Archeologia e Calcolatori
37.1, 2026, 295-312
doi 10.19282/ac.37.1.2026.15

ARCHER: ARTIFACT CLASSIFICATION AND HANDLING WITH ENHANCED ROBOTICS

1. Introduction and state-of-the-art

In recent years, increasing attention has been paid to the use of methodologies related to robotics and AI in the Humanities, specifically, archaeology, which has long been regarded as the border between the Humanities and technical-scientific disciplines (Roosevelt *et al.* 2015; Bogdani 2019; Fiorucci *et al.* 2020). A whole series of technological tools have been developed to help researchers to collect the data as automatically as possible – possibly preprocessing it before human intervention – in order to provide archaeologists with more time for analyses and reflections beyond the capability of machines. Most of these technological aids focused on wide-field surveys of excavation sites by using uncrewed vehicles to carry out exploration of places otherwise difficult to investigate (Jeansoulin, Papini 2007; Eisenbeiss *et al.* 2008; Fernandez-Hernandez *et al.* 2014; Campana 2017). These tools have significantly expanded the volume of data collected for subsequent analysis. Some attempts at automatic surveys were also tried (Barazzetti *et al.* 2012). Unmanned Aerial Vehicles (UAVs), piloted remotely or using semi-automatic flight plans, were successfully employed for photogrammetric surveys of sites during excavation or research campaigns. This approach enabled the 3D or virtual reality reproduction of the places, allowing archaeologists to simulate real-world observation and tour (Olivito, Taccola 2014; Cassanelli *et al.* 2016; Olivito, Emanuele 2016). In the underwater field, progress has been made in the investigation of maritime sites (Warren *et al.* 2007; Bingham *et al.* 2010; Conte *et al.* 2010; Kuhn *et al.* 2014); despite the challenging environment in which the devices operate. Unmanned underwater vehicles (UUVs) have achieved great success, particularly on the seabed, where many artefacts had previously been inaccessible due to the difficulties in reaching them. UUVs are now unlocking new avenues of investigation that were previously closed off.

Robotics, when applied to the study of archaeological artefacts, is still a marginal field. The recent RePAIR (RePAIR 2021) and AUTOMATA (AUTOMATA 2024) research projects indicate that robotics will pave the way for the application of technological tools to the archaeological sector also for lab activities. A significant contribution has been made by the ArchAIDE project (Gattiglia 2018; Anichini, Gualandi 2019; Wright 2019; Gualandi *et al.* 2021; Anichini *et al.* 2022) through the development of a mobile app for the automatic recognition of ceramic sherds. This app, which utilises AI based on convolutional neural networks trained on thousands of images, allows researchers to quickly and accurately obtain detailed information about pottery fragments from a single image. The classification of ceramic fragments, which are among the most common finds in archaeological excavations, is typically a time-consuming and repetitive process, whether conducted manually by experts or with the assistance of AI. As a result, many sherds remain unclassified and stored in the warehouses of universities and other institutions; so far, there have been few attempts to automate the process of fragment recognition and sorting using automatic tools.

Among them, the RASCAL system (Wang *et al.* 2021) is noteworthy for its large-scale digital documentation of ceramic sherds: it captures metadata (weight, colour) and high-resolution images in a structured database, using a robotic arm with an integrated depth camera to handle sherds efficiently while minimising user intervention. This paper outlines the design and implementation of a system that integrates a robotic arm with a deep learning-based classification framework to identify and collect fragments, subsequently depositing them into designated containers. The primary objective of this study is to develop a system that automates the recognition and classification of archaeological potsherds using AI, in combination with a robotic arm. By leveraging an improved version of the neural network developed for appearance-based recognition of Majolica di Montelupo developed within the ArchAIDE project (Gualandi *et al.* 2021) and integrating it with a robotic arm, this work addresses the classification and sorting of archaeological ceramic fragments lying on a working surface. These tasks are performed in a simulation environment where, first, an image of the fragments – either captured from the real world or stored on the PC – is processed using state-of-the-art machine learning algorithms to decompose it, isolating and measuring each of the sherds (instances). These instances are then classified through the AI algorithms, assigning them a class with a confidence level. In the second phase, the classified fragments are spawned in the simulation environment on a working surface; then, a pick-and-place task is performed by the robotic arm in order to sort the fragments by picking them up and putting them in boxes. This automation aims to address a significant challenge in post-excavation activities, where a large quantity of ceramic material remains unclassified due to time and resource constraints: in fact, archaeologists typically should spend hours manually classifying and sorting ceramic fragments, forcing them to do repetitive, time-consuming tasks, preventing them from focusing more on excavation, in-field analysis, or rest.

It is often argued that the primary challenge in ceramic analysis is typological interpretation, but this perspective overlooks the significant logistical bottleneck known as the 'quantification crisis' in post-excavation workflows. Before high-level interpretation can occur, thousands of fragments must be physically sorted, measured, and counted to meet standard excavation reporting requirements. This process is inherently repetitive and consumes hundreds of person-hours that could otherwise be dedicated to stratigraphic analysis or field investigation. From the point of view of the practical workflow, the ARCHER system is designed for bulk batch-processing rather than the handling of individual isolated fragments. In a real-world archaeological scenario, the intended use case involves taking all fragments recovered from a single stratigraphic unit – typically stored together in a bulk container (i.e.: a box, a bag) – and placing them simultaneously on the system's working surface.

