Machine Learning for UAV and Ground-Captured Imagery: Toward Standard Practices

Authors Anonymous

ABSTRACT

Our collaborative work began in 2019 with the intent to overcome obstacles that had arisen from the inability to access curated artifact collections from remote locations. It was our specific aim to not only create digital twins of excavated objects that could not be exported out of their country of origin, but also to emphasize the contextual associations of objects residing in hidden museum collections using a range of digital techniques. As part of a growing field project in 2022, machine learning (ML) with YOLOv5, a family of compound-scaled object detection models trained on the COCO dataset was used to classify visual data and advance our understanding of in situ archaeological phenomena prior to destructive fieldwork. While not the sole contribution, the use of object-based machine learning improved quality and range of information obtained in non-destructive site surveys and improved data sharing capacity. Despite challenges encountered while training the algorithm and classifying objects, combining ML with drone data collection will continue as part of our long-term spatial data recording procedure. Despite both success and failures reported here, this work contributes to greater standardization of ML techniques in archaeological practice.

Keywords: machine learning, archaeology, UAV & ground-based imagery classification, YOLOv5, non-destructive site survey

Introduction

As part of our growing field project in 2022, we applied machine learning with YOLOv5 (DeepLearning 2020; Jocher2020), a family of compound-scaled object detection models trained on the COCO dataset (see PyTorch 2023) to classify visual data and advance our understanding of *in situ* phenomena on the ground prior to conducting destructive archaeological fieldwork. Beyond what have become well-established digital field and laboratory methods in archaeology (e.g. high-resolution data capture, instant digital recording and transfer with tablets and other small devices) over the past few decades, the application of Machine Learning (ML) for Unmanned Aerial Vehicle (UAV) and ground-captured imagery was expected to greatly expand our ability to select areas for intensified investigation or excavation in expedient, non-destructive ways. Expanding on what we had learned while analyzing physical collections in previous field and laboratory analyses (see Sharp 2019), the development of a neural network that was aimed directly at survey presented several unforeseen challenges.

Although some initial challenges while learning to train the algorithm and classify objects left much work to be done, ML on archaeological survey data captured with UAV will continue to be more tightly integrated into our standard data collection procedure in the field. With intensifying efforts to integrate ML into various disciplines and industry, however, we find ourselves in relatively uncharted territory with few applied archaeological models to follow (but see Bonhage et al. 2021; Caspari and Crespo 2019; Orengo et al. 2021; Sakai et al. 2023).

Upon our return to the field in 2022 after a three-year hiatus, we discovered that some areas of our study area had become dangerous and inaccessible, necessitating the incorporation of remote sensing technology in our initial surveys. The application of ML object/structure detection was expected to facilitate not only precision recording of all areas of interest but also to help determine various environmental risk factors like El Niño/La Niña climate fluctuations. Additionally, we were hopeful that our non-destructive site monitoring data would assist with heritage preservation issues that would arise in the future (e.g., Caster et al. 2022). This paper discusses our successes and failures and future directions of our work in northern Peru, with the primary aim to contribute to the dialogue at this transformational moment in digital archaeology. Looking toward the future and aiming to foreshadow challenges that our project and others will face, we urge for intensified efforts toward standardization of practice and integration of streamlined methods in this fast-expanding branch of computational archaeology.

Implementing YOLO for UAV and Ground-based Image Classification and Analysis

In 2022, we applied UAV-remote sensing and deep learning in our survey season to increase the amount of terrain our small team could study over the ten-day period. In anticipation of training a dataset for automated object detection and classification, we used digital cameras and "LiDAR" enabled smartphones with GPS to augment the range of data that our survey teams encountered on the ground (and which were expected to be encountered in the future). As we have witnessed in other cases, we anticipated the quality and coverage of data collected to improve dramatically using these cost-effective techniques and highly portable tools. Additionally, we aimed to test the feasibility of applying UAV in our proposed long-term research; one of several non-destructive survey techniques that we are hopeful will become standardized for surveys in challenging topographical, environmental, or biodiverse terrain.

