Digitization and Reconstruction of Archaeological Artifacts

Helena Cristina da Gama Leitão
Instituto de Computação
Universidade Federal Fluminense
hcgl@ic.uff.br

Jorge Stolfi
Instituto de Computação
Universidade Estadual de Campinas
stolfi@ic.unicamp.br

The concrete problem

Plain ceramic fragments are often found by the thousands in achaeological digs. Those finds probably include many pieces that could be reassembled into complete artifacts. However, that task is hardly ever attempted, due to the large cost of trying millions of possible pairs by hand.

The computational problem

The input data are $N \approx 10^4$–10^5 irregular curves, each defined by $L \approx 10^3$ sample points. The goal is to find any pairs of curves that have long sections with sufficiently similar shapes.

Prior work

At present, computerization of this task is limited to sorting fragment images according to general attributes such as color and texture. For plain pottery, this approach is not effective: one needs algorithms that can find matching pieces by comparing their shapes.

Is there is enough data?

Matching fragments typically fit together with sub-millimeter accuracy, therefore false matches are fairly rare. Statistical analysis of a sample of freshly fractured ceramic shows that one centimeter of fragment outline gives at least 18 bits of useful information about the shape of the matching fragment.

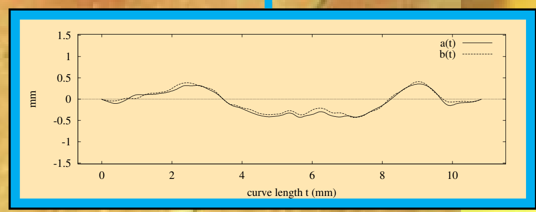

Multi-scale matching

We find matching fragment outlines by multi-scale comparison of the curvature graphs. We filter and resample the outlines at various scales of detail, and look for promising candidate pairs at the coarsest scale, where each outline has only a few tens of sample points. We then repeatedly select the most promising pairs, and compare them again at the next finer scale of detail. In the end, we are left with a small set of fragment pairs that are most likely to be adjacent in the original objects.

Incremental matching

To compute the optimum pairing of samples for a given candidate, we use an incremental variant of dynamic programming (DP), that considers only pairings similar to the one found at the previous (coarser) scale. This modified version reduces the cost of DP from quadratic to linear in the number of samples, and removes a factor of L from the total cost.

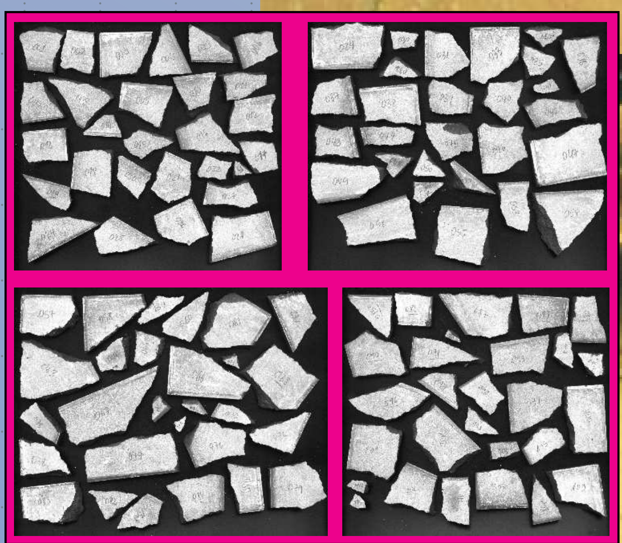

Experiment

Five unglazed ceramic tiles, each measuring about 25×6 cm, were shattered with a steel bar into 112 fragments measuring 1 to 5 cm (bottom right) which were digitized on a 300 dpi document scanner.

Results

In a typical run, the program listed 42 correct pairs (blue lines), among the 54 most promising candidates; and failed to find 28 correct pairs.

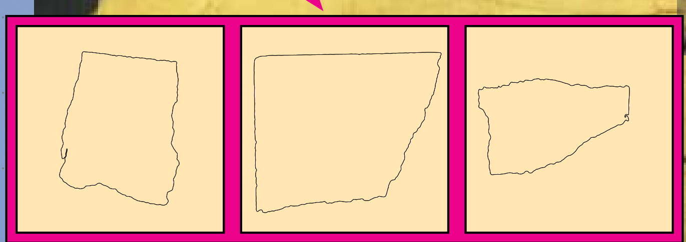

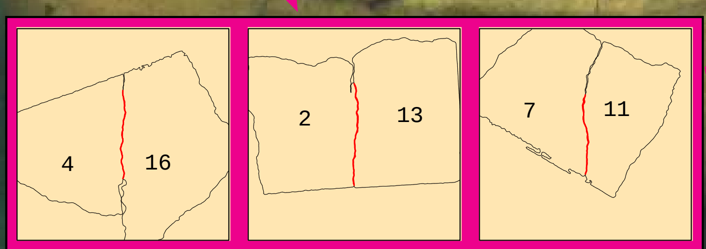

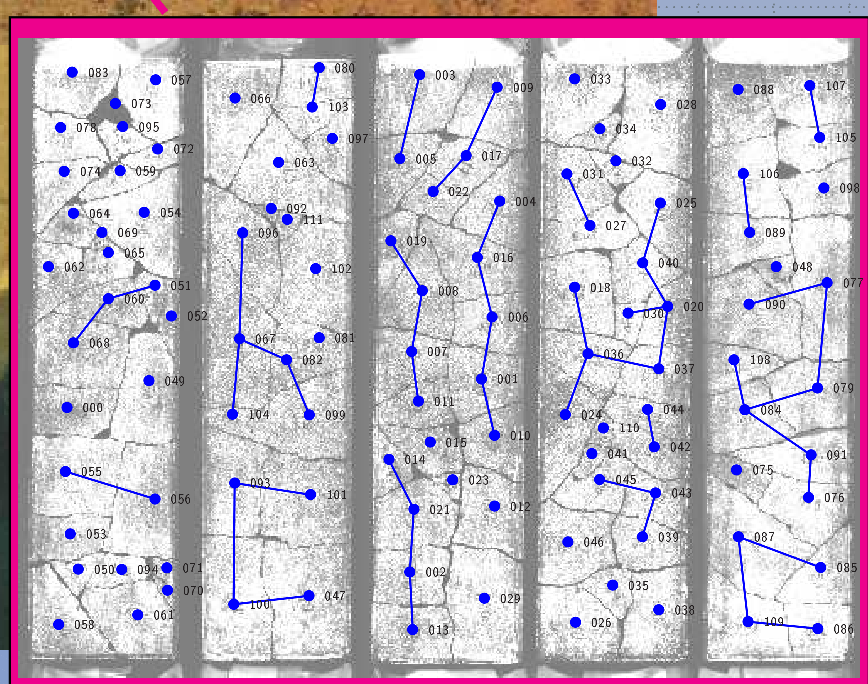

What next?

We are now extending our tools to curved pottery fragments. We plan to use structured-light stereo vision to detemine the true (view-corrected) curvature at each point along the fracture line.

Seeking collaboration

We seek collaboration with potential users, especially archaeological museums with substantial collections of pottery fragments.

Support: CAPES, CNPq, FAPERJ, UFF, Unicamp.