The robotic system then autonomously segments, identifies, and physically sorts the entire group into specialized containers based on their identified classes. This allows the archaeologist to return the material to its original storage unit not as a chaotic mass, but as a structured and ordered set, with each fragment already linked to its corresponding digital metadata in the ArchAIDE database. Consequently, the system acts as a force-multiplier for the researcher, automating the tedious physical reorganization that usually follows a major excavation campaign.

2. System design

The simulated robotic system has been realised by combining a deep learning-based classification system with a ROS-based software framework for control and simulation, which implements a robotic platform with a Tinkerkit robotic arm and an Intel RealSense camera.

2.1 *The classification system*

The classification system was based on the ArchAIDE deep learning algorithm for appearance-based recognition (Anichini *et al.* 2022), which provides automatic recognition of archaeological potsherds decoration, specifically Majolica di Montelupo decoration, from pictures. The ArchAIDE system returned a top-5 results list, allowing users to choose the correct option from the five suggestions. In this respect, there are two major issues when applying this system to this work setting. The original ArchAIDE classification system was developed using Tensorflow 1.x with a pre-trained ResNet101 CNN, which has since become obsolete and difficult to maintain. Moreover, ArchAIDE required active human interaction to identify the correct classification from the list of five options. To enable the present system to function autonomously, the deep learning framework was updated to Pytorch (Paszke *et al.* 2019), testing different neural networks in order to obtain higher reliability and accuracy in recognition. Additionally, the system was modified to utilise top-1 accuracy, allowing the robotic arm to operate independently without human intervention.

A ResNet50 CNN was chosen to test the recognition workflow of the robotic system. It is a convolutional neural network that is 50 layers deep which utilises a bottleneck design consisting of 48 convolutional layers, one

MaxPool layer, and one Average Pool layer. With approximately 25.5 million trainable parameters, the model was fine-tuned using transfer learning on 66 decoration genres of 'Majolica di Montelupo'. This specific configuration was selected after testing indicated it offered the most efficient balance between validation accuracy and inference speed compared to deeper models like ResNet101. Specifically, the ResNet50 achieved an inference time of 0.7 seconds, ensuring the robotic arm could operate without significant computational delays. The specific *corpus* of Majolica di Montelupo has been chosen to leverage a pre-existing dataset of thousands of training images, thereby focusing research efforts on the novel integration of the CNN with the ROS-based robotic control framework. This ensured the prototype could be tested against a variety of decoration genres (66 selected classes) to validate the effectiveness of the autonomous sorting workflow.

To ensure the statistical validity of the classification results, the dataset was partitioned into two sets for training and validation. The validation set, which constitutes the testing data used to evaluate model performance and establish the confidence threshold th_{conf}, as we will discuss later, contains images that were excluded from the training phase. This prevents data leakage and ensures that the accuracy metrics reported in Tab. 1 accurately represent the network's capacity for inference on previously unseen ceramic materials. It is important to say that, while the current study utilises this *corpus* as a specific case study for training and testing, the underlying architecture – built on the ROS framework – is designed to be material-agnostic. The objective is to demonstrate a functional robotic work-cell that can be recalibrated for various ceramic or lithic classes by swapping the neural network weights, rather than presenting a finished general-purpose solution for all archaeological materials.

2.2 *Robotic platform and software*

A 6-degree-of-freedom Arduino Tinkerkit Robot Arm was integrated into the system. This arm is capable of full 360-degree rotation at the base and has a maximum operating distance of 80 cm. The working payload at 32 cm operating distance is 150 g, and the maximum payload at the minimal arm configuration is 400 g, well within the limits of the vast majority of ceramic sherds collected during excavation. The interface is provided by an Arduino board (i.e. Arduino Uno) with a shield that allows the servomotors to be powered on directly from the power supply shared with the board and controlled via Arduino PWM pins. For image capture, we use an Intel RealSense D415 camera with a depth sensor, which allows us to obtain an RGB image for feeding the ArchAIDE Convolutional Neural Network (CNN) for recognition and a point cloud to calculate the grasp pose and perform the pick-and-place operation. The robotic arm and the Intel RealSense camera are mounted on a custom-designed platform.

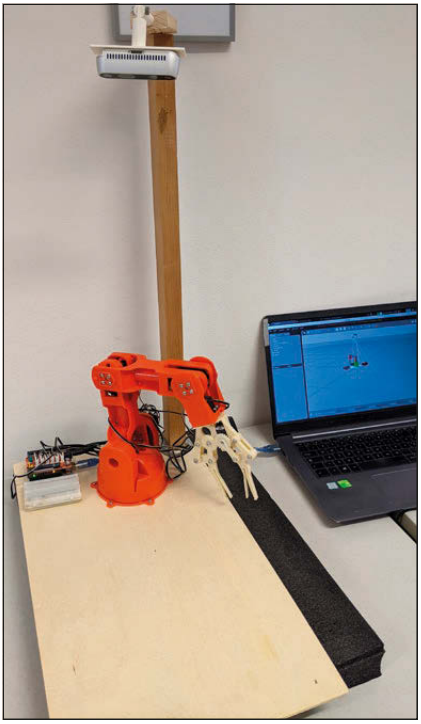

Fig. 1 – Robotic system.

The base measures 25 cm per side and has two circular notches with a radius of 12.5 cm centred on the corners of the plinth. This design ensures that the fragments closest to the arm rest directly on the table and not on the base. The pylon which holds the camera is 60 cm high and has a square cross-section of 2.5 cm; at the top, it extends 17.5 cm towards the centre of the base, ending with a cylindrical pin that provides one degree of rotational freedom. This pin supports a cylindrical hinge at the end of which the camera is mounted. This configuration allows the camera to be oriented conveniently to catch all the fragments on the table (Fig. 1).

