Quantifying Object Similarity: Applying Locality Sensitive Hashing for Comparing Material Culture

Abstract: We present a novel technique that compares and quantifies images used here to compare similarities between material cultures. This method is based on locality sensitive hashing (LSH), which uses a relatively fast and flexible algorithm to compare image data and determine their level of similarity. This technique is applied to a dataset of sculpture faces from the Aegean, Anatolia, Cyprus, Egypt, Iran, Indus/Gandhara, the Levant, and Mesopotamia. Results indicate that the objects can be differentiated based on regional differences and show similarities to other locations that share specific material culture traits. Images from known locations enable a network of compared objects to be constructed, where inverse closeness centrality and link weights are used to indicate areas that have a greater or less cultural similarity to other regions. Different periods are assessed, and the results demonstrate that objects from earlier than the 9th century BCE show greater similarity to other local and Egyptian items. Objects from between the 9th and 4th centuries BCE increasingly show interregional similarity, with the eastern Mediterranean, including the Aegean, Anatolia, Egypt, and Cyprus, having close similarity to multiple regions. After the 4th century BCE, greater sculptural similarity is found across a wide area, including the Aegean, Cyprus, Egypt, Mesopotamia, and Gandhara. In general, sculptures from more distant areas increase in similarity in later periods, that is starting from the 9th century BCE. The results demonstrate that the technique can be applied to quantifying object similarity and extended to a broad range of archaeological objects, while also being a tool for rapid analysis that requires minimal data compared to some machine learning techniques. The code and data are provided as part of the outputs.

Keywords: locality sensitive hashing, sculpture, similarity, material culture, network, centrality

1. Introduction

Comparing material culture has been a long-standing endeavor for art historians, archaeologists, and other scholars interested in how material culture changes across time and space as well as how they compare. While this has been an enduring interest, many methods have traditionally applied qualitative assessments to changes in material culture and when judging similarity between objects (e.g., Winter 2005). Such methods work well when sufficient expertise is present or data are limited to allow an expert to make an assessment. Quantitative and computer vision approaches, including machine-learning techniques, have been developed in archaeology to automate comparisons and classification of material cultures that compare imagery and form in objects, which allow classification and identification of different material culture types (Forrester-Sellers et al., 2017; Gansell et al., 2014; Bevan et al., 2014; Wang et al., 2010). While a lot of progress has been made in classification in particular, measuring levels of similarity in objects is still poorly developed. Few methods exist that are easily transferable to different types of objects while also having flexibility in the dimensionality of data, such as images with different sizes. Computational methods deployed are also often slow, requiring a high percentage of sub-sampling within images. Some techniques are limited by their need for training data, which sometimes requires a lot of examples, to make them applicable in archaeology. This is not always feasible, particularly in cases where relatively rare objects are analyzed. On the other hand, computational methods that provide metrics for similarity between objects could be useful in determining cultural relationships and provide an automated way for practitioners to compare objects for given material culture.

The primary goal of this work is to introduce and demonstrate the utility of a novel method in archaeology used to compare material culture: it applies a form of locality sensitive hashing (LSH;

Kulis et al., 2012; Paulevé et al., 2010). This is a family of algorithms that assess data similarity, including images and other forms of visual data. Among different visual media, the technique can be used to compare photographs to determine their similarities, including similarities between specific objects within images. By focusing on common and specific traits, such as faces on sculpture in this paper, then items from different regions and periods can be compared. The benefit of LSH is that it creates a similarity value when comparing two objects. If the locations of objects are known, then similarity scores can be mapped spatially and in relation to other comparable objects from different regions, allowing a relationship network to be created and analyzed. This helps to compare objects across regions, suggesting where common stylistic traits might be similar or shared. The LSH algorithm is also relatively fast, allowing many images to be compared. Another benefit is the minimal data requirement, as this technique can be used with two or more images. Overall, the method creates a quantitative-based output for object similarity. The technique can be extended to many types of object comparisons, including video and 3D images.

This work begins by providing background in the quantitative and computer vision techniques comparable to the approach taken here. Data and the cases applied in the analysis are also presented, including relevant background on the objects discussed. The deployed LSH method is discussed along with the graph analysis used to compare object similarity across networked regions in different periods. A test is done to demonstrate LSH's utility in addressing its intended design. Results from the approaches are then given, showing overall comparisons between the different regions and periods as well as outputs for the entire dataset. Discussions on the insights gained from the analysis and wider utility of the approach for material culture comparison is given. Future research potential is also discussed. As part of the output of this work, and given the utility of the approach to other forms of data, the code and sample data are provided for readers.

2.0 Background

2.1 *Quantitative Techniques for Object Similarity*

A variety of techniques have been created to measure object or image similarity, often falling under the wider field of computer vision. One family of techniques include scale-invariant feature transformation (SIFT), which can be used to recognize objects or features in images (Lowe, 1999). Methods related to this technique generally determine local key points where data from given base images are stored. Images can then be compared to these locally stored points of data to determine the level of similarity between two or more objects. Matching areas between images can be measured using Euclidean distances between points, where location, scale, and orientation of given points that closely match are considered to be good fits and are more likely to suggest similarity in images. The approach is a form of image descriptor, which compares base images or reference images to new images given in an assessment. Geometric blur is another technique that can be used to compare image similarity by looking at signal averages, including areas with sparse signals, to produce a signal discriminative descriptor (Berg et al., 2005). The image is blurred around given points in a picture, and samples are taken across these blurred regions which can then be compared to other images across comparable areas. Another approach is to use the silhouettes of image objects, by graphing the object edges (to represent the contours of the objects) to aid comparisons (Sebastian et al., 2002; Siddiqi et al., 1999). Scene gist is a method that represents a picture as a weighted set of dimensions and treats the scenes within it spatially; scene signatures are compared to different objects where similar spatial structure in images indicates objects are similar (Oliva and Torralba, 2001). Other approaches have emphasized speed in searching and identifying similarities. For instance, mean squared error is a simple technique that allows users to compare pixel values and return objects that have the most similar pixel values based on minimal error between objects (Lehmann and Casella, 2011). This generally requires scaling the images to make them comparable, while also placing greater emphasis on images that produce outliers in results, which bias closely matched data.

Increasingly, deep learning techniques (Foster, 2019) have been employed to classify imagery, including 2D and 3D images, using such techniques as convolutional neural networks (Rawat and Wang, 2017) or restricted Boltzmann machines (Srivastava and Salakhutdinov, 2014). Both methods are a set of techniques that translate images through multiple analytical layers applying different weighted and transformational functions that reduce dimensions of data into informative sets that differentiate and, therefore classify, objects based on similarity levels. Ideally, the algorithms are trained using a subset or sample data, before they are then used to classify a full or larger number of images.

