



Context-adaptive navigation of 3D model collections

Silvia Biasotti^a, Elia Moscoso Thompson^{a,*}, Michela Spagnuolo^a

^aIstituto di Matematica Applicata e Tecnologie Informatiche 'E. Magenes' - CNR

ARTICLE INFO

Article history:

Received January 15, 2019

Keywords: 3D similarity exploration, similarity measures combination, 3D object classification
2000 MSC: 68U05, 65D18

ABSTRACT

When reasoning about similarity in a collection of objects with heterogeneous qualities, there are several aspects of interest that can be followed to explore the collection. Indeed, the notion of similarity among 3D models is not only grounded on the geometric shape but also, for instance, on the style, material, color, decorations, common parts. These are all important factors that concur to the concept of similarity. Search engines for visual content are expected to address similarity assessment in collections, providing a higher degree of flexibility with respect to the traditional 3D object retrieval operations. In this work, we describe the design and functioning of a search engine working on multiple factors and discuss the results on a number of collections, which challenge existing 3D object retrieval engines.

© 2019 Elsevier B.V. All rights reserved.

1. Introduction

The digital era makes available an increasing number of scans of objects. The availability of this wealth of data and information opens-up the possibility to build new tools to study and explore collections, leveraging on the digital accessibility to large sets of exemplars that can provide insights difficult to achieve by manual inspection of the datasets. Visual search engines are the natural tools for supporting comparative studies of object collections and they are generally based on *content-based* approaches [1]. They are based on the combination of *shape descriptors* as signatures synthesizing the geometric content of 3D models [2, 3], and of *similarity measures* for matching descriptors [4, 5]. The first prototype of 3D search engine [6] permits to select models from sketches, text searches or selecting a number of global shape descriptors. Query results are ordered according to their distance from the query model. The implicit assumption behind this kind of search is that there exists a unique key to interpret the dataset and that

is possible to codify the intuition behind human similarity in a single measure. Unfortunately, a single shape similarity measure is not sufficient to characterize the many different flavors of shape comparison when dealing with complex and heterogeneous datasets [7, 8]. For this reason, we think that search engines should be flexible, multi-modal and able address similarity assessment in collections of different nature. This is particularly important when the assets to be studied are only parts of objects, and therefore their overall shape does not have, in most cases, a meaning, as for fractured archaeological findings. Moving from these considerations, *facets* were proposed in [9] as quantitative or categorical aspects that are relevant for the user. Therefore, faceted browsing, i.e. the possibility for the user to specify the desired shape properties, has been identified, altogether with query refinement, as one of the challenges that search engines for 3D models must be tackled for effectively exploring 3D shape collections [8].

In this paper, we present the settings and experimental results of a search engine for 3D shape collections able to address faceted queries and automatic search relaxation. This search approach was originally presented in [10], with the setting and initial experiments carried out in the context of the GRAVITATE project [11], targeting similarity search to support the analysis of fractured archaeological objects of various collections. In

*Corresponding author: Tel.: +39-010-6475-697; fax: +39-010-6475-660;
 e-mail: silvia.biasotti@ge.imati.cnr.it (Silvia Biasotti),
elia.moscoso@ge.imati.cnr.it (Elia Moscoso Thompson),
michela.spagnuolo@ge.imati.cnr.it (Michela Spagnuolo)

this paper, we generalize the approach to a wider context of object types and signatures. In particular, the novel content with respect to [10] is:

- larger set of descriptors: ten descriptors are included and tested in the current search framework, with respect to the five ones of the original paper;
- new collections for the validation: while the original paper was focusing only on archaeology, in this paper we tested and evaluated the efficacy of the search methodology on a totally different collection of objects, targeting the application of the search to a broad and generic search context. We believe that this testing is particularly significant to the development of visual content-based search engines;
- extended evaluation: the performance of the search framework has been discussed in more details, up to the inclusion of a quantitative, yet limited, evaluation session on a ground truth that we were able to collect for the archaeological domain.

