Surface Partial Matching & Application to Archaeology

Arik Itskovich and Ayellet Tal

Technion

Abstract

Partial matching is a fundamental problem in shape analysis, a field that is recently gaining an increasing importance in computer graphics. This paper proposes a novel approach to performing partial matching of surfaces. Given two surfaces $\mathcal{M}_{\mathcal{A}}$ and $\mathcal{M}_{\mathcal{B}}$, our goal is to find the best match to $\mathcal{M}_{\mathcal{A}}$ within $\mathcal{M}_{\mathcal{B}}$. The key idea of our approach is to integrate feature-point similarity and segment similarity. Specifically, we introduce a probabilistic framework in which the segmentation and the correspondences of neighboring feature points allow us to enhance or moderate our certainty of a feature-point similarity. The utility of our algorithm is demonstrated in the domain of archaeology, where digital archiving is becoming ever more widespread. In this domain, automatic matching can serve as a worthy alternative to the expensive and time-consuming manual procedure that is used today.

Keywords: Matching, similarity, archaeology, cultural heritage

1. Introduction

Finding correspondences between objects is a fundamental problem in computer vision, computer graphics, and geometry processing, with applications in a variety of areas, such as medicine, biology and archaeology [1, 2]. This paper focuses on matching surfaces, which are represented as polygonal meshes. Previous work mainly concentrated on determining the similarity of whole surfaces. There exists a large amount of work related to this topic; see [1, 3, 4] for comprehensive surveys. In certain applications, however, this does not suffice – we wish to find only similar sub-surfaces. The added difficulty stems from the fact that helpful global techniques, such as scaling, alignment, or symmetry cannot be utilized.

A reasonable avenue to address partial matching is to examine local feature points, e.g., [5, 6, 7, 8]. In contrast to global descriptors, which attempt to represent the overall shape of the surface, local descriptors are defined only for salient points of the surface. However, there are often many points on the surface with similar local descriptors. Which of them should be selected?

The term partial matching has been used to describe three different variants: part-in-whole matching, part-to-part matching, and partial

Preprint submitted to Computer & Graphics

matching. Part-in-whole matching attempts to match whole objects, but with special emphasis on a selected part, thus taking into account the context within which the part is defined [9, 10]. Part-to-part matching attempts to match subregions of two input surfaces, where the parts to be matched are not known beforehand [11, 12, 13, 14, 15, 16, 17]. This is important for applications such as alignment and registration of partial views of 3D objects acquired by a 3D scanner, as well as for self-similarity and other shape analysis tasks [18]. The third category of *partial matching* is the problem addressed in this paper. Given a specific part of an unknown surface, our goal is to detect similar parts on other surfaces, regardless of the global surface this part belongs to. Since the surface from which the query is extracted is unknown, neither global information nor the scale and the orientation of this part relative to the whole object, can be utilized.

We propose an algorithm for detecting partial matching of surfaces. In [19], an interesting approach is presented, in which the objects in the database are processed and decomposed into segments. We avoid such pre-processing and maintain only the whole objects. Our key observation is that though isolated feature points often do not suffice, their aggregation provides adequate information re-



Figure 1: (a): Matching a dancer from a broken Hellenistic (first century BCE) amphora (top left) to dancers on a vase (right). (b): Detecting cupids on Hellenistic oil lamps

garding similarity. We introduce a probabilistic framework in which segmentation and neighboring feature points allow us to enhance or moderate the certainty of feature similarity. We show that this scheme manages to detect partial similarity automatically, even when the data is very noisy and the scales are unknown.

We demonstrate the usefulness of our algorithm for archaeology (Figure 1). This domain is an intriguing application of computer graphics techniques [20] and a variety of topics has been investigated including matching fragments [21, 22], shape illustration [23, 24], relief extraction [25], and more.

We focus on partial similarity, which is a central task in archaeology. After finding a new artifact, the archaeologist aims at locating it in time and space. Knowing where else in the world similar artifacts were found, reflects on commerce, mutual influence, etc.

Traditionally, archaeological artifacts are drawn by hand and printed in the reports of archaeological excavations – an extremely expensive and timeconsuming procedure. This aids the archaeologist to compare artifacts without actually holding them in her hand. Such drawings are often inaccurate, since the precision of the drawing depends on the qualifications of the artist. In addition, this technique does not always suffice due to space limitations that force the archaeologist to choose which objects will be drawn and decide on a small fixed set of viewing directions. Digitizing the findings by a high resolution scanner and creating archaeological databases will be a welcome alternative.

Since most artifacts are found broken, whole similarity and part-in-whole similarity are not helpful. Part-to-part matching is also insufficient, since the archaeologists usually wish to mark the "meaningful" parts they are looking for. Instead, partial similarity is the only alternative. An added difficulty in this domain is that the objects are often eroded or noisy, so the algorithm should be fundamentally robust.

The contribution of this paper is hence twofold. First, we propose a novel technique for partial similarity, which handles noisy data (Sections 2-5). The major novelty of this technique is a probabilistic method for integrating point similarity and segment similarity. In addition, both point similarity and segment similarity are discussed and new algorithms for addressing them are introduced. Second, we demonstrate its applicability in a challenging domain – archaeology.