In this regard, the range of circumstances we confronted in regional surveys are both long-standing and recent. In northern coastal Peru, concerns that have become more prominent are the result of increased frequency of extreme climate events, as well as reduced site monitoring and protection capacity over the past five years. Given the challenging environment and sometimes perilous terrain, the ability of UAV to ensure the safety of field crews makes it an ideal choice.

Nepal and Eslamiat (2022) provided the context and justification for using UAV drones as one of this season's survey method. UAVs have already been employed in numerous domains, including traffic

Auteur
Verwijderd: machine learning (

Auteur
Verwijderd:)

Auteur
Verwijderd: unmanned aerial vehicle (

Auteur
Verwijderd:)

Auteur
Verwijderd: machine learning (

Auteur
Verwijderd:)

Auteur
Opmerking [1]: Why are these specific references chosen? There are actually many articles on the use of ML and DL on remotely sensed data for archaeology. See for instance, Lambers et al. 2019 for an (already slightly outdated) overview.

Auteur
Opmerking [2]: This parts reads as a conclusion not as an introduction, please consider rephrasing / restructuring.

Auteur
Opmerking [3]: Machine learning and deep learning are used inconsistently, I would stick to one of the terms.

monitoring, surveillance, inspection, surveys, etc. UAV use in contemporary fieldwork is transforming the way the world is viewed as real-time deep learning algorithms increase both speed and accuracy of information capture. With the deployment of deep neural networks in recent years, UAVs have taken over aerial sensing research in the urban, environmental, and agricultural sectors (see Nepal and Eslamiat 2022).

Deep Learning is a subset of Machine Learning that uses many layers to extract data features. Its application has been used in a wide range of industries, from creating autonomous UAV trajectories to now making significant strides in accurate object classification. Their methods provide real-time object detection, which makes them appropriate for autonomous robotics and UAV applications in archaeology as well. The development of computer vision and deep learning methods has also been aided by the usage of graphics processing units (GPUs) for deep learning algorithms (Nepal and Eslamiat 2022). This made it possible for us to incorporate object detection methods appropriate for our real-time application in archaeology.

In this study, we tested the utility of both UAVs (both DJI Mavik and DJI Phantom 4 Pro) and ground-based image capture (iPhone Pro MAX and various Android smartphones) in our archaeological fieldwork. When combined with increasingly established object detection capabilities of deep learning programing, we found these methods to be both economical and accessible, a clear advantage recognized in other academic and industry applications (see Nepal and Eslamiat 2022). Although several deep learning and CNN (Convolutional Neural Network) algorithms have been presented since 2012, YOLO is a single stage deep learning system that detects objects using a convolution neural network. Several deep learning algorithms now available like recurrent neural networks (RNN, see Mittal 2020), cannot detect an object in a single run, but YOLO enables the detection across a neural network in forward propagation, making it appropriate for real-time application. For these reasons and because of its accelerated rate of detection, YOLO was the deep learning program used for our study (Nepal and Eslamiat 2022).

Analytical Challenges

The primary dataset for training the YOLO algorithm consisted of imagery captured during 2022 fieldwork (Figure 1 and 2). Images ranged from pictures of individual pottery fragments, pottery in a local workshop visited by the survey team in 2019, vessels fragments photographed in the lab during 2015 field season (Sharp 2019), and images downloaded from the internet that were representative of the range of vessel forms and decorative styles on pottery observed in the field.

Auteur
Opmerking [4]: Do you mean object detection or image classification?

Auteur
Verwijderd: in

Auteur
Opmerking [5]: I think a CNN is a Deep Learning algorithm…Deep Learning is the method and CNNs are one of the algorithms.

Auteur
Opmerking [6]: Reference?

Auteur
Opmerking [7]: I think there might be confusion here with R-CNN. The difference lies in that YOLO is a single stage detector, while R-CNN is a two stage detector. Basically, the difference is whether localization and classification is separated or not.