With respect to the current state of the art, the novelty brought by the paper is the development of a search framework able to adapt the navigation to the properties of the objects, which may differ for different collections, and adapt also to the goals of the search session, which may differ from session to session even if the user is the same. The search framework defined is flexible with respect to the choice of descriptors, and also to the inclusion of different similarity distances, depending on the properties one might be interested to capture. Another advantage of the proposed solution is related to the ease of formulating the query, which allows to specify the search intent by combining multiple descriptors in a single query. We believe it is important to explore the possibility to move from traditional object retrieval methods to more general search by similarity, or better, by similarities. An easy approach to formulate the search intent is therefore useful given the intrinsic difficulty to map the search intent, inherently multi-faceted, into a fixed combination of descriptors.

The reminder of the paper is organized as follows. In Section 2, we overview the state of the art on the exploration of 3D object collections, also discussing the most popular shape descriptors and similarity measures. In Section 3, we describe how the search engine was conceived as a combination of several properties and how they act as independent filters in the conceptual dimensions that describe the collection of 3D models. Section 4 details the mechanism behind the properties combination and how to automatically set the search engine parameters according to the variations emerging in the collection. In Section 5 we describe the validation scenario: first, we present the collections used for the validation, in the archaeological and generic domains; we then discuss ground truths and evaluation methodology, given the peculiarity of the search addressed which can be hardly captured by quantitative performance measures. It will be shown how the criteria elaborated for the archaeological domain and the results obtained generalize to a wider domain. Experimental results on these datasets are discussed in Section 6. Finally, in Section 7, we draw some conclusions and sketch plans for future development.

2. Related work

In the last fifteen years, search engines have been addressed with a large number of content-based techniques aimed at detecting global shape similarity [6, 12] or part-based search and retrieval in 3D object collections [13, 14] or specific contexts, like product design [15, 16], parametric shape collections [17] and cultural heritage [18, 10]. As similarity is a cognitive process, the user's intent during the search should be included in the loop [19]: methods for relevance feedback were introduced for 3D object search [20, 21], while the first 3D search engine able to support user interaction appeared quite recently [22] under the paradigm of the faceted meta-data.

2.1. 3D shape similarity

Most of the methods for shape analysis and description focus in the 3D geometric information [2, 3, 23, 14, 5]. Shape can be stored in a wide variety of descriptions (histograms, matrices, graphs, etc.) and can be based on many different types of information (punctual, normal vectors, surface or volumetric, possibly with attributes). All methods extract salient geometric information, and use that to derive a concise description. Over the years, descriptors and similarity measures were refined to deal with geometric invariances, from rigid transformations to isometries, according to the application at hand [24, 5]. A further characterization of the descriptors as local or global depends on their locality with respect to the whole model. Appearance was also taken into account beside geometry, resulting in techniques for the retrieval of *textured* 3D models, which are common for example in cultural heritage applications [25, 18]. Finally, other aspects can concur to the concept of similarity, such as functionality [26] or semantics [27].

Shape descriptors. For 3D objects exploration, only a sub-part of the shape descriptors actually available in the literature is adopted, being the selection a tuning between computational efficiency and efficacy of the description.

The Lighfield descriptor (LFD) [28] is one of the most popular global descriptors [29, 30, 31], possibly with some ad-hoc variation [32] like the inner-distance [33]. Other choices of global descriptors are the Shape Diameter Function (SDF) [34], the Spherical Harmonics (SH) [35], the Heat Kernel Signature (HKS) [36], the histograms of Gaussian Curvature (hGC) [29], the Shape Distributions (SD) [37], the Voxel Shape Histogram (VSH) [38], the 3D Shape Context [39], altogether with local geometric properties of facets [40], and representative landmarks [41]. In general, not one single descriptor is considered and a combination is adopted. For instance, [32] considers four different shape descriptors: LFD (with Inner-Distance), SFD, SH and HKSM [29] uses hGC, LFD and SD; [17] adopts SD, VSH and LFD. Other approaches encode parts as a box, [42, 43]. In this case, the 3D model comes with multiple disconnected components and its representation is a probabilistic, part-based deformable model, which encode clusters of parts, correspondence across clusters and alignment of a template to each model. At a higher level of abstraction, [44] encodes each

shape with a spatial relation graph. Each node in the graph corresponds to a component, while each edge denotes a contact or symmetry relation.