2. General approach

Given two surfaces $\mathcal{M}_{\mathcal{A}}$ and $\mathcal{M}_{\mathcal{B}}$, the goal is to find the best match of $\mathcal{M}_{\mathcal{A}}$ within $\mathcal{M}_{\mathcal{B}}$. The key idea of our method is to base the search not only on feature similarity, but also on segment similarity. We show that combining these similarity measures is beneficial. This general scheme is not restricted to any specific feature-point similarity or

Algorithm 1 Partial matching

Input: Meshes $\mathcal{M}_{\mathcal{A}}$ and $\mathcal{M}_{\mathcal{B}}$.

Output: The region of $\mathcal{M}_{\mathcal{B}}$ that best matches $\mathcal{M}_{\mathcal{A}}$.

- 1: Point similarity:
 - 1.1 Detect a set of salient points (vertices) both on $\mathcal{M}_{\mathcal{A}}$ and on $\mathcal{M}_{\mathcal{B}}$.
 - 1.2 Compute the local descriptors for the salient points.
 - 1.3 Measure the similarity between the salient points of $\mathcal{M}_{\mathcal{A}}$ and those of $\mathcal{M}_{\mathcal{B}}$.
- 2: Segment similarity:
 - 2.1 Segment $\mathcal{M}_{\mathcal{A}}$ and $\mathcal{M}_{\mathcal{B}}$.
 - 2.2 Compute descriptors for the segments.
- 2.3 Measure the similarity between the segment descriptors of $\mathcal{M}_{\mathcal{A}}$ and $\mathcal{M}_{\mathcal{B}}$.
- 3: Integrate point similarity (Stage 1) and segment similarity (Stage 2).
- 4: Obtain the region(s) in $\mathcal{M}_{\mathcal{B}}$ that is most similar to $\mathcal{M}_{\mathcal{A}}$.

segment similarity, and can successfully integrate other types of similarity. Nevertheless, we propose specific methods for performing feature-point matching and segment matching. In practice, the surfaces are given as polygonal meshes, represented by sets of vertices and faces.

Our algorithm is outlined in Algorithm 1 and illustrated in Figure 2. First, the salient points are detected and their similarity is computed. Avoiding the consideration of all the vertices of the mesh not only improves the performance, but also enhances the results, since non-distinctive vertices are ignored. Then, both $\mathcal{M}_{\mathcal{A}}$ and $\mathcal{M}_{\mathcal{B}}$ are segmented into meaningful components and their segments are matched. Next, given the above similarity measures, they are integrated. The goal is to compute consistent correspondences between the salient vertices. Finally (Step 4), the similar region(s) in $\mathcal{M}_{\mathcal{B}}$ is determined according to the correspondence established in the previous stage. The bounding ellipsoid of the partial model is aligned with the detected region in $\mathcal{M}_{\mathcal{B}}$ using Principal Component Analysis (PCA). If a more precise alignment, which follows the meaningful boundaries, is sought, it is determined by employing the ICP algorithm [26]. We elaborate on Stages 1-3 in Sections 3-5.

3. Salient points – detection & similarity

This section defines the saliency of points (vertices in the polygonal mesh representation) and proposes algorithms for detecting them, computing their descriptors, and measuring their similarity (Stage 1 of Algorithm 1).

Salient vertices: Our definition of salient vertices is inspired by SIFT [27] and is similar to [15]. Recall that in scale-space a signal is represented as a family of smoothed signals, the *scale-space representation*, parameterized by the size s of the smoothing kernel used for suppressing fine-scale structures [28]. Given an input surface, the scale-space representation is a series of surface approximations, at coarser and coarser scales of detail, constructed using geometric low-pass filtering. These surfaces are then encoded relative to each other by expressing each level as a scalar displacement of its predecessor. Intuitively, as the scale grows, smoother versions of the surface are obtained. Over time the small disturbances thin out, smoothing the high frequencies, while the general shape is only slightly degraded. The difference between two versions of the surface, smoothed at two scales s_1 and s_2 , occurs at regions that contain features of scales between s_1 and s_2 .

Let the *normal displacement* of a point be the difference between its locations in the direction of the normal in two adjacent levels. (In practice we use 10 scales.)

Definition 3..1. A salient vertex is a vertex for which the normal displacement between consecutive levels of scale space is a local extremum.

Vertex descriptor: We are seeking a local descriptor that is invariant to similarity transformations (translation, rotation, and scaling). We propose a combination of two normalized signatures – the shape index S_I and the Willmore energy W.