We used ROS (Robotic Operating System) with MoveIt! (Coleman *et al.* 2014) motion planning framework to control the robot and the D415 camera, Gazebo 1.11 with RViz 1.14 to visually control motion planning for simulation. The integration of RViz and Gazebo with MoveIt! provides a high-fidelity environment that ensures excellent sim-to-real transferability. This framework operates as a realistic, physics-based digital twin, utilizing sensor modeling and world-building that treat data as it would appear in a real-world scenario. By employing the same ROS-based control nodes and motion planning for both the virtual model and the physical arm, the system minimizes the transition gap. While this article presents a proof-of-concept validation, the architecture is designed to be deployed directly onto physical hardware. The bridge between ROS/MoveIt! and the robotic arm is provided by the Arduino ROS library, which converts instructions coming from ROS nodes into commands for Arduino.

3. AI recognition workflow

The system is designed to allow the user to visualise the pick-and-place operation in simulation, enabling verification of its success and ensuring that the arm does not enter singularity – a condition where certain configurations cause a loss of degrees of freedom, i.e. the arm can no longer move independently in all intended directions, leading to unpredictable motion or instability. If so, the system proceeds with the grasp operation of the fragments by adding only an additional parameter when launching the simulation. The simulation has been conducted in Gazebo, i.e. a robust, open-source platform for robotics simulation, widely used for testing robotics algorithms, autonomous systems, and machine learning in a controlled, dynamic environment that provides realistic physics-based simulations, detailed 3D environments, supporting sensor modelling and custom world-building (Koenig, Howard 2004).

3.1 *World setting, image caption and fragment recognition*

We use a ROS launch file to spawn the simulation in the Gazebo environment. Additionally, the Rviz visualisation tool (Kam *et al.* 2015) is used to obtain additional information on each spawned element. The arm, the platform and camera are positioned at the centre of the world reference system; three cylindrical containers are spawned behind the platform, each spaced 60 degrees apart from the other (Fig. 2). As described later, the fragments of the three selected classes will be deposited in these boxes.

Images of the fragments can be obtained either by using the Intel RealSense camera mounted on the real-world system or a pre-existing image stored on the PC (for demonstration purposes only). It is crucial to include a 10 cm chessboard scale on the table alongside the fragments. OpenCV will use this scale to determine the dimensions of the fragment on the working surface. The image is processed via the Python library OpenCV, with this workflow (Fig. 3, original on the left, segmented on the right):

– the image is converted into black and white;
– a Gaussian filter is applied to facilitate segmentation;
– the fragments are divided into single instances through edge recognition using a binary and an Otsu filter for thresholding;

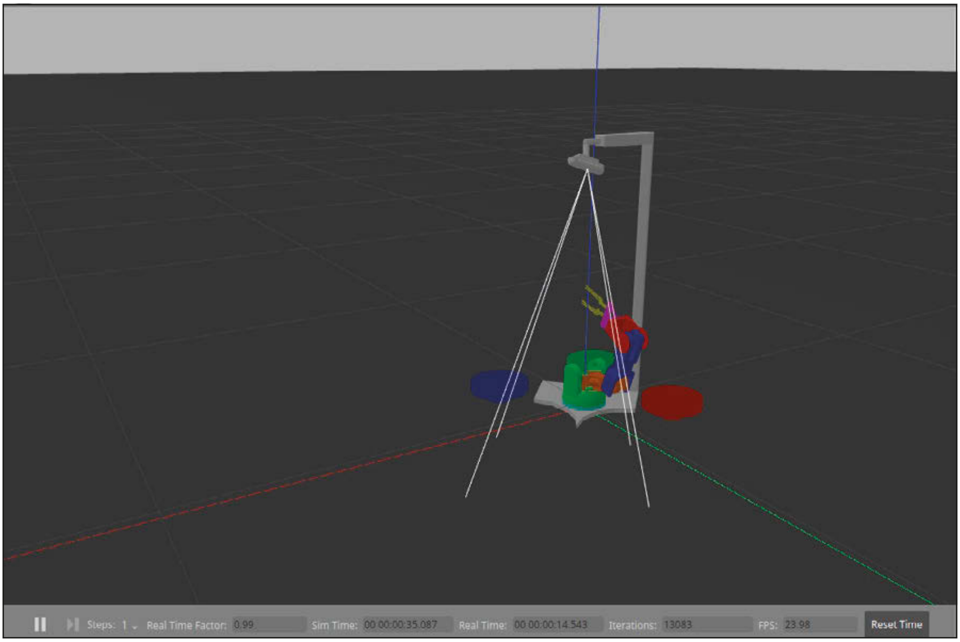

Fig. 2 – The robotic manipulation system in the Gazebo simulation environment.

– the background is blanked out;
– instances are measured (by comparison with the chessboard scale at the bottom of the image), establishing their bounding box size and coordinates of the centre of each fragment.

The single instances are then passed to the trained deep learning model, which assigns each a class and a confidence level for recognition. Each instance is then spawned into the simulation according to its size and coordinates of the centroid. The coordinates of the centroid are measured relative to the origin of the world axes; this origin coincides with the centre of the robotic arm's base and plinth. Then, colour is applied to the fragments depending on the assigned class (Fig. 4). Once all the fragments have been spawned, the system uses a ROS topic to let other scripts identify the names assigned to them, the coordinates of the centre of each fragment, the container in which each fragment should be deposited, and its average dimension.

3.2 *Class and fragment selection*

As said, the CNN used for the classification is trained on 66 decoration genres (ceramic classes) of 'Majolica di Montelupo' out of the original 83

Fig. 3 – Segmentation and image recognition.