Finally, it is worth mentioning, that there exist methods suitable for the exploration of model collections made of uniform families of objects that do not adopt descriptors in the usual conception, like [45] and [46]. In [45] a meta-representation of a family of objects is the set of PDFs (Probability Density Functions) of relations defined between shape parts. The PDFs can be used to estimate the probability of a specific value for any relation (e.g., the angle between wing and fuselage). In [46] a collection of shapes is represented as a collection of shape differences. In this formulation, shapes are compared by the differences of two inner products: the area-based and the conformal-based inner product.

Distances. In the standard approach to similarity assessment distances between 3D shapes are computed as distances between shape descriptors. For instance, when the shape is described as a feature vector or a bag of features, a common choice is to select the *Minkowski* L^p family of distances, for instance in [30] used the L^1 or Manhattan distance, [44, 43, 13, 42] adopted the L^2 or Euclidean distance, etc. Sometimes, the L^p distance is applied to an embedding of the feature vectors in the 3D Euclidean space, as done in [31] through the Multi-dimensional Scaling (MDS), or to a reduction via MDS of the similarity matrix between an object and a set of templates as done in [44] for a human model and a set of user-specified human poses. Other distances, often selected when the objective is to deal with shape or part classification, are the Mahalanobis distance [47], the Bhattacharya distance [48] and the Earth Mover's Distance (EMD) [49].

When dealing with global/local descriptors also the similarity distance may reflect this characteristic, with a hierarchical computation of global and local similarity scores, [40, 30, 13]. For instance, the overall idea in [30] is to compute these scores only for models that are sufficiently similar, namely in the sense of k -nearest neighbours with respect to the 2D views of each model used to compute the LFD. In this way, it is possible to reduce the computational complexity from quadratic to linear in the number of models. Then, the sparse distance metric is propagate to all pairs of models via a heat diffusion embedding and the local similarity is derived from the initial co-alignment given by the light field computation. Sub-part, fuzzy correspondences are first evaluated only for sufficiently similar models and again propagated to all pairs of models, by using N -order graphs.

2.2. Dataset exploration

The exploration of large collections of 3D models is a challenging task and only a few methods tackle such a problem. Most of methods make several hypotheses on the classes of the dataset, for instance focusing on man-made objects [50, 45, 44, 13, 42]. Again, [45] assumes that the input shapes are coming from the same family and are pre-segmented and consistently labeled, with labels from a pre-defined, family-specific set of labels, while [44] assumes that during experiments all models belong to the same category (only chairs, only

bikes...). Others, like [50], assume that the shapes in the collection all share a common global part structure (e.g., a set of four-legged chairs or two-winged airplanes) and that most of the shape variability within the collection can be explained in terms of the relative sizes and positions of the shape parts (e.g., changing size and positions of the wings in the airplanes, but the number of wings is fixed). Similarly, [31] focuses on the evolution of an entire population of 3D models which belong to a certain class (e.g., lamps, chairs, candelabra and teapots). Here, the assumption is that the initial set of shapes was pre-analyzed to possess part correspondence and built-in structural information such as inter-part symmetries. The set should be reasonably rich and diverse and the shapes therein sufficiently developed.

Other methods consider classes of 3D object made of variations of a set of basis shapes, e.g. [46] proposes a human dataset of 64 shapes, each with 6.5K vertices and a "bunny dataset" containing 218 (approximately conformal) deformations of the Stanford bunny, each with 14K vertices. In similar settings, [17] focuses on the exploration of parametric shape collections, i.e. shapes that can span the large set of possible geometries resulting from the variability of a fixed set of parameters, like it may happen for CAD models. In this case, the crucial point is to identify the shape descriptors that reflect the continuous variation of the parameters into the shape descriptor space.

The use of heterogeneous shape collections, like the Princeton Shape Benchmark¹ and 3D Warehouse², is acknowledged by [32, 29, 43, 30] and, at the moment, the collection used in [30] is the largest one adopted for content-based exploration, with its 103738 3D models. These methods do not explicitly declare any particular assumptions/requirements on the type of the 3D shapes but do not permit a multi-facet exploration by a user-driven combination of different descriptors.

