BLOG POST #8 (KCL & MIN, June 2026) - Automating Specialists' Work - Without Replacing Specialists

Decisions, decisions

The digitisation of cultural heritage artefacts involves many decisions. Each individual decision may not seem overly complex to the human specialists who have to make them, but when combined, the whole becomes very complex. Humans get used to new skills very quickly, naturalising what was once unfamiliar into self-evident acts that are carried out without needing to think about them. Consider, for example, riding a bicycle. Think about how many movements, decisions and adjustments need to be made constantly, even before going into traffic. There was a time that you had to think about each of these and it was very complex. But after having learned it, it becomes simple and does not require thought. Now imagine explaining how to ride a bicycle. Next, imagine having to explain how to ride it expertly. Finally, consider having to explain this to something (not someone!) that does not have any experience of moving through space, does not have a body, and does not have a natural way of interacting with the world. You would need to write a small (and very dull!) book to do it.

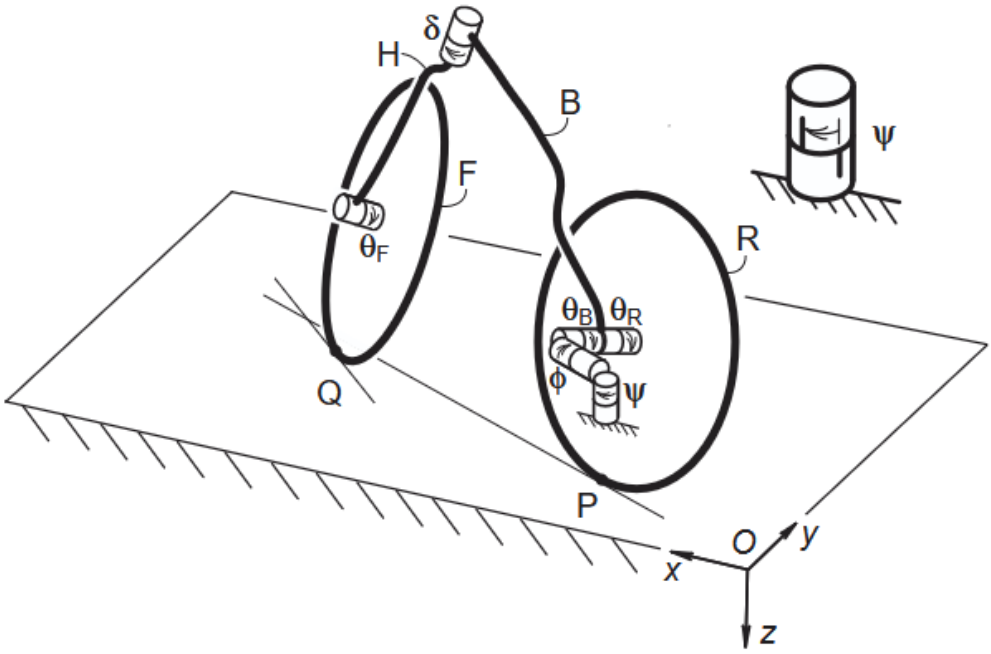


Figure 1: Riding a bicycle is not nearly as easy as experience tells us. (Image from J.P. Meijaard, J.M. Papadopoulos, A. Ruina, and A.L. Schwab 2007. Linearized dynamics equations for the balance and steer of a bicycle: a benchmark and review. Proceedings of the Royal Society A 463: 1955-1982)

AUTOMATA

Funded by the European Union

UK Research and Innovation

This project has received funding from the European Union's HORIZON RIA research and innovation programme under grant agreement N. 101158046

Now think about all the complex decisions that human specialists need to make when digitising archaeological artefacts and generating usable data from them. In a project like AUTOMATA, where large numbers of artefacts must be processed, it quickly becomes clear that manually encoding rules for every possible scenario is simply not feasible. Yet the system still needs to make decisions about instrument calibration, assess the quality of the acquired data, and determine the appropriate course of action.