Generally, for all of these computer vision methods, algorithms that sample a greater percentage of images are more accurate, but they require more time and computational resources. For some of these methods even analyzing only a few images at megabyte scale can take considerable computational time. Additionally, some computer vision and machine learning techniques require base data or training sets, which may not always be available, particularly for rarer objects, as often found in archaeology. Many of these methods require images to have the same dimensions, which either may not be possible or requires even more processing time before the analysis can be completed. One set of methods that enable fast, generally accurate, comparisons, without requiring any training images, is LSH (Tsai and Yang, 2014; Jegou et al. 2008). LSH uses a sampling method, similar to some of the other computer vision techniques discussed above, but one that hashes datasets: it translates pixel values into fixed-size values called hashes. Hashes are then organized into “buckets” or containers based on similarity. Buckets are effective ways to organize data which can then be compared to other image data. Nearest-neighbor, clustering methods, or direct comparisons can be applied to allow similarities in data (or hashes) to be measured; a similarity score resulting from the comparison of any two images across a given defined scene is provided. The advantage of this technique is that it is effective with multi-dimensional data, including 3D images, and it does not require the scaling of images, as methods like the mean squared error do. This is particularly beneficial when comparing legacy image data as well as newly-obtained images. LSH can also be used to segment images and be applied to object recognition and image indexing (Kacem et al., 2015). This has an advantage for image databases because it allows images to be organized based on similarities. Finally, compared to the methods described above, it is relatively fast at comparing a small or large set of images. These advantages make LSH useful for searching and comparing both large and small datasets when analyzing objects and their similarity in different images. These advantages have, in fact, led to similar LSH algorithms being used in image database searches, creating the possibility for fast comparisons of data (Paulevé et al., 2010).

To our knowledge, and despite its advantages, LSH methods have not been developed for archaeology. In this work, we demonstrate the advantages of LSH for comparing material culture using available images, while also making the code available so that other researchers can deploy this method in their own areas of interest.

2.2 *Applied Data*

There are conditions in which LSH is best used, where objects that are most comparable should be used in the method. For this paper, we attempt to show the utility of the approach by using data we are familiar with given our regional archaeological expertise. This has led us to develop several selective criteria in choosing our case study. First, while LSH can work with many types of images given to it, the utility of the outputs is best realized by having a dataset that we have some understanding of so that accuracy can be evaluated. This enables a form of validation to be done in that we can check to see if similar objects appear as similar. Secondly, we wanted to have a dataset that was sufficiently spatially diffuse and numerically great enough to show the utility of the networking outputs developed from LSH that are discussed in the methods. Thirdly, we wanted to be sure that objects are comparable in their type, such as male statues, and regions (e.g., faces) within objects compared. This also required photographs to be taken from generally the same angle or perspective. Fourthly, there is a minimal data quality, that is pictures should not be too low of a resolution so that identification features (e.g., eyes, noses, cheeks) are present. Finally, we also wanted to work with samples that were legally acquired. Based on the criteria stated above, the statue data we found fit the qualities required. However, the method is intended to be general and can be applied to a variety of other objects.

While objects such as pottery, paintings, jewelry, and others could potentially be used with the approach discussed, we found that statues provided a sufficient number of examples while other artifacts we searched, including cylinder seals and ceramics, were not sufficient, at least for regions and periods we are familiar with. In the discussion, we summarize criteria, based on the results of this effort, that users can be used to guide the data selection process. We briefly discuss the dataset used in this work, which acts as a test case for the method's utility. We chose stone sculptures depicting human males, common to many Old World societies, to create the dataset applied here because they met our selection criteria discussed above. Stone sculptures are also common through many different periods, allowing comparisons to be measured both temporally and morphologically between given regions. Such sculptures have also been studied previously using alternative approaches to the one advanced here, which allows these results to be compared to other works. Our data, as well as the code used for this investigation, are provided here for download (see Supplementary Material).

The data covers the eastern Mediterranean and South-West Asia, with images focusing on the central sections of sculptured faces. The sculptures derive from the wider Aegean, Anatolia, Cyprus, Egypt, Levant, Mesopotamia, Iran, and Indus and Gandhara regions (Figure 1). A total of 233 images of statue faces are included. The regional breakdown of the statues is as follows: 31 Aegean, 10 Anatolia, 41 Cyprus, 52 Egypt, 24 Levant, 44 Mesopotamia, 10 Iran, 3 Indus, and 18 Gandhara. Ideally, we would use statues from well-defined archaeological sites, but not all examples derive from archaeological sites that are known, as statues were obtained in different periods. The periods covered by these sculptures range from the 3rd millennium BCE to the Late Roman period or about 300 CE. This wide spatial and temporal coverage allows comparisons to be made between these cultures and to see how cultural relationships changed over time. Part of the data are categorized as "pre-9th century BCE," which corresponds to the Bronze and Early Iron Ages in the Near East until about 900 BCE. The long period was marked by political fragmentation and city-states or regional states were common and persistent (Van de Mieroop, 2004). Interaction between political entities was frequent, and various levels of interaction occurred (e.g., commercial, diplomatic, and military), including the movement of merchants and artisans who helped spread specific material culture traits and stylistic motifs across wider areas (e.g., see Pfälzner, 2015; Stockhammer, 2013; Zaccagnini, 1983). Material culture of this period often retained markedly regional characteristics in comparison with later periods. Regions nearer to each other tended to display more similarities in their material cultures than more distant regions, although even between nearby regions clear physical differences are evident in sculpture and in other cultural items (Stiebing and Helft, 2018).

The following period in the sample data, referred to as the 9th-4th centuries BCE period, spans the 9th century to about 330 BCE. This is a period characterized by a succession of larger empires that controlled most of the Near East and Egypt (Cline and Graham, 2011). This is the start of the so-called

Age of Empire period (AoE) in the Near East and Egypt. Amongst other empires, the Assyrian and Babylonian Empires brought together most areas of the Near East under their control, while the Achaemenid Empire extended from the Aegean to Central Asia, incorporating many cultures within its borders (Waters, 2014). This was a period of increased interregional similarities in material culture, due to the intensification of long-distance trade from the Mediterranean to Central Asia, and the movement of artisans across the wider Near East (Altaweel and Squitieri, 2018).

The following period is from the late 4th century BCE and into Late Antiquity, referred to as the post-4th century BCE period here. This was the period when, after Alexander the Great's conquests, the Near East from the Levant to eastern Iran was included under the Seleucid Empire, which promoted the establishment of colonies with Greek inhabitants across the empire (Wiesehöfer, 1996). Contact between the Greek and Near Eastern cultures was not new; however, interaction increased with the presence of Greek settlers throughout the Near East and Egypt, which boosted the blending of Greek and Near Eastern material cultural elements (Invernizzi, 2012; Burn, 2004). Egypt, too, participated in this phenomenon during the Ptolemaic dynasty, which greatly promoted the mixing of Greek and traditional Egyptian artistic motifs (Hill, 2016). Some stylistic features, such as naturalistic facial and hair details, as well as Greek-inspired style of folding cloth, made their way from the Mediterranean deep into Central Asia, including into India, and persisted there for centuries as witnessed by Gandhara culture (Behrendt, 2007).