3. Similarity for collections with uneven property variations

When reasoning about similarity in a collection of objects with heterogeneous properties, it is useful to use multi-modal information in combination (e.g., geometry and texture) and also proceed someone in an iterative manner, by grouping items, or their parts, into meaningful clusters. This scenario is still quite far from the scenario depicted by the current state-of-the-art: most methodologies for comparing, retrieving, or classifying objects in repositories are based on a single analysis of the geometric 3D shape, possibly building on specific invariants, such as the presence of axis of symmetry [51] or appendages [52]. Even if several recent approaches for similarity assessment aim at identifying (dense or sparse) correspondences among the model elements (e.g., [53, 13, 54]) or combining texture and geometry information (e.g., [18, 55, 56]), they actually pursue a shape matching at a global level rather than evaluating similarities based on the comparison of specific

¹<http://shape.cs.princeton.edu/benchmark/>

²<https://3dwarehouse.sketchup.com/>

features. With respect to these strategies, we present an approach that addresses similarity on the basis of a set of aspects that describe the models, following an approach similar in spirit to the faceted meta-data proposed for images in [57].

This work has been pushed by a challenging scenario for visual search, archaeology, where traditional approaches fail and where trendy approaches based on learning are not really effective due to the lack of large training sets and to the inherent complexity of capturing the search intent of users.

3.1. Similarity characterization in an archaeological context

Archaeology is a challenging context for similarity assessment. First of all, similarity is the guiding principle behind reasoning for the analysis of artifacts and it usually involves all aspects which may contribute to the shape: geometry, color, textures, chemical properties of color pigments, and many more. Beside data related to the tangible essence of assets, artifacts come with all the contextual information, primarily in textual form, related to the archaeological description and documentation of the findings themselves.

What complicates the computational model of similarity is the fact that, typically, assets are fractured, eroded, and their color faded. Moreover, pieces of fractured objects may be mixed-up. At the same time, there are a lot of visual cues that may suggest that fragments belong to the same object, and these cues are used jointly or in sequence to answer typical questions related either to re-assemble (*may these fragment match together?*) or to re-association (*could this fragment belong to that collection?*) or simply to the stylistic analysis of the fragment itself (*do these fragments share the same decoration?*).

In the initial phase of the GRAVITATE project, we worked with the archaeologists to understand what were the basic similarity criteria used in traditional studies, within the scope of the project case studies. In the following, we enumerate the properties identified as useful by the users and the motivations that concurred to build our first search engine [10].

1. *Overall fragment size*: the size of the fragment, in terms of the overall space it occupies. Alone, such a criterion is generally too rough to assess similarity, but it becomes useful as a further filter to help restricting results obtained in queries.
2. *Thickness*: the thickness of a fragment. This property indicates, for instance, if two shards could come from the same object.
3. *Roughness*: the grain and finish of the fragment surface. Again, similarity according to this property may be significant for assessing if two fragment belong to the same object or where manufactured in the same workshop.
4. *Shape continuity*: if two fragments fit close together on a same object, it is likely that they exhibit a similar, overall curvature of the outer skin surface.
5. *Color*: colorimetric information provide a rich set of visual cues for assessing similarity, and color distribution is an important one. This property indicates that users may want to cluster fragments with similar colors, independently of other characteristics.

6. *Decorations*: some fragments exhibit parts of the same decorations, usually as local relief or colored/painted patterns, providing again an important insight for similarity assessment.
7. *Overall shape*: the global shape of a model. This criterion is the closest to the classical global similarity criterion. This property is used to reason on similarity aiming fragments that contain a relevant and semantically significant part of an object, such as a head or a feet.

The criteria listed above represent the basic similarity types, or better, comparison axis, that are used as single searches or as concatenated filters by the archaeologists in their current practices manual. The interesting fact is that all the properties above are used, with different weight and relevance, in all reasoning processes underlining re-assembling, re-association or re-unification problems.

3.2. Matching similarity criteria to shape signatures

The second important step is to translate the similarity criteria identified into computational tools, able to perform according to the criteria expressed. This can be done either resorting to state of the art descriptors methods, or by implementing new ones if none can provide the behaviour sought. We translated all the similarity criteria except for the shape continuity and the decorations, which are still work in progress.