The shape index is a local signature that provides a quantitative measure of the shape [5]. It is invariant to scale and Euclidean transformations. Let k_1 and k_2 be the maximal and the minimal principal curvatures respectively. At point p the shape index



Figure 2: Algorithm Outline. Given a cupid-shaped relief extracted from a Hellenistic oil lamp $\mathcal{M}_{\mathcal{A}}$ (top) and a different Hellenistic oil lamp $\mathcal{M}_{\mathcal{B}}$ (bottom), our algorithm detects the part within $\mathcal{M}_{\mathcal{B}}$ that is most similar to $\mathcal{M}_{\mathcal{A}}$. The two cupids (query and result) differ both in shape and in pose. (The other cupid of $\mathcal{M}_{\mathcal{B}}$ will be found as a second match.)

is defined as

$$S_I(p) = \frac{1}{2} - \frac{1}{\pi} \arctan \frac{k_1(p) + k_2(p)}{k_1(p) - k_2(p)}.$$
 (1)

The Willmore energy measures the amount of deviation of a surface S from a sphere [29]. Let H be the mean curvature and K be the Gaussian curvature. The Willmore energy is defined as:

$$\mathcal{W} = \int_{S} H^2 \,\mathrm{dA} - \int_{S} K \,\mathrm{dA},\tag{2}$$

where dA is the area. Since the integrand $(H^2 - K)$ dA is invariant to Möbius transformations, W is a conformal invariant of the surface.

While the shape index represents the shape well, it falls short of discriminating between certain shapes, such as shapes having $k_2 \approx 0$. The Willmore energy does not distinguish between convex and concave shapes. Their combination into a single 2D vector, manages to better distinguish between shapes. This is illustrated in Figure 3. Given the query point in magenta (left), using only the Willmore energy, the points in magenta and green (right) equally match. Using only the shape index signature, the points in magenta and blue equally match. The combined signature retrieves only the magenta point, as expected.

Vertex similarity & confidence: For each salient vertex $v \in \mathcal{M}_{\mathcal{A}}$, the Euclidean distances dist between its descriptor to those of the vertices of $\mathcal{M}_{\mathcal{B}}$ are computed. A pair whose distance is below a certain threshold is considered similar. In practice, this threshold is set to 0.01.

We associate with each corresponding pair an initial probability value, which expresses our confi-



Figure 3: The Willmore energy matches the magenta point also to the green point, whereas the shape index matches it also to the blue point. The combined signature retrieves only the matching magenta point.

dence in this correspondence. A pair will get a high probability if the distance between the descriptors of its vertices is small. This is realized by setting

$$\mathcal{P}\left((v,u)\right) = \frac{1}{\varepsilon + dist(descriptor(v), descriptor(u))}$$

where in our implementation $\varepsilon = 1e - 5$.

4. Segmentation – computation & similarity

This section introduces our segmentation technique and our segment descriptors and then proposes a novel algorithm for measuring segment similarity (Stage 2 of Algorithm 1).

Segmentation: Given a mesh (either $\mathcal{M}_{\mathcal{A}}$ or $\mathcal{M}_{\mathcal{B}}$), the goal is to efficiently segment it into meaningful components. We base our algorithm on the general linear region growing scheme of [30]. This is an incremental strategy of starting from a random face and collecting faces along the way until no adjacent faces can be found whose distance to the current patch is sufficiently small. The benefit of this scheme is its efficiency.

The major consideration is how to define the distances between faces. Differently from [30], our distance is defined as follows. For each face f, we consider the mean curvature of the face H(f) as the average curvature of its vertices. Then, the distance between two faces is defined as:

$$dist(f_1, f_2) = \eta |H(f_1) - H(f_2)|,$$

where η is a small positive constant (0.1 in our implementation) if the sign of $H(f_1)$ is the same as the sign of $H(f_2)$ and 1 otherwise.

Other segmentation algorithms can be utilized as well, depending on the application at hand. For instance, for a few objects (Figures 1(a), 10(a), 11), we use the relief segmentation algorithm of [25].

Segment descriptor: The descriptor of a segment is the 3D Shape Spectrum Descriptor (SSD) of [31]. It is defined as the distribution of the shape index over the entire segment. In practice, the distribution is represented by a histogram whose bins uniformly quantize the range of the shape index values.

Similarity between segments: We are given the sets of segments $S_A = \{A_1, A_2, \ldots, A_n\}$ and $S_B = \{B_1, B_2, \dots, B_m\}$ of $\mathcal{M}_{\mathcal{A}}$ and $\mathcal{M}_{\mathcal{B}}$ and their descriptors. The goal is to establish a match between every segment in S_A to a subset of S_B . Our algorithm consists of two stages. First, we define a single Earth Mover's Distance (EMD) [32] optimization problem that determines the correspondences between segments of $\mathcal{M}_{\mathcal{A}}$ to segments of $\mathcal{M}_{\mathcal{B}}$. Then, to improve the obtained similarity measure, we employ the Random Sample Consensus (RANSAC) algorithm [33] that distinguishes between the true correspondences (inliers) and the false correspondences (*outliers*). As discussed below, RANSAC estimates the similarity transformation between the inliers and hence guarantees invariance to similarity transformation, and in particular to scale.

1. Computing the Earth Mover's Distance: The EMD evaluates the dissimilarity between two multidimensional distributions. It has a couple of benefits that make it appropriate for our problem. First and foremost, it enables partial matching in a very natural way. Second, it can be applied to general variable-size descriptors, which subsume histograms.

