Proposal for CS109B projects: Visual recognition of images on ancient Greek vases

Participants:
Natasha Bershadsky, Lecturer in the Department of the Classics
nbershadsky@fas.harvard.edu

The project proposes to develop a model for visual recognition of images on ancient Greek black-figure and/or red-figure vases. The goal of the project is to enable intelligent sorting of large quantities of images on vases into meaningful categories, according to objects depicted and compositions of the scenes. Such an application will propel forward research related to the vase-painting imagery, giving us new tools for parsing the pictorial language of the vases.

This is an ambitious goal, but it would be valuable to make progress in providing a partial solution. To explain the problem, let me quickly describe relevant traits of the ancient Greek vase painting.

Greek vase painting is a highly uniform field, where representations are extremely consistent, and multiple visual stable expressions ("formulas") exist on the level of the individual objects represented or the composition of the scenes (see fig. 1). To understand a particular scene on a vase, it is most useful to find its "siblings," similar scenes that differ in a few details. In this way we start to understand the system of representations to which a given scene belongs.

Massive quantities of images of Greek vases are already available in open access databases (our archive was built on their basis). These image collections are searchable; however, the search operates according to verbal tags: for example, if a verbal description of a vase contains the word "shield," such a vase will appear in a search for "shield (as does Image A);" if the verbal tag is lacking, the vase will be lost to the search (Image B). Further, there is no simple way to find all scenes that would feature, for example, a warrior in the middle, surrounded by figures that face him: such a search needs to be performed 'by hand.'

Fig. 1: Images A and B

The project invites you to work on creating models that would sort images on vases on the basis of their visual characteristics. Different approaches can be explored: for example, one can focus on training a classifier to identify a round shield. We are very interested in such object-based

image recognition studies. Alternatively, it would be most useful to build a model that could identify ‘siblings’ of a given vase, even without identifying any particular images in the scene.

Technical notes:

Images in both black-figure and red-figure pots were created through sharp contrasts (black images on light backgrounds for black figure, and light images on black background for red figure). Details are added by incisions (black figure) or lines (red figure). The images lack three-dimensional shading cues on which some of the visual recognition tools rely.

One particular problem is that the images in the database are not standardized: they contain both closeups of the painted scenes, and photographs of whole vases.

Database:

Our database is the armsandarmor archive, a database of more than 100,000 images of ancient Greek vases indexed and organized using machine learning to search by color, feature recognition, and verbal tags. It's based on software developed by Jeff Steward at the Harvard Art Museums and deployed on an autoscaling postgres instance on Google Kubernetes Engine.

The interface can be explored at https://armsandarmor.orphe.us/ or crawl the API at https://armsandarmor.orphe.us/api/items; the data can be also exported.