc.) of the lands around each site in the form of shapefiles were acquired from open source repositories (HDX, 2020). The following paragraphs will discuss the methodology of the proposed approach in detail which can be divided into 4 main steps: data preprocessing, data training, prediction and assessment (see [Figure 2](#)).

Figure 2 Processing workflow from data acquisition to result assessment.

2.1 DATA PREPROCESSING

The data preprocessing runs along three distinct paths. The first is the SAR path which will generate a continuous supply of coherence maps over the monitored area, as well as provide us with accurate representation of urban settlements. The second is the Optical path which will yield information on the ground cover type by using vegetation and soil indices. The last and third one is the auxiliary data path which will provide valuable extra information to the feature vector in the form of streets, buildings, and other relevant geographic information in the vicinity of each site (Figure 2).

2.1.1 SAR path

The consecutive acquisitions from the Sentinel-1 satellites are obtained through the ESA Copernicus Open Access Hub (Copernicus Open Access Hub, 2019) database. Sentinel-1 SLC IW products are acquired using the Terrain Observation with Progressive Scan (TOPSAR) technique (De Zan et al., 2006), meaning each SLC product features numerous sub-swaths, in this case 3 sub-swath with 9 bursts each. The appropriate sub-swaths are selected and the needed bursts are extracted from both image pairs. The orbit files are then applied to the resulting products which are geocoded using the 3" Shuttle Radar Topography Mission (SRTM) Digital Elevation Model (DEM). The above-mentioned steps conclude the coregistration of the two SAR images, and allows us to start estimating the coherence between the two SAR acquisitions.

Interferometric Coherence (γ) estimation is generated from two SAR acquisitions of the same area from geometrically close orbital position, and relies on the phase information between them to estimate surface changes. It is calculated using the normalized cross-correlation between two selected acquisitions as shown in Equation 1, where s_1 and s_2 represent the signals acquired at two different times, t_1 and t_2 .

$$\gamma = \left| \frac{\langle s_{t1} \cdot s_{t2}^* \rangle}{\sqrt{\langle s_{t1} \cdot s_{t1}^* \rangle \langle s_{t2} \cdot s_{t2}^* \rangle}} \right| \quad (1)$$

Coherence is limited to the interval [0,1]. A strong coherence, γ close to 1, indicates that the two acquisitions are almost identical (including static speckle distribution) which can only occur if no ground change occurred between t_1 and t_2 , and the geometric location of the satellites at the time of acquisitions was very close. On the other end of the spectrum are the low coherence pixels which indicate that there was a change between the two acquisitions which can be attributed to many sources (Adragna et al., 2008: 293). These fall into several categories such as volumetric coherence which is caused mainly by leafy vegetation (γ_{volume}), orbital coherence caused by the satellite orbit position differences (γ_{orbit}), noise and processing errors ($\gamma_{processing}$) and most importantly when it comes to this approach, temporal coherence which is caused by ground changes between the two acquisitions ($\gamma_{temporal}$). This is better portrayed in Equation 2.

$$\gamma = \gamma_{volume} \times \gamma_{orbit} \times \gamma_{processing} \times \gamma_{temporal} \quad (2)$$

Volumetric coherence (γ_{volume}) loss due to vegetation is addressed by adding a vegetation index to the feature vector (see Section 2.1.2), while orbital coherence (γ_{orbit}) loss is considered to be negligible due to the well controlled orbital parameters of the Sentinel-1 constellation (ESA, 2020a). Processing coherence ($\gamma_{processing}$) loss which stems from noise in the system is neglected due to high signal to noise (SNR) ratio the Sentinel-1 system. Finally, temporal coherence ($\gamma_{temporal}$) loss is caused by changes on the ground between the two SAR acquisitions. This can be due to an infinite amount of reasons which can cause a change in the distance of the return signal from the satellite to the ground pixel between t_1 and t_2 . One of those reasons is archaeological looting where looters move soil by digging holes and creating heaps in order to get to archaeological layers. Additionally, total destruction of archaeological monuments contributes to temporal coherence loss by eliminating the building in question and thus altering the signal return distance. Once a coherence estimation map is created, multiple TOPS bursts are merged to create a continuous image.

Once created, coherence maps are thresholded in order to extract appropriate patches to be classified. The threshold value is chosen by inspecting the training data which includes cases of known looting and their appropriate coherence loss signatures. The threshold is then chosen based on the weighed mean of looting training patches as shown in Equation 3, where i represent pixel value, $\bar{x}_w$ is the weighed mean, x_i is pixel value within looting patches, and w_i is the weight given to each pixel. The latter is calculated using Equation 4 where p is the count of pixels in looting patch j .

$$\bar{x}_w = \frac{\sum_{i=1}^n w_i \cdot x_i}{\sum_{i=1}^n w_i} \quad (3)$$

$$w_i = \frac{1}{p_j} \quad (4)$$

Finally, the thresholded patches are vectorized and patches with very high coherence are discarded since they signify that no change occurred at that particular patch between t_1 and t_2 while the ones with lower coherence, i.e. falling below the calculated threshold are kept to be processed. One of the acquired SAR image is also used to generate an urban area mask by using speckle divergence which can be simply done using SNAP's radar processing features (Esch et al., 2013). A water body mask is also generated by taking advantage of the specular reflection of the signal off such surfaces. In this case, water bodies are masked out from any future processing.