This is where AUTOMATA's approach comes in. The goal is not to automate specialists out of the process, but to automate the repetitive quality-control tasks they perform every day. By taking over these routine checks, the system allows specialists to focus on interpretation, decision-making, and the more challenging cases where their expertise is most valuable.

Machine Learning

Thankfully, techniques exist to have a machine learn from examples, given a dataset: Machine Learning. Rather than write an exhaustive list of explicit rules, we present the machine with data and the conclusions we want it to draw and set it to work. In principle, this sounds simple enough. In practice, however, it requires careful thought about how to represent information in a way that a machine learning model can process, how to gather and prepare the necessary data, and, crucially, how to enable the model to evaluate its own performance and improve from experience. In AUTOMATA, the model learns to identify the numerical patterns that distinguish high-quality acquisitions from problematic ones using examples that have been labelled by specialists.

In our case, the machine needs to learn how to judge, for each artefact, whether the 3D model associated with it has been generated properly, and whether all the archaeometric data that have been gathered from the artefact are good enough for future analyses and have come from the correct areas on the artefact. In order for the machine to learn what characterises good and bad cases, it needs to be presented with the available information in a way that it can work with. For example, if we take a look at the below image of a 3D model of a ceramic sherd, humans can instantly see that there are problems with it. This is because humans know that you cannot see through pottery like this, and that when it breaks, it does not get weird knobbly textures at the edges.

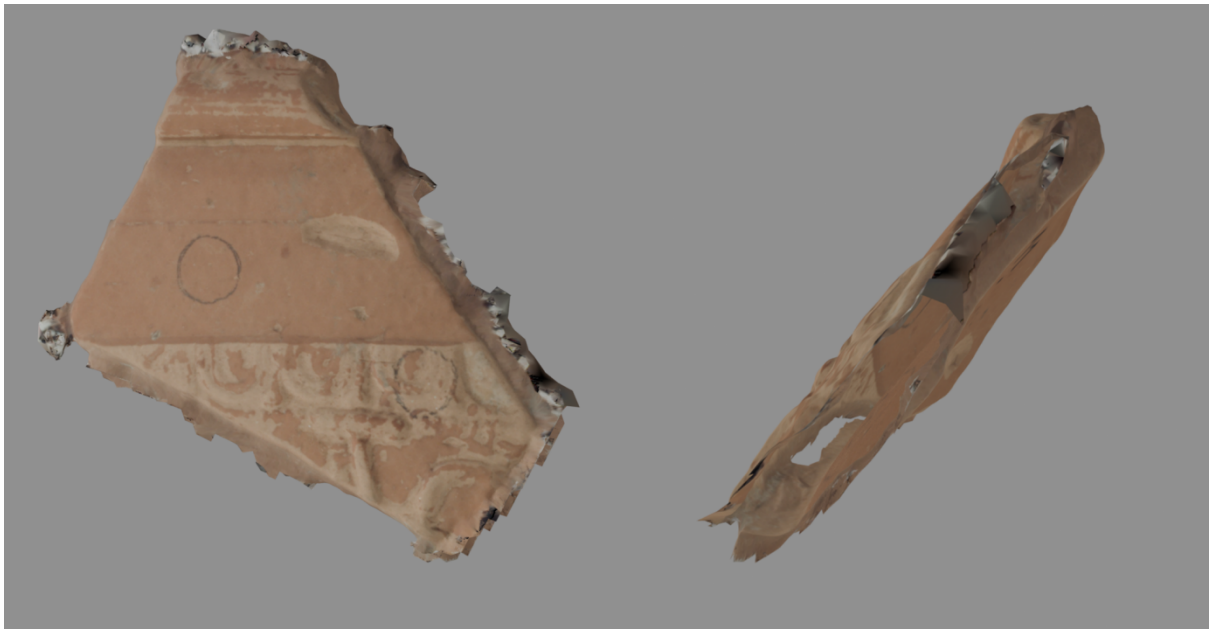


Figure 2: Two views of a 3D model of a ceramic sherd showing problematic holes and knobbly textures at the edges.

