Ancient Ceramics: Computer aided Classification

Dorrit Porter[1], Peter Werner[2] und Sven Utcke[1]

[1] Arbeitsbereich Kognitive Systeme, Fachbereich Informatik, Universität Hamburg, Germany
{6porter,utcke}@informatik.uni-hamburg.de
[2] Arbeitsbereich Stadtbaugeschichte
Technische Universität Hamburg-Harburg, Germany
p.werner@tu-harburg.de

Zusammenfassung. Die Klassifikation von Gebrauchskeramiken, die in fast allen Ausgrabungen rund um die Welt zu den zahlreichsten Funden gehören, ist eine zeitaufwendige und mühevolle Arbeit. Viele der dem Anschein nach uninteressanteren Funde wandern deshalb unklassifiziert ins Archiv. Die automatische Klassifizierung von Gebrauchskeramiken wäre von daher eine willkommene Hilfe für viele Archäologen, die sie von Routinetätigkeiten entlasten könnte. Darüberhinaus würde die automatische Klassifikation von Funden über mehrere Ausgrabungsstätten hinweg eventuell zu ganz neuen Einsichten führen — Schlussfolgerungen, die mit dem traditionellem Ansatz der manuellen Klassifikation nur schwer zu gewinnen gewesen wären. In dieser Mitteilung beschreiben wir einen ersten Prototypen für die Klassifikation innerhalb eines solchen Systems und zeigen, dass bereits mit einem recht willkürlichem Satz an Merkmalen und einem untrainiertem System eine sinnvolle Klassifikation möglich ist.

Keywords: Archäologie; Mustererkennung; rotationssymmetrische Objekte

Ancient Ceramics: Computer aided Classification

Dorrit Porter[1], Peter Werner[2], and Sven Utcke[1]

[1] Arbeitsbereich Kognitive Systeme, Fachbereich Informatik,
Universität Hamburg, Germany
{6porter,utcke}@informatik.uni-hamburg.de
[2] Arbeitsbereich Stadtbaugeschichte
Technische Universität Hamburg-Harburg, Germany
p.werner@tu-harburg.de

Abstract. The classification of pottery as found in many archaeological sites around the world is a time consuming and often tedious task which can take months and even years for a single site. Many of the seemingly less interesting artifacts simply disappear into storage unclassified. Automating the classification of pottery would hence be a welcome help for many archaeologists, freeing them from routine work. In addition the routine classification of all finds across several sites might also yield new and interesting results which would not have been possible using traditional approaches. This paper describes a first proof of concept for the classification machinery inside such a system and demonstrates that even with a rather haphazardly chosen set of features and an untrained system reasonable classification results are possible.

Keywords: archaeology; pattern classification; rotationally symmetric objects

1 Introduction

Ceramics usually have a short period of life, while at the same time being subject to changes in fashion and style. They are also among the most numerous artefacts in most excavations. All this makes them uniquely suited when it comes to dating strata, that is to distinguish between chronological and ethnic groups and to create a chronological frame of reference.

The classification of pottery as found at excavations the world over is traditionally done by archaeologists right at the site and, all too often, according to a system applicable only to this site, and by this archaeologist. This has not remained unnoticed by neither archaeologists nor computer scientists, and both have come up with a number of systems to remedy the situation. The archaeologists have done so as early as in the mid-fifties, and have mostly concentrated on the creation of a general vocabulary to be used in classification [1,2,3,4]; however, no definite scheme has surfaced so far. Computer scientists, on the other hand, have come to the field comparatively late and have mainly concentrated on the automatic documentation [5,6], skirting the difficult (and often subjective) issues of classification (but see also [7,8]).

In our work (but not this article) we will in contrast concentrate on classification. We are directly building on the work of archaeologists, trying to unify the existing classification schemes, which rely on a set of non-orthogonal and quite possibly highly redundant features, using the tools of traditional pattern recognition. To this end a wide variety of "traditional" features will be implemented and then reduced to a minimum set of features which is still able to classify pottery as good as an average archaeologist. The rational behind this effort is to free archaeologists from routine work, but also to allow easier comparisons across different sites and cultures, which might well allow new insights, e. g. into trading routes, cross-cultural relationships, and patterns of migration.