Patterns of increasing interaction between cultures from the Mediterranean to Central Asia during the period covered by our study have been previously discussed in several archaeological works; however, our method proposes a quantitative way to demonstrate the increasing similarity in material culture that previous studies have tackled mostly through qualitative methods. By doing so, we intend to show that quantitative methods such as that we propose here can effectively help researchers interested in studying levels of material culture similarity across time and space.

2.3 *Data Organization*

The data are organized by the periods designated above (pre-9th century BCE, 9th-4th centuries BCE, and post-4th century BCE). A metadata file provides information about each compared image. The data columns used to provide information on images are: *file*, *time*, *period*, *culture*, *region*, *modern country*, and *source*. *File* is the name of the file; *time* is the time range of a given object with negative numbers indicating BCE dates. *Period* is the period in which the object originated, using local or regional chronologies; *culture* is the object's cultural affiliation. *Region* is the archaeological region, such as Levant; *modern country* represents not only where or nearly where the object is found but here acts as a representative country for the archaeological region the object is from (e.g., Iraq representing Mesopotamia). The *modern country* column is used to enable network relationships to be created in sculpture similarity using a representative modern country. These outputs enable the network relationships discussed below. Finally, *source* is the source of the data and where the object is currently located. Table 1 provides an example of the metadata file used; the repository (see Supplementary Material) provides the full dataset as well as the metadata used, and includes additional explanations and documentation on how to use the data and code.

Based on what is stated above, it is important for this implementation to make object images as comparable as possible for LSH so that the differences within images mostly reflect morphological differences. This means finding the same type of object, and in this study statue faces were used. Images were found using museum databases and downloaded from Creative Commons (2020) where possible for the given regions analyzed (see Supplementary Material for data information). Using existing images demonstrates the utility of this approach for legacy data, although for other use cases users may want to apply their own images. First, all images were rendered in grayscale. While algorithms within LSH should already minimize color differences in the evaluation, to focus more on the shapes within the images, this is still done as an initial precaution to flatten color variation. Second, the images were cropped from just above the eyebrows, to the eyebrow corners, and then down to the chins; all images were standardized in this way (see Figure 2 for examples). This was also done because virtually all human faces analyzed included this region of the face, making the images generally comparable. Similarity can still be identified when images show only part of an object; however, the same area of the face is included in images being compared to keep images standardized. We also apply face recognition and automated face cropping software that we incorporated in the code and provide to users such that it enables another option to standardize compared areas. Third, the images were frontal views and at eye level. At times, images deviated from this slightly, but this general rule was applied as closely as possible. Features in faces should still be evident even in slightly off-center images. Fourth, only sculptures depicting male figures were used, providing a common gender for comparison. Fifth, background areas away from the body or face of a sculpture were given a solid black color in order to standardize comparisons. This is a type of masking that, although it does not nullify the background, makes the proportion of background to the face part of the differences between sculptures. Finally, sculptures with minimal facial damage were preferred. Many sculptures had at least some minor damage, but generally large breaks were avoided and sculptures should be more than 95% complete. While all of these criteria are not required for our method, and raw images without any modification can be used, these criteria help to maintain a relative consistency across the image areas to be analyzed.

Insert figure here
Figure 1. Locations where sculptures derive from and example images of sculptures.

Insert table here
Table 1. Example metadata used for image data.

3.0 Methods

The applied methods discussed below are made accessible by the repository given, where further instructions and descriptions on the methods are provided. The discussion below summarizes the key methods used.

3.1 *Locality Sensitive Hashing*

The first step is to apply LSH to compare images from different periods. What LSH attempts to achieve is a nearest-neighbor analysis using multidimensional datasets. The algorithm's main advantage is its fast-search capabilities, allowing similar data results to be found within a potentially large dataset quickly (Slaney and Casey, 2008). The algorithm can be applied to different types of data, including 2D and 3D images, videos, and texts; the examples used in this work are 2D images. Images, and more specifically pixels, are multidimensional representations that show objects in given scales, depending on the resolution. Such pixels display values that can be informative in any given image. The nearest-neighbor search tries to match similar pixels so that images that are similar receive higher scores. In addition to pixel dimensionality, images have distinct features that represent groups of pixels that can be divided. These groups are called hashes, and they represent values for image pixel data; this is a type of grouping or representation of data that simplifies the pixel values from an image into a data representation of fixed-size values (that is a hash). In the algorithm, the results of similar hashes are placed in the same category, or “bucket”. These buckets assume that there is a reasonable probability these hashes belong together or are very similar. The hash can then be used to represent a given simplified value set for an image that is compared to other values, that is other hashes, used to represent other images. Values representing two similar images placed in the same bucket divisions would likely return as comparable or even the same. This has the benefit of reducing the multidimensionality of image pixels to smaller datasets, allowing a quicker comparison to be made while also analyzing similarities based on hash values. The higher the number of similar hashes and the more they share buckets increases the similarity score for any two images. There are variations of LSH: the algorithm used here is a near duplicate detection LSH that scores images based on similarity, with exact duplicates scoring a value of 1.0 and others scoring lower values to near 0 (Yang et al., 2009). This is based on the hash and bucket alignment between images. Figure 2 summarizes this process, displaying example images that are processed using what is described and the notation given. Effectively, the figure shows hashes derived from rows and bands of pixels, that are then divided into buckets to be used for image comparisons and subsequent similarity scores. The function can be summarized using the following notation for images processed:

$$\begin{gathered} S_i = dh(f, n) \\ \forall b_i \in B: \\ S'_i = S[(b_i + 1) * r] \\ S'_i \in H \end{gathered} \tag{1}$$

where the image signature (S) of each image (i) is obtained by the image input (f) and hash size (n) using a difference hash (dh) function (Fei et al., 2017), which looks at gradient differences in adjacent pixels. Then, for each band (b) number (t), defined by the user as to how to divide an image signature, hashing is portioned for each signature (S') column based on rows (r) in the matrix (defined as n^2/number of bands). The resulting S', that is the modified image signature, is then placed as part of the hash bucket (H) container used in image comparisons for similarity among candidate image pairs. This process effectively divides an image into relevant buckets that can be compared in the next step. The similarity measure is then conducted using the following:

$$\begin{gathered} H \leftarrow H(i, j) \\ \forall i, j \in H: \\ v_{ij} = S'_i + S'_j \\ m_{ij} = \left(n^2 - v_{ij}\right)/n^2 \\ m_{ij} \in D \end{gathered} \tag{2}$$

where candidate image pairs (i, j) are assessed from the hash buckets. The signatures (S') for the candidates are then summed (v) and a similarity score (m) is created using the input hash size (n). The near duplicate set (D) then contains the output similarity score. There is also a similarity threshold, which is a minimum value that can be used to enable m to belong in D, with the threshold left at 0 and allowing all possible candidates for similarity to be placed together even if m is low. Images lacking any similarity were not scored. The Supplementary Material has relevant instructions on how to use the Python code that implements the algorithm discussed here. Overall, the idea is that images that are similar should have at least some shared hash space and score between 0 and 1, with images closer to 1 being near duplicates or very similar.