Auteur
Opmerking [8]: For information in the difference in speed you can have a look at Olivier & Verschoof-van der Vaart 2021, 10.5334/jcaa.78.

Auteur
Opmerking [9]: I would state which YOLO version is used as there is quite a difference between them.

Figure 1: Example of field images of pottery and other materials in physical setting that comprise original training data.

Figure 2: Additional pottery fragments in association with stones and twigs.

Downloading and installing programs to begin training revealed many compatibility issues amongst pre-installed programs on our Windows PCs — as finding out which version was compatible with another program and/or our own computers proved difficult at times — which necessitated a bit of experimentation, downloading and redownloading and frequent program updates. Equally so, offline training of the algorithm with installed programs was bulky and required several steps. Considering this limitation of offline training on a Windows PC, we decided to train YOLOv5 for pottery sherd detection

Auteur
Verwijderd: -

Auteur
Verwijderd: –

Auteur
Opmerking [10]: Is this relevant for the current article or would a single sentence suffice?

using Google Colab, a partly free computing application with a paid tier, that allows Python commands to be implemented and stored in the Google cloud.

Once we began working in the cloud, we were heavily reliant on the aid of Deep Learning (DL) [1] computing tools, which included AI engine, makesense.ai, to label our datasets categorically through bounding boxes. The archaeologists on our team found this technique to be easy to learn and apply successfully, requiring only basic programming skills. Conveniently, the collaboration between Github repository and Ultralytics (Jocher 2020) supplied a Google Colab forum prefixed with a YOLOv5 training notebook[2]. This provided the foundational base where our dataset (taken from 2022 fieldwork imagery) was input into Colab's YOLOv5 training applications.

Our first attempts at training the algorithm offline with installed software were not as easy as expected. For example, it was essentially impossible to "plug-and-play," so to speak. We also discovered that using the lab imagery from 2015, which had been processed at a lower resolution and was free of 'background noise', did not work as well as (considerably more complex) field imagery. In fact, while training the algorithm, Christofis discovered that the prepared diagnostic imagery and images of whole-vessel pottery specimens downloaded from the web worked exactly 0% of the time (Figure 3). Although such images are certainly useful for typological identification of pottery and other artifacts, they were unsuitable for training the algorithm alongside field-captured imagery.

Importantly, YOLOv5 was never able to identify any pottery sherd images taken in a lab setting once we used the complex images captured in the field in initial training sessions. In all lab-based datasets, YOLOv5 could not detect the differences between stone and sherd. These occurrences continued through approximately 50 images and three training runs, where all sherds were incorrectly detected as stones in our output. These results continued to come back just as inconclusive, even with aid provided by comparative training of lab-based datasets against stock images of stones. This means that, given the current capabilities of YOLOv5, the algorithm cannot be utilized for legacy assemblages if training multiple datasets (like survey imagery) subsequently and/or simultaneously. We believe, however, that laboratory collections and other prepared diagnostics can be used if trained separately from the objects photographed in the field. We plan to test this hypothesis in the future.

---

[1] The standardization and training methods discussed could not have been constructed without the help of DeepLearning's tutorial, "YOLOv5 training with custom data" (DeepLearning 2020).

[2] We were thankful to have support from Weights and Biases' online YOLOv5 tutorial (Davies 2022)– a step-by-step guide on how to download and kickstart YOLO's object detection features.

Auteur
Opmerking [11]: Reference?

Auteur
Verwijderd: ary-

Auteur
Opmerking [12]: These are also results, I would move them to the result section.

Figure 3: Results using images processed in the lab to train the algorithm after initial attempts at training using field photos.