This article describes the preliminary steps of such a system, implementing 21 of the most commonly used (and most easily programmed) features from the pertinent archaeological literature. At present, all of the implemented features are based on the assumption of a rotationally symmetric object's profile, which is a cross-section of the object in the direction of the axis of symmetry; other important features such as decoration, material, colour or the manufacturing process will be added at a later stage. The 21 currently implemented features are described in Section 2. Classification is currently done using a very simple, untrained, k-means classifier described in Section 3. For testing purposes we then applied this classification to 30 vessels from the Late Bronze Age settlement of Tall Munbaqa/Syria, the ancient Ekalte. The results, which are described in Section 4, were surprisingly good and prompted this early publication. However, they also show the way for further improvements, and these, together with our intended research direction, are discussed in Section 5, which concludes this article.

2 Features used

In the following we will give a short description of the 21 features currently implemented — 15 general features and 6 for the vessel's neck — and cite the relevant publications from which they were taken. Since most of these features are rather self-explanatory we will describe in more detail only those features where the meaning or implementation isn't straightforward. Since all features are based on the object's profile, we do not distinguish between individual feet or a ring supporting the object. The 15 general features, many of which are depicted in Figure 1, are:

1. Height in centimetres (including a possible stand or pedestal support).
2. Width in centimetres.
3. Height to width ratio [1,2,8,9].
4. Relative height of the widest point [1,2,9].
5. Relative width at top (highest point) [1,2,3,8,9]. Also called orifice or mouth.
6. Relative width at base. This is the maximum width at which the object touches the ground relative to the object's maximum width.
7. Relative average width [1].

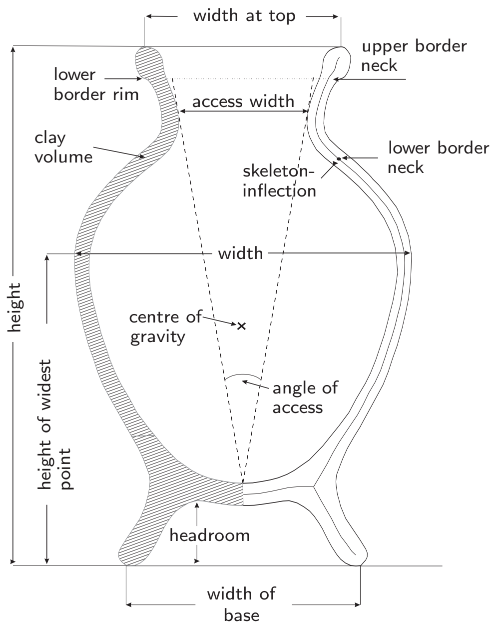

Fig. 1. Some of the features currently implemented

8. Relative headroom, similar to [2]. This is the maximum clearance of a concave base relative to the object's height.
9. Clay efficiency [8]. This is the ratio of a vessel's capacity relative to the volume of the clay.
10. Relative centre of gravity [1]. This is based on the assumption of a homogeneous density of the clay and set in relation to the object's height.
11. Relative access width, similar to [1]. This is the width of the inner point with the biggest (inner) height-to-width ratio relative to the maximum width, see Figure 1.
12. Angle of access. This is the maximum angle under which one can directly reach the middle of the vessel's bottom, calculated from the access width.
13. Mean relative wall thickness. Computed along the dominant skeleton arm (by doubling the minimum distance to the next contour point), and relative to the diagonal of the smallest enclosing rectangle (to incorporate both very wide as well as very tall vessels). Feet and other decorations and ornaments lead to a slightly overestimated value.
14. Type of rim, similar to [3]. 0 for a simple rim, 1 for a thickened rim, 2 for a more complex rim.
15. Skeleton-complexity. The number of additional skeleton arms (belonging to at least 5 profile points).

The 6 features specific to the neck, which is the section between the body and the rim (see Figure 1) are:

16. Existence. Similar to [2] the main criterion is a (significant) inflection in the top 50 % of the skeleton. Also one of the main features in [1].
17. Relative length of neck. This is the length of the skeleton between the beginning of the neck and the beginning of the rim (which ends the neck) relative to the height of the vessel.
18. Relative height of neck I. This is really the height of the inflection relative to the overall height.
19. Relative height of neck II. The height of the inflection relative to the height of the widest point below the neck.
20. Relative width of neck I. This is the minimum width of the neck relative to the maximum width of the vessel.
21. Relative width of neck II. This is the minimum width of the neck relative to the width at top.