Insert figure here

Figure 2. Process using LSH on images allowing image comparisons. Hashed rows and bands of images are placed in buckets which are then used to compare to the hashed outputs of other images. Similarity scores are derived based on comparisons of hashed values.

3.2 *Network Outputs*

The next step is to create a network space for all compared images using their modern location information provided in the metadata, which is under the *modern country* column. While the above output is a quantitative metric that measures similarity between two images, two images can also represent a potential link, based on the ancient region where the item was produced (e.g., Aegean, Mesopotamia). Graph theory has been utilized as a means to compare complex decorative patterns along with spatial attributes on ceramics, as an example (Huet, 2018), although here we use a network based approach for assessing regional similarity scores. The images compared from each of the originating regions allow a network output to be created. Link values (in this case two-way links) are median similarity scores for all images within the group of regions under comparison, and can include a region linked to itself in cases where images are from the same region. Median value is used since a single image could be relatively similar or dissimilar, whereas the median value can be representative of a given temporal and regional dataset, helping to minimize values from outliers. Then, the inverse closeness centrality, which measures the sum of the edge similarity values for the compared images, was considered an appropriate measure for establishing centrality, and was used here to determine which nodes, that is which regions, are more central in this network (Hofstad, 2017; Newman, 2010). Median centrality scores were also used, since the dataset is uneven and some regions had more sculpture representation. The centrality metric was used because it allows direct comparisons with multiple regions; node images having the highest centrality value will generally indicate that the node has more images similar to it from different nodes. Areas that show greater similarity to other areas demonstrate that they likely share stylistic traits with a higher number of regions. The network output is used to show regional material culture comparisons utilizing LSH.

3.3 Verification of LSH Method

As the LSH method is new in how it is applied to archaeological problems, and as a key stage in software development (Craig and Jaskiel, 2002), we conducted a verification test to ensure the algorithm works to design. That is, we designed a test to see if the similarity measures from LSH reflect outputs that suggest the LSH approach and outputs at least reflect the intended design. The algorithm provided here is tested with sample data, some of which reflect comparability to the case study data presented, while others are different. The test data and metadata for these data are provided in the Supplementary Material. There are ten test cases that reflect more and less extreme variations of the types of data LSH is used for here. Figure 3 reflects outputs of this test, showing which test examples matched more similarly and which ones did not register any results. Two images (2, 6) are identical, which score as 1 as expected. Another two images (4, 5) are slight variations from each other where the face angles are rotated in one of the images (5), with the score reflecting high image similarity (0.79). Another test includes a non-facial object (1) as an extreme case example that deviates from the others. However, the remaining examples represent statue faces, with many being similar to each other and, in fact, from the same region in cases, which we expected to show some similarity to each other. Two examples are from Cyprus (8, 9) and two are from Egypt (7,1 0), with the remainder being Classical (4, 5) or recent copies of Renaissance works (2, 6). One statue face (3) is from a modern statue example. Factors such as the angle the face is set (e.g., 5, 9), shape of lips (e.g., 2, 5), and chin/cheek curvature (10, 7) are some features that influence results. Overall, the results reflect the types of cases the algorithm should detect, with close matching between statues from similar periods or cultures, while more extreme (i.e., corner cases) examples either reflect an identical statue face or examples that did not compare easily, which should not produce any similarity output. The full results of the verification step are provided in the data outputs.

Insert figure here

Figure 3. Test images and output similarity scores (between 0 and 1) used to verify the LSH algorithm functions to design. The numbers reflect the samples with the metadata given in the Supplementary Material.

4.0 Results

4.1 *Scenario Data*

The results are organized by the range of periods described earlier. For the networks created, center points of regions are used as nodes in results given. While the similarity threshold is left at 0, the analysis applies different parameter value tests for *n*, that is hash size, and *b*, the number of bands. Intervals of 10 were tested between 10-80 for *n* and *b* on all the images. This was done to find potential outputs sensitive to both similar and dissimilar images. Generally, lower *n* values result in more images registering similarity scores in the algorithm's outputs, as there are fewer hashes to place pixel values. A greater number of hashes lowers the likelihood that images will show any similarity result. For *b*, lower values lead to fewer image similarity outputs, as bands are larger and require more pixels to be similar to register comparable scores. High *b* values result in many images registering similarity scores, where images are divided into more bands. While it is difficult to create a clear quantitative procedure to choose the best *b* and *n* values, we chose outputs that clearly registered qualitatively similar images and had low similarity scores for dissimilar items. Consequently, we used *n*=20 and *b*=30 for the analysis. This selection balanced the two types of input parameters while generally producing a lower number of images registering similarity scores. Alternative parameters, however, may also be appropriate for this and other datasets.

4.2 *Pre-9th Century BCE*

We included 62 objects for analysis from this period. The number of sculptures measured per region are: the Aegean (7), Anatolia (2), Egypt (22), Indus (3), Iran (3), Levant (9), and Mesopotamia (16). The unevenness of the samples indicates that not all regions have sculpture images that could easily be found using online sources. Nevertheless, the overall sample set is large enough to determine variation between images.