After our YOLOv5 repository was established and all complimentary programs downloaded, we began the process of setting up a standardized workspace. This started with creating a master folder to contain training datasets (images) and their allocated labels — a compartment for each titled "images" and "labels." Inside these two folders, two more spaces were created to separate "train" and "validation" images and/or labeling. Validation datasets were used as a precautionary addition to the training datasets, meant to fine-tune and evaluate detection accuracy. After setting up our master folder, we imported the datasets from our 2022 fieldwork into the train and validation containers of the images folder (our validation images were always a smaller sample size than training). We tried to choose collections that replicated the natural 'noise' of the landscape, where pottery sherds were scattered along rocks, found on dirt, and in association with sticks and leaves.

Auteur
Verwijderd: –

Auteur
Verwijderd: ,

Auteur
Opmerking [13]: A validation dataset is generally used to check what the performance is of a model during training. Maybe rephrase this sentence or please explain more how the dataset was used in this research.

With the master folder uniform, we began the process of bounded-box labeling for both the training and validation datasets. To do so, the images were imported into makesense.ai's object detection browser. This free-to-use GPLv3 website was quite user-friendly (if not a tedious affair), that allowed us to upload photos, set up personalized labels, and transfer our new labels into various components of our images. For this study, we chose detection of four classes — sherd, stone, leaf, and stick — due to the repetition of appearance that these elements demonstrated within our images. As makesense.ai's manual object detection process would be used to cross-analyze YOLOv5's own capabilities to detect parallel sets, being thorough and detailed (while labeling) was crucial for a standardized accuracy reading. After labeling all sherds, stones, leaves, and sticks within each image of each dataset, we exported all labels into a YOLO compatible zip file, which was then extracted into the designated label file.

Following the exportation of our labels, our master training folder (now consisting of full datasets and labels) was ready to be uploaded for YOLOv5 training. Here, we transitioned to the next part of our research with the help of Google Colab. Google Colab's YOLOv5 workspace[3] is conveniently set up as a functioning tutorial – much of the code needed to train, validate, and detect with YOLOv5 has already been provided. After cloning YOLO, we activated the first line of code. This allowed us access to all applications in an individual Colab session. After cloning finalized, we uploaded and unzipped our training data into Colab. Here we found that YOLOv5's pre-existing classes have no applicability to our own study. For example, before re-classification images of pottery specimens had been misidentified as modern objects like toilets, hammers and books (Figure 4). Ergo, we edited the code on this file to match the classes/labels that we had created and implemented during the beginning of this research process. Ensuring that our project-specific constraints were now recognized by YOLOv5's algorithm, we changed the number of classes to 4, and the names (labels) to "stone", "sherd", "stick", and "leaf" which improved detection results substantially (Figure 5). With these adjustments in place, the data were re-imported back into Colab's YOLOv5, where we could finally test detection capabilities by advancing to "train" in our workspace. Activating the second block of prewritten code and experimenting with different epoch bounds we achieved the following results.

[3] https://colab.research.google.com/github/ultralytics/yolov5/blob/master/tutorial.ipynb

Auteur
Verwijderd:

Auteur
Verwijderd: –

Auteur
Verwijderd: –

Auteur
Verwijderd: run,

Auteur
Verwijderd: , and train through

Auteur
Opmerking [14]: I think it is not necessary to be this detailed about the code on Colab. A sentence that the classes were adjusted and a reference to the code is enough.

Auteur
Opmerking [15]: What I am missing is the parameters that were used for training: number of epochs, learning rate, etc.

Figure 4: Image depicting a shell and pottery fragment misidentified as a toilet and book.

Figure 5: Improved object detection using smaller set of relevant classes.

Preliminary Results

Precision and recall are two commonly used evaluation metrics in ML and information retrieval. Precision tells us how often the model is correct when it predicts a positive label. A high precision score indicates that the model is good at correctly identifying positive instances, while a low precision score indicates that the model is making a lot of false positive predictions.

Recall tells us how often the model correctly identifies all the positive instances. A high recall score indicates that the model is good at finding all the positive instances, while a low recall score indicates that the model is missing a lot of positive instances.