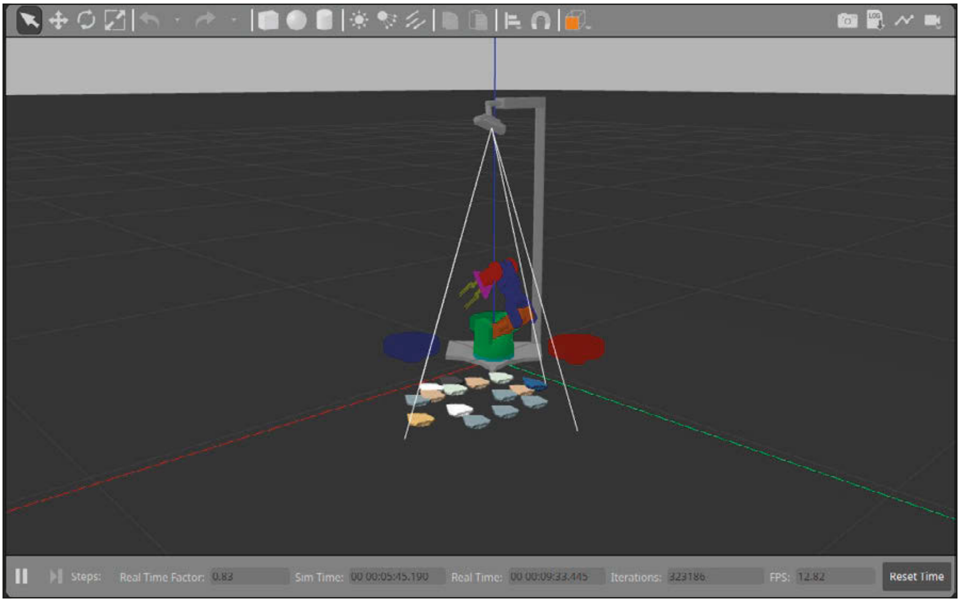

Fig. 4 – Fragments are spawned into the world.

of the ArchAIDE project, each containing at least 15 images: this selection was performed to get a reliable top-1 accuracy during training and, then, inference. As mentioned, the system includes only three cylindrical boxes for fragment sorting; therefore, it is necessary to select exactly three classes from the K (≤ 66) classes to which the N fragments on the table may belong. For this reason, after recognition, the system groups the fragments into their respective classes. For each class, the average confidence level, the standard deviation and the number of fragments belonging to that class are calculated. The three classes are then selected, following these prioritised criteria:

– the highest average confidence level;
– the lowest standard deviation;
– the highest number of fragments.

In any case, to be selected, the average confidence level of a class must be above a certain threshold $\underline{P_i} \geq th_{conf}$ and the standard deviation under a certain threshold $\sigma_i \leq th_\sigma$, calculated as in equations below, where $\underline{P_i}$ is the mean confidence level of each class, σ_i is the mean standard deviation of the class, $P_j(i)$ is the confidence level of each fragment belonging to the class, and M is the number of fragments in the class.

$$\forall C_i{:}\ \underline{P_i} = \frac{\sum_{j=1}^{M} P_j^{(i)}}{M} \geq th_{conf}$$

$$\forall C_i{:}\ \sigma_i = \sqrt{\frac{1}{M}\sum_{j=1}^{M} (P_j^{(i)} - \underline{P_i})^2} \leq th_\sigma$$

The value of th_{conf} is established based on the global accuracy of the trained CNN on the testing dataset. The fine-tuning of this threshold followed a heuristic approach: by setting th_{conf} equal to or slightly higher than the model's global accuracy, we ensure that the robotic arm only attempts to sort fragments where the AI's confidence exceeds its own average performance. This serves as a critical safeguard to minimize misclassification errors in an autonomous setting; any fragment with a prediction confidence below this baseline is discarded rather than being incorrectly placed in a container.

Despite the high top-1 accuracy of the trained model, some fragments may be incorrectly classified. The same confidence threshold used above for class selection is used to limit misclassification errors. If the confidence score is below it (regardless of the predicted class), the fragment is discarded. The threshold is set to be greater than or equal to the global accuracy of the trained model.

3.3 *Pick-and-place operation*

Once the fragments have been created in the simulation and the three classes have been selected, as previously described, pick-and-place operations can be performed. The script creates a transform coordinate (TF) system centred on the fragment. Then, using the MoveIt! libraries and KDL plugin,

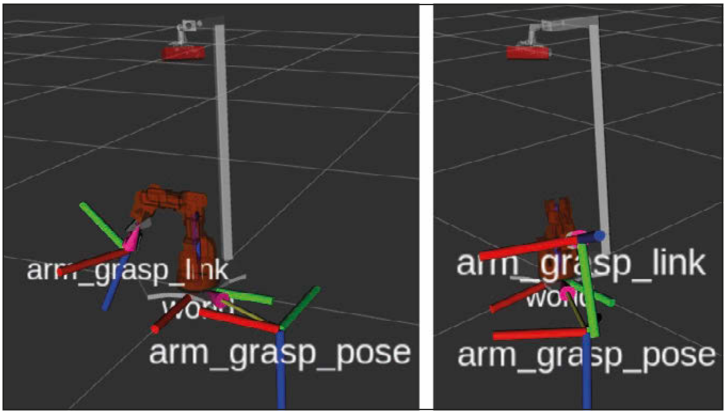

Fig. 5 – Double-TF scenario.

