CS486C – Senior Capstone Design in Computer Science

Project Description

Project Title: Automation of Prehistoric Ceramic Recording and Classification	
Sponsor Information: NAU NORTHERN ARIZONA UNIVERSITY Department of Anthropology	Dr. Leszek Pawlowicz Department of Anthropology Northern Arizona University leszek.pawlowicz@nau.edu

Project Overview:

For archaeologists working in the American Southwest, painted ceramic sherds (bits of pottery with painted designs) are one of the most important types of artifacts. Based on the varieties of designs present, they can assign a "type" to the ceramic, which in turn can yield information about who made it, when they made it, and how far did it travel from where it was made. Often, this is the only way to gather that information.

But there's a significant problem with this method. While the definitions of what designs define a ceramic type have been well-established for close to a century, the interpretation of these designs is a highly-subjective process, and archaeologists with decades of experience can disagree on what type should be assigned to a sherd. Our work has found levels of disagreement of up to 48% between experienced archaeologists; this in turn can results in errors in assigning dates and cultural associations to an archaeological site.

In 2021, with my colleagues in the NAU Department Of Anthropology and the Coconino National Forest, I published a paper on using the deep learning method "convolutional neural networks" (CNNs) to classify sherds by type; it's available online, open-access, at:

https://www.sciencedirect.com/science/article/pii/S0305440321000455

In brief, we found that CNNs have the potential to identify painted ceramics by type with an accuracy comparable to, and in some cases superior to, highly-experienced archaeologists. This method offers the promise of bringing objectivity and consistency to the classification of painted ceramics. And we are continuing work to improve the accuracy of CNN models, and expand it to multiple types of ceramics.

The next step is to make this approach accessible and usable to the archaeological community at large, and this is where the Capstone project would be a huge help. There are two major applications for this technology:

1. Archaeologists often do "field surveys", where they walk across large areas looking for signs of archaeological sites. When they find one, they will record their findings, including what types of ceramics are present. However, these field workers often do not have experience with the types of painted ceramics present in an area, leading to classification errors. A self-contained smartphone app that would allow classification and recording of painted ceramics would be a huge help for this. It would have to be self-contained, as this work is often done in remote areas with no Internet connectivity.

2. Archaeological excavations typically produce huge numbers of ceramic sherds, often in numbers so large that some level of "triage" has to be performed, analyzing only a few of the decorated sherds, and not photographing any of them. An automated system that could photograph and analyze hundreds to thousands of sherds would be a boon to archaeology. Even a system that only records the artifacts without

analyzing them would be a terrific help; we have several projects in preparation that would benefit enormously from an automated photography system.

Solution overview

Since the publication of our paper, I have been intermittently looking around for someone to work on a basic smartphone app that uses our CNN model to classify ceramics by type. Not having found anyone, I gritted my teeth, and learned enough Dart/Flutter to create a very simple Android app that seems to work well. It lets the user take a photograph of a sherd, and comes back with a predicted type, along with a confidence factor. If the confidence factor for a type is greater than 0.1 out of 1, it will list all types with confidence > 0.1. The user can also load in a sherd photo from the phone's image gallery, and do a prediction on it.

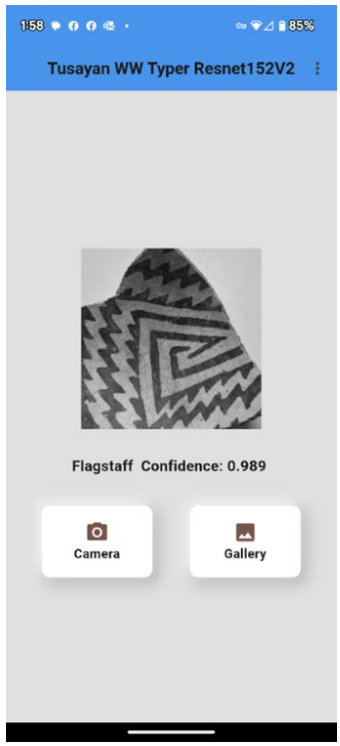

While fine as a test app, the app both lacks a number of desirable features, and also includes some clumsy hacks I needed to get it to work. So I would like to see a number of additions/refinements to the app:

1. The ability to save the sherd photos in full resolution.

2. The saving of metadata associated with the sherd, preferably in CSV format:

    a. The ID of the site where the sherd was found.

    b. Any subsidiary information (e.g. like a specific location within the site).

    c. An abstracted GPS location (not highly accurate, for security reasons).

    d. The filename of the sherd photo.

    e. The classification results from the model, including confidence levels.

    f. The ability to override the top prediction; while the model is as accurate as experts, it does sometimes make obvious mistakes. This override would let the user choose an alternate type, but keep the original predicted type and confidence level in the CSV file.

g. Ultimately, the ability to upload the image files and CSV files to a remote server when the user is back in Internet range

3. Creation of an iOS version of the app. Dart/Flutter is a multiplatform framework, and says that the same code base can be adapted for Android, iOS, Windows, Mac, Linux and Web applications. However, some modifications need to be made to the interface to support these apps.

Optionally, creation of web and OS-based apps; Dart/Flutter also supposedly supports this using a common codebase. I can provide the Dart/Flutter app code I created as an example, but the participants would be free to rip this to shreds to create a better app.

For the automated photography and classification system, I visualize this as a conveyor belt and webcam controlled by a PC. The archaeologist puts sherds on the conveyor belt, the PC/webcam recognizes the sherd, captures an image, and saves that image. Optionally, it could run a CNN to classify that sherd, and save the metadata in a separate file.

Definitely optional: The CNN model used is a standard image classification model, with transfer learning and retraining used to improve its suitability for sherd classification. If someone wanted to investigate whether alternate or custom CNN models might produce better results, I wouldn't object.

Impact of successful product

The impact of this work on Southwest archaeology could be massive. Archaeologists would have a consistent and objective method for classifying ceramics by design type, making data taken by different archaeologists fully comparable. Automating the acquisition of sherd data would increase the amount and accuracy of data that could be recorded. Peter Pilles, chief archaeologist of the Coconino National Forest, has this to say about the app:

"… during a field survey where we compared my identifications with those proposed by the app, I was impressed with the agreement of the app's identifications with my own. I also learned that my concept of where to make a distinction between two types was too restrictive. Upon reflection, I realized that this may explain why sites dating to the periods represented by this type are under-represented in our survey files. … In sum, I highly support the continued development and improvement of this app."

Knowledge, skills, and expertise required for this project:

The skills required for this project are fairly straight forward:

- Programming experience in Python, Dart/Flutter, or other multiplatform mobile development tools, or acquiring those skills for the project
- Interfacing computers with mechanical and visual systems

Equipment Requirements:

- For the smartphone app, the development platform and software/tools are freely available online; students would need to provide their own smartphones, and PCs with the required software platforms (Mac for iOS, PCs with Windows or Linux for Android) for development and testing.
- For the automated acquisition and classification of large numbers of sherds, a stand-alone webcam and conveyor belt would be needed.

Software and other Deliverables:

- A strong as-built report detailing the design and implementation of the product in a complete, clear and professional manner. This document should provide a strong basis for future development of the product.
- Complete professionally-documented codebase, delivered both as a repository in GitHub, BitBucket, or some other version control repository; and as a physical archive on a USB drive.

- Upon completion of the project, I anticipate writing a paper on the project, to be submitted to a peer-reviewed journal (most likely Advances In Archaeological Practices). All of the project participants would be included as co-authors, and their roles explicitly singled out in the publication.