To apply the EMD distance, the distance between a "pile of earth" *i* to a "hole" *j* is the L_1 distance between the SSD histogram of segment $B_i \subset \mathcal{M}_{\mathcal{B}}$ to the histogram of segment $A_j \subset \mathcal{M}_{\mathcal{A}}$. The weight of B_i is a scalar defined as $w_{B_i} = Area(B_i)/4\pi R_i^2$, where R_i is the radius of the bounding sphere of segment B_i .

2. Improving the similarity measure with RANSAC: The correspondences obtained in the previous stage may be inaccurate due to two reasons. First, depending on the segmentation algorithm used, the segmentations of $\mathcal{M}_{\mathcal{A}}$ and $\mathcal{M}_{\mathcal{B}}$ might differ even if the models look similar to the naked eye. Second, since $\mathcal{M}_{\mathcal{A}}$ is a sub-surface, some meaningful parts may be missing or be adjacent to the boundaries, as illustrated in Figure 4.

To improve the obtained correspondence, we employ the RANSAC algorithm, which fits a model to experimental data that might contain outliers. In



Figure 4: The cyan leg in (a) is very different from its corresponding segment in (b), since $\mathcal{M}_{\mathcal{A}}$ is a broken piece of $\mathcal{M}_{\mathcal{B}}$.

our case, the model is a similarity transformation between $\mathcal{M}_{\mathcal{A}}$ and $\mathcal{M}_{\mathcal{B}}$, the inliers are the segments of S_A whose correspondences are correct (i.e, they are compatible with the estimated transformation), and the outliers are the other segments.

At each iteration four segment pairs are chosen, from those found by the EMD optimization to match. The similarity transformation between the centers of mass of the segments is estimated [34]. The transformation that yields the maximum number of inliers is considered correct and the final set of inliers is determined accordingly.

The matching probability of an inlier segment in S_A and its matched segment in S_B is set to 1. To determine the matching probabilities of an outlier segment in S_A and segments in S_B , the similarity transformation is applied to all the vertices of S_A 's segment. The matching probabilities are then set to be proportional to the number of transformed vertices that fall close to those of the segment in S_B .

Results: Figure 5 demonstrates the matching results of the two stages described above. It can be seen that after Stage 1, one segment is matched incorrectly. After applying RANSAC, the correct correspondence is detected.



Figure 5: Top: Matching the segments of a flower to those on the fragment of a Hellenistic vase. After applying the EMD the green segment is matched to a different flower on $\mathcal{M}_{\mathcal{B}}$. RANSAC fixes it. Bottom: Matching cupid-shaped reliefs of two Hellenistic oil lamps. After applying the EMD the heads and the torsos are matched correctly, but the legs are not. RANSAC fixes the matching of the green legs.

5. Integrating feature-point similarity and segment similarity

We are given the similarity results of the salient vertices (Stage 1 of Algorithm 1) and the similarity results of the segments (Stage 2). Our goal is to integrate these similarity measures, so as to obtain more reliable and consistent correspondences between the vertices (Stage 3).

Intuitively, if a pair of vertices resides in segments having a high matching probability, the correspondence probability of the pair should increase. Conversely, if the pair resides in segments having a low matching probability, its correspondence probability should decrease. Moreover, the correspondences of the vertices in the neighborhood of the pair provide additional information regarding our confidence in the correspondence of this pair.

Problem definition: Let $S_A = \{A_1, A_2 \cdots, A_n\}$ be the set of segments of $\mathcal{M}_A, \mathcal{M}_A = \bigcup_{i=1}^n A_i$ and $S_B = \{B_1, B_2 \cdots, B_m\}$ be the set of segments of $\mathcal{M}_B, \mathcal{M}_B = \bigcup_{j=1}^m B_j$. We are given a salient vertex, $v \in A_i$, its corresponding salient vertex $u \in B_k$, and their initial correspondence probability $\mathcal{P}((v, u))$. We are also given the segment matching probabilities $\mathcal{P}(A_iB_j) \forall B_j \subset \mathcal{M}_B$. Our goal is to modify the known initial correspondence probability $\mathcal{P}((v, u))$, so as to better reflect the matching probability of segment A_i and the correspondence probabilities of the other salient vertices in A_i .

The naive approach: Intuitively, in order to update the correspondence probability $\mathcal{P}((v, u))$, incorporating the matching probability of the segment, we could simply multiply the two:

$$\mathcal{P}_{new}((v, u)) = \mathcal{P}\left((v, u)\right) \cdot \mathcal{P}\left(A_i B_k\right).$$

Indeed, if the probability $\mathcal{P}(A_iB_k)$ is small, the correspondence probability of (v, u) will decrease as expected. However, this technique is not as robust as we may wish. As illustrated in Figure 6, our confidence in the probability should take into account not only the probability of the segment correspondence, but also those of the neighboring vertices. The figure shows a case in which even if the correspondence of (v, u) agrees with that of A_iB_k , the correspondence should not increase but rather decrease, since the other vertices point to a different correspondence.

Our approach: To overcome the drawback of the naive approach, we take into account the correspondence probabilities of the other vertices. Intuitively,



Figure 6: The naive correspondence is not sufficiently robust. v should be matched to a vertex in B_j even if $\mathcal{P}(A_iB_k)$ is high.

if we consider the correspondences of the neighboring vertices, we achieve a "voting procedure", in which similar correspondences increase the probability of a pair. There are various ways to perform voting procedures that integrate different measures. Our proposed approach is inspired by [35].

