A Framework for Design Identification on Heritage Objects

Jun Zhou*
Yuhang Lu*
zhouj@mailbox.sc.edu
yuhang@cec.sc.edu
University of South Carolina
Columbia, South Carolina

Karen Smith
S.C. Department of Natural Resources
Columbia, South Carolina
smithky@drn.sc.gov

Colin Wilder
University of South Carolina
Columbia, South Carolina
wildercf@mailbox.sc.gov

Song Wang
University of South Carolina
Columbia, South Carolina
songwang@cec.sc.edu

Paul Sagona
University of South Carolina
Columbia, South Carolina
sagona@sc.edu

Ben Torkian
University of South Carolina
Columbia, South Carolina
torkian@email.sc.edu

ABSTRACT

A challenging problem in modern archaeology is to automatically identify fragmented heritage objects by their decorative full designs, such as the pottery sherds from Southeastern America. The difficulties of this problem lie in: 1) these pottery sherds are usually fragmented so that each sherd only covers a small portion of its underlying full design; 2) these sherds can be so highly degraded that curves may contain missing segments or become very shallow; and 3) curve patterns may overlap with each other from the making of these potteries. This paper presents a deep-learning based framework for matching a sherd with a database of known designs to find its underlying design. This framework contains three steps: 1) extracting curve pattern using an FCN-based curve pattern segmentation method from the digitized sherd's depth map, 2) matching a sherd with a non-composite (single copy of a design) pattern combining template matching algorithm with a dual-source CNN re-ranking method to find its underlying design, and 3) matching a sherd with a composite (multiple copies of a design) pattern using a Chamfer Matching based method. The framework was evaluated on a set of sherds from the heartland of the paddle-stamping tradition with a subset of known paddle-stamped designs of Pre-colonial southeastern North America. Extensive experimental results show the effectiveness of the proposed framework and algorithms.

KEYWORDS

Cultural heritage protection, curve pattern extraction, curve pattern matching, composite curve patterns, neural networks

1 INTRODUCTION

All around the world, the archaeological record is filled with fragmentary objects of bone, pottery, shell, stone, wood, and cloth variously embellished with realistic and abstract designs. These designs may include figural imagery like that seen on ancient Maya [24] or the carved marine shell gorgets of late prehistory in North America [22]. Humanities and social science scholars have put these designs to many uses, such as building chronologies, tracking trade networks, reconstructing aspects of style and the creative process, and understanding the creation and expression of identity.

Without question, most of these topics are best addressed using complete designs rather than design fragments. Traditionally, complete designs are composed using whole artifacts; fragments of designs are then identified as belonging to complete compositions manually by visual assessment. The task of matching design fragments to whole designs can be highly time-consuming. As a result, millions of broken cultural heritage objects stored in museums around the world remain unstudied from a design perspective, and large numbers of decorated objects found in the archaeological record contribute little to our understanding of style, production and use, and meaning.

Computer-aided identification of the designs from fragmented cultural objects has attracted great interest among archaeologists and computer scientists in recent years [10, 12]. In this paper, we take pottery sherds found on archaeological sites in the heartland of the paddle-stamping tradition of southeastern North America as our case study, and develop a framework to identify the underlying carved wooden paddles impressed on pottery from the Carolinas to the Gulf Coast. The elaborately carved wooden paddles of the Southeastern Woodlands, a small fraction of which are shown in Figure 1. And the ornate curvilinear paddle impressions on countless pottery sherds of the Swift Creek style tradition made ca. AD 350 to AD 650, shown in Figure 2 frames our study case.

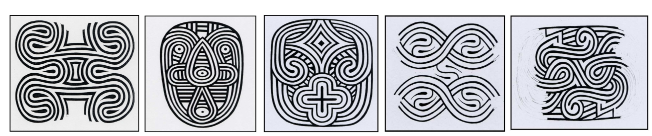

Figure 1: Five paddle designs reconstructed by Frankie Snow. Original design reproduced with permission, courtesy of Frankie Snow, South Georgia State College.

Identifying the full curvilinear paddle design from fragmentary sherds is a highly challenging problem. First, each sherd only contains a small portion of the underlying full paddle design. Second, the available sherds rarely come from the same vessel, and it is difficult to assemble them into large pieces for more complete curve patterns. Third, one carved paddle may be applied multiple times on

*Both authors contributed equally to this research.

PEARC '19, July 28-August 1, 2019, Chicago, IL, USA

DOI: 10.1145/3332186.3332190

the pottery surface with spatial overlap. As a result, curve patterns detected on sherds may be incomplete or very noisy due to the gap when applying a planar carved paddle onto a curved pottery surface and to the erosion of sherd surfaces over thousands of years. Furthermore, a sherd may contain a *composite* pattern, i.e. 2.(b), a small fragment of multiple, partially overlapping copies of the same design. Such a composite pattern is not simply a portion of the full design. Therefore, matching it to the underlying full design is not a simple partial-to-global matching problem [5].