While the data cover a long period, and certainly there is evidence for regional interaction, statue face elements are very distinct, often even within a region. Link similarity scores are indicated in Figure 4a along with inverse closeness centrality measures that show Egypt as being by far the most central (+23) in having similarities with regional statue styles. This indicates it generally has more similarity with others relative to the regions measured. Outside of Egypt, all other nodes are less than 5, indicating limited similarity of these regions to other areas compared to Egypt. Given that Egypt had the most sculpture images, the centrality result is not surprising. However, the link similarity values, which use median similarity between regions, also suggest more regions have similarity to Egypt than other regions. The median inverse centrality output also shows Egypt as the most central region, demonstrating it is likely to be more central even with a higher relative sample than other regions (Figure 4b). Egypt averaged 0.56 similarity with other regions using the median link values. Median similarity for sculptures within the Aegean are 0.68, which is the highest median score, indicating why the overall median centrality value for the Aegean is also relatively high despite the fact the region has few regional links. Sculptures from the Aegean in this period mainly come from the Cyclades, a limited geographical area. The Indus and Anatolia showed no similarity to any region. Overall, the highest median link similarity values are for those that compare images from the same region. The Levant, Egypt, and Aegean all have median similarity values over 0.6 for links to themselves, that is comparing images within these regions, while Mesopotamia and Iran have a median similarity link value over 0.6. For Mesopotamia and Iran, the examples come from southern Mesopotamia (southern Iraq) and southwestern Iran, suggesting the similarity score is relatively higher due to the fact these are very close areas. Although neighboring regions more commonly have similarity values with some of their neighbors, more distant regions, with the exception of Iran and Egypt and Mesopotamia and Egypt, mostly did not register any similarity.

Insert figure here
Figure 4. Networks showing median link similarity and inverse closeness centrality for nodes (a) and median inverse closeness centrality values (b) for sculptures in the pre-9th century BCE.

4.3 *9th-4th Centuries BCE*

From around the 9th century BCE and into the post-Iron Age period before the rise of Alexander the Great in the late 4th century BCE, increased frequency of empires and larger states in the Near East and neighboring regions became evident (Cline and Graham, 2011). Statues from this period (89 total) are compared similarly to the pre-9th century BCE case. Statues derived from: the Aegean (6), Anatolia (3), Cyprus (28), Egypt (15), Iran (6), Levant (15), and Mesopotamia (16). Figure 5 demonstrates the networked results, where results show generally more regions, including those beyond immediate neighbors, having similarity scores. Cyprus is the most central with a value of 1242 (Figure 5a), but it also has the most samples. Using the median scores for centrality (Figure 5b), the Aegean and Egypt are comparable at about 0.52. This demonstrates that relative to their sample numbers, these areas displayed relatively high centrality. This time, we see that most regions have some similarity even with more distant regions, although more distant regions are generally less similar. Most of the centrality can be seen in the eastern Mediterranean, including the Aegean, Cyprus, Egypt and Anatolia, while weaker centrality values are evident in eastern regions (Mesopotamia and Iran). Link similarity showed median values between 0.5-0.54, which is lower than the range of similarity in Figure 4. Overall, the results highlight that generally more regions are comparable than the previous case. In part, this is because there is more data for this period. However, the overall number of linkages and higher median centrality scores across the region demonstrates that more regions begin to have more similar characteristics in their sculptures.

Insert figure here
Figure 5. Networks showing median link similarity and inverse closeness centrality for nodes (a) and median inverse closeness centrality values (b) for sculptures between the 9th and late 4th centuries BCE.

4.4 *Post-4th Century BCE*

Data for this period represent 82 objects, covering from the end of the 4th century BCE to about 300 CE or slightly later. The regions compared include: the Aegean (18), Anatolia (5), Cyprus (13), Egypt (15), Gandhara (18), Iran (1), and Mesopotamia (12). Overall, there are more relatively central nodes, with the Aegean (75) having the highest inverse closeness centrality (Figure 6a), indicating multiple areas showing regional similarity to each other. Looking at median centrality scores, Cyprus, Mesopotamia, and Gandhara showed the highest scores at around 0.53 (Figure 6b). Even in distant regions, such as Gandhara and the Aegean (0.525), relatively high link similarity is evident in Figure 6. While Iran only has one example in this case, this example has some similarity, indicating many sculptures showed similar traits. The multiple connections across neighboring regions indicates broader similarity over wider distances relative to previous examples. Furthermore, centrality scores are now higher in areas more distant than the eastern Mediterranean. The Roman examples in the dataset, which spanned across the eastern Mediterranean region including from the Aegean, Cyprus, Egypt, and Mesopotamia, do stand out, where the median link values for all these regions in that cultural period is over 0.57. In fact, many individual images scored well over 0.7, which are among the highest similarity values in this dataset.

Insert figure here
Figure 6. Networks showing median link similarity and inverse closeness centrality for nodes (a) and median inverse closeness centrality values (b) for sculptures from the late 4th century BCE to about 300 CE.

4.5 *All Periods and Regions*

In looking at all of the dataset and images, regions can be compared to each other to see how similar they are to other regions over long periods. This shows which regions generally link strongly. While Egypt and Mesopotamia have the highest centrality (Figure 7a), mainly because they have a large number of samples, using only median outputs (Figure 7b) show values appear to converge across the broader regions investigated between 0.51-0.53. Both Egypt (0.542) and Gandhara (0.54) have the highest overall mean link similarity for images. In this case, both these values are for images having similarity to other images from the same region. Particularly for Egypt, this result demonstrates a long-term consistency in sculptures, as sculptures dated from the third millennium BCE to the late Roman period, where many of the earlier and later images often showed some similarity. In the case of Gandhara, the temporal range is more limited.

Insert figure here
Figure 7. Networks showing median link similarity and inverse closeness centrality for nodes (a) and median inverse closeness centrality values (b) for sculptures from all the periods analyzed.

4.6 *Sculpture Similarity Across Time and Regions*

Image data can be summarized by looking at all the different periods statues date to and seeing what similarity scores are achieved for images (Figure 8). The results show earlier statue faces generally have higher median similarity scores, often over 0.55, although there are far fewer examples prior to 1000 BCE. On the other hand, those after 1000 BCE have median similarity values ranging closer to 0.51-0.53, but there are greater ranges in individual values with images scoring well over 0.7 in some cases and well below 0.45 in others. To further look into how the data change in the periods analyzed, and using kernel density estimation to smooth the statue image data, it is evident that similarity values for all images compared demonstrate greater convergence later in time when comparing within a region and between two regions, that is within a node and between two nodes (Figure 9). Figure 9a-b represents images from the pre-9th century BCE for within and between node similarity respectively. The kernel density measurements show a higher portion of within node (that is within a region) comparisons having similarity values at or greater than 0.654, which is the median value. This is the highest median value achieved for any type of node comparisons for period categories compared. It shows that the within region similarity scores higher than in other periods. On the other hand, Figure 9b shows comparisons between different nodes, that is between different regions, which have a median around 0.565. This means that the within region similarity is higher than the between region similarity for this period. For measurements in the 9th-4th centuries BCE (Figure 9c-d) and 4th century BCE and later (Figure 9e-f), the similarity values are more similar when looking at within (c and e) and between node (d and f) comparisons. In this case, the distributions for similarity values are much closer to normal distributions. For the 9th-4th centuries BCE images, median similarity for both within and between node similarity is around 0.52. For the 4th century BCE and later, the image similarity median values are 0.54 and 0.53 for within and between node comparisons respectively. Overall, this suggests from the 9th century BCE and later, sculptures within regions as well as sculptures between regions that are similar show generally less differences or score more comparably. The pre-9th century BCE figures have the highest median similarity scores, although the samples are the smallest set from the three datasets.