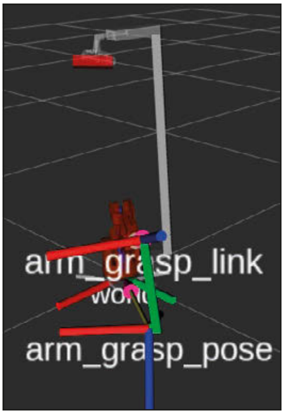

Fig. 6 – Single-TF scenario.

it positions the end effector (EEF) on the fragment, i.e. the system attempts to align the EEF's transform coordinate system with the one created on the fragment, allowing the execution of the grasping.

Two different Inverse Kinematics solvers were employed to address possible unintended movements dangerously close to the working surface that could interfere with the positioning of the other fragments. The first, based on https://github.com/studywolf/blog/tree/master/InvKin, a simplified system with 3 degrees of freedom (DOF) (shoulder-elbow-wrist), keeps the end effector raised from the working surface, ensuring a vertical positioning of the EEF on the selected fragment. Then, once the positioning has been

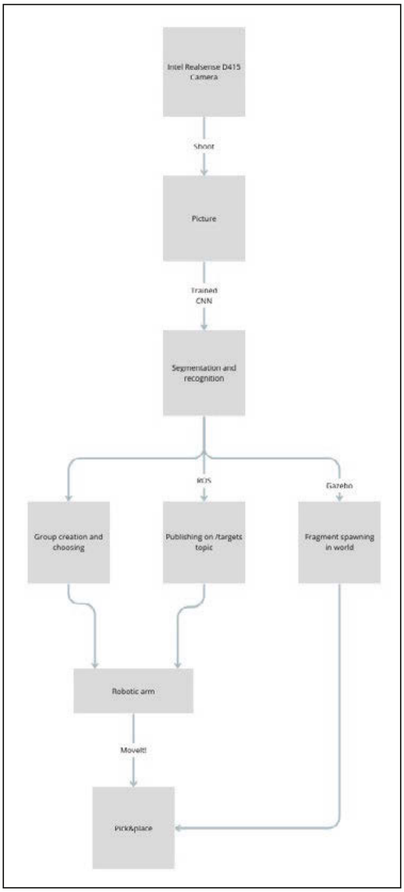

Fig. 7 – System workflow.

carried out, the second, a coordinate system, is created over the fragment and MoveIt! and solves the pose for the pick-and-place operation.

To prevent a too tight grasp, which could damage the fragment, the average size of the fragment – derived from the bounding box size calculated during the OpenCV workflow – is used to determine the extent of closure of the gripper jaws. After various attempts, two different scenarios for creating the TF coordinate system on the selected fragment through two distinct scripts were chosen: in the first scenario (double-TF scenario), a TF coordinate system, rotated by approximately 90 degrees around the z–axis relative to the world axes orientation, is initially created (Fig. 5, left). In the event that MoveIt! fails to resolve the kinematic chain with this TF within the 5-second timeout,

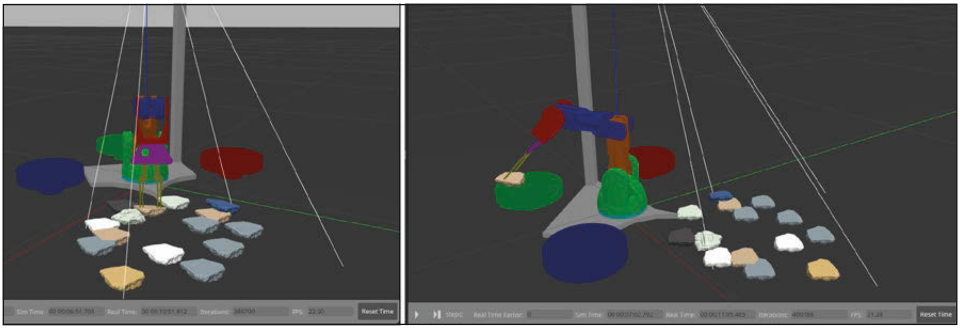

Fig. 8 – Pick-and-place operation.

it is replaced by one rotated approximately 45 degrees around the z-axis with respect to the world coordinate system (Fig. 5, right).

In the second scenario (the single-TF scenario), a coordinate system rotated approximately 45 degrees around the z-axis with respect to the world coordinate system is created, which reduces the calculation time for the grasp pose (Fig. 6).

The overall recognition and pick-and-place workflow can be summarised in the flowchart in Fig. 7: in the first phase, the fragments are classified, and the three classes to be sorted are chosen using machine and deep learning methods. In the second phase, the sorting takes place using a robotic arm that grasps the fragments and deposits them in the three cylindrical boxes. The pick-and-place operation (Fig. 8) takes an average of 15 seconds for each fragment, from the moment the system receives the ROS message on the topic to when the fragment is released in the box.

4. Results

As described above, tests were carried out on the images in Fig. 9 using the two possible scenarios for creating TFs. A ResNet50 trained on the 66 selected classes of the dataset has been used since this network achieved the best trade-off between accuracy and inference speed among the CNNs tested for the project (ResNet50, ResNet101, Inception_v3), as shown in Tab. 1.

The tests showed some variability in the success rates depending on the TF configuration used. Generally, the single-TF configuration had a higher success rate in grasping and correctly placing the fragments in the corresponding bowls (Tab. 2). The success rate across various trials ranged between approximately 66% and 80%, with only a few fragments missed or dropped.

Meanwhile, the double-TF configuration resulted in lower success rates in some cases, with certain tests showing rates as low as 33%, though some trials achieved 100% success (Tab. 3).