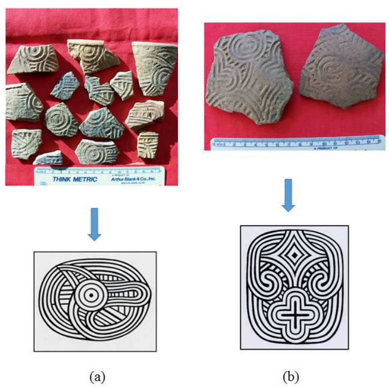

Figure 2: Sample pottery sherds (top) and their underlying wooden paddle designs (bottom). Two pottery sherds in (b) contain a composite pattern, resulting from the multiple applications of the carved paddle with partial spatial overlaps. Original designs reproduced with permission, courtesy of Frankie Snow, South Georgia State College.

In this paper, we develop a new framework for identifying carved paddle designs from pottery sherds by addressing these challenges. More specifically, we extract the curve pattern from a sherd and then match it to each known design in a database and return the best matched design. As shown in Fig 3, this framework contains three steps: 1) extract a curve pattern from a sherd, 2) identify the underlying design for a sherd with a non-composite pattern, or 3) identify underlying design for a sherd with a composite pattern. In particular, we extract curve pattern using an FCN-based curve pattern segmentation method from a digitized sherd's depth map, and then match the sherd with a non-composite pattern by combining template matching algorithm with a dual-source CNN re-ranking algorithm to find its underlying design, or match the sherd with a composite pattern using a new Chamfer Matching algorithm.

In our experiments, we evaluate the proposed framework on a set of sherds from the heartland of the paddle-stamping tradition with a subset of known paddle-stamped designs of pre-Colonial southeastern North America. Our result shows our algorithms are much better than several other state-of-art matching algorithms.

The remainder of this paper is organized as follows. Section 2 reviews the related work. Section 3 introduces the proposed framework including new algorithms for curve pattern extraction and matching the curve pattern on a sherd with known designs. Section 4 introduces the collected test data and the experiment results. Section 5 discusses the future work that will improve our algorithms and followed by a brief conclusion in Section 6.

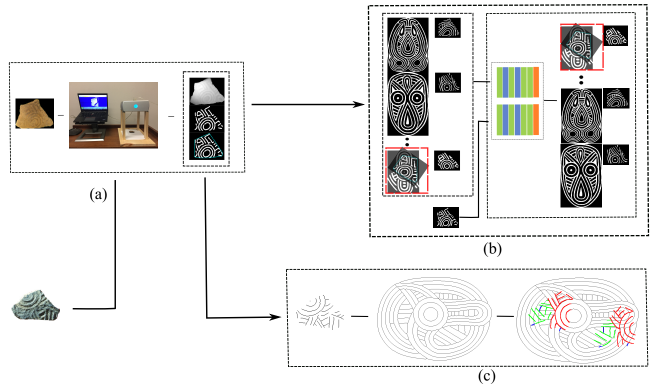

Figure 3: An illustration of a framework on identifying the underlying design for a sherd. (a) Extract a curve pattern from a sherd. (b) Identify the underlying design for a sherd with a non-composite pattern. (c) Identify the underlying design for a sherd with a composite pattern. Original design reproduced with permission, courtesy of Frankie Snow, South Georgia State College.

2 PREVIOUS WORK

Many previous works on computer-aided processing of archaeological fragments, such as pottery sherds, were focused on classifying whether different fragments come from the same vessel and the classification results are then used to aid the 3D reconstruction of the underlying whole object, such as a full vessel [8, 9]. Color and texture information have been widely used for fragment classification [20, 23]. Many geometric features were also used to classify archaeological fragments [13, 14, 25]. Several previous works [7, 8] are focused on developing algorithms to assemble sherds into larger pottery pieces, or the whole vessel, by fitting the boundary shape of sherds.

However, the sherds with the same design are rarely from the same vessel and it is impossible to reconstruct an entire vessel using available sherds. As a result, we could not use the color, texture, and geometric information in this work as in previous fragment classification researches.

From the algorithm perspective, this paper aims to extract curve pattern from a sherd's depth image and to find a match between a partial curve pattern (on a sherd) and a full curve pattern (design).

Curve-pattern extraction from RGB or gray-scale images is closely related to curve-based image segmentation. Seveal algorithms have been proposed in many specific applications. For example, [18] utilized an energy criterion based on intensity and local boundary smoothness to extract blood vessels in medical images. [27] constructed a statistical shape model to extract sulcal curves on the outer cortex of the human brain. [33] proposed a tree-based algorithm to detect curve like cracks from pavement images. These methods all rely on specific assumptions in respective applications and it is not easy to extend the segmentation algorithm developed for one application to another application.

Curve-pattern matching has been a long-studied problem in computer vision and image processing. By thinning all of the curve pattern to one-pixel wide, many shape matching algorithms have been developed for matching curve patterns. For example, Belongie et al. [3] proposed a shape context approach, which uses a log-polar histogram as the feature to match two curve patterns. Roman-Rangel [26] extended the shape context algorithm to a Histogram of Orientation Shape Context (HOOSC) algorithm. Barrow et al. [2] proposed the widely used Chamfer matching algorithm by pre-computing a distance map for efficiently locating one curve pattern on another curve pattern. Brunelli [4] developed a template matching method. In Perceptual Hash (pHash) [30], each image was coded into a 64-bit fingerprint number, which was then used as features to compare and match images. Many of these existing matching algorithms are sensitive to noise and deformation present in the curve patterns segmented from sherds.