2.1.2 Optical path

Optical data is obtained from two sensors. The Landsat data, used for observations prior to 2015, is downloaded through the United States Geological Survey database (USGS, 2020), while data from Sentinel-2, used for observations after 2015, is downloaded from the ESA Copernicus Open Access Hub (Copernicus Open Access Hub, 2019). This data is used to generate two main data items, the brightness index and the vegetation index which are later added to the feature vector in order to better understand the terrain that lies at each particular [Figure 2](#), which in turn can help to better understand the origin of the ground disturbance recovered from the InSAR Coherence patch analysis.

Vegetation cover is computed using the Normalized Vegetation Difference Index (NDVI), which can yield accurate masks highlighting pixels with healthy vegetation (Campbell et al., 2011: 518). NDVI takes advantage of the difference between the Near Infra-red and the Red band reflectance that occurs in healthy vegetation, and is calculated using Equation 5 below,

where the NIR is represented by Sentinel-2 Band 8 and RED by Band 4.

$$NDVI = \frac{NIR - RED}{NIR + RED} \quad (5)$$

Pixels with high healthy vegetation tend to generate a high loss of coherence due to volumetric decorrelation, as well as to the fact that the leaves of such trees are never in a constant position and thus are never in the same place between two consecutive satellite passes. Although numerous vegetation indices are available, NDVI was chosen due to the fact that it performed the best at matching the volumetric decorrelation generating from plants in our test cases.

The Brightness Index (Escadafal, 1989), BI, is used to estimate the brightness of soil and served as a good indicator as to where the surface is not covered by any vegetation. This is calculated using the equation below (Equation 6) where the Red band is represented by Sentinel-2 Band 4 and the Green by Band 3.

$$BI = \sqrt{\frac{RED^2 + GREEN^2}{2}} \quad (6)$$

2.1.3 Auxiliary data

Relative geographic data is added to better describe the location of the site with regards to its surroundings, rather than its actual geographic location. In this case, data describing the patch's accessibility via local roads, its proximity to settlement as well as known archaeological sites are inserted. These are derived from open access urban shapefiles for the countries in question (HDX, 2020).

2.2 TRAINING SITES

Collecting data on looting activities proved to be one of the most challenging parts of this research. The lack of dedicated databases recording archaeological looting means that it is almost impossible to collect a coherent body of data from a large region. While the EAMENA database is a good place to start, the data provided on each site does not always feature the date of disturbance (EAMENA, 2021), which hinders our ability to find temporally close SAR image pairs. Instead, data on looted archaeological sites is spread over numerous platforms, including numerous social networks such as Facebook and other platforms. This situation makes it difficult to collect coherent data, especially data regarding the looting and destruction dates of the sites in question. For this reason, this paper relies on a small dataset collected by the author from publicly known looting occurrences (i.e. publicized in news outlets), looting activities mentioned in research papers, as well as using Google Earth's historical imagery to assess

and date damage to known sites. The resulting set is a group of sites that were destroyed or looted distributed between Syria and Iraq as seen in [Figure 3](#).

3 DATA TRAINING

3.1 DATA PREPARATION

With the above data processing, we are able to construct a feature vector that describes each coherence patch. This can be divided into three sections originating from SAR, Optical, and Auxiliary data. SAR data provides coherence statistics and patch area, as well as urban area masks, while Optical data provides NDVI and BI statistics, finally Auxiliary data adds relative location information in the form of proximity to streets, urban centres, and known archaeological sites are added (see [Figure 2](#)).

Once the feature vector is constructed from the above data, we are then able to prepare it for training. The first step in this process is normalization. This is important due to the high scale variability between the different entries of the vector (Géron, 2017: 72), for example the NDVI Mean entry in the feature vector is a value between -1 and 1 while the Patch Area entry varies between 225 m^2 and 10000 m^2 . In order to achieve this, we scale each feature vector entry to a unit vector. Once the data is normalized, we reduce the dimension of the feature vector using Principal Component Analysis (PCA) to a total

of 8 principal components, this value showed the most promising results on the available data and preserved 95 per cent of the data variance ([Figure 2](#)). The dimension reduction occurs when the data is projected onto the hyperplane of defined by the 8 principal components.

At this stage, the data is divided into three main subsets: a holdout test set, kept on the side only to be used to test the effectiveness of the approach once the models are trained (see Section 4), i.e. not affected by the processing steps discussed in the following sections. The two other subsets are a training set and a validation set, split according to an 8:2 ratio from a balanced dataset (see [Figure 2](#) and Section 3.2) which will be further processed and used to train the models to classify the data. This is further discussed in the following sections.

3.2 BALANCING THE DATASET

Looking at the data at this stage of the analysis, it is very clear that the known looting patches are heavily outnumbered by those of coherence loss patches of non-looting origin by a ratio of 100 to 1 (see [Figure 4](#)), creating an imbalanced dataset (Hoens et al., 2013: 44). This is mainly related to the nature of the coherence maps which are littered with coherence loss patches due to the general nature of the terrain, as well as constant human activities between consecutive satellite acquisitions which can create a difference in the signal between

Figure 3 Map of the Near East showing the location of the training sites whose coherence patch signature was used to train the models.

two successive acquisitions. Such a heavy skew in data severely hinders the classification of minority classes (Weiss and Provost, 2003). Although many datasets are not balanced, an imbalance rate such as the one presented in this paper can be classified between modest and extreme imbalance (Weiss, 2013: 15).

This class imbalance problem is a rather common issue faced in data mining and data science. As a result, numerous approaches have been developed to overcome such imbalance. At the most basic level are over-sampling and under-sampling the minority and majority class respectively, which entails randomly replicating a number of the minority class or randomly deleting some of the majority class until a balanced dataset is obtained (Hoens et al., 2013: 45). Such naïve approaches to over and under sampling come with numerous disadvantages, in this case it is clear that by removing a large chunk of data entries in order to balance the dataset one loses a lot of information; on the other hand, simply copying the same data entry in order to balance the dataset results in numerous identical entries severely increases the risk of overfitting. To combat such disadvantages, numerous more sophisticated approaches to over and under sampling were developed, a detailed review of which can be found in Hoens and Chawla (2013).