CNN	Loss		Accuracy (%)		Inference time
	Training	Validation	Training	Validation	
ResNet50	1.0410	2.2867	75.5629	46.8077	0.7
ResNet101	1.6763	2.3342	54.1715	39.2473	1.6
Inception_v3	1.0982	2.4161	71.6801	42.5067	1.1

Tab. 1 – CNNs results.

Filename	No. of Sherds	Sherds assigned to a chosen class	Sherds correctly placed	Success (%)
img005DSC02101_1	6	4	3	75.0%
img005DSC02101_2	5	5	4	80.0%
img043DSC02180_1	4	3	2	66.7%
img148SG101877_1	6	5	4	80.0%
img148SG101877_2	6	6	4	66.7%
img183SG101921_2	5	3	2	66.7%

Tab. 2 – Results of single-TF pick & place experiments.

Filename	No. of Sherds	Sherds assigned to a chosen class	Sherds correctly placed	Success (%)
img005DSC02101_1	6	4	4	100.0%
img005DSC02101_2	5	5	2	40.0%
img043DSC02180_1	4	3	1	33.3%
img148SG101877_1	6	5	1	20.0%
img148SG101877_2	6	6	3	50.0%
img183SG101921_2	5	3	2	66.7%

Tab. 3 – Results of double-TF pick & place experiments.

The difference in success rates could be attributed to the cluttered scene in which the robotic arm had to perform grasps, which likely led to fewer accurate grasps, especially in more complex fragment shapes. Missed fragments, or those that fell after a grasp attempt, were more common in the double-TF configuration. The highest observed percentage of missed fragments reached 80%, reflecting the challenges in reliably grasping this configuration. However, even with the single-TF, some trials still exhibited missed fragments, albeit at a lower rate. Overall, the single-TF configuration appeared more reliable, especially in terms of grasp accuracy and reduced fragment drop rates. However, the double-TF configuration may still be beneficial in non-cluttered scenarios where fragments are well-spaced from each other.

The simulation results serve as a systematic validation of the system's control logic and kinematic chain. The success rates, ranging between 66% and 80% for the single-TF configuration, demonstrate that the AI-driven grasping logic can effectively manage the complex geometries of archaeological fragments within a realistic physical simulation. We acknowledge that these

Fig. 9 – Test images.

results represent a successful proof-of-concept; consequently, the primary objective of this phase was to refine the software-hardware interface. The validation of these algorithms on physical archaeological materials, accounting for real-world surface friction and material fragility, constitutes the immediate subsequent step of this project's development.

5. Discussion

The current setup achieves a grasping and placement success rate of approximately 66-80%, but lacks the ability to verify grasp success or detect fragment loss during transport to the sorting box, limiting its practical application in real-world archaeological scenarios without continuous human intervention. To address this, the next development will implement current feedback monitoring for the servomotors. When the end effector (EEF) successfully grasps a fragment, the absorbed current increases significantly as the gripper's internal surfaces collide with the fragment's edges, preventing complete closure. Conversely, the current remains low when the EEF is empty due to the lack of collisions. This feedback mechanism would allow the system to identify and retry unsuccessful grasps, improving efficiency and reducing operational downtime, as well as human intervention, making it more useful in real-world archaeological scenarios. Furthermore, employing a small, cost-effective robotic arm demonstrates the system's feasibility with minimal economic resources, highlighting its portability and potential use during excavation campaigns.

However, the system's design also supports the integration of larger, more advanced robotic arms – such as the Franka Emika or Cobot UR5 – which offer greater payload capacities and extended reach. These alternatives are particularly suited for in–lab activities requiring enhanced performance for more demanding tasks. A comparison with the similar RASCAL system (Wang *et al.* 2021) highlights key differences in focus and implementation, despite shared features such as the use of simulation environments like Gazebo for testing and refinement and the integration of the Robot Operating System (ROS) framework to ensure adaptability and seamless hardware integration. Both systems employ 6-degree-of-freedom robotic arms, but their hardware differs: RASCAL uses a WidowX 250 robotic arm integrated with a depth camera for object detection and manipulation, while our system utilises an Arduino Tinkerkit robotic arm and an Intel RealSense D415 camera for imaging and classification. RASCAL is primarily designed for the large-scale digital documentation of ceramic sherds, capturing detailed metadata such as weight, colour, and high-resolution images, which are stored in a structured database. In contrast, our system focuses on classifying ceramic fragments using convolutional neural networks (CNNs), physically sorting them into predefined categories. While both systems aim to automate processes with minimal user intervention, their objectives diverge. RASCAL prioritises the creation of a comprehensive and scalable digital archive. Meanwhile, our system addresses the demands of post-excavation classification and sorting, emphasising precision in recognising and organising artefacts and also providing data storage in the ArchAIDE database (Gualandi *et al.* 2021).

6. Conclusions

While many once-disruptive technologies have become integral to archaeological practice, robotics remains underutilised. Innovative projects like RePAIR (RePAIR 2021), which focuses on reconstructing Roman frescoes using AI-supported robotics, and the newly funded AUTOMATA (AUTOMATA 2024), aimed at developing a robotic AI-enhanced work cell for digitising ceramics and lithics, signal a burgeoning interest in this area. Robotic systems, particularly robotic arms, offer solutions to laborious tasks such as quantification, puzzle-solving, and large-scale digitisation – activities often hindered by limited time and human resources. The prototype we presented addresses the common post-excavation challenge of pottery quantification, a time-intensive process that frequently leaves numerous ceramic fragments unclassified. To mitigate this issue, we propose automating the classification process through a robotic arm capable of physically sorting fragments based on their identified classes by being integrated with an AI system trained for recognising and categorising ceramic fragments. The convolutional neural network (CNN) (recently presented at conferences and will be discussed in another paper) was developed using the framework established for the Majolica di Montelupo in the ArchAIDE project, with the primary objective to optimise accuracy, enhancing the reliability of a system operating without human intervention.