While we trained initially at 150 epochs, 60 was identified as the ideal constraint for our research, as it allowed us to load the number of times that YOLOv5 trained our datasets without overfitting them. When YOLOv5 was run through other epoch modes, such as its default setting 3, and results were still promising.

Auteur
Opmerking [16]: Reference, and maybe just the formulae?

Auteur
Verwijderd: bounds

Auteur
Opmerking [17]: How was this determined?

After training for 150 epochs, we managed to obtain a precision score of 0.28, which meant out of 100, only 28 predictions were correct. Even though this was the highest precision score out of all our training runs, this was still quite low for our application. This was due to the limited number of examples that were available, as we only used 56 images for training this model (on this specific run). On the second run, we obtained a 0.23 recall score– still low for our application. Although both of these outputs exhibit low precision and recall, it is possible to make significant improve detection when new data are added. A precision score of 0.28 and a recall score of 0.23 supports that YOLOv5 is able to correctly detect sherds against the complexities upon the natural landscape (as seen in our collation of stones, leaves, and sticks).

Average Precision (AP) is a performance metric that measures how well a model can detect and localize objects in an image. When analyzing results, mAP simply refers to the mean of Average Precision.

We achieved a mAP value of 0.19, which meant that the average precision across all the classes in the dataset was 0.19. In other words, the model's performance in detecting and localizing objects in the images was not ideal (The mAP score ranges from 0 to 1, where a score of 1 indicates perfect performance, and a score of 0 indicates that the model is not able to detect any objects correctly). A mAP score of 0.19 is again quite low, and it suggests that the YOLOv5 model needs to be improved in order to achieve better performance. It is important to note, however, that the interpretation of a mAP value depends on the dataset and the specific task being evaluated. Thus, a score of 0.19 may be considered 'good' or 'bad' depending on the difficulty of the task. In other words, archaeological object detection may be more difficult than other utilities, and therefore, what is considered a 'good' mAP value may be lower than in other applications.

A common way to monitor a model's performance during a training session is by plotting mAP as a function of the number of iterations. This graph below provides information of the YOLOv5 model's performance during our training process. This gives us guidance on how well the model is fitting the training data. During the initial stage of the training, mAP increases rapidly, and then slows down as the number of iterations increases. After some time, it remains constant. As we can see in the chart below, the model is constant after 60 iterations, which means that the model begins to overfit the data.

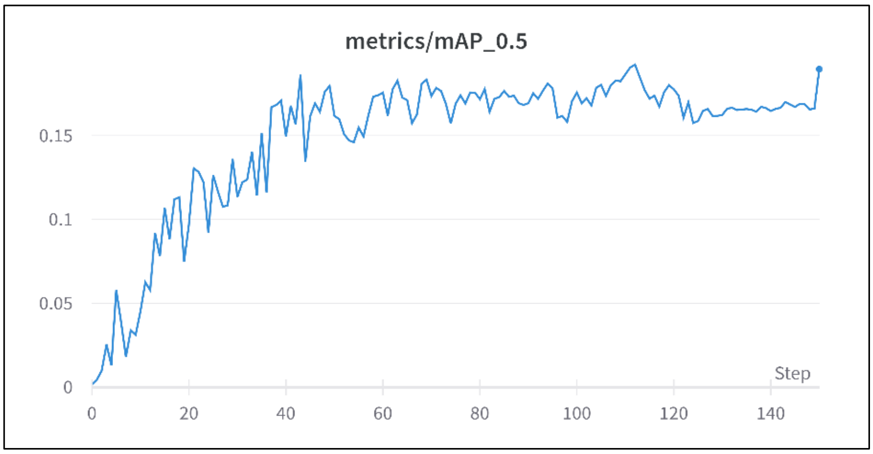

Figure 6: Graphical illustration of results after 150 iterations.

Discussion and Concluding Remarks

Auteur
Verwijderd: at

Auteur
Verwijderd: datasets

Auteur
Opmerking [18]: What was the precision on this second run and the recall on the first?