Given the nature of the dataset presented in this paper, it is beneficial to use a mix of over-sampling the minority class and under-sampling the majority class. This is mainly due to the fact that, in addition to having an imbalance ratio of 100 to 1, the minority class is represented by a small number of entries, which means that under-sampling the majority class will result in a very small dataset prone to overfitting. In this case, Synthetic Minority Oversampling Technique + Edited Nearest Neighbour Rule (SMOTE+ENN) (Gustavo, Ronaldo & Carolina, 2004) is used in order to oversample the minority class while at the same time reducing some of the majority class; this in turn reduces the risk of overfitting. It must be noted that the Test Set is not affected by the SMOTEENN algorithm. In order to oversample the data using SMOTE, the below steps are applied (Hoens et al., 2013: 47):

- Locate random minority data instance a .
- Locate k nearest instances to a ; in this case $k = 5$.
- Consider instance b as one of the randomly chosen k neighbours.
- Add synthetic point c as a convex combination of the chosen instances a and b .

To reduce the majority class, the algorithm uses the Edited Nearest Neighbour approach (ENN) to remove majority class instances that differ from at least two of their neighbours (Wilson, 1972). This is accomplished by applying the following steps:

- For each instance i in the majority class, choose the k nearest neighbours; in this case $k = 3$.
- Find the classes associated with each of the k neighbours.
- Remove the instance i from the data when its class is not in accordance with the dominant class of its k neighbours, determined in the previous step.

This combined approach balances the data and reduces the effects of noise that could arise from only oversampling the minority class. Using this method, the data was manipulated to create 44.4 per cent minority class and 55.6 per cent majority class, a ratio that was chosen as it resulted in the best performance in this case (Figure 4).

3.3 MODEL ENSEMBLE

Having balanced the dataset by adding the synthetic data, training models are expected to perform much better at classifying the minority class. In this case, we resort to *ensemble learning* in order to better predict the patch classes from the over-sampled dataset (Liu et al., 2013: 62). Several approaches were undertaken in order to reach an optimal solution. In this case, SMOTEENN followed by AdaBoost (Freund et al., 1997), SMOTEBoost (Chawla et al., 2003), as well regular Random Forests (Breiman, 2001) to classify the data after SMOTEENN. In the end, a voting approach that combines the results of all the above approach proved to be the most efficient (Figure 5).

Figure 4 Pie charts showing the SMOTEENN effect on the training and validation datasets.

The three models mentioned above (AdaBoost, SMOTEBoost, and Random Forests) were trained on all possible PCA pair combinations, which ensured that the trained models remained very simple and resulted in a significant decrease in the number of false positives. With a total of 8 PCA component representing each data instance, this resulted in 28 PCA pair combinations trained using the three chosen models, yielding a total of 84 predictions per instance.

This approach was favoured over training a single model due to the fact that each resulting model in this approach could be very simple, and very general, thus ensuring that the decision boundaries can be applied efficiently to the test set. The hyperparameters of each model were tuned using RandomSearchCV, with the used parameters discussed in the next paragraphs.

AdaBoost uses sequential base learners, in this case 250 decisions trees with a maximum tree depth of 15, and a minimum sample of leaves of 25 (which is the minimum number of samples that is allowed on a leaf node of the tree) and a learning rate of 0.95, in order to create a single strong learner. Each base learner is trained separately and used to make predictions on the dataset. Based on the results of the preceding learner, the subsequent model learns from the dataset after changing the weights of the misclassified instances in the preceding model. In this fashion, each learner slightly improves on the results of its preceding one, in order to finally reach a strong model (Géron, 2017: 200).

SMOTEBoost combines the two steps, SMOTE and AdaBoost, into a single algorithm. In this case, instead of simply applying SMOTE to the entire dataset prior to the process, SMOTE is applied to the minority class at every single 'boost' of the AdaBoost algorithm (Freund et al., 1997: 120; Liu et al., 2013: 66). In this case, a minority to majority ratio of 0.9 was chosen at every step of the algorithm. The total number of estimators here was again chosen to be 50 and a learning rate of 0.9 was set.

Random Forests were also used due to their ability

to generalize well and avoid overfitting. A total of 300 estimators were used in this case with a maximum tree depth of 10, and a minimum sample of leaves totalling 15.

Passing the training features to this system, we get a total of 84 prediction per instance (resulting from 28 PCA pair combinations trained with three different models), which consists of both good and bad predictions. We aggregate the predictions by simply adding the prediction vectors resulting from the 84 predictions, which means that instances that are often predicted as looting and destruction tend to accumulate a higher aggregate prediction score, while those instances that are rarely classified as looting and destruction tend to accumulate a very low score (Figure 5). This is apparent in the aggregate prediction plot shown in Figure 6, where blue instances are those of positive instances of coherence loss patches (due to archaeological destruction or looting), while the red instances are those of negative instances of coherence loss patches. It is obvious that there is a clear split between those instances that are rarely classified as positive instances, and thus have an aggregate score close to zero, and those that have a higher aggregate score (Figure 6). It is important to note that in this case, the validation data has been passed through the SMOTEENN algorithm, and the data is balanced by a ratio of 0.8 minority to majority classes. However, the plot remains noisy, and in order to get the final results in this case, a threshold was set in which instances with an aggregate vote value above 70 per cent of the maximal aggregate value are chosen as confirmed positive instances, the rest are relegated to negative instances. This value was chosen based on a trial and error approach, with such value yielding the best compromise between false positives and false negatives (see Figure 5).