The use of a relatively specialized dataset (66 classes with 15+ images each) was a deliberate choice to test the system within a controlled typological environment. We acknowledge that such a sample size presents risks of overfitting in a production environment; however, in this prototype, we mitigated these risks through transfer learning and the implementation of stringent confidence thresholds. Future iterations will leverage larger, more diverse datasets – such as those being developed by the AUTOMATA project – to transition from this validated proof-of-concept to a more robust, general-purpose archaeological tool.

This dual approach facilitates the classification of archaeological pottery – preventing artefacts from remaining unsorted and neglected in archaeological warehouses – and generates classification data, which can be archived for future analysis. The system is designed to accommodate alternative neural networks trained on diverse datasets, allowing for various classification objectives and applicability across different archaeological contexts and material types. Future development could incorporate a human-in-the-loop approach to enhance overall performance further, as evidenced by projects like Pithia (Assael *et al.* 2022). By combining automation, adaptability, and scalability, the project successfully balances technological innovation with the practical realities of archaeological practice. While the system provides a versatile solution for post-excavation tasks, its implementation is envisioned as a support tool for authorized institutional laboratories. By integrating robotics with established digital archives, the ARCHER system demonstrates significant potential to address the 'quantification crisis' while adhering to the rigorous conservation and administrative standards required by the international archaeological community.

Quirino Saraceni, Gabriele Gattiglia
Università di Pisa
quirino.saraceni@phd.unipi.it, gabriele.gattiglia@unipi.it

Nevio Dubbini
Miningful Srls, Pisa
nevio.dubbini@gmail.com

REFERENCES

Anichini F., Gattiglia G., Gualandi M. 2022, *Classificare in modo automatico i frammenti ceramici: ArchAIDE, un'utopia realizzata*, «Agoghè», 275-294 (https://doi.org/10.12871/978883339612518).

ANICHINI F., GUALANDI M.L. 2019, *Una "rete neurale" per il riconoscimento automatico della ceramica. Il Progetto ArchAIDE*, in *Atti del LI Convegno internazionale della ceramica (Savona 2018)*, Firenze, All'Insegna del Giglio, 73-81.

ASSAEL Y., SOMMERSCHIELD T., SHILLINGFORD B., BORDBAR M., PAVLOPOULOS J., CHATZIPANAGIOTOU M., ANDROUTSOPOULOS I., PRAG J., FREITAS N. 2022, *Restoring and attributing ancient texts using deep neural networks*, «Nature», 603, 280-283 (https://doi.org/10.1038/s41586-022-04448-z).

AUTOMATA 2024, *AUTOMATA: Automated enriched digitisation of archaeological lithics and ceramics* (https://cordis.europa.eu/project/id/101158046).

BARAZZETTI L., REMONDINO F., SCAIONI M., BRUMANA R. 2012, *Fully automatic UAV image-based sensor orientation*, in *International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences*, XXXVIII5/W16, 277-284 (https://doi.org/10.5194/isprsarchives-XXXVIII-5-W16-277-2011).

BINGHAM B., FOLEY B., SINGH H., CAMILLI R., DELAPORTA K., EUSTICE R., MALLIOS A., MINDELL D., ROMAN C., SAKELLARIOU D. 2010, *Robotic tools for deep water archaeology: Surveying an ancient shipwreck with an autonomous underwater vehicle*, «Journal of Field Robotics», 27, 702-717 (https://doi.org/10.1002/rob.20350).

BOGDANI J. 2019, *Archeologia e tecnologie di rete: metodi, strumenti e risorse digitali*, Roma, BraDypUS.

CAMPANA S. 2017, *Drones in archaeology. State-of-the-art and future perspectives*, «Archaeological Prospection», 24, 4, 275-296 (https://doi.org/10.1002/arp.1569).

CASSANELLI C., CORRETTI A., TACCOLA E. 2016, *Rilevamento aerofotogrammetrico da UAV e restituzione 3D dell'edificio fortificato medievale di Entella*, «Annali della Scuola Normale Superiore di Pisa. Classe di Lettere e Filosofia», 8, 2, 104-168.

COLEMAN D., ŞUCAN I.A., CHITTA S., CORRELL N. 2014, *Reducing the barrier to entry of complex robotic software: A MoveIt! case study*, «Journal of Software Engineering for Robotics», 5, 1, 3-16 (https://doi.org/10.6092/JOSER_2014_05_01_p3).

CONTE G., ZANOLI S.M., SCARADOZZI D., GAMBELLA L. 2010, *Robotics tools for underwater archaeology*, in *Science for Cultural Heritage. Technological Innovation and Case Studies in Marine and Land Archaeology in the Adriatic Region and Inland. Proceedings of the VII International Conference on Science, Arts and Culture (Veli Lošinj 2007)*, Singapore, World Scientific Publishing, 187-194 (https://doi.org/10.1142/9789814307079_0016).

EISENBEISS H. *et al.* 2008, *The autonomous mini helicopter: a powerful platform for mobile mapping*, in *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XXXVII, Part B1, ISPRS Congress, Beijing, China, Gotha, Copernicus GmbH, 977-983.