Insert figure here
Figure 8. Box and whiskers plots for similarity scores in different periods. Negative values represent BCE dates and positive are CE dates for given centuries. Dates represent the earliest date given for statues.

Insert figure here
Figure 9. Kernel density measures for similarity link values with median (dotted lines) values for all links. Graphs a, c, and e indicate images compared to other images from the same node (region); graphs b, d, and f demonstrate comparisons between images from different nodes. Graphs a and b are pre-9th century BCE images, graphs c and d are 9th-4th centuries BCE, and graphs e and f are for those from the late 4th century BCE and later.

5.0 Discussion

5.1 *Benefits of the Methodology*

The LSH method presents several benefits that are applicable to the test case as well as potentially other material culture case studies. First, the algorithm enables rapid analysis, where 233 image samples can be analyzed in seconds on a simple desktop computer. While a much larger dataset could slow the analysis, the benefit is the number of images could be in the thousands and overall size of images on the order of gigabytes before the analysis would need to be distributed to clusters or high performance systems. This would facilitate analysis for those who cannot easily access distributed resources. Second, comparisons did not require training data. Only two images are required for the analysis to work. While deep learning classification could be powerful for analysis, often simply finding similarity is all that is needed to generate informative results. Third, images compared have very different dimensionality, meaning that no adjustment or standardization of image dimensions is required in this analysis. This is useful in helping to facilitate the image processing stage of analysis, as no adjustments are required. Fourth, and more importantly, the analysis produced results that are informative and matched our qualitative expectation for the given material cultures analyzed. Qualitatively, the statue faces in the pre-9th century BCE are far more diverse than those from later periods. Although similarity scores averaged higher values in this period, far fewer statues recorded any similarity values, with only 30 outputs from a sample of 62. On the other hand, the other periods each registered over 3000 similarity scores. The results suggest that the convergence in material culture similarity in the centuries after the 9th century, also evident qualitatively, are apparent in results. Even though the results are expected, they demonstrate the analysis can generate a quantitative output for object similarity. This quantitative approach can potentially be applied alongside with traditional approaches to help researchers processing large image datasets and obtain quantitative evaluations of object similarity that can then be compared with evaluations obtained through traditional methods. Fifth, extending this approach to potentially other forms of objects, including objects that are less known or common objects such as ceramics and coins, is a possibility with simple 2D photographs. Ideally, 3D data with high quality images would be the best approach, but the case study deliberately used simple photographs to demonstrate that legacy datasets are also effective for analysis. As legacy data make-up the vast majority of available data, analyzes should make use of these data where possible. Outside of having to select hash size and bands, the method is automated in the analysis phase. Data pre-processing or adjustments made prior to analysis, although executed here, might not be necessary, as hashes within an image would inform if there is some similarity between images that are not cropped or masked in any way. In other words, if part of a scene is similar to another image, then some similarity score should return. Data organization, by time and region in this case, might, however, be needed to enable easier archaeological or material culture interpretation of results. We provide the code in this work in the hope it can be applied by others to different cultural objects or even parts of objects compared for similarity.

5.2 *Insights to Material Culture*

To further demonstrate how the approach taken here benefits material culture interpretation, we briefly discuss further insights obtained for the cases analyzed. Object similarities, material culture hybridization and the blending of different styles have long been the focus for archaeologists and art historians alike, who have often linked these phenomena to intercultural interactions (e.g., Hodos, 2012; Burkert, 2007). Concepts such as 'international', 'intercultural' or 'hybrid' styles as opposed to 'regional' or 'local' styles have been frequently used to label the blending of stylistic motifs developed in a particular region with motifs borrowed from foreign cultural milieus (Pfälzner, 2015; Feldman, 2006). The concept of style in both archaeology and history of art has received many theoretical definitions, though key elements of these definitions concern the presence of shared and recurrent motifs on objects that, taken together, represent an ensemble suitable to assess continuity (or discontinuity) in cultural traditions as well as intercultural interactions visually expressed (Sanz and Fiore, 2014; Conkey and Hastorf, 1990). The results of LSH and using the network approach advanced demonstrate that quantitative methods can effectively aid the identification of cross-cultural material culture similarities and show the strength of those similarities. Nevertheless, it should be noted that quantitative methods such as that one we propose are not intended to substitute traditional methods of stylistic evaluations. In the case proposed here where we apply LSH method to statue faces, we are limiting the definition of "style" to the presence of recurrent features across images and their digital similarity. However, we believe that the quantitative assessment of material culture similarity can offer valid insights on the interactions among different cultures and how these are reflected visually. The increased convergence and more objects showing similarity over time from the 9th century BCE can be interpreted to be the result of long-term empire formation and continuity in western Asia, which increased close interaction among people of different cultural backgrounds, including artisans and merchants, leading to syncretistic developments in various cultural traditions (Altaweel and Squitieri, 2018). This is part of a process termed as 'universalism' or 'globalization' as some have called it, which saw states increasingly share common languages, art, religions, and other cultural phenomena as large scale empires emerged. In other words, for the sculptures analyzed, the results reflect a syncretistic development that is comparable to other forms of material culture and wider cultural phenomena during the 1st millennium BCE and 1st millennium CE.

5.3 Selection Criteria for Users

We recognize that the methods deployed in this work are new and, therefore, to best use them users may require to deploy selection criteria to select appropriate use case samples. First, we do think users should select cases they are reasonably comfortable in understanding or at least work with regional and cultural experts who have an understanding of the samples so results can be qualitatively checked. Material cultural experience means results can be evaluated or at least investigated with given expertise. For instance, in this work, similarity scores did align with qualitative discussion in other works (e.g., Altaweel and Squitieri, 2018; Behrendt, 2007; Burn, 2004). Secondly, while a strength of the method is it works with a minimal dataset, clearly having more samples improves confidence in results. We also suggest that samples come from spatially diffuse areas so that comparisons have greater meaning. Only comparing results from within a site, for instance, may not show sufficient differences among comparable objects, although comparing those objects to others from different sites might. This leads to the third point, which is objects should be comparable. Users of the applied methods should be sure they are comparing the same general type of object and that areas within objects that are compared are similar. Ideally, the images should be taken from the same angle or similar angles to minimize the effects of perspective on the outputs. Damage on the objects should be limited in areas compared. For instance, if comparing ceramics then we suggest selecting pottery types such as plates or bowls and comparing to other plates and bowls with images from the same angle and the same areas within objects. If patterns of decoration are compared, then only the decoration areas on artifacts should be used in comparisons. We also noticed that selecting samples with similar levels of preservation worked the best in results. Potentially some artifacts could be altered over time based on natural or human-made alterations conducted in the past, including preservation affecting objects. From our examples, we see that alterations or changes to the primary object or area compared should be minimal, or at least controlled for, to increase effectiveness of LSH, although some damage or changes could be tolerated as not all of our statues were perfectly preserved. Fourthly, while we could have enhanced our dataset by selecting more samples, it was also evident that not all potential samples have sufficient data quality. Users should be sure that key distinctive features compared, such as noses, eyes, and cheeks in our case, are evident and of reasonable quality. This means pixels need to be informative and capture features to sufficient detail. There is a potential that varied image quality may produce uneven results. We did find that samples lower than 5 kB in size tended to have less comparability to larger images (e.g., greater than 50 kB), which reflects unevenness in image quality and pixel values (see below). If users create their own data, we suggest comparing images with comparable image quality. While the algorithm is flexible to allow variation in quality, the interpretive value is likely to be better with closer quality data.