3.4 EVALUATION METRICS

In order to evaluate the trained models and their effectiveness on the test set, it is important to define

Figure 5 Workflow showing the aggregate voting and thresholding process.

Figure 6 Validation data instance prediction plot.

what are the metrics that can serve our objective. In this particular approach, the aim is to detect the positive minority class of looting and destruction instances, often represented by 0.1 to 2 per cent of the data (discussed above in Section 3.2), while at the same time reducing the amount of False Positives (FP) which would render our results unreliable. A reliable system in this case must ensure that first, a very high percentage of the positive, looting and destruction, instances are correctly classified. Misclassifying these instances is detrimental to the models' functional purpose, since, if the model does not reliably detect positive instances, its purpose of detecting looting and destruction comes into question. Secondly, a reliable system must ensure that none, or a minimal number of negative non-looting and destruction instances are mistaken for positive ones. A large number of negative instances being classified as positive (i.e. non-looting patches being classified as looting and destruction) will decrease the confidence in the system.

Two metrics help us quantify the above observations. The first being the *sensitivity* metric (often referred to as recall) which returns the ratio of correctly classified positive instances to the total number of positive instances in the dataset. The second is *specificity* which returns the ratio of the correctly classified negative instances to the total number of negative instances in the dataset (Japkowicz, 2013: 192). The combination of *sensitivity* and *specificity* is often used in the field of medical imaging, where data is often imbalanced and the risk of misclassifying a positive instance as negative (disease as non-disease) could be life threatening, while the risk of classifying many negative instances as positive (non-disease as

disease) could lead to unnecessary surgery (Szyk et al., 2019). These metrics are calculated using Equations 7 and 8, where TP refers to True Positive instances, FP to False Positives, TN to True Negatives and FN to False Negatives.

$$Sensitivity = \frac{TP}{TP + FN} \quad (7)$$

$$Specificity = \frac{TN}{FP + TN} \quad (8)$$

Having two metrics makes direct model comparisons challenging. In order to alleviate this difficulty, especially with the presence of numerous models as is the case here, we rely on the G-mean. This metric was first introduced by Kubat et al. (1998) with the aim of better evaluating model predictions on imbalanced datasets, and combines both *sensitivity* and *specificity*. It can be easily calculated using Equation 9. In addition, given the fact that our data is imbalanced with a minority positive class, it is naturally expected that a missed positive instance will have a larger effect on lowering the G-mean score, which in turn will guide us to train models that prioritize correctly labelling positive looting and destruction instances.

$$G_{mean} = \sqrt{sensitivity \times specificity} \quad (9)$$

We use the above metrics to fine tune our model ensemble, aiming to increase the results without overfitting the models. The resulting models, described in Section 3.3, show promising results individually with relatively high G-mean scores as seen in [Table 1](#). However, the model ensemble predictions result in the highest G-mean and

confirms that the *ensemble* is a better classifier than any of the individual learners within it.

The prediction results on the validation set are shown in [Figure 6](#), where the maximum aggregate score achieved is 63, which leads to a threshold value of 44 which represents 70 per cent of the maximum (see Section 3.3 and [Figure 5](#)). By applying this approach, we end up with a total of 77 TP instances and 7 FN instances (total positive instances in the validation set is 84). We also were able to classify all TN instances in their correct class. From these values we obtain a G-mean score of 0.957, indicating that our model is able to classify both Positive and Negative instances correctly to a good degree. We rely on the ensemble to classify the coherence patches from our test sets, which are presented in the following section.

MODEL	G-MEAN SCORE
AdaBoost	0.809
SMOTEBoost	0.938
Random Forest	0.945
Model Ensemble	0.957

Table 1 G-mean scores of the different models used.

4 RESULTS

In this section, the above approach is applied to two different test cases of known destruction activities in the Near East. The first being the known case of the destruction of Palmyra, Syria which yielded good results and limited amount of FP. The second case is that of Ain Dara, Syria, the temple of which was targeted by an air strike leading to its destruction in early 2018, and proved more challenging when it comes to reducing the number of FPs. The validation data in these cases comes from available of VHR optical images released on Google Earth some time after each incident as seen in [Figures 8 and 10](#).

The first test case comes from Palmyra, Syria, where ISIS destroyed the Temple of Bel along with several other tower tombs (above ground tower like structures) in 2015. The high degree of publicity that accompanied the destruction, as well as the relevance of the destroyed buildings, are the reasons that this site was chosen to be a representative case in this article. [Figure 7](#) clearly shows the coherence loss associated with the destruction event (within green boxes) along with countless other coherence losses with various decorrelation origins that are not of interest in this case. The chosen model ensemble can clearly eliminate the majority of the unwanted coherence

Figure 7 Top: VHR optical image of Palmyra showing polygons classified as coherence loss induced by looting/destruction, Below: Coherence Map – SAR Pair: 22-08-2015/03-09-2015.

loss patches, resulting in a single FP, and highlight the ones that are of interest to a good degree with a single FN. **Figure 7** shows the patches that were classified as having destruction at their origin (red polygons), while the rest was discarded. The single FP instance is located in a vegetated area and shown within the orange box, it is safe to assume that vegetation was at the origin of the decorrelation in this case. The single FN instance is shown in a red box. The G-mean score of the model ensemble in this case is 0.863, and a summary of the results is shown in **Table 2**.