Deep-learning based algorithms, particularly the CNN-based algorithms, have been recently used for image segmentation [1, 31] and image matching[11, 29]. For image segmentation, the most influential one is the Fully Convolutional Network proposed by [6, 17]. However, if we directly apply these deep-learning based segmentation algorithms to our problem of extracting curve patterns, it may produce non-curve segments because the CNNs are trained directly on the color/intensity images. For image matching, several algorithms have been proposed. For example, a patch-based local image matching network, called MatchNet [11], was developed to jointly learn the feature representation and the matching function from image data. In [29], after exploring multiple neural networks, a CNN-based model called DeepCompare was developed to match a variety of images based on their appearance. In [16], Lin et al. designed a CNN-based model for fast image retrieval using binary hash codes. These methods are mainly developed for matching color or gray-scale images and show degraded performance in matching binary images of curve patterns with noise and deformation.

3 PROPOSED METHOD

The full pipeline of the proposed framework is illustrated in Figure 3.

3.1 Curve-pattern extraction from a sherd

Given a pottery sherd, the first step of our framework is to extract curve pattern from this sherd. Generally speaking, extracting curve pattern from the surface of a sherd is a typical low-level image segmentation problem. However, the erosion and sediment usually make the visibility of curve patterns on the sherd very weak and blurred, which substantially increases the difficulty in accurately segmenting them. In this step, we use the excavated pottery sherds associated with the Woodland period for experiments and we found that it is very difficult to extract these curve patterns from the camera-taken RGB images of these sherds. Given that these curve patterns are stamped on the surfaces of pottery vessels by carved paddles, curve patterns usually show larger depth than the adjacent non-curve surface. Therefore, in archeology, 3D scanners are usually applied to achieve the 3D depth image of the sherd surface, as illustrated in Figure 4, and the curve patterns are then segmented directly from the depth image.

However, after buried under the earth for thousands of years, together with possible shallow stamping in making the vessel, the curve patterns can still be difficult to segment even from the scanned

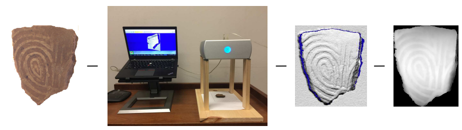

(a) (b) (c) (d)

Figure 4: An illustration of scanning sherds for depth images. (a) RGB image of a sherd. (b) Setup of a 3D scanner. (c) 3D point cloud of the sherd surface obtained by the 3D scanner. (d) Depth image of the sherd surface: pixel intensity represents the depth value at a location.

(a) (b) (c) (d)

Figure 5: An illustration of segmenting curve patterns from sample sherds. (a) Depth images of sherds, where darker pixels have larger depths. (b) FCN-extracted curve skeletons. (c) Refined curve skeletons by using a dense prediction CNN. (d) Final segmented curve patterns with recovered curve width, masked by the sherd boundaries (indicated by red contours).

high-resolution depth images. In our previous work [19], we developed a CNN-based algorithm to more accurately and reliably segment the stamped curve patterns from the depth images of the sherds, by learning and incorporating the implied curve geometry, such as curve smoothness and parallelism, in the underlying designs. Specifically, we train a Fully Convolutional Network (FCN) to detect the skeletons of the curve patterns in the depth images. Then, we train a dense prediction convolutional network to identify and prune false positive skeleton pixels. Finally, we recover the curve width by a scale-adaptive thresholding algorithm to get the final segmentation of curve patterns. Figure 5 shows the sample results after each step of this algorithm. We also extract the boundary of a sherd, indicated by red contours in Figure 5. The sherd boundary provides a mask to exclude all the information outside the sherd boundary from matching in the later steps.

It has been shown in [19] that this CNN-based algorithm can segment the curve pattern from a sherd much more accurately than other low-level and high-level image segmentation algorithms.

3.2 Identify the underlying design for a sherd with a non-composite pattern

Most of the sherds contain non-composite patterns, i.e. only one copy of partial designs are presented on these sherds. The second step of our framework is to identify underlying designs for sherds with non-composite patterns. The segmented curves from above step can be far from perfect because of the strong noise and shallow stampings on the unearthed sherds. In particular, the curve pattern segmented from a sherd may show deformation from its underlying design due to the drying process in making the object. In this step, we elaborate on a two-stage matching algorithm that is robust

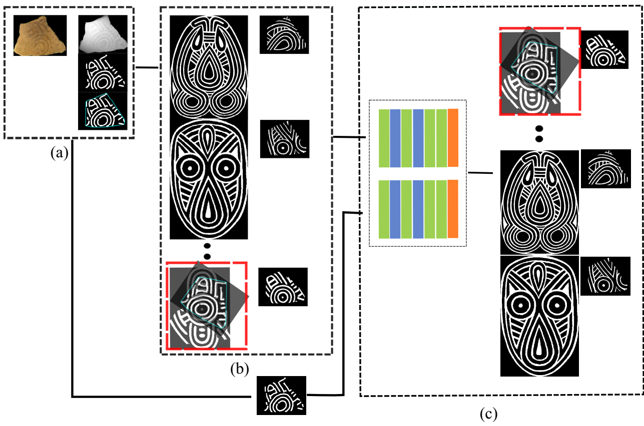

Figure 6: An illustration of the full pipeline of identifying design for a sherd with non-composite patterns. (a) Curve pattern segmentation from a sherd. (b) Stage 1: template matching with all the designs for selecting a small set of candidate matchings of the input sherd. (c) Stage 2: CNN-based re-ranking of the candidate matchings. Correctly matching design is shown in red box, which is ranked low in Stage 1 but ranked at the top in Stage 2. Original design reproduced with permission, courtesy of Frankie Snow, South Georgia State College.