3 Classification

As a very first test we used essentially unweighted data and a simple k-means algorithm. So as not to give any measurement unwarranted weight we normalised the standard deviation of all measurements to 1. We used Euclidean distance, random placement of seeds and between 50 and 500 runs for each k. For a set of 240 different vessels which can be hand-segmented into approximately $k = 80$ subsets, many of which are formed by only one vessel, we found no straightforward way to automatically determine this k.

However, when we compared the results of segmenting 30 vessels into a given number of $k = 9$ sets with the results from a previous hand-classification of the same vessels, we found that the results were mostly surprisingly accurate and even partly surpassed the manual classification. This was all the more surprising as no training was involved; the classification was done on an unmodified, i. e. non-orthogonal, redundant and arbitrarily scaled, set of measurements. The results are given in the next section and provide a lower bound for what can be expected after orthonormalisation and training.

4 Results

In this test, 30 out of 240 vessels were manually classified by an expert, in a similar manner, with similar accentuation on the more interesting vessels and under similar time constraints as exist for real excavations. This classification resulted in 9 separate clusters. We then ran a simple k-means algorithm with $k = 9$ over the same 30 vessels. This resulted in 2 clusters which were identically grouped by both manually as well as the machine segmentation (comprising 6 vessels, or 20 % of all objects, see Figure 2), 2 archaeologically untenable classifications and 1 archaeologically questionable classification (also affecting 6 vessel, 20 %, see Figures 5 and 6), but also 4 clusters where the machine classification actually improved on the manual classification (affecting 18 vessels, 60 %, see Figures 3

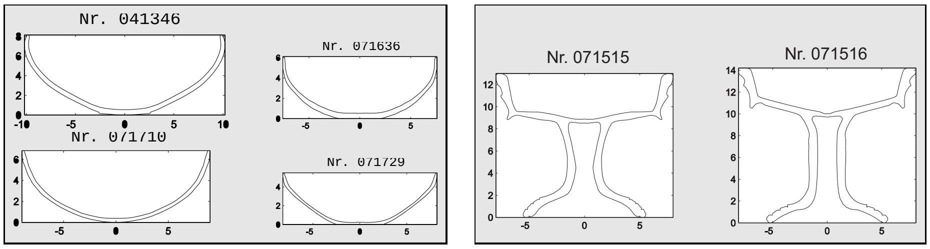

(a) Man. and mach. Cluster 2. (b) Man. and mach. Cluster 7.

Fig. 2. Clusters shared by the manual and machine classification.

(a) Manual cluster 3.

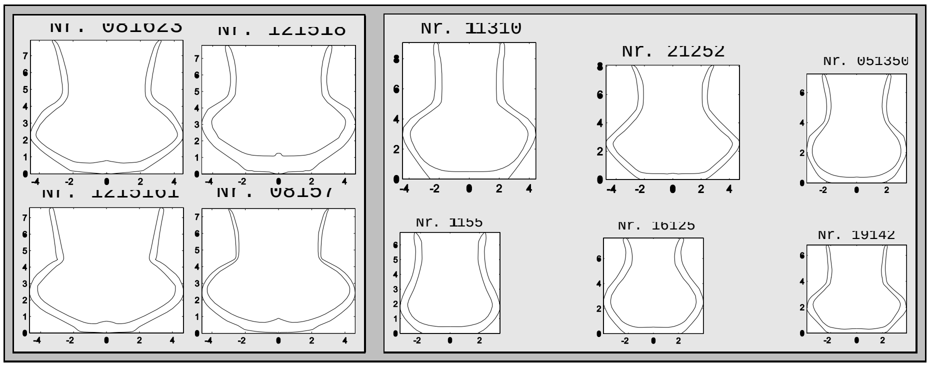

(b) Machine cluster 1. (c) Machine cluster 8.

Fig. 3. Machine-Clusters 1 and 8 comprise the manual cluster 3.

(a) Manual cluster 4. (b) Man. cluster 8.