The second example presented in this article is that of Ain Dara, Syria. The site was first targeted by an air strike damaging its main temple at the beginning of 2018, and

was almost completely bulldozed at the end of the same year. In this case, we are concerned with the first case of destruction. While the damage in this case is extensive, the main difference between this case study and that of Palmyra is the surrounding terrain. Palmyra features an almost exclusively un-vegetated area, which meant that coherence loss due to vegetation was minimal and highly contained. On the other hand, the case of Ain Dara presented a different picture. The site, especially in January, is located within a well vegetated area rich in agricultural fields (see **Figure 9**). In addition, the area of the archaeological site itself was not cleared of vegetation and appears to be overgrown, as was apparent from the NDVI image, as well as from the validation VHR

Figure 8 VHR validation images from Google Earth showing the effects of the destruction on the Temple of Bel in Palmyra.

Figure 9 Top: VHR optical image of Ain Dara showing the polygons classified as coherence loss induced by looting/destruction, Below: Coherence Map – SAR Pair: 06-01-2018/30-01-2018.

imagery shown in [Figure 10](#). This resulted in a very noisy coherence estimation map which features a significantly higher number of coherence patches ([Table 2](#)), and thus made it difficult to isolate the only positive instance of coherence loss patch resulting from the destruction (the noise persists in coherence maps with 12 and 24 day temporal resolution). However, the algorithm succeeded in clearing the majority of the patches (FP polygons are indicated by a red box) and correctly indicating the coherence loss patch resulting from the air strike (green box) as seen in [Figure 9](#) and [Table 2](#).

A deeper look at how the models performed revealed that the aggregate predictions ([Figure 11](#)) resulted in a maximum score of 39, which yielded a cut off score of 27 according to the 70 per cent threshold set in Section 3.3. With this threshold, we obtain a total of 1 TP, in this case the only correct coherence patch loss due to the air strike.

We also obtain 15 other FP instances, 496 TN instances and 0 FN instances ([Table 2](#)). By calculating G-mean we again get a very good value of 0.985 which indicates that our model has done a decent job over this test set.

5 DISCUSSION

The test cases above, obtained exclusively from open source data, are promising. Three out of four target destruction patches in Palmyra were correctly labelled with minimal False Positives, and the single destruction patch in Ain Dara was correctly labelled, albeit with numerous False Positives present in the resulting set. While this approach is not meant to replace the use of optical VHR imagery nor the reliance on human operators to confirm the presence of destruction; it allows the

Figure 10 VHR validation images from Google Earth showing the effects of the Air Strike.

Figure 11 Instance prediction plot for Ain Data test set data.

SITE	TOTAL PATCHES	TP	TN	FP	FN	G-MEAN SCORE
Palmyra	150	3	145	1	1	0.863
Ain Dara	512	1	496	15	0	0.984

Table 2 Classification results from the test sites.

interested parties to monitor a very large region using minimal resources. In this case, VHR data could be used only as validation data when positive patches are reported in archaeological sites of interest.

One of the main limitations to this approach is its failure over vegetated areas. This of course stems from the fact that such areas will generate coherence loss which will be indistinguishable from any human induced change in their close proximity. This fact limits the target areas to those with a majority soil land cover zones, and excludes a large chunk of landscape. In addition, the spatial resolution, which was resampled to 15×15 m ground resolution for the coherence maps, means that the small scale changes will be almost impossible to distinguish. Such small scale changes could come in the form of exploratory looting holes which looters often dig in order to check the prospect of the underlying layers (Tapete and Cigna, 2019: 13).

There is much to be improved however, starting with the collection of a much larger sample of looted sites which can be included in the training data. The dataset as it stands is relatively small, and thus the trained models might not be able to generalize well for positive patches outside the region, as well as positive patches that do not display the same statistical distribution as those from this small training set. In order to expand this database, data concerning destruction and looting in archaeological sites should be recorded along with precise dates of ground activities, which enables the acquisition of temporally close SAR image pair, and reduces the risk of temporal decorrelation unrelated to the events of interest.

While the use of open source data dictates the resolutions of what can be detected, and limits this approach in some ways; it is the essence of this article to present a purely open source methodology, where all the used data is openly available through ESA Copernicus Open Access Hub, as well as through the USGS Earth Explorer.

6 CONCLUSION

The use of Remote Sensing to monitor archaeological sites is on the rise with the ever increasing number of earth observation satellites in orbit. More recently, researchers have been exploring numerous ways to automatically detect these changes using computer algorithms, thus opening the door to monitoring large

swaths of land instead of focusing on single site cases. However, the majority of such approaches rely on VHR images, whether Optical (e.g. Masini et al., 2020) or SAR (e.g. Tapete, Cigna & Donoghue, 2016).

While this article aims to build the frameworks to an automatic detection of archaeological looting and destruction, the focus here is to rely solely on Open Source data by exploiting the synergy between Sentinel-1 and Sentinel-2. Sentinel-1 data yields urban maps on the one hand, and the most important data for this research, coherence maps on the other. Sentinel-2 provides the complementary data, NDVI and BI, which allows us to better understand the ground cover at the base of the coherence loss. Using the provided data from the above-mentioned satellites, we construct a feature vector that describes each coherence loss patch within the coherence map. We then resort to Machine Learning to train a model *ensemble* to classify the patches and automatically identify the ones associated with looting and destruction. The reliance on an *ensemble* stems from the fact that our data is relatively small (see [Figure 3](#)) and highly imbalanced, thus an *ensemble* will help avoid overfitting and ensures that our models are able to generalize well despite the previously mentioned issues. The *ensemble* includes a Random Forest, AdaBoost, and SMOTEBoost, who already work well on a standalone basis as attested by their G-mean scores, but result in a better performance when their aggregated results are considered (see [Table 1](#)).