to noise, errors and deformation present in the segmented curve patterns. We formulate this problem by matching the curve pattern segmented from a sherd against each location, with each possible orientation, of each known design, and then select the design with the lowest matching cost as the matched design. This exhaustive matching procedure identifies not only the matched design, but also the matched location and orientation on the matched design. Based on this problem formulation, the key issue is then the definition of an appropriate cost in matching the curve patterns segmented from a sherd to a location of a full design, with a specified orientation. This problem is nontrivial in the proposed archaeology application for two reasons. First, the exhaustive matching against each possible location and orientation of each design leads to a very large search space. To prevent from an overly slow algorithm, we require the matching cost to be very efficient to compute for each possible solution in the search space. Second, compared with the underlying design, the curve patterns segmented from the sherd usually contain strong noise and deformations in the drying process in making many of these objects, many years of erosion and sediment under the earth, and the imperfectness of the curve-segmentation algorithms.

To address this problem, we developed a new two-stage matching algorithm, with a different matching cost in each stage, as shown in Figure 6. In Stage 1, we propose to use a classical **template matching** method, which is highly computationally efficient, over the whole search space to identify a small set of candidate matchings on all the known designs. This simple matching cost can help efficiently reduce the search space of solutions. In Stage 2, we further derive a new matching cost by training a dual-source Convolutional Neural Network (CNN) and apply this more computationally-intensive matching to re-rank the candidate matchings identified in Stage 1.

CNN architecture is shown in Figure 7(a). These two sub-networks take candidate matchings and sherd curve patterns as the inputs, respectively. Each sub-network consists of a sequence of convolution, max pooling layers and a global average pooling layer (GAP) for feature learning, as detailed in Table 1. We implement this dual-source CNN by truncating AlexNet [15], as shown in Fig. 7(b), to "conv4" layer and replacing all layers after "conv4" layer with a GAP layer. Both inputs, i.e., candidate matchings and sherd curve patterns are re-sized to 227×227 pixels, before being fed to the sub-network. Parameters are listed in Table 1.

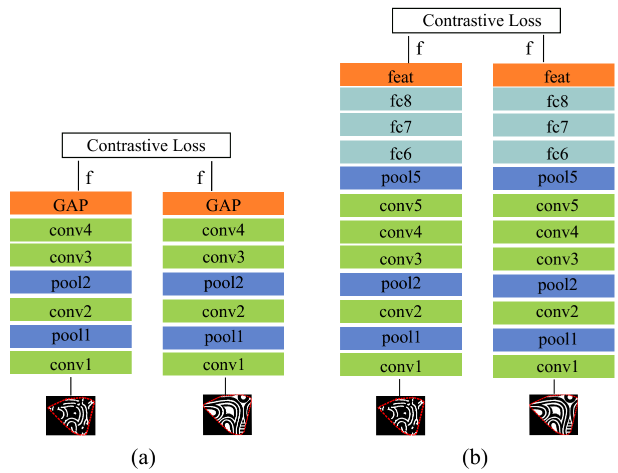

Figure 7: An illustration of the dual-source CNN architectures: (a) the proposed CNN and (b) AlexNet.

Table 1: Configuration of each sub-network; n, k, s, p stand for the number of outputs, kernel size, stride and padding size respectively.

Name	Type	Configuration
GAP	GAP	$k : 13 \times 13, s : 1$
conv4	Convolution	$n : 384, k : 3 \times 3, s : 1, p : 1$
conv3	Convolution	$n : 384, k : 3 \times 3, s : 1, p : 1$
pool2	MaxPooling	$k : 3 \times 3, s : 2$
conv2	Convolution	$n : 256, k : 5 \times 5, s : 1, p : 2$
pool1	MaxPooling	$k : 3 \times 3, s : 2$
conv1	Convolution	$n : 96, k : 11 \times 11, s : 4, p : 0$
data	Input	227×227 binary image

Through this supervised learning, various kinds of noise and deformations in the segmented curve patterns can be implicitly identified and suppressed in computing the CNN-based matching cost.

3.3 Identify the underlying design for a sherd with a composite pattern

In the making of the pottery, the carved paddle usually stamped on the pottery multiple times to ensure full coverage of the surface.

As a result, a large number of the pottery sherds contain composite patterns, i.e. each pattern is a part of its underlying design, and these patterns can overlap with each other. Classical matching methods, such as Chamfer matching require one pattern to be a portion of the other. This is not true in this case as the curve pattern on the sherd is a composite one. In our previous work [32], we developed a new algorithm that can automatically identify multiple components of the composite pattern extracted from the sherd.