Auteur
Verwijderd: has means to

Auteur
Opmerking [19]: You are using mAP 0.5 (see figure 6), this should be explained.

Auteur
Opmerking [20]: Reference.

Auteur
Opmerking [21]: There has been quite some research on the added complexity of archaeological versus normal object detection. See for instance the PhD thesis of Verschoof-van der Vaart: https://hdl.handle.net/1887/3256824 Especially chapter 7.

With the semi-successful training of the algorithm, we hope to standardize YOLOv5's object-based methodology and expand it to our low altitude drone-based survey data. How does Google Colab fit into this picture? Certainly, it provided us with a convenient tool when software downloads and versioning presented critical problems. A critical question that emerges, however, is whether it is necessary to work independent of the online system. In particular, we look towards the limitations of Google Colab: its potential for data loss and even its potential untimely discontinuation in the future (as has been the case with other Google products like Google Poly). These challenges notwithstanding, the ability to conduct high-resolution classification of objects on the ground without collecting them or using expensive GPS units are some of the key advantages of using YOLOv5 technique.

Auteur
Opmerking [22]: An option would be to use Roboflow – Colab and then run the trained model locally.

With the training of the YOLO object detection algorithm on material culture collected on the ground and with UAVs (see Orengo et al. 2021), we have entered a new era of digital archaeology. This affordable approach produces high-resolution spatial data as well as material culture records that are informative, experiential, immersive, and easily accessible (i.e. limited post-processing) and permit us to leave objects *in situ* without destroying site components. Among others, the ability to weed out non-essential data and focus on the essential, visually navigate and detect objects and features non-destructively are proving to be key advantages (see also Mittal 2020). While fundamentally successful at detecting pottery sherds on the ground, the ability to transfer data into our GIS for spatial analysis (based on accurately classified images captured with location enabled smartphones and machine learning with YOLOv5) has the capacity to transform traditional time-consuming survey, object recording, and collection.

Auteur
Verwijderd: s

Auteur
Opmerking [23]: Has this already been developed? Because generally images lose their geospatial information when going through a CNN.

For archaeologists, a critical advantage of incorporating object-based machine learning techniques is never having to pick up an artifact while at the same time being able to identify it stylistically, quantitatively, and locationally. Applying such methods contributes to the speed and accuracy with which new survey data can be captured. Through high-resolution *in situ* recording, using images captured on the ground or remotely using UAVs, the preservation of fragile and at-risk archaeological sites is significantly improved. Furthermore, through the advancement of non-destructive techniques, we gain a high-resolution view of the landscape and surrounding environment while actively working towards its conservation.

In this experiment, YOLOv5 has exhibited potential for consistent detection of archaeological phenomena, although statistically marginal. In this case, it was successful at detecting and classifying the presence of pottery sherds when used with images captured on the ground rather than the UAV imagery. Presently, images captured using our UAV camera were too vague to produce the necessary detail for distinguishing pottery from rocks. During the next step of this project, we plan to re-capture drone data at lower altitudes (approximately 15-20 meters) and use higher-resolution cameras to test our methodology.

We also believe that, given more time to enlarge the dataset and retrain the model, we support ideas that implementing programs like YOLOv5 into UAV surveys could considerably supplement the way that today's digital archaeological fieldwork and surveying is done. UAV-captured surveying and diagnostic evaluation with machine learning tools will save time, money and resources, while allocating more space to actively research, excavate, and connect with communities. The immense amount of data provided by AI-commanded drones as well as its "no-touch" approach also leaves opportunities for a better (and more ethical) digital recording of both archaeological and contemporary phenomena; thus, assisting in the current revolution to advance a more decolonized archaeology and emphasize preservation.