The model *ensemble* is fine-tuned and tested, with satisfying results, on two sites: Palmyra and Ain Dara, both of which are located in Syria. In Palmyra, the model succeeded in highlighting the coherence loss resulting from the destruction of the Temple of Bel, as well as the coherence loss from the destruction of numerous tower tombs while only missing a single positive instance, and misclassifying another negative instance as positive. In Ain Dara, the model was able to correctly detect the single positive instance; but misclassified 15 negative instances in the process. The results presented in Section 4 are encouraging, however, the presented approach is still a work in progress. A larger dataset must be collected, which in turn will help us create a more generalizing model, as well as predict a wider array of coherence loss patterns than the ones present in the training set. On the other hand, no matter how large the dataset grows, this approach is still limited by the type of satellite data we are using; in this case, we are limited by the spatial resolution which means that small disturbances to archaeological sites are unlikely to be detected. This is counterbalanced however by the fact that the data is freely available, which will undoubtedly lead to a higher uptake of such model within research communities and local authorities.

Finally, no matter how efficient a detection system is, merely detecting archaeological looting and destruction

will not deter individuals from looting archaeological artefacts, or destroying archaeological structures. Instead, one should think of ensuring that the approach is implemented within the relevant authorities who are in charge of protecting the cultural heritage sites, who can then act upon the presence of a looting and destruction instance. In the case of the proposed approach, the open source nature of data ensures that the operation costs remain minimal when it comes to data acquisitions, as well as when it comes to the human operators monitoring the system. In this regards, I believe that it is only through the use of such open source data that we can encourage a larger engagement within the cultural heritage community in the MENA region and elsewhere, as well as increase their reliance on satellite data in order to effectively protect cultural heritage sites from further damage.

COMPETING INTERESTS

The author has no competing interests to declare.

AUTHOR AFFILIATION

Hassan El-Hajj, PhD Candidate orcid.org/0000-0001-6931-7709
Classical Archaeology Department, Philipps-Universität
Marburg, Biegenstrasse 10, 35037 Marburg, Germany

REFERENCES

- Adragna, F** and **Nicolas, J-M**. 2008. Interferometry. In: Maitre, H (ed.), *Processing of Synthetic Aperture Radar Images*. Hoboken, NJ: John Wiley and Sons Inc. pp. 279–300. DOI: <https://doi.org/10.1002/9780470611111.ch12>
- Agapiou, A**. 2020. Detecting Looting Activity Through Earth Observation Multi-temporal Analysis over the Archaeological Site of Apamea (Syria) during 2011–2012. *Journal of Computer Applications in Archaeology*, 3(1): 219–237. DOI: <https://doi.org/10.5334/jcaa.56>
- Agapiou, A, Lysandrou, V** and **Hadjimitsis, D**. 2017. Optical Remote Sensing Potentials for Looting Detection. *Geosciences*, 7. DOI: <https://doi.org/10.3390/geosciences7040098>
- Bewley, R, Wislon, A, Mattingly, D, Kennedy, D, Banks, R, Bishop, M, Bradbury, J, Cunliffe, E, Fradley, M, Jennings, R, Mason, R, Rayne, L, Sterry, M, Sheldrick, N** and **Zerbini, A**. 2016. Endangered Archaeology in the Middle East and North Africa: Introducing the EAMENA Project. In: Campana, S, Scopigou, R, Carpentiero, G and Cirillo, M (eds.), CAA2015. *Keep the Revolution Going: Proceedings of the 43rd Annual Conference on Computer Applications and Quantitative Methods in Archaeology*. Oxford: Archaeopress, pp. 919–932.
- Breiman, L**. 2001. Random Forests. *Machine learning*, 45: 5–32. DOI: <https://doi.org/10.1023/A:1010933404324>
- Campbell, J** and **Wynne, R**. 2011. *Introduction to Remote Sensing*. 5th ed. New York: The Guilford Press.
- Casana, J**. 2015. Satellite Imagery-Based Analysis of Archaeological Looting in Syria. *Near Eastern Archaeology*, 78(3): 142–152. DOI: <https://doi.org/10.5615/neareastarch.78.3.0142>
- Casana, J** and **Laugier, EJ**. 2017. Satellite Imagery-Based Monitoring of Archaeological Site Damage in the Syrian Civil War. *PLOS One*, 12(11). DOI: <https://doi.org/10.1371/journal.pone.0188589>
- Casana, J** and **Panahipour, M**. 2014. Satellite-Based Monitoring of Looting and Damage to Archaeological Sites in Syria. *Journal of Eastern Mediterranean Archaeology and Heritage Studies*, 2(2): 129–152. DOI: <https://doi.org/10.5325/jeasmedarcherstu.2.2.0128>
- Cerra, D** and **Plank, S**. 2020. Towards Early Warnings for Damages to Cultural Heritage Sites: the Case of Palmyra. In: Hadjimitsis, D, Themistocleous, K, Cuca, B, Agapiou, A, Lysan- drou, V, Lasaponara, R, Masini, N and Schreier, G (eds.), *Remote Sensing for Archaeology and Cultural Landscapes: Best Practices and Perspectives across Europe and the Middle East*. Switzerland: Springer, pp. 221–239. DOI: https://doi.org/10.1007/978-3-030-10979-0_13
- Cerra, D, Plank, S, Lysandrou, V** and **Tan, J**. 2016. Cultural Heritage Sites in Danger – Towards Automatic Damage Detection from Space. *Remote Sensing*, 8(781). DOI: <https://doi.org/10.3390/rs8090781>
- Chawla, N, Lazarevic, A, Hall, LO** and **Bowyer, KW**. 2003. SMOTEBoost: Improving Prediction of the Minority Class in Boosting. In: Lavrač, N, Gamberger, D, Todorovski, L and Blockeel, H (eds.), *Knowledge Discovery in Databases: PKDD2003*. Berlin, Heidelberg: Springer, pp. 107–119. DOI: https://doi.org/10.1007/978-3-540-39804-2_12
- Cigna, F** and **Tapete, D**. 2018. Tracking Human-Induced Landscape Disturbance at the NASCA Lines UNESCO World Heritage site in Peru with COSMO-SkyMed InSAR. *Remote Sensing*, 10(572). DOI: <https://doi.org/10.3390/rs10040572>
- Copernicus Open Access Hub**. 2019. Available at <https://scihub.copernicus.eu/twiki/do/view/SciHubWebPortal/WebHome#dias-box> [Last Accessed 04 May 2019].
- Cunliffe, E**. 2014. Remote Assessments of Site Damage: a New Ontology. *International Journal of Heritage in the Digital Era*, 3(3): 453–473. DOI: <https://doi.org/10.1260/2047-4970.3.3.453>
- De Zan, F** and **Monti Guarnieri, A**. 2006. TOPSAR: Terrain Observation by Progressive Scans. *IEEE Transactions in Geosciences and Remote Sensing*, 44(9): 2352–2360. DOI: <https://doi.org/10.1109/TGRS.2006.873853>
- EAMENA**. 2021. EAMENA: Endangered Archaeology in the Middle East and North Africa. Available at <https://database.eamena.org/index.html> [Last Accessed 29 January 2021].