Formally, given the state of correspondences $\mathcal{L}_{\mathcal{M}_{\mathcal{A}}}$ and a vertex $v \in A_i$, our goal is to calculate

$$\mathcal{P}_{new}((v,u)) = \mathcal{P}\left((v,u) \mid \mathcal{L}_{\mathcal{M}_{\mathcal{A}}}\right),$$

where $\mathcal{L}_{\mathcal{M}_{\mathcal{A}}} = \{(v', u', B_j) | v' \in \mathcal{M}_{\mathcal{A}}, u' \in B_j \subset \mathcal{M}_{\mathcal{B}}, \mathcal{P}((v', u')) > 0\}$ (where $\mathcal{P}((v', u')) > 0$ indicates that u' matches v').

That is to say, we wish to calculate the correspondence probability of the pair (v, u), given the correspondences and the *locations* of the corresponding vertices of all the salient vertices of $\mathcal{M}_{\mathcal{A}}$. By locations we refer to the segments on which the corresponding vertices reside, and not their exact location within the segment. For now we assume that v and u, each resides in a single segment.

In essence, it suffices to consider only the correspondences of the vertices that lie in A_i . Therefore, given the state of correspondences $\mathcal{L}_{A_i} = \{(v', u', B_j) | v' \in A_i, u' \in B_j, \mathcal{P}((v', u')) > 0\}$, our goal is to calculate

$$\mathcal{P}_{new}((v, u)) = \mathcal{P}\left((v, u) \mid \mathcal{L}_{A_i}\right),$$

which is the correspondence probability of the pair $(v \in A_i, u)$, given the *locations* of the correspondences of all the salient vertices of A_i .

By Bayes' rule:

$$\mathcal{P}\left((v,u) \mid \mathcal{L}_{A_i}\right) = \frac{\mathcal{P}\left(\mathcal{L}_{A_i} \mid (v,u)\right) \mathcal{P}\left((v,u)\right)}{\mathcal{P}\left(\mathcal{L}_{A_i}\right)}.$$
 (3)

Since $\mathcal{P}((v, u))$ is given, all we need to calculate are $\mathcal{P}(\mathcal{L}_{A_i} \mid (v, u))$ and $\mathcal{P}(\mathcal{L}_{A_i})$. We next explain how they are computed.

By the law of total probability,

$$\mathcal{P}(\mathcal{L}_{A_{i}}) = \mathcal{P}(\mathcal{L}_{A_{i}} \mid (v, u)) \mathcal{P}((v, u))$$

+
$$\mathcal{P}\left(\mathcal{L}_{A_{i}} \mid \overline{(v, u)}\right) (1 - \mathcal{P}((v, u))).$$

So, to compute $\mathcal{P}(\mathcal{L}_{A_i})$, we need to compute $\mathcal{P}(\mathcal{L}_{A_i} \mid (v, u))$ and $\mathcal{P}(\mathcal{L}_{A_i} \mid \overline{(v, u)})$. Below, we elaborate on the calculation of the first term; the second one is computed analogically.

By construction, for segment A_i , $\sum_{B_j \in \mathcal{M}_{\mathcal{B}}} \mathcal{P}(A_i B_j) = 1$. Therefore, using the law of total probability:

$$\mathcal{P}\left(\mathcal{L}_{A_{i}} \mid (v, u)\right) = \sum_{B_{j} \in \mathcal{M}_{\mathcal{B}}} \quad \mathcal{P}\left(\mathcal{L}_{A_{i}} \mid (v, u), A_{i}B_{j}\right) \cdot \mathcal{P}\left(A_{i}B_{j} \mid (v, u)\right). \quad (4)$$

In order to break the cycle of dependencies between the vertex similarity and the segment similarity, we set

$$\mathcal{P}(A_i B_j \mid (v, u)) = \mathcal{P}(A_i B_j).$$
(5)

Moreover, we assume that the correspondences of the salient vertices of A_i are independent (and they were computed this way). Hence,

$$\mathcal{P}\left(\mathcal{L}_{A_{i}} \mid (v, u), A_{i}B_{j}\right) = \prod_{v' \in A_{i}} \mathcal{P}\left(\mathcal{L}_{v'} \mid (v, u), A_{i}B_{j}\right),$$
(6)

where $\mathcal{L}_{v'}$ is the state of correspondence of a salient vertex $v' \in A_i$ (i.e., its corresponding vertex u' and the segment in $\mathcal{M}_{\mathcal{B}}$ in which u' resides).