5.4 *Limitations*

From conducting the analysis, some limitations are evident. One limitation is that our work was applied to images we can find using online resources and those from reputable institutions. The Indus had the least samples, which likely makes the region's result the least clear, but it also demonstrates that the method can work with a small set of samples and still demonstrate comparability with other regional samples. Ideally, we would also have more data from regions such as the Levant and Anatolia to demonstrate regional differences, but it proved difficult to find more samples. Many sculptures found are often those sold in the antiquities market, which we did not use. Additionally, some of the data used are from unclear provenance found within a region. While location could be generalized to a region, specifying the site objects come from would have enabled a more complex network to be built since an analysis across sites, rather than regions, would have been possible. This could potentially show how some sites are more or less similar to other sites rather than having to generalize at a regional level as we did. It is a strength that the approach could use legacy and limited data, but this is not always ideal given image quality might be poor among samples present and exceptions in cases could more heavily influence results. We found, for instance, images that are lower than 5 kB did not compare as well to higher quality images (more than 50 kB). Ideally, image quality should be more than 10 kB and image quality similar for all images to make suitable comparisons, although variations in quality can be tolerated and useful outputs still obtained. Images are not always taken from appropriate angles, which are angles that are comparable for all images, with the results potentially being less clear, but, as stated above, areas showing similarity within objects still have potential to register similarity so long as similar areas are present in images even if taken from different angles. Nevertheless, the results are likely to register a lower similarity score than what would be expected, which would limit the approach. Additionally, the methods advanced are likely to work best with objects that have only minor damage, which is often not possible in archaeology. Heavily damaged objects produce very different and inconclusive results, but if large parts of an object are preserved the analysis may still be possible.

6.0 Conclusion

6.1 *Future Work*

The versatility of the proposed method makes it suitable for several applications in the field of image comparison in archaeology and affiliated disciplines. It is a valid aid to the comparison of objects and the study of their cultural similarity. More samples of objects and images could make results more accurate, as few samples could simply produce outliers or exceptional cases. Having good quality images also can improve results. Future applications could target larger datasets with visual data such as ceramics, cylinder seals, and coins, which makeup a large number of items for some cultures in the Near East and neighboring regions. However, other very different forms of material culture could be deployed, including paintings and jewelery. Different media can be made such as 3D images; new databases for objects could be analyzed using such data. The approach could also be adapted to mobile devices by creating an application with the code, enabling it to be a tool that can be used in the field or by those processing archaeological data. Archaeological projects could deploy the code to compare objects excavated with existing databases or images to see how similar objects are. Other potential work using this effort includes developing approaches that reconstruct broken objects and then comparing items using comparable methods to what was advanced here.

7.0 Acknowledgements

Left out for the purpose of review.

8.0 References

Altaweel, M., Squitieri, A., 2018. Revolutionizing a world: from small states to universalism in the Pre-Islamic Near East. UCL Press, London.

Behrendt, K. A., 2007. The Art of Gandhara in the Metropolitan Museum of Art. Yale University Press, New Haven, CT.

Berg, A.C., Berg, T.L., Malik, J., 2005. Shape matching and object recognition using Low Distortion Correspondences. In: 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’05). Presented at the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’05), IEEE. San Diego, CA, pp. 26–33. https://doi.org/10.1109/CVPR.2005.320.

Bevan, A., Li, X., Martinón-Torres, M., Green, S., Xia, Y., Zhao, K., Zhao, Z., Ma, S., Cao, W., Rehren, T., 2014. Computer vision, archaeological classification and China's terracotta warriors. Journal of Archaeological Science 49, 249–254. https://doi.org/10.1016/j.jas.2014.05.014.

Burkert, W., 2007. Babylon, Memphis, Persepolis: eastern contexts of Greek culture. Harvard University Press, Cambridge, Mass.; London.

Burn, L., 2004. Hellenistic Art: from Alexander the Great to Augustus. British Museum Press, London.

Cline, E.H., Graham, M.W., 2011. Ancient empires: from Mesopotamia to the rise of Islam. University Press, Cambridge; New York.

Conkey, M., Hastorf, C. 1990. Introduction, In: Conkey, M. and Hastorf, C. (Eds.). The uses of style in archaeology: 1-4. Cambridge University Press, Cambridge.

Craig, R.D., Jaskiel, S.P., 2002. Systematic software testing, Artech House computing library. Artech House, Boston.

Creative Commons, 2020. https://search.creativecommons.org/. Accessed 5/1/2020.

Fei, M., Ju, Z., Zhen, X., Li, J., 2017. Real-time visual tracking based on improved perceptual hashing. Multimed Tools Appl 76, 4617–4634. https://doi.org/10.1007/s11042-016-3723-5

Feldman, M. H. 2006. Diplomacy by design. Chicago University Press, Chicago.

Forrester-Sellers, J., Velasco, A., Pickering, R., Diaz, J.C., 2017. Classifying ancient west Mexican ceramic figures using three-dimensional modelling and machine learning. Presented at the Electronic Visualisation and the Arts (EVA 2017). https://doi.org/10.14236/ewic/EVA2017.3.

Foster, D., 2019. Generative deep learning: teaching machines to paint, write, compose, and play. O'Reilly Media, Sebastopol, CA.

Gansell, A.R., Meent, J.-W. van de, Zairis, S., Wiggins, C.H., 2014. Stylistic clusters and the Syrian/South Syrian tradition of first-millennium BCE Levantine ivory carving: a machine learning approach. Journal of Archaeological Science 44, 194–205. https://doi.org/10.1016/j.jas.2013.11.005.

Hill, M., 2016. Egypt in the Ptolemaic Period. In: Heilbrunn Timeline of Art History: The Metropolitan Museum of Art, 2000–, New York. http://www.metmuseum.org/toah/hd/ptol/hd_ptol.htm.