The machine does not know this. In fact, the machine has no idea that the image (or the 3D model) is of a ceramic sherd. It has no concept of 'sherd', or even that an image can have something in it. All it sees when presented with an image is a grid of numbers. Rather than train the machine to visually inspect images of 3D models, we try to describe the models numerically. For example, the fact that we can see through the sherd means that there are holes in the mesh that makes up the 3D model's surface. For a normal 3D model of a solid object, there are no holes in the surface mesh (of course this assumes that we are not modelling an object with a hole in it, in which case, we would have a hole in the modelled object, but not in the mesh we model the object with). In fact, every edge is shared by exactly two faces. For there to be a hole in the mesh, an edge will have a single adjacent face. So if we count the number of faces per edge, if there are any that have only one, this means that we have a hole in the mesh.

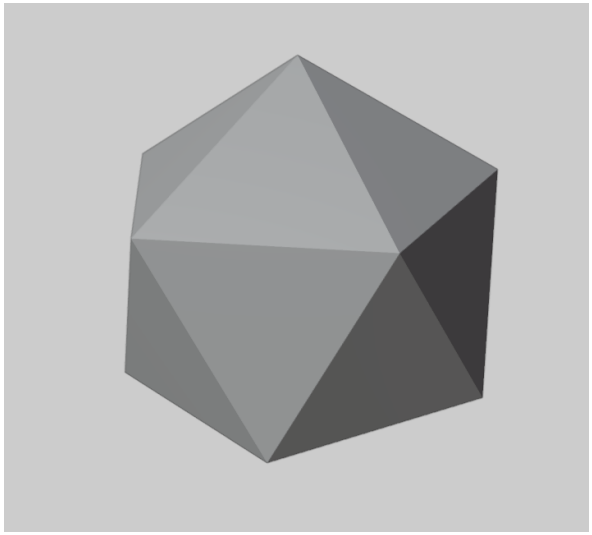


Figure 3: Icosahedron, showing that each edge is shared by exactly two faces

Similar ways of expressing potential flaws with 3D models numerically had to be formulated and extracted for each 3D model. Not every flaw is as straightforward as the presence of holes in the mesh. Therefore, rather than tell the machine 'this is a problem' for each kind of flaw, we defined a number of numerical indicators, called features, and scored each model's on them. This set of features is called a feature vector and each of these describes the 3D model in the way that is relevant to our task. The machine is presented with each model's feature vector together with the information of whether the model in question is a good one, or a bad one. Equipped with this information, the algorithm identifies combinations of features that are consistently associated with good or problematic models. For example: 'if there are many very flat areas on the model and there are holes in the mesh it is a bad model', or 'if transitions between areas are smooth and there are roughly half as many vertices as faces, it is a good model'.

Once a working classification model has been constructed, we can use it to classify data that have not been seen or labelled by a human expert. The model is presented with the features generated from the 3D model and based on these features it generates a prediction of whether this 3D model should have the label 'good' or 'bad'. This predicted label is then used to make automatic decisions about the acquisition pipeline. If the 3D model is considered good, then acquisition can proceed as is. If the 3D model is problematic, the machine can flag this, either to automatically correct the flaws, or to a human operator to intervene. Of course, this is a simplified version of what is actually happening. In reality, we have multiple classification algorithms operating at different stages of the process. Also, while automatic correction sounds simple, in actual fact this may turn out to be very complex.

Clearly, such an approach is not limited to the 3D modelling stage of the acquisition pipeline. For Hyperspectral Imaging, X-Ray Fluorescence, and Raman Spectrometry, similar approaches have been defined. The generated features are different, of course, but the approach is the same. In this way, complex decision processes, that rely on specialist knowledge, can be modelled and automated relatively straightforwardly. This will allow the AUTOMATA project to process large amounts of artefacts while requiring relatively little specialist time (once we have taught the model). Importantly, AUTOMATA is not designed to replace specialists. Instead, it enables them to focus their expertise where it has the greatest impact. The system automates routine checks, flags uncertain or problematic cases, and defers to human judgement whenever intervention is needed. By combining machine learning with expert oversight, AUTOMATA can process larger collections more efficiently while maintaining the high standards required for cultural heritage research.