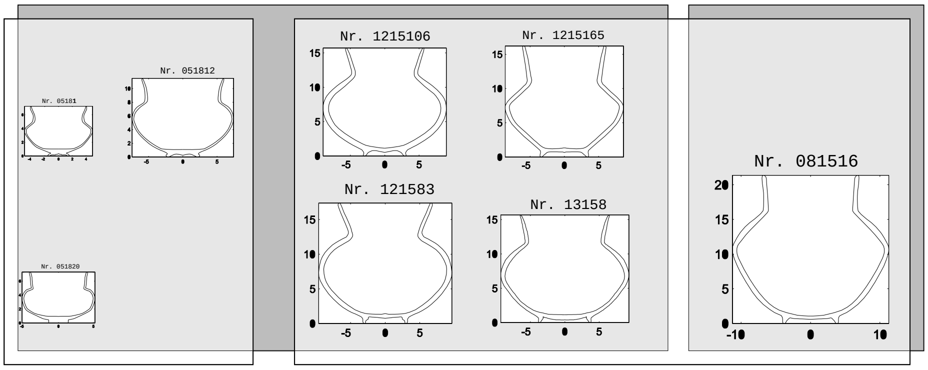

(c) Mach. cluster 5. (d) Machine cluster 6.

Fig. 4. Manual clusters 4 and 8 (dark) are comprised of the same objects as the machine clusters 5 and 6 (light).

(a) Manual cluster 9.

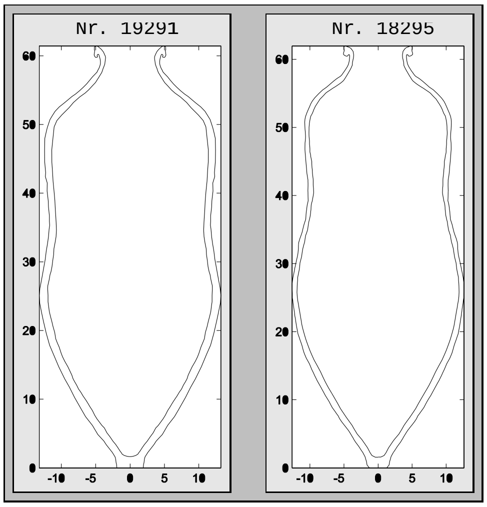

(b) Machine 4. (c) Machine 9.

Fig. 5. Manual cluster 9 is erroneously split into two separate clusters.

and 4). This result was consistently reached in 3 trials with 500 random placements of seeds each. We will now analyse these results in more detail.

The two clusters in Figure 2, which were grouped identical by both manual as well as machine classification, are both visually very distinctive, at least within this particular set of 30 vessels, were they are the only bowl-shaped vessels or vessels with a particularly high stand respectively. It is therefore not surprising that these objects were similarly grouped by both the human expert as well as the machine.

More interesting are the manual clusters 3 (Figure 3), and 4 and 8 (Figure 4), which get redistributed into the 4 machine-generated clusters 1 and 8 (Figure 3), and 5 and 6 (Figure 4). Here the redistributed clusters have been judged more appropriate or at least as good as the manual classification by our human expert. This shows that even with such an extremely simple classifier as the k-means algorithm we are already able to support archaeologists in their daily practice (at least with such a small number of objects as were tested here).

However, for the remaining 4 manual clusters (cluster 9 and 1, 5, 6 in Figures 5 and 6) or 4 machine-clusters (cluster 9, 4, and 3), afflicting 6 vessels altogether, we get results which are clearly unacceptable from the point of view of an archaeologist. This is immediately obvious even for a layperson in the case of the erroneously subdivided manual cluster 9 in Figure 5; the two vessels are indeed quite close in feature-space and get combined as soon as we set $k = 8$. It is therefore all the more interesting that the archaeologically quite distinct man-

(a) Manual cluster 1. (b) Manual 5. (c) Manual 6.

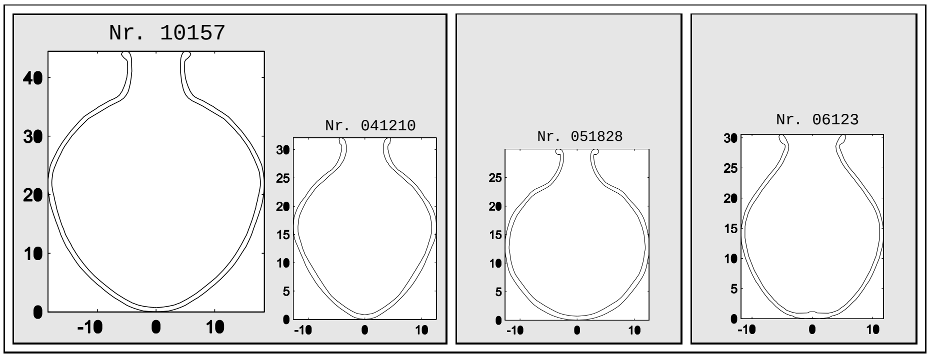

(d) Machine cluster 3.