FERNANDEZ-HERNANDEZ J., GONZÁLEZ-AGUILERA D., RODRÍGUEZ-GONZÁLVEZ P., MANCERA-TABOADA J. 2014, *Image-based modelling from Unmanned Aerial Vehicle (UAV) photogrammetry: An effective, low-cost tool for archaeological applications*, «Archaeometry», 57, 128-145 (https://doi.org/10.1111/arcm.12078).

FIORUCCI M., KHOROSHILTSEVA M., PONTIL M., TRAVIGLIA A., DEL BUE A., JAMES S. 2020, *Machine Learning for Cultural Heritage: A survey*, «Pattern Recognition Letters», 133 (https://doi.org/10.1016/j.patrec.2020.02.017).

GATTIGLIA G. 2018, *Classificare le ceramiche: dai metodi tradizionali all'intelligenza artificiale. L'esperienza del progetto ArchAIDE*, in D. MALFITANA (ed.), *Archeologia, Quo vadis? Riflessioni metodologiche sul futuro di una disciplina. Atti del Workshop internazionale (Catania 2018*), Monografie dell'istituto per i Beni Archeologici e Monumentali IBAM, Catania, CNR, 285-298.

GUALANDI M., GATTIGLIA G., ANICHINI F. 2021, *An open system for collection and automatic recognition of pottery through neural network algorithms*, «Heritage», 4, 140-159 (https://doi.org/10.3390/heritage4010008).

JEANSOULIN R., PAPINI O. 2007, *Underwater archaeological knowledge analysis and representation in the VENUS project*, in *AntiCIPAting the Future of the Cultural Past.*

Proceedings of the XXI International CIPA Symposium (Athens 2007), Athens, CIPA. The ICOMOS/ISPRS Committee for Documentation of Cultural Heritage, 388-394.

Kam H.R., Lee S.-H., Park T., Kim C.-H. 2015, *RViz: A toolkit for real domain data visualization*, «Telecommunication Systems», 60, 2, 337-345 (https://doi.org/10.1007/s11235-015-0034-5).

Koenig N., Howard A. 2004, *Design and use paradigms for Gazebo, an open-source multi-robot simulator*, in *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (Sendai 2004)*, Piscataway, NJ, IEEE, 2149-2154 (https://doi.org/10.1109/IROS.2004.1389727).

Kuhn V., Drews-Jr P., Gomes S., Cunha M., Botelho S. 2014, *Automatic control of a ROV for inspection of underwater structures using a low-cost sensing*, «Journal of the Brazilian Society of Mechanical Sciences and Engineering», 37, 1, 361-374.

Olivito R., Emanuele T. 2016, *Tecniche fotogrammetriche e applicazioni di Digital-, Virtual-e Cyber-Archaeology: le Digital Technologies e gli scavi del laboratorio LSA/SAET*, «Annali della Scuola Normale Superiore di Pisa. Classe di Lettere e Filosofia», 8, 97-103.

Olivito R., Taccola E. 2014, *3D Modelling in the agora of Segesta: Techniques and data interpretation*, «Archeologia e Calcolatori», 25, 175-188 (https://doi.org/10.19282/ac.25.2014.09).

Paszke A., Gross S., Massa F., Lerer A., Bradbury J., Chanan G., Killeen T., Lin Z., Gimelshein N., Antiga L., Desmaison A., Köpf A., Yang E., DeVito Z., Raison M., Tejani A., Chilamkurthy S., Steiner B., Fang L., Bai J., Chintala S. 2019, *PyTorch: An imperative style, high-performance deep learning library*, «ArXiv», abs/1912.01703 (https://api.semanticscholar.org/CorpusID:202786778).

RePAIR 2021, *RePAIR: Reconstructing the Past. Artificial Intelligence and Robotics meet Cultural Heritage* (https://www.repairproject.eu/).

Roosevelt C., Cobb P., Moss E., Olson B., Ünlüsoy S. 2015, *Excavation is destruction-digitization: Advances in archaeological practice*, «Journal of Field Archaeology», 40, 325-346 (https://doi.org/10.1179/2042458215Y.0000000004).

Wang D., Lutz B., Cobb P.J., Dames P. 2021, *RASCAL: Robotic Arm for Sherds and Ceramics Automated Locomotion*, in *2021 IEEE International Conference on Robotics and Automation (ICRA) (Xi'an 2021)*, Piscataway, NJ, IEEE, 6378-6384 (https://doi.org/10.1109/ICRA48506.2021.9561057).

Warren D.J., Church R.A., Eslinger K.L. *et al.* 2007, *Deepwater archaeology with autonomous underwater vehicle technology*, in *Offshore Technology Conference (Houston 2007)*, Houston, TX, Offshore Technology Conference, 1329-1339.

Wright H. *et al.* 2019, *Developing the ArchAIDE application: A digital workflow for identifying, organising and sharing archaeological pottery using automated image recognition*, «Internet Archaeology», 52 (https://doi.org/10.11141/ia.52.7).

ABSTRACT

Ceramics, one of the most important sources for reconstructing the past, constitute one of the most common finds in archaeological excavations. The classification of ceramic sherds based on origin, dating, and typology requires a long and sometimes tedious process carried out by specialized archaeologists, who could instead dedicate their time to fieldwork. For this reason, these fragments often remain unclassified in the warehouses of universities and heritage authorities for long periods. AI and robotics can help automate this task: this project proposes a possible solution by presenting the development of a robotic arm that performs pick-and-place operations on fragments classified using ArchAIDE, an AI-based ceramic classification application. The contribution of this article is an automated system which, after recognizing the fragments using the ArchAIDE AI trained on 66 ceramic classes, collects and sorts them into appropriate containers based on the assigned class.