University of Oxford Introduces a Sector-Leading Image Recognition ML Prototype to Augment Digitization in Numismatics

The **University of Oxford** houses 21 million objects in the collections of its **Gardens, Libraries & Museums** (GLAM)—artifacts and specimens that are among the world's most significant. One aspect of GLAM's mission is to preserve these assets and make them accessible to the world for education and research. However, the organization has only enough space to display about 10 percent of its holdings at a time, and there's an enormous backlog of artifacts still waiting to be cataloged. To tackle that obstacle, GLAM used Amazon Web Services (AWS) to build an enhanced image recognition system that would help accelerate the process of cataloging artifacts.

The Gardens and Museums IT team used **Amazon SageMaker**, a fully managed service that provides developers and data scientists with the ability to build, train, and deploy machine learning (ML) models quickly. Powered by **Amazon Elastic Compute Cloud** (Amazon EC2) instances, the models were trained and deployed at low cost to automatically catalog the extensive coin collection of the Ashmolean Museum—which is the United Kingdom's first public museum and the world's first university museum. On AWS, the image recognition system identifies and catalogs coins in a fraction of the time it would take human volunteers to complete the same task.

Updating Infrastructure and Creating an ML-Based Cataloging System on AWS

GLAM comprises four museums—the Ashmolean Museum of Art and Archaeology, the Oxford University Museum of Natural History, the Pitt Rivers Museum, and the History of Science Museum—as well as the Bodleian Libraries and the Oxford Botanic Garden & Harcourt Arboretum. In 2019, GLAM saw 900,000 visitors to its digital collections. Its 21 million objects include live specimens and plants, historical artifacts, and even images of objects that were damaged, lost, or returned to collectors. "For many years, the museums were not overly active at investing in and managing the information technology infrastructure that underpins all our digital services," says Anjanesh Babu, systems architect and network manager at Gardens and Museums IT. "After years of underinvestment in outdated infrastructure, the University of Oxford brought together a strategic focus on digital transformation through the GLAM Digital program." As part of this program, the Digital Estate Improvement project was the foundational part to deliver root and branch improvements to the infrastructure to make it fit to meet the digital aspirations laid out in the objectives. In 2017, the project uploaded 60 TB of digital records to **Amazon Simple Storage Service** (Amazon S3), an object storage service that offers industry-leading scalability, data availability, security, and performance.

To optimize access to the collections for digital teaching and research, GLAM set its sights on an ML solution that would reduce the time that a research department needed to identify and catalog an object. For this, Anjanesh had to identify a suitable well-cataloged collection that would become the prototyping candidate. **Jerome Mairat**, curator of the Heberden Coin Room in the Ashmolean Museum, had previous experience of developing digital collections from the ground up and offered to support this exploration. The solution would first work with the Roman Provincial Coinage online, a world-renowned research project in numismatics. "I wanted to have a practical example of what we could do to demonstrate to stakeholders," explains Anjanesh. "Coins are a natural carrier of that message and have huge engagement power, so that took us on the journey to engage the data science team at AWS." The first step in machine learning is to decide what you want to predict—in this case, Anjanesh wanted to predict a simple outcome: heads or tails—that is, obverse or reverse. Given a set of known training data, could an ML solution predict the right side of a coin with a high degree of accuracy? This was the much needed segue into the ML world. The outcome went beyond a simple "heads or tails" to much enriched data classification.

Saving Time and Money While Automating Image Processing on AWS

GLAM used AWS to build and deploy 11 ML models on Amazon SageMaker in about 10 weeks. Within that time, research and experimentation took about a month. "The prototyping was rapid and went beyond my expectations," says Anjanesh. Using an initial dataset of 100,000 images, GLAM first used Amazon SageMaker Notebooks to build, train, and experiment with models. Later, it switched to Amazon SageMaker training jobs because many training jobs could be launched simultaneously on Amazon EC2 P3 Instances (powered by NVIDIA V100 Tensor Core GPUs) as Amazon EC2 Spot Instances. Amazon SageMaker managed the training jobs so that they would run when compute capacity became available. By using Spot Instance pricing, GLAM could train the models at 10 percent of the cost of Amazon EC2 On-Demand Instance pricing and in 50 percent less time. And because Amazon SageMaker is framework agnostic, GLAM was able to train the models on PyTorch and implement its own algorithms and scripts.