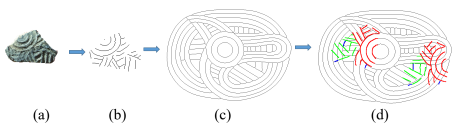

Figure 8: An illustration of identify the underlying design for a sherd with a composite pattern. (a) A sherd image. (b) Curve extraction from a sherd. (c) Curve extraction from a design. (d) A sherd matching to two locations on a design. Original design reproduced with permission, courtesy of Frankie Snow, South Georgia State College.

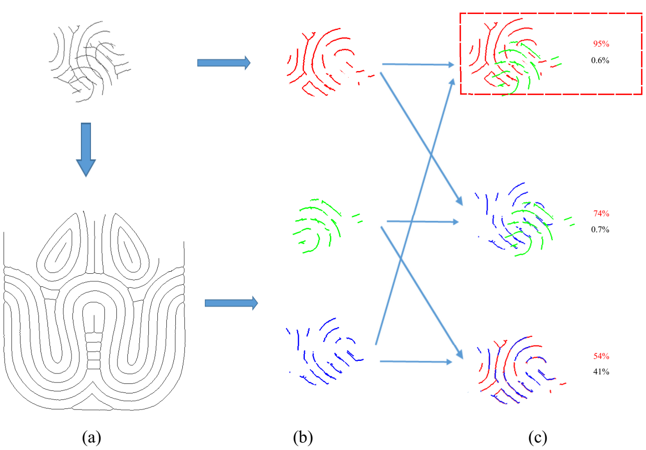

Figure 9: The process of combining candidate components for matching to a design. The optimal result is indicated in the red box. (a) Matching a sherd pattern (top) to a design pattern (bottom). (b) Candidate Components. (c) Combining candidate components (best matching is shown in red box). Original design reproduced with permission, courtesy of Frankie Snow, South Georgia State College.

Taking sherd curve pattern images and design images, we first use standard edge-thinning algorithm to reduce the curve width to one pixel as illustrated in Figure 8 (b).

Although the width of curves presents important cue in matching a sherd and a design, we try not to use the curve width information because it is very difficult to accurately measure the curve width from a deteriorated sherd surface. Second, we extended classical Chamfer matching method to match the one-pixel-wide curve patterns from a sherd against each location, with each possible orientation, of each known design. Different from classical Chamfer matching algorithm, we do not pick the design with the lowest matching cost. Instead, for each design, we select a number of matchings as candidates as long as a threshold percentage of total pixel matches. Shown in Fig 9, these candidates are then combined and reconstructed. The combination with the most matching pixels (defined as completeness) and least overlapping pixels (defined as disjointness) is taken as the best matching, and its normalized completeness is taken as its matching score. The design with the highest matching score is selected as the sherd's underlying design. Note that, matching score is the higher the better.

4 EXPERIMENTS

For our study, we collected a set of 1000 pottery sherds that were excavated in various archaeological sites located in Southeastern North America. 900 of these pottery sherds contain non-composite patterns represent 98 unique paddle designs, and the rest 100 pottery sherds contain composite patterns representing 20 unique paddle designs. Each sherd in the set only displays one design, while the same design may be applied to the surfaces of multiple sherds. We divide these 900 sherds with non-composite curve patterns into two groups of equal sizes, one group is for CNN training and the other is for CNN testing.

In our experiment, we use the Cumulative Matching Characteristics (CMC) ranking metric to evaluate the matching performance. To identify the underlying design of a sherd pattern U, we match it against all 98 or 20 designs depending on whether it is a sherd with a non-composite pattern or a composite pattern. We then sort these 98/20 designs in terms of the matching scores and pick the top L designs with the highest scores. If the ground-truth design of this sherd is among the identified top L designs, we treat it a correct design identification under rank L. We repeat this identification for all 450 sherds with non-composite patterns in CNN testing set, and 100 sherds with composite patterns respectively, and calculate the accuracy, i.e., the percentage of the correctly identified sherds, under each rank L, $L = 1, 2, \cdots, 20, \cdots, 98$. This way, we can obtain a CMC curve in terms of rank L to evaluate the performance of a matching algorithm, shown in Figure 10 and Figure 12 for sherds with non-composite patterns and sherds with composite patterns respectively. The higher value in this curve, the better the matching performance.

Since matching methods applied to sherds with non-composite patterns and sherds with composite patterns are different, we conducted two sets of experiments to evaluate our proposed framework. First, to evaluate the effectiveness of the proposed framework on sherds with non-composite curve patterns, we select eight existing matching algorithms for comparison: Template Matching [4], Chamfer Matching [2], Shape Context [3], Nearest Neighbor [28], pHash [30], Gabor [21], DeepCompare [29] and MatchNet [11]. Experiments are conducted on the testing dataset with 450 sherds with non-composite patterns.