Fig. 6. Manual clusters 1, 5, and 6 get erroneously combined to machine cluster 3.

ual clusters 1, 5, and 6 in Figure 6 get combined into just one machine-cluster. Obviously these 4 objects are considered closer in appearance to each other than the two (very similar) objects in Figure 5. This is indication both of a weakness of our simple k-means algorithm (obviously some archaeologically uninteresting differences between the two objects in Figure 5 get assigned undue weight), but also indicative of the fact that our current set of 21 features still misses some important differentiations. The main difference between the 3 manual clusters in Figure 6 are the form of the base (conoidal (pointed) base versus globular (rounded) or flat base) but in particular the overall form of the objects, which are sort of heart-shaped versus spherical versus oval. None of this (except for the flat base in Figure 6(c)) can currently be captured by our set of features, which apparently need to be extended by some curvature based properties, or possibly by approximation of the vessel-shape with known geometric objects as was, e. g., done in [4].

5 Conclusion and outlook

In this paper we presented a selection of 21 different features for the classification of rotationally symmetric ceramics which we selected from the pertinent archaeological literature. We have demonstrated that based on these features, and for a small sample of 30 different vessels from the Late Bronze Age, even a very simple, untrained classifier already achieved classification results which in many cases rival the results from a manual classification by an expert. However, we have also seen that for some shapes the classification is still suboptimal.

This suggests a number of different courses for our future work. As a very first step we certainly need to apply our algorithm to many more vessels; the main constraint here is the time it takes a human expert to come up with a usable manual classification. Next we need to train our classifier in order to reach a more satisfactory weighting of the individual features. A particular problem in this context are the Features 14, 15, and 16, which are discrete, while all other features are continuous.

Once a trained classifier can be applied to more objects, we will be better able to assess what other features are needed; the results in Section 4 already suggested that we will probably need features which can capture the overall form of the object, be it curvature based or by fitting geometric entities.

The ultimate goal then will be a detailed analysis of parameter-space — maybe there is a natural division of parameter-space according to basic form (amphora versus beaker and so on)? Or even a natural partition of the parameter-space that allows (block) diagonalisation of (derived) features? Maybe a PCA will come up with new, meaningful features? All this should provide new insights for both computer scientists as well as archaeologists.

References

1. Shepard, A.O.: Ceramics for the Archaeologist. 5 edn. Carnegie Institution, Washington D. C. (1965)
2. Karstens, K.: Allgemeine Systematik der einfachen Gefäßformen. Volume 16 of Münchener vorderasiatische Studien. Profil Verlag, München, Wien (1994)
3. Bauer, I., Endres, W., Kerkhoff-Hader, B., Koch, R., Stephan, H.G.: Leitfaden zur Keramikbeschreibung (Mittelalter–Neuzeit). 2 edn. Kataloge der prähistorischen Staatssammlung. Verlag Michael Lassleben Kallmünz/Opf., München (1993)
4. Ericson, J.E., Stickel, E.G.: A proposed classification system for ceramics. World Archaeology **4** (1973) 357–367
5. Adler, K., Kampel, M., Kastler, R., Penz, M., Sablatnig, R., Schindler, K., Tosovic, S.: Computer aided classification of ceramics - achievements and problems. In: Proc. of 6th Intl. Workshop on Archaeology and Computers, Vienna, Austria (2001)
6. Kampel, M., Sablatnig, R., Costa, E.: Classification of archaeological fragments using profile primitives. In Scherer, S., ed.: Computer Vision, Computer Graphics and Photogrammetry - A Common Viewpoint, Proc. of the 25th Workshop of the Austrian Association for Pattern Recognition (OEAGM). Volume 147 of Schriftenreihe der OCG., Oldenburg, Wien, München (2001) 151–158
7. Kampffmeyer, U.: Untersuchungen zur rechnergestützen Klassifikation der Form von Keramik. Volume 11 of Arbeiten zur Urgeschichte des Menschen. Peter Lang, Frankfurt am Main, Bern, New York, Paris (1988)
8. Streckner, C.: Das SAMOS-Projekt: Neue Wege der Informatikanwendung in der Archäologie. Archäologie in Deutschland (1989) 16–21
9. Steuer, H.: Die Südsiedlung von Haithabu. Volume 6 of Die Ausgrabungen in Haithabu. Karl Wachholz Verlag, Neumünster (1974)