- Escadafal, R.** 1989. Remote Sensing of Arid Soil Surface Color with Landsat Thematic Mapper. *Advances in Space Research*, 9(1): 159–163. DOI: [https://doi.org/10.1016/0273-1177\(89\)90481-X](https://doi.org/10.1016/0273-1177(89)90481-X)
- Esch, T, Marconcini, M, Felbier, A, Roth, A, Heldens, W, Huber, M, Schwinger, M, Taubenbock, H, Muller, A and Dech, S.** 2013. Urban Footprint Processor—Fully Automated Processing Chain Generating Settlement Masks from Global Data of the TanDEM-X Mission. *IEEE Geoscience and Remote Sensing Letters*, 10(6): 1617–1621. DOI: <https://doi.org/10.1109/LGRS.2013.2272953>
- European Space Agency.** 2020a. Sentinel-1 orbit. Available at <https://sentinel.esa.int/web/sentinel/missions/sentinel-1/satellite-description/orbit> [Last Accessed 15 September 2020].
- European Space Agency.** 2020b. Sentinel-1 SAR user guide introduction. Available at <https://sentinel.esa.int/web/sentinel/user-guides/sentinel-1-sar> [Last Accessed 05 March 2020].
- Fisk, R.** 1991. The Biggest Supermarket in Lebanon. *Berytus*, 39: 243–252.
- Freund, Y and Schapire, RE.** 1997. A Decision-Theoretic Generalization of Online Learning and an Application to Boosting. *Journal of Computer and System Sciences*, 55(1): 119–139. DOI: <https://doi.org/10.1006/jcss.1997.1504>
- Géron, A.** 2017. *Hands-on Machine Learning with Scikit-learn and Tensorflow: Concepts, Tools, and Techniques to Build Intelligent Systems*. Sebastopol, CA: O'Reilly Media.
- Gustavo, B, Ronaldo, P and Carolina, M.** 2004. A Study of the Behavior of Several Methods for Balancing Machine Learning Training Data. *SIGKDD explor. newsl.*, 6(1): 20–29. DOI: <https://doi.org/10.1145/1007730.1007735>
- Hoens, R and Chawla, N.** 2013. Imbalanced Datasets: From Samplings to Classifiers. In: He, H and Ma, Y (eds.), *Imbalanced Learning: Foundations, Algorithms, and Applications*. Hoboken, New Jersey: John Wiley and Sons, pp. 43–60. DOI: <https://doi.org/10.1002/9781118646106.ch3>
- Humanitarian Data Exchange.** 2020. *The humanitarian data exchange*. Available at <https://data.humdata.org/> [Last Accessed on 05 September 2020].
- Japkowicz, N.** 2013. Assessment Metrics for Imbalanced Learning. In: He, H and Ma, Y (eds.), *Imbalanced Learning: Foundations, Algorithms, and Applications*. Hoboken, New Jersey: John Wiley and Sons, pp. 187–206. DOI: <https://doi.org/10.1002/9781118646106.ch8>
- Kubat, M, Holte, R and Matwin, S.** 1998. Machine Learning for the Detection of Oil Spills in Satellite Radar Images. *Machine learning*, 30: 195–215. DOI: <https://doi.org/10.1023/A:100745223027>
- Liu, X and Zhou, Z.** 2013. Ensemble Methods for Class Imbalance Learning. In: He, H and Ma, Y (eds.), *Imbalanced Learning: Foundations, Algorithms, and Applications*. Hoboken, New Jersey: John Wiley and Sons, pp. 61–82. DOI: <https://doi.org/10.1002/9781118646106.ch4>
- Manjunath, BS and Ma, WY.** 1996. Texture Features for Browsing and Retrieval of Image Data. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(8): 837–842. DOI: <https://doi.org/10.1109/34.531803>
- Masini, N and Lasaponara, R.** 2020. Recent and Past Archaeological Looting by Satellite Remote Sensing: Approach and Application in Syria. In: Hadjimitsis, D, Themistocleous, K, Cuca, B, Agapiou, A, Lysandrou, V, Lasaponara, R, Masini, N and Schreier, G (eds.), *Remote Sensing for Archaeology and Cultural Landscapes: Best Practices and Perspectives Across Europe and the Middle East*. Switzerland: Springer, pp. 123–137. DOI: https://doi.org/10.1007/978-3-030-10979-0_8
- Parcak, S.** 2015. Archaeological Looting in Egypt: A Geospatial View (Case Studies from Saqqara, Lisht, and El Hibeh). *Near Eastern Archaeology*, 78(3): 196–203. DOI: <https://doi.org/10.5615/neareastarch.78.3.0196>
- Rayne, L, Bradbury, J, Mattingly, D, Philip, G, Bewley, R, and Wilson, A.** 2017. From Above and on the Ground: Geospatial Methods for Recording Endangered Archaeology in the Middle East and North Africa. *Geosciences*, 7(100). DOI: <https://doi.org/10.3390/geosciences7040100>
- Rayne, L, Gatto, M, Abdulaati, L, Al-Haddad, M, Sterry, M, Sheldrick, N and Mattingly, D.** 2020. Detecting Change at Archaeological Sites in North Africa Using Open-Source Satellite Imagery. *Remote Sensing*, 12(3694). DOI: <https://doi.org/10.3390/rs12223694>
- Savage, S, Johnson, A and Levy, T.** 2017. Terrawatchers, Crowdsourcing, and At-Risk World Heritage in the Middle East. In: Vincent, M, Lopez-Menchero, VB, Ionnides, M and Levy, T (eds.), *Heritage and Archaeology in the Digital Age: Quantitative Methods in the Humanities and Social Sciences*. Springer. DOI: https://doi.org/10.1007/978-3-319-65370-9_4
- Silverman, H and Fairchild Ruggles, D.** 2007. Cultural Heritage and Human Rights. In: Silverman, H and Fairchild Ruggles, D (eds.), *Cultural Heritage and Human Rights*. Springer. pp. 3–22. DOI: https://doi.org/10.1007/978-0-387-71313-7_1
- Stewart, C.** 2017. Detection of Archaeological Residues in Vegetated Areas Using Satellite Synthetic Aperture Radar. *Remote Sensing*, 9(118). DOI: <https://doi.org/10.3390/rs9020118>
- Stewart, C, Montanaro, R, Sala, M and Riccardi, P.** 2016. Feature Extraction in the North Sinai Desert Using Space-born Synthetic Aperture Radar: Potential Archaeological Applications. *Remote Sensing*, 8(825). DOI: <https://doi.org/10.3390/rs8100825>
- Stone, E.** 2004. Archaeological Site Looting: The Destruction of Cultural Heritage in Iraq. In: Emberling, G and Hanson, K (eds.), *Catastrophe! the Looting and Destruction of Iraq's Past*. Chicago: The Oriental Institute Museum of the University of Chicago, pp. 65–80.
- Szyc, L, Hillen, U, Scharlach, C, Kauer, F and Garbe, C.** 2019. Diagnostic Performance of a Support Vector Machine for Dermatofluoroscopic Melanoma Recognition: The Results of the Retrospective Clinical Study on 214 Pigmented Skin