Hodos, T., 2012. Cyprus and the Near East, In: Smith, T.J., Plantzos, D. (Eds.), A Companion to Greek Art. Wiley, pp. 312–329. https://doi.org/10.1002/9781118273289.ch16

Hofstad, R. van der, 2017. Random graphs and complex networks. Cambridge series in statistical and probabilistic mathematics. Cambridge University Press, Cambridge.

Huet, T., 2018. Geometric graphs to study ceramics. In: Matsumoto, M., Uleberg, E. (Eds.), CAA2016: Oceans of Data Proceedings of the 44th Conference on Computer Applications and Quantitative Methods in Archaeology. Archaeopress, pp. 311-324.

Invernizzi, A., 2012. Remarks on the intercultural encounters in the Hellenized Orient. Parthica 14, 89–108.

Jegou, H., Amsaleg, L., Schmid, C., Gros, P., 2008. Query adaptative locality sensitive hashing. In: 2008 IEEE International Conference on Acoustics, Speech and Signal Processing. Presented at the ICASSP 2008 - 2008 IEEE International Conference on Acoustics, Speech and Signal Processing, IEEE, Las Vegas, NV, USA, pp. 825–828. https://doi.org/10.1109/ICASSP.2008.4517737.

Kacem, N., Zagrouba, E., Barhoumi, W., Gallas, A., 2015. Locality-sensitive hashing for region-based large-scale image indexing. IET Image Processing 9, 804–810. https://doi.org/10.1049/iet-ipr.2014.0910.

Kulis, B., Grauman, K., 2012. Kernelized Locality-Sensitive Hashing. *I*EEE Trans. Pattern Anal. Mach. Intell. 34, 1092–1104. https://doi.org/10.1109/TPAMI.2011.219.

Lehmann, E.L., Casella, G., 2001. Theory of Point Estimation {Springer Texts in Statistics}. Springer-Verlag New York Inc., Dordrecht.

Lowe, D.G., 1999. Object recognition from local scale-invariant features. In: Proceedings of the Seventh IEEE International Conference on Computer Vision. Presented at the Proceedings of the Seventh IEEE International Conference on Computer Vision, IEEE, Kerkyra, Greece, pp. 1150–1157 vol.2. https://doi.org/10.1109/ICCV.1999.790410.

Newman, M.E.J., 2010. Networks: an introduction. Oxford University Press, Oxford; New York.

Oliva, A., Torralba, A., 2001. Modeling the shape of the scene: A holistic representation of the spatial envelope. International Journal of Computer Vision 42, 145-175.

Paulevé, L., Jégou, H., Amsaleg, L., 2010. Locality sensitive hashing: A comparison of hash function types and querying mechanisms. Pattern Recognition Letters 31, 1348–1358. https://doi.org/10.1016/j.patrec.2010.04.004.

Pfälzner, P. 2015, The Art of Qatna and the question of "international style". In: Pfäalzner, P. and al-Maqdissi, M. (Eds.), Qatna and the networks of Bronze Age Globalism. Harrassowitz, Wiesbaden, pp. 33-41.

Rawat, W., Wang, Z., 2017. Deep Convolutional Neural Networks for Image Classification: A Comprehensive Review. Neural Computation 29, 2352–2449. https://doi.org/10.1162/neco_a_00990.

Sanz, I. D. and Fiore, D. 2014. Style: its role in the Archaeology of Art. In: Smith, C. (ed.), Encyclopedia of Global Archaeology. Springer-Verlag New York, pp. 7104-7111

Sebastian, T.B., Klein, P.N., Kimia, B.B., 2002. Shock-Based Indexing into Large Shape Databases. In: Heyden, A., Sparr, G., Nielsen, M., Johansen, P. (Eds.). Computer Vision — ECCV 2002. Springer, Berlin; Heidelberg, pp. 731–746. https://doi.org/10.1007/3-540-47977-5_48

Siddiqi, K., Shokoufandeh, A., Dickinson, S.J., Zucker, S.W., 1999. Shock graphs and shape matching. International Journal of Computer Vision 35, 13–32. https://doi.org/10.1023/A:1008102926703

Slaney, M., Casey, M., 2008. Locality-Sensitive Hashing for finding Nearest Neighbors [Lecture Notes]. IEEE Signal Processing Magazine 25(2), 128–131. https://doi.org/10.1109/MSP.2007.914237

Srivastava, N., Salakhutdinov, R., 2014. Multimodal learning with deep boltzmann machines. Journal of Machine Learning Research 15, 2949-2980.

Stiebing, W.H., Helft, S.N., 2018. Ancient near eastern history and culture. Third edition. Routledge, New York.

Stockhammer, P. 2013. From hybdridity to entanglement, from essentialism to practice. Archaeology Review from Cambridge 28 (1), 11-28.

Tsai, Y.-H., Yang, M.-H., 2014. Locality preserving hashing. In: 2014 IEEE International Conference on Image Processing (ICIP). Presented at the 2014 IEEE International Conference on Image Processing (ICIP), IEEE, Paris, pp. 2988–2992. https://doi.org/10.1109/ICIP.2014.7025604.

Van De Mieroop, M., 2004. A History of the Ancient Near East. Blackwell, Oxford.

Wang, J.Z., Ge, W., Snow, D.R., Mitra, P., Giles, C.L., 2010. Determining the sexual identities of prehistoric cave artists using digitized handprints: a machine learning approach. In: Proceedings of the International Conference on Multimedia - MM '10. Presented at the international conference, ACM Press, Firenze, p. 1325. https://doi.org/10.1145/1873951.1874214.

Waters, M., 2014. Ancient Persia: a concise history of the Achaemenid Empire, 550– 330 BCE. Cambridge University Press, New York.

Wiesehöfer, J., 1996. Ancient Persia: From 550 BC to 650 AD. Trans. Azizeh Azodi. I.B. Tauris, London.

Winter, I. J., 2005, Establishing group boundaries: toward methodological refinement in the determination of sets as a prior condition to the analysis of cultural contact and/or innovation in the 1st millennium BCE ivory carving. In: Suter, C. E. and Uehlinger, C. (Eds.), Craft and Images in Contact. Academic Press, Fribourg, pp. 23-42.

Yang, X., Zhu, Q., Cheng, K.-T., 2009. Near-duplicate detection for images and videos. In: Proceedings of the First ACM Workshop on Large-Scale Multimedia Retrieval and Mining - LS-MMRM '09. Presented at the First ACM workshop, ACM Press, Beijing, China, p. 73. https://doi.org/10.1145/1631058.1631073.

Zaccagnini, C., 1983. Patterns of mobility among ancient Near Eastern craftsmen. Journal of Near Eastern Studies 42(4), 245– 64.