The resulting image recognition system involves a series of models. First, several ML models conduct image preprocessing so that the image recognition models can produce the best results. For example, to digitize a coin, volunteers—usually university students—photograph each side and then catalog it. If the coin is offset by even 20 degrees, the ML models can't process it as well, so convolutional neural networks are used to turn each image to 90 degrees—a process that Jerome says on its own would save the Ashmolean Museum up to 3 years of work. Another model removes the background from the image of the coin, and yet another uses state-of-the-art adversarial networks to deblur, denoise, and upscale an image to make it high quality. This process is so effective that even pictures from phones are usable.

Once the image is ready, more ML models extract the features of the coin—such as whether it has a cost, is made of metal, or portrays a person—and use this information to find coins with similar features in GLAM's index. A transformer model then generates image captioning or metadata, which is tagged to each image. All models are deployed on Amazon EC2 G4 Instances (powered by NVIDIA T4 Tensor Core GPUs), which cut inference time from minutes to seconds.

The image recognition system is expected to save up to 3 years of work on a collection of 300,000 coins. "It's replacing our production line with an ML model that will improve the curation steps," says Anjanesh. Jerome adds, "Now we can focus our volunteers on other steps that add value. The ML process improves the workflow and productivity and adds value for the public." ML models are expected to lead the way in categorizing GLAM's bulk image datasets for the future.

Analyzing a coin, which previously took volunteers anywhere from 10 minutes to hours, is expected to take just a few minutes once the image recognition system is in place. "If we have ML models against datasets, and if they're responsible for validation and image enhancements, that would not only save staff time but also enable us to educate volunteers and perhaps upskill them to work with such models. The knowledge students can gain from this is another value that we can add to the whole process," explains Jerome.

The image recognition system may also be used to spur visual engagement with visitors. For example, the Ashmolean Museum used to hold object identification sessions, during which people could bring objects and seek curatorial advice to identify these objects and learn their stories. Now, through AWS-powered image recognition, a person could photograph an object and upload that image from home to learn details about it, leading to possibilities of virtual object identification.

Continuing to Make GLAM's Collections More Accessible on AWS

GLAM plans to apply its image recognition system to objects beyond coins, including more complex items like plants, gemstones, butterflies, and other collections. It also wants to share an open-source version of the system with other universities and research institutions on GitHub to foster sharing of datasets as part of the larger strategic road map.

Looking ahead, Anjanesh imagines additional ways ML could improve accessibility to GLAM's collections and streamline internal processes. Currently, website visitors can't use a single search box to look for a specific object across GLAM's collections; they have to visit a separate website for each museum or library. There are aspirations to deliver a cross-collection search function for all of GLAM. Ultimately, Anjanesh envisions an even grander search function that would comb the collections of a multitude of participating universities and heritage partners across the world. Additionally, on a more localized front, ML solutions could be used to monitor and adjust temperatures of galleries to best preserve objects—which would augment the work by collections care teams.

By building an image recognition system on AWS, GLAM has greatly increased access to its collections for students, researchers, and public visitors while saving its staff and volunteers a massive amount of work. "I appreciate how much effort AWS put into this," says Anjanesh. "I thought this project would be complex and time consuming, but using AWS made it easy through common off-the-shelf tools that are portable as well as quick to deliver."

University of Oxford's Gardens, Libraries & Museums

The University of Oxford's Gardens, Libraries & Museums house some of the world's most significant collections. They provide important places of scholarly enquiry and serve as the front door to the wealth of knowledge and research generated at Oxford, welcoming over 3 million visitors each year.

ABOUT

Benefits of AWS

- Built and deployed 11 ML models in about 10 weeks
- Expects to save up to 3 years of work cataloging a collection of 300,000 coins
- Expects coin analysis to take a few minutes versus times ranging from 10 minutes to hours
- Cut inference time from minutes to seconds
- Complements the work already being carried out by volunteers

AWS Services Used

Amazon EC2

Amazon Elastic Compute Cloud (Amazon EC2) is a web service that provides secure, resizable compute capacity in the cloud. It is designed to make web-scale cloud computing easier for developers.

Learn more

Amazon EC2 Spot Instances

Amazon EC2 Spot Instances let you take advantage of unused EC2 capacity in the AWS cloud. Spot Instances are available at up to a 90% discount compared to On-Demand prices.

Learn more

Amazon SageMaker

Amazon SageMaker is a machine learning service that you can use to build, train, and deploy ML models for virtually any use case.

Learn more

Amazon Simple Storage Service

Amazon Simple Storage Service (Amazon S3) is an object storage service that offers industry-leading scalability, data availability, security, and performance.

Learn more

aws