In Template Matching, we directly use OpenCV implementation of Template Matching for finding the best matched designs, as well as locations and orientations. In Chamfer Matching, sherd curve pattern I_T and each design are first thinned to one-pixel-wide skeleton U and V, respectively. Then U is translated and rotated to match V in terms of Chamfer distance. The Chamfer matching cost is then defined as the minimal distance, including all translations

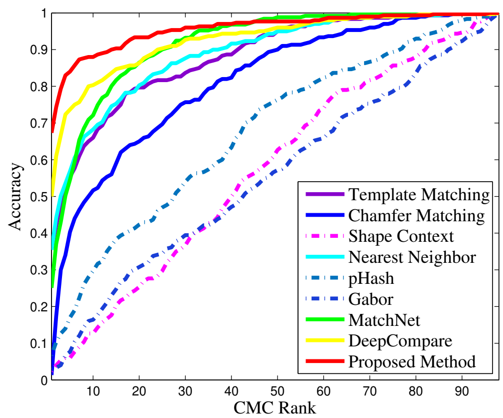

Figure 10: CMC curves of the proposed method and the eight comparison matching methods.

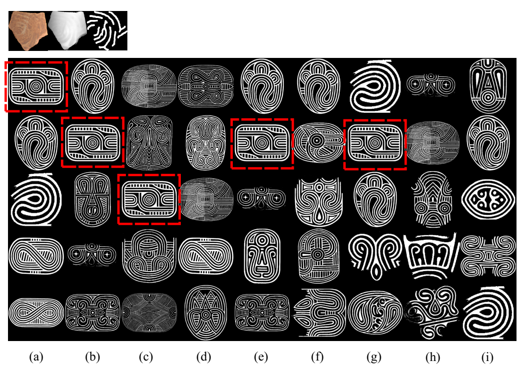

Figure 11: The top 5 matched designs (from top to the bottom) identified by (a) the proposed method, (b) Template Matching, (c) Nearest Neighbor, (d) MatchNet, (e) DeepCompare, (f) Shape Context, (g) Chamfer matching, (h) Gabor and (i) pHash, respectively. True designs are highlighted in the red box. Original designs reproduced with permission, courtesy of Frankie Snow, South Georgia State College.

and rotations. Shape Context follows the same formation of Chamfer Matching by sliding U over V and calculate the shape-context matching at each location of sliding for best matching locations. We directly use the Shape Context implementation, as well as its matching cost, from the OpenCV package[1]. Nearest Neighbor, pHash, Gabor, DeepCompare and MatchNet are used to re-rank the candidate matchings that are selected by the proposed Stage 1 Template Matching method. The same CMC ranking metric is then computed for each of them for performance evaluation. Specifically, for Nearest Neighbor, we directly calculate the intensity difference between a pair of inputs as their matching cost. pHash was implemented using pHash library[2]. For Gabor, we construct gabor features using Gabor filter from OpenCV package[3]. For the MatchNet, we employ its original network architecture and training parameters, then fine-tune with the above training dataset on the model trained on "Yosemite" dataset[4]. For DeepCompare, we choose the 2-channel deep network introduced in [29], and fine-tune the model with the canny edge images generated from the "Yosemite" dataset.

CMC curves of the proposed method and the eight comparison methods are shown in Figure 10. We can see that our method achieves the best CMC performance, and outperforms the second-best matching method by 17.3% on Rank-1 CMC value. Figure 11 shows the identification result of the proposed method and the eight comparison methods on a sherd with non-composite pattern segmented from a degraded sherd. We can see that, the proposed method matches the true design (in red box) at CMC Rank 1, while the other comparison methods do not.

Second, to justify the effectiveness of our framework for sherds with composite patterns, we pick four classical matching algorithm for performance comparison in the experiments: Template Matching [4], Chamfer Matching [2], Shape Context [3] and Histogram of Orientation Shape Context (HOOSC) [26]. We follow the same setup for the first three as that described for sherds with non-composite patterns and for HOOSC, we use the same setup as Shape Context, but incorporating the orientation measure into the log-polar histograms. HOOSC was implemented by HG Zhao[5] using MATLAB. Shown in Figure 12, the top-1 CMC rank of the proposed method

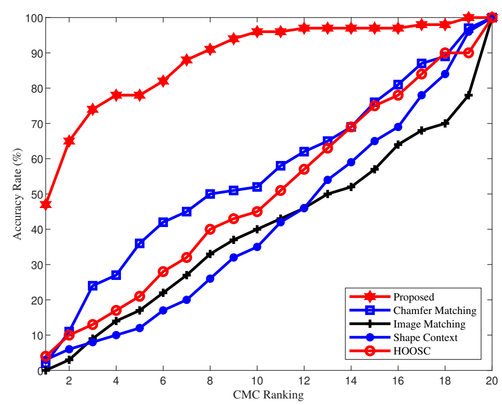

Figure 12: CMC curves of the proposed method and the four comparison methods.

is %46, while all of the four comparison methods show very poor performance by having top-1 CMC ranks below %5 and the CMC curves along the diagonal line. The major reason for their poor performance is that they do not consider and cannot well handle the composite patterns present on the sherds. By explicitly considering the possible composite patterns, the proposed method achieves much better CMC performance.

Figure 13 shows the identification result of two sample sherds with composite patterns on the proposed method and four comparison methods. We can see that, the proposed method can identify

[1] https://docs.opencv.org/3.0-beta/modules/shape/doc/shape_distances.html
[2] https://www.phash.org/
[3] https://docs.opencv.org/3.0-beta/modules/imgproc/doc/filtering.html
[4] https://github.com/hanxf/matchnet
[5] https://github.com/CyberZHG/Sketch-Based/tree/master/HOOSC

the correct designs under CMC rank 1, while the four comparison methods can only identify the correct designs under much higher CMC ranks.