Auteur
Verwijderd: grow and train

Yet with our new-found ability to partially detect, analyze, interpret and share, we find ourselves ever at the cross-roads of experimentation, trial, and error. Although creative and thought provoking, we aim for standardized methodologies to lend credibility, longevity, and repeatability to our approach. With this paper and our shared insights, it is our aim to contribute to the dialogue on combining such approaches, helping to further integrate this advancing technique into our trans-disciplinary community of computational archaeology practitioners.

Acknowledgements

The authors wish to thank the field crew from the 2022 field season of Proyecto de Investigación Arqueológica Gallinazo del Norte (or PIAGN) for their help in recording the data used in this research. We are grateful for the loan of the DJI Mavic 2 drone used to capture data in 2022 PIAGN survey by Aerospace Engineering Lab of Dr. Hossein Eslamiat, Southern Illinois University, Carbondale.

Funding

A portion of this project was funded through NSF Award # 1519048 Doctoral Dissertation Research Award: Social Interaction as Determined Through Spatial And Technological Analysis.

Conflict of interest disclosure

The authors declare that they comply with the PCI rule of having no financial conflicts of interest in relation to the content of the article.

References

Bonhage, Alexander, Mahmoud Eltaher, Thomas Raab, Michael Breuß, Alexandra Raab, and Anna Schneider
2021 A modified Mask region-based convolutional neural network approach for the automated detection of archaeological sites on high-resolution light detection and ranging-derived digital elevation models in the North German Lowland. Archaeological Prospection 28(2):177–186. DOI:https://doi.org/10.1002/arp.1806.

Caspari, Gino, and Pablo Crespo
2019 Convolutional neural networks for archaeological site detection–Finding “princely” tombs. Journal of Archaeological Science 110:104998.

Caster, Joshua, Joel B. Sankey, Helen Fairley, and Alan Kasprak
2022 Terrestrial lidar monitoring of the effects of Glen Canyon Dam operations on the geomorphic condition of archaeological sites in Grand Canyon National Park, 2010–2020. US Geological Survey.

Davies, David.
2022 “ YOLOv5 Object Detection on Windows (Step-By-Step Tutorial).” https://wandb.ai/onlineinference/YOLO/reports/YOLOv5-Object-Detection-on-Windows-Step-By-Step-Tutorial---VmlldzoxMDQwNzk4.

DeepLearning.
2020 “YOLOv5 training with custom data.” YouTube.
https://www.youtube.com/watch?v=GRtgLlwxpc4&t=780s

Gómez, Francisco Díaz, Josué Jiménez Peiró, Amparo Barreda Benavent, Bárbara Asensi Recuenco, and Juan Hervás Juan
2015 Modelado 3D para la generación de patrimonio virtual. Virtual Archaeology Review 6(12):29–37.

Mittal, Payal, Raman Singh, and Akashdeep Sharma
2020 Deep learning-based object detection in low-altitude UAV datasets: A survey. Image and Vision computing 104:104046.

Nepal, Upesh, and Hossein Eslamiat

2022 Comparing YOLOv3, YOLOv4 and YOLOv5 for Autonomous Landing Spot Detection in Faulty UAVs. Sensors 22(2):464. DOI:10.3390/s22020464.

Orengo, Hector A., Arnau Garcia-Molsosa, Iban Berganzo-Besga, Juergen Landauer, Paloma Aliende, and Sergi Tres-Martínez
2021 New developments in drone-based automated surface survey: Towards a functional and effective survey system. Archaeological Prospection 28(4):519–526.

PyTorch
2023 https://www.pytorch.org, accessed August 31, 2023.

Sakai, Masato, Yiru Lai, Jorge Olano Canales, Masao Hayashi, and Kohhei Nomura
2023 Accelerating the discovery of new Nasca geoglyphs using deep learning. Journal of Archaeological Science 155:105777. DOI:https://doi.org/10.1016/j.jas.2023.105777.

Sharp, Kayeleigh
2019 Rethinking the Gallinazo: A Northern Perspective from the Mid-Zaña Valley, Peru. Southern Illinois University at Carbondale.