Thus, we need to compute $\mathcal{P}(\mathcal{L}_{v'} \mid (v, u), A_i B_j)$. It is done by considering the correspondence of $v' \in A_i$ as a random variable and distinguishing between two cases. If the correspondence of (v', u') is correct (with probability $\mathcal{P}((v', u'))$), it must reside in B_j since $v' \in A_i$ and this correspondence should suit that of (v, u). If the correspondence is incorrect, u' could lie anywhere, and thus the segment on which it lies is chosen randomly among the segments of $\mathcal{M}_{\mathcal{B}}$. In this case, the probability that u' resides in a segment is set to the relative area of this segment and is denoted as β' . Formally,

$$\mathcal{P}\left(\mathcal{L}_{v'} \mid (v, u), A_i B_j\right) = 1_{\left\{u' \in B_j\right\}} \cdot \mathcal{P}\left((v', u')\right) + (1 - \mathcal{P}\left((v', u')\right)) \cdot \beta', (7)$$

where $1_{\{u' \in B_j\}}$ is the *indicator function*, defined as:

$$1_{\{u' \in B_j\}} = \begin{cases} 1, \ u' \in B_j \\ 0, \ u' \notin B_j \end{cases}$$

Putting it all together, Equation (4) becomes:

$$\mathcal{P}\left(\mathcal{L}_{A_{i}} \mid (v, u)\right) = \sum_{B_{j} \in \mathcal{M}_{\mathcal{B}}} \mathcal{P}\left(A_{i}B_{j}\right) \prod_{v' \in A_{i}} \left[1_{\left\{u' \in B_{j}\right\}} \cdot \mathcal{P}\left(\left(v', u'\right)\right) + \left(1 - \mathcal{P}\left(\left(v', u'\right)\right)\right) \cdot \beta'\right].$$

Up till now we assumed that a vertex $v \in \mathcal{M}_{\mathcal{A}}$ belongs to a single segment. If v belongs to multiple segments Equation (4) can be rewritten as:

$$\sum_{A_i \in \operatorname{Seg}(v)} \mathcal{P}\left(\mathcal{L}_{A_i} \mid (v, u), v \in A_i\right) \mathcal{P}\left(v \in A_i\right) = \sum_{A_i \in \operatorname{Seg}(v)} \sum_{B_j \in \mathcal{M}_{\mathcal{B}}} \mathcal{P}\left(A_i B_j\right) \mathcal{P}\left(v \in A_i\right) \cdot \prod_{v' \in A_i} \mathcal{P}\left(\mathcal{L}_{v'} \mid (v, u), v \in A_i, A_i B_j\right).$$
(8)

Finally, after calculating the correspondence probabilities for every pair (v, u) separately, we normalize the probabilities of all the matches of v.

Examining Equation (8), it can be seen that as required, the similarity probability of a pair of vertices depends on the correspondence probability of the segment(s) on which the vertices reside, on the probabilities of the neighboring vertices, and on the initial correspondence probability of the pair.

Results: Figure 7 demonstrates a result of our algorithm. We wish to match the sub-surface in Figure 7(a) to that in Figure 7(b). We focus on the salient vertex marked in red. The matching segments are colored in corresponding colors. The initial correspondences of the red vertex are displayed in Figure 7(c) in different colors and their initial correspondence probabilities are: p_{red} = $0.05; p_{green} = 0.28; p_{blue} = 0.48; p_{magenta} =$ 0.19. After applying our algorithm, the final probabilities are: $p_{red} = 0.91; p_{green} = 0; p_{blue} =$ 0.09; $p_{magenta} = 0$. Evidently, the correct matching probability increased to 0.91, whereas the other correspondences were significantly decreased. For comparison, we also applied the scheme using the naive approach. Here, the correct matching probability (after normalization) increased to 0.48, whereas the incorrect matching to the blue point received an even higher value and increased to 0.52. Hence, the uncertainty remained, with a slight preference to the incorrect matching.

Figure 8 compares our final result to the one that would be obtained if the naive approach were utilized, given a dancer as a query. It can be seen that



Figure 7: The probability of the red vertex in (c) increased from 0.05 to 0.91, whereas the probability of the blue vertex decreased from 0.48 to 0.09.



Figure 8: Comparison of the final results of our matching and the naive approach

while our result retrieved a single dancer, the naive approach distributed the matched vertices between the three dancers.

6. Results: Partial similarity in Archaeology

Archaeological artifacts are inherently noisy and eroded, not only due to the scanning process, but also because of their very nature, found after spending thousands of years underground. This makes the difficult problem of comparing objects even harder. Our results demonstrate not only the ability of the algorithm to detect non-identical sub-meshes, but also the algorithm's robustness to noise.

Figure 1 shows two examples. In Figure 1(a) a dancer from a fragment of a vase is matched to the dancers on a significantly-noisy Amphora. (The dancers were embedded on the amphora synthetically by [25].) In Figure 1(b) a cupid from a Hellenistic oil lamp is the query. Our algorithm matches this query to the cupids on a different oil lamp. The poses as well as the shapes of the matched cupids differ, i.e., the query cupid has hair while the matched cupids do not, the matched cupids have wings while the query does not, etc.

In Figure 9 the input queries are Greek letters $(\Upsilon, M A)$ extracted from Hellenistic stamps (either from the same stamp (a) or from different stamps (b)). In the latter case, not only the letters may differ in shape, but also the scale ratio is unknown. Our algorithm manages to detect the letters, even though it can be seen that the M on the stamp is immersed in the stamp's background.

In Figure 10, given an instance of a pattern (blue), the other instances are detected (magenta). This can be viewed as the extraction of 3D textures.