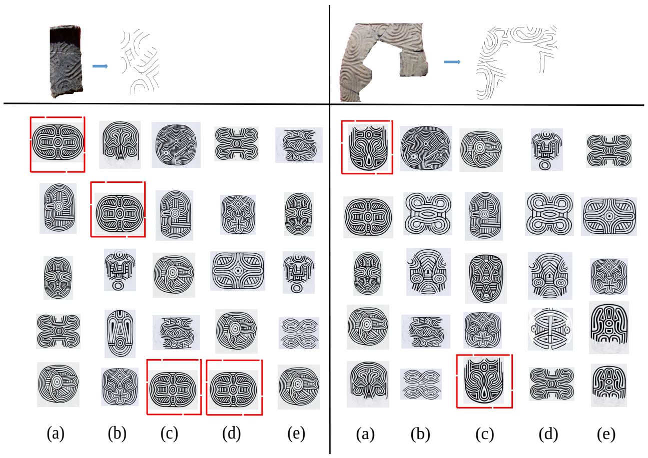

Figure 13: The top 5 matched designs (from top to the bottom) identified by (a) the proposed method, (b) Chamfer Matching, (c) Shape Context, (d) HOOSC, respectively. True designs are highlighted in the red box. Original designs reproduced with permission, courtesy of Frankie Snow, South Georgia State College.

5 FUTURE WORK

Many different adaptations, tests, and experiments have been left for the future due to lack of time and data (i.e. the experiments with real data on sherds with composite patterns are usually very time consuming, requiring even weeks to finish a single run). Concerning different methods applied to sherds with non-composite patterns and composite patterns in our framework, our future work would also involve combining these two cases into one single method or developing an algorithm to automatically identify whether a sherd contains a non-composite pattern or a composite pattern.

6 CONCLUSION

In this paper, we explored an important and challenging task in archaeology: identifying the curve design on the surfaces of highly fragmented and degraded pottery sherds. We developed a new framework to match the curve patterns segmented from sherds to a set of known designs. First, we extract curve patterns from a sherd using an FCN-based image segmentation method. Then, we used a 2-stage matching algorithm to match a sherd with non-composite pattern to a set of known designs. Alternatively if the sherd is a composite pattern, we developed a new Chamfer matching algorithm to match the sherd to find its underlying design. In the experiment, we validated the proposed framework by using a set of real sherds together with their corresponding designs from the Woodland Period in Southeastern North America. Comparison to several existing matching methods verified that the proposed framework can achieve a new state-of-the-art performance .

7 ACKNOWLEDGMENT

This research was supported by the National Science Foundation Archaeology and Archaeometry Grant Program (1658987), the National Center for Preservation Technology and Training Grants Program (P16AP00373) and University of South Carolina Social Sciences Grant Program. We would like to show our gratitude to Professor Frankie Snow at South Georgia State College for sharing his pearls of wisdom and design images with us during the course of this research. We also thank Dr. Matthew Compton, Curator of the R. M. Bogan Repository at Georgia Southern University for generously sharing his collection, and our colleague Professor Scot Keith for encouraging the pursuit of this research.