- Lesions. *Diagnostics*, 9(103). DOI: <https://doi.org/10.3390/diagnostics9030103>
- Tapete, D and Cigna, F.** 2019. Detection of Archaeological Looting from Space: Methods, Achievements and Challenges. *Remote Sensing*, 11(2389). DOI: <https://doi.org/10.3390/rs11202389>
- Tapete, D, Cigna, F and Donoghue, DN.** 2016. 'Looting Marks' in Spaceborn SAR Imagery: Measuring Rates of Archaeological Looting in Apamea (Syria) with TerraSAR-X Staring Spotlight. *Remote Sensing of Environment*, 17: 42–58. DOI: <https://doi.org/10.1016/j.rse.2016.02.055>
- United States Geological Survey.** 2020. Available at <https://earthexplorer.usgs.gov/> [Last Accessed 17 November 2020].
- Weiss, G.** 2013. Foundation of Imbalanced Learning. In: He, H and Ma, Y (eds.), *Imbalanced Learning: Foundations, Algorithms, and Applications*. Hoboken, New Jersey: John Wiley and Sons, pp. 13–41. DOI: <https://doi.org/10.1002/9781118646106.ch2>
- Weiss, G and Provost, F.** 2003. Learning when Training Data are Costly: The Effect of Class Distribution on Tree Induction. *Journal of Artificial Intelligence Research*, 19: 315–354. DOI: <https://doi.org/10.1613/jair.1199>
- Wilson, D.** 1972. Asymptotic Properties of Nearest Neighbor Rules Using Edited Data. *IEEE trans. syst. man cybern.*, 2: 408–421. DOI: <https://doi.org/10.1109/TSMC.1972.4309137>

TO CITE THIS ARTICLE:

El-Hajj, H. 2021. Interferometric SAR and Machine Learning: Using Open Source Data to Detect Archaeological Looting and Destruction. *Journal of Computer Applications in Archaeology*, 4(1), pp. 47–62. DOI: <https://doi.org/10.5334/jcaa.70>

Submitted: 06 January 2021 Accepted: 10 February 2021 Published: 16 March 2021

COPYRIGHT:

© 2021 The Author(s). This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC-BY 4.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited. See <http://creativecommons.org/licenses/by/4.0/>.

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