Finally, Figure 11 shows several coins. This is a very interesting example, since these coins are considered a major discovery, which brings up many questions. Were similar coins found elsewhere? Were they stamped using the same die? The query is a goblet extracted from a coin dated to the first century CE (blue). This goblet is matched by our algorithm to reliefs of goblets on four other coins (magenta). These coins are not identical – they differ in the shape of the goblet, in its location within the coin, in the level of noise, in the state of the burnout, and in the location of the cracks. In one of them, the goblet is reflected. Our algorithm copes with these conditions well.



Figure 9: Typography of Hellenistic stamps



Figure 10: Detecting repeating patterns on a variety of Hellenistic artifacts. The shapes of the similar reliefs are not identical. Nevertheless, our algorithm finds them.

Comparison to previous work: We are not aware of related work in the domain or archaeology. We therefore compare our results to those presented in [17] for general, noiseless objects. Though [17] is the closest work to ours, it addresses a different problem, where the parts to be matched are not predefined. This problem is important for applications such as alignment and registration, but it is less suitable for archaeology, since the archaeologists usually have a clear idea regarding the "characteristic" parts of the objects they wish to match.

In Figure 12(a), we mark one flower on the Buddha (1M faces) and the other flowers are detected. The query time in our Matlab implementation is 20-25% faster than [17]'s on similar PC configurations. In Figure 12(b), given a single star, all the stars on the knot are detected, although their sizes and shapes differ across the model. This result can be compared to the result of [17] in which stars were placed on the plane. It can be seen that our matching is more precise. The major benefit of our method, however, is its ability to handle noise.

Implementation: We ran the algorithm on an Intel Core i7 machine with 3GB of memory. On a model of 275K faces (the coin in Figure 11), the online query takes only 1.2 seconds. The preprocessing takes 2.9 minutes, out of which the salient vertex processing takes 2.2 minutes and the segment processing takes 42 seconds. Most of this time (1.9 minutes) is devoted to curvature estimation.

Limitations: If the relief is extremely eroded and the segmentation algorithm fails to segment it, our algorithm might fail. This is illustrated in Figure 9(middle), where the circled M was not de-



Figure 11: Coins from the first century CE. The coins differ in the shape of the goblet, in its location within the coin, in the level of noise, in the state of the burnout, and in the location of the cracks. Our algorithm finds the magenta goblets given the blue query.



Figure 12: Comparisons to [17]: Given a flower (star), the others are detected.

tected.

7. Conclusion

This paper proposed a novel algorithm for partial matching, which is based on the observation that local feature similarity by itself does not suffice, yet the aggregation of neighboring locally-similar features does. To realize this observation, our algorithm consists of three steps. First, locally-similar vertices are detected and our confidence in their correspondence is formulated as a matching probability. Second, the meshes are segmented and matched. Finally, vertex similarity and segment similarity are combined, to yield a more reliable match.

The utility of the algorithm was evaluated in a challenging application – archaeology, in which partial matching is a key task. The (fragments of) artifacts found in the excavation should be matched to those previously found in this site or elsewhere in the world. We showed that our technique is useful for the type of models prevalent in this domain. It can handle noisy and broken objects.

In the future it will be interesting to examine this algorithm in the context of global consistency. The use of feature points is often insufficient due to many ambiguities. Our point/segment integration can be viewed as a consistency criteria for filtering out false positives.

Acknowledgements: We are grateful to Ayelet Gilboa from the Zinman Institute of Archaeology, University of Haifa, for her invaluable advice and for letting us use the models. This research is supported in part by the Google Research Award Program and by the Israeli Ministry of Science and Technology, grant 3-3421.