REFERENCES

[1] Vijay Badrinarayanan, Alex Kendall, and Roberto Cipolla. 2017. Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE transactions on pattern analysis and machine intelligence* 39, 12 (2017), 2481–2495.
[2] Harry Barrow, Jay Tenenbaum, Robert Bolles, and Helen Wolf. 1977. Parametric correspondence and Chamfer matching: Two new techniques for image matching. In *International Joint Conference on Artificial Intelligence*, Vol. 2. 659–663.
[3] Serge Belongie, Jitendra Malik, and Jan Puzicha. 2001. Shape Context: A new descriptor for shape matching and object recognition. In *Advances in Neural Information Processing Systems*. 831–837.
[4] Roberto Brunelli. 2009. *Template Matching Techniques in Computer Vision: Theory and Practice*. John Wiley & Sons.
[5] Y. Cao, Z. Zhang, I. Czogiel, et al. 2011. 2D nonrigid partial shape matching using MCMC and contour subdivision. In *IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2345–2352.
[6] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. 2018. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence* 40, 4 (2018), 834–848.
[7] Fernand Cohen, Ezgi Taslidere, Zexi Liu, et al. 2010. Virtual reconstruction of archaeological vessels using expert priors & surface markings. In *IEEE Conference on Computer Vision and Pattern Recognition - Workshops*. IEEE, 7–14.
[8] David Cooper, Andrew Willis, Stuart Andrews, et al. 2001. Assembling virtual pots from 3D measurements of their fragments. In *Conference on Virtual Reality, Archeology, and Cultural Heritage*. ACM, 241–254.
[9] Ayelet Gilboa, Avshalom Karasik, Ilan Sharon, et al. 2004. Towards computerized typology and classification of ceramics. *Journal of Archaeological Science* 31, 6 (2004), 681–694.
[10] Radim Halíř. 1999. An automatic estimation of the axis of rotation of fragments of archaeological pottery: A multi-step model-based approach. In *International Conference in Central Europe on Computer Graphics, Visualization and Interactive Digital Media*. WSCG.
[11] Xufeng Han, Thomas Leung, Yangqing Jia, Rahul Sukthankar, and Alexander Berg. 2015. MatchNet: Unifying feature and metric learning for patch-based matching. In *IEEE Conference on Computer Vision and Pattern Recognition*. 3279–3286.
[12] Martin Kampel and Robert Sablatnig. 2007. Rule based system for archaeological pottery classification. *Pattern Recognition Letters* 28, 6 (2007), 740–747.
[13] Avshalom Karasik and Uzy Smilansky. 2008. 3D scanning technology as a standard archaeological tool for pottery analysis: practice and theory. *Journal of Archaeological Science* 35, 5 (2008), 1148–1168.
[14] Avshalom Karasik and Uzy Smilansky. 2011. Computerized morphological classification of ceramics. *Journal of Archaeological Science* 38, 10 (2011), 2644–2657.
[15] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. ImageNet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*. 1097–1105.
[16] Kevin Lin, Hueifang Yang, Jenhao Hsiao, and Chusong Chen. 2015. Deep learning of binary hash codes for fast image retrieval. In *IEEE Conference on Computer Vision and Pattern Recognition - Workshops*. 27–35.
[17] Jonathan Long, Evan Shelhamer, and Trevor Darrell. 2015. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 3431–3440.
[18] Liana M Lorigo, Olivier D Faugeras, W Eric L Grimson, Renaud Keriven, Ron Kikinis, Arya Nabavi, and C-F Westin. 2001. Curves: Curve evolution for vessel segmentation. *Medical image analysis* 5, 3 (2001), 195–206.
[19] Yuhang Lu, Jun Zhou, Jun Chen, Jing Wang, Karen Smith, Wilder Colin, and Song Wang. 2018. Curve-structure segmentation from depth maps: A CNN-based approach and its application to exploring cultural heritage objects. *AAAI Conference on Artificial Intelligence* (2018).

[20] Michael Makridis and Petros Daras. 2012. Automatic classification of archaeological pottery sherds. *Journal on Computing and Cultural Heritage* 5, 4, Article 15 (2012), 21 pages.

[21] Bangalore Manjunath and Weiying Ma. 1996. Texture features for browsing and retrieval of image data. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 18, 8 (1996), 837–842.

[22] Philip Phillips and James Brown. 1978. *Pre-Columbian Shell Engravings from the Craig Mound at Spiro, Oklahoma, Part 1.* Peabody Museum Press, Cambridge.

[23] Li-Ying Qi and Ke-Gang Wang. 2010. Kernel fuzzy clustering based classification of Ancient-Ceramic fragments. In *International Conference on Information Management and Engineering*. IEEE, 348–350.

[24] Dorie Reents-Budet. 1994. *Painting the Maya Universe: Royal Ceramics of the Classic Period (Duke University Museum of Art).* Duke University Press, Durham.

[25] Edgar Roman-Rangel, Diego Jimenez-Badillo, and Estibaliz Aguayo-Ortiz. 2014. Categorization of Aztec potsherds using 3D local descriptors. In *Asian Conference on Computer Vision - Workshops*. Springer, 567–582.

[26] Edgar Roman-Rangel, Carlos Pallan, Jean-Marc Odobez, et al. 2011. Analyzing ancient Maya glyph collections with contextual shape descriptors. *International Journal of Computer Vision* 94, 1 (2011), 101–117.

[27] Xiaodong Tao, Jerry L Prince, and Christos Davatzikos. 2002. Using a statistical shape model to extract sulcal curves on the outer cortex of the human brain. *IEEE Transactions on Medical Imaging* 21, 5 (2002), 513–524.

[28] Roger Weber, HansJörg Schek, and Stephen Blott. 1998. A quantitative analysis and performance study for similarity-search methods in high-dimensional spaces. In *International Conference on Very Large Data Bases*, Vol. 98. 194–205.

[29] Sergey Zagoruyko and Nikos Komodakis. 2015. Learning to compare image patches via convolutional neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition*. 4353–4361.

[30] Christoph Zauner. 2010. *Implementation and Benchmarking of Perceptual Image Hash Functions.* Upper Austria University of Applied Sciences.

[31] Shuai Zheng, Sadeep Jayasumana, Bernardino Romera-Paredes, Vibhav Vineet, Zhizhong Su, Dalong Du, Chang Huang, and Philip HS Torr. 2015. Conditional random fields as recurrent neural networks. In *Proceedings of the IEEE international conference on computer vision*. 1529–1537.

[32] Jun Zhou, Haozhou Yu, Karen Smith, Colin Wilder, Hongkai Yu, and Song Wang. 2017. Identifying designs from incomplete, fragmented cultural heritage objects by curve-pattern matching. *Journal of Electronic Imaging* 26, 1 (2017), 011022–011022.

[33] Qin Zou, Yu Cao, Qingquan Li, Qingzhou Mao, and Song Wang. 2012. CrackTree: Automatic crack detection from pavement images. *Pattern Recognition Letters* 33, 3 (2012), 227–238.