References

- J. Tangelder, R. Veltkamp, A survey of content based 3D shape retrieval methods, Multimedia Tools and Applications 39 (3) (2008) 441–471.
- [2] B. Zitova, J. Flusser, Image registration methods: a survey, Image and vision computing 21 (11) (2003) 977– 1000.
- [3] B. Bustos, D. Keim, D. Saupe, T. Schreck, D. Vranić, Feature-based similarity search in 3D object databases, ACM Computing Surveys 37 (4) (2005) 345–387.
- [4] O. van Kaick, H. Zhang, G. Hamarneh, D. Cohen-Or, A Survey on Shape Correspondence, in: Proc. of Eurographics State-of-the-art Report, 2010.
- [5] C. Dorai, A. Jain, COSMOS-A representation scheme for 3D free-form objects, PAMI 19 (10) (1997) 1115– 1130.
- [6] A. Johnson, M. Hebert, Using spin images for efficient object recognition in cluttered 3D scenes, PAMI 21 (5) (1999) 433–449.
- [7] M. Elad, A. Tal, S. Ar, Content Based Retrieval of VRML Objects An Iterative and Interactive Approach, in: Eurographics Workshop on Multimedia, 97–108, 2001.
- [8] P. Shilane, T. Funkhouser, Selecting Distinctive 3D Shape Descriptors for Similarity Retrieval, in: International Conference on Shape Modeling and Applications, 18, 2006.
- [9] T. Funkhouser, M. Kazhdan, P. Shilane, P. Min, W. Kiefer, A. Tal, S. Rusinkiewicz, D. Dobkin, Modeling by example, ACM Trans. on Graphics 22 (3) (2004) 652–663.
- [10] L. Shapira, S. Shalom, A. Shamir, D. Cohen-Or, H. Zhang, Contextual Part Analogies in 3D Objects, International Journal of Computer Vision 89 (2) (2010) 309–326.
- [11] S. Belongie, J. Malik, J. Puzicha, Matching shapes, in: IEEE International Conference on Computer Vision, vol. 1, 454–461, 2001.
- [12] U. Castellani, M. Cristani, S. Fantoni, V. Murino, Sparse points matching by combining 3D mesh saliency with statistical descriptors, in: Computer Graphics Forum, vol. 27, 643–652, 2008.
- [13] T. Funkhouser, P. Shilane, Partial matching of 3D shapes with priority-driven search, in: Symp. on Geometry Processing, 131, 2006.
- [14] N. Gelfand, N. Mitra, L. Guibas, H. Pottmann, Robust global registration, in: Symp. on Geometry Processing, vol. 2, 5, 2005.
- [15] X. Li, I. Guskov, Multi-scale features for approximate alignment of point-based surfaces, in: Symp. on Geometry processing, 217–226, 2005.
- [16] M. Novotni, P. Degener, R. Klein, Correspondence generation and matching of 3D shape subparts, Tech. Rep., CG-2005-2, Universitat Bonn, 2005.
- [17] R. Gal, D. Cohen-Or, Salient geometric features for partial shape matching and similarity, ACM Trans. on Graphics 25 (1) (2006) 130–150.
- [18] S. Biasotti, S. Marini, M. Spagnuolo, B. Falcidieno, Sub-part correspondence by structural descriptors of 3D shapes, Computer-Aided Design 38 (9) (2006) 1002– 1019.
- [19] A. Ferreira, S. Marini, M. Attene, M. Fonseca, M. Spagnuolo, J. Jorge, B. Falcidieno, Thesaurus-based 3D Ob-

ject Retrieval with Part-in-Whole Matching, International Journal of Computer Vision 89 (2010) 327–347.

- [20] H. Rushmeier, Eternal Egypt: experiences and research directions, in: Modeling and Visualization of Cultural Heritage, 22–27, 2005.
- [21] B. Brown, C. Toler-Franklin, D. Nehab, M. Burns, D. Dobkin, A. Vlachopoulos, C. Doumas, S. Rusinkiewicz, T. Weyric, A System for High-Volume Acquisition and Matching of Fresco Fragments: Reassembling Theran Wall Paintings, ACM Trans. Graph. 27 (3) (2008) 84.
- [22] D. Koller, J. Trimble, T. Najbjerg, N. Gelfand, M. Levoy, Fragments of the City: Stanford's Digital Forma Urbis Romae Project, in: the Third Williams Symposium on Classical Architecture, vol. 61, 237–252, 2006.
- [23] M. Kolomenkin, I. Shimshoni, A. Tal, Demarcating Curves for Shape Illustration, ACM Transactions on Graphics 27 (5) (2008) 157:1–9.
- [24] M. Kolomenkin, I. Shimshoni, A. Tal, On Edge detection on Surfaces, in: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2767–2774, 2009.
- [25] R. Zatzarinni, A. Tal, A. Shamir, Relief Analysis and Extraction, ACM Trans. on Graphics 28 (5) (2009) 136:1–9.
- [26] P. Besl, H. McKay, A method for registration of 3-D shapes, IEEE Trans. on pattern analysis and machine intelligence 14 (2) (1992) 239–256.
- [27] D. Lowe, Distinctive image features from scale-invariant keypoints, Int. Journal of Computer Vision 60 (2) (2004) 91–110.
- [28] R. Young, The Gaussian derivative model for machine vision: I. Retinal mechanisms, Spatial Vision 2 (4) (1987) 273–293.
- [29] A. Bobenko, P. Schröder, Discrete Willmore flow, in: Symp. on Geometry Processing, vol. 255, 101–110, 2005.
- [30] B. Chazelle, D. Dobkin, N. Shouraboura, A. Tal, Strategies for polyhedral surface decomposition: an experimental study, in: SOCG, 297–305, 1995.
- [31] T. Zaharia, F. Preteux, 3D shape-based retrieval within the MPEG-7 framework, in: SPIE Nonlinear Image Processing and Pattern Analysis, 133–145, 2001.
- [32] Y. Rubner, L. Guibas, C. Tomasi, The earth mover's distance, multi-dimensional scaling, and color-based image retrieval, in: DARPA Image Understanding Workshop, 661–668, 1997.
- [33] M. Fischler, R. Bolles, Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography, Communications of the ACM 24 (6) (1981) 381–395.
- [34] B. Horn, H. Hilden, S. Negahdaripour, Closed-form solution of absolute orientation using orthonormal matrices, J. Optical Society of America 5 (7) (1988) 1127– 1135.
- [35] M. Kolomenkin, I. Shimshoni, Image Matching Using Photometric Information, in: IEEE Conference on Computer Vision and Pattern Recognition, 2506–2514, 2006.