As a driving guideline for associating a shape descriptor to a similarity criterion we followed the following rules:

- *simplicity* of the descriptor. If two or more descriptors are suitable for a task, we select the simplest one, in terms of output storage and matching complexity (e.g., real numbers or feature vectors);
- *scalability*, depending on the details we have to focus on (e.g. a chiselled decoration) we need to compare high-resolution 3D models with millions of vertices and therefore, we prefer the shape signatures that are able to deal with the larger meshes (in term of vertices and faces);
- *coherence*, meaning that the shape signatures should rank the query results with respect to a model as close as possible to the similarity types identified by the archaeologists.

Moreover, it can happen that a single descriptor is not enough to fully capture the complexity of a similarity criterion or is able to capture it only partially; in these case, we propose either a combination of descriptors or present multiple choices to the user.

In the description below, we use the term *compatibility* instead of *similarity* to discuss the mapping from criteria to descriptors, to emphasize that each similarity measure contributes to a reasoning rather than retrieving a *crisp* result. In other words, each criterion acts as a filter with respect to a specific property and it is not meant to return the whole similarity assessment. To give an example of this effect, when selecting a colorimetric property when the the query model is mainly red, in the results we expect models with a predominance of the red texture, independently of their shape.

In the following, we define the implemented descriptors.

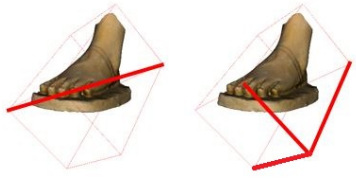


Fig. 1. Left: the diagonal and Right: the three edges of the MMB.

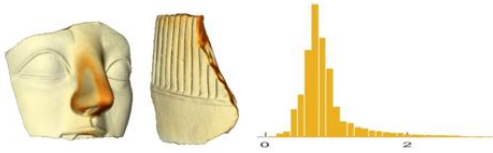


Fig. 2. Left: Two fragments are colored according to the local SDF value, ranging from white-yellow (low values) to red (high values). Right: the distribution of the SDF function; the abscissa of the maximum value of the histogram corresponds the thickness descriptor.

1. *Compatibility in terms of the overall size.* We identify the object size with the volume occupied by the object, therefore we use the oriented minimal bounding box as an approximation of such a volume. Two descriptors are used to map this criterion: the diagonal of the minimal bounding box (MBB) and the vector of the length of the three sorted edges of the MBB (MMB(edge)). We also considered the hull packing value (i.e., the ratio between the minimum and maximum edge of the minimal bounding box) as proposed in [12] but, in our experiments, the MMB(edge) value seems to provide a higher discriminating capability. Figure 1 depicts the meaning of the descriptors based on the minimal bounding box on a GRAVITATE fragment. The distance between the descriptors is the L^1 distance, which works nicely as the descriptors are scalar values or three-dimensional vectors.
2. *Compatibility in terms of the thickness.* The thickness is computed as the shape diameter function (SDF) [34] that has been shown to provide a stable approximation of the diameter of a 3D object with respect to a view cone centered on the object surface and with the axis aligned to the surface normal. After computing the SDF function on each point of the fragment surface, we defined as thickness descriptor the average of the most frequent value. Other possible choices for the thickness descriptor are possible, for instance, we could consider considering the thickness around fragment profiles or fragment fractures [58, 59]. However, these methods generally require a heavy pre-processing step, often not completely automatic, either in terms of the detection of fragmentation orientation or identification of the fracture creeks. With reference to Figure 2, the thickness descriptor correspond to the bin that scores the maximum value of the histogram. As distance between two thickness descriptors, we adopt the usual L^1 distance.
3. *Compatibility in terms of the roughness.* Often, roughness is quantified by the deviations, in the surface normal direction, of a real surface from its ideal form. It is often asso-

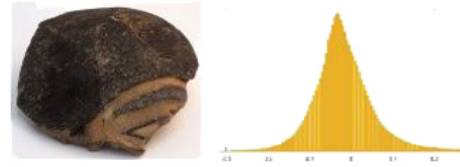


Fig. 3. Left: The hairs of the statue are represented on the fragment by a regular relief pattern. Right: the corresponding mean curvature distribution histogram.

ciated to local, shape texture-like properties [60]. Among the possible descriptions, we identified the mean curvature and the shape index as two possible geometric properties for representing these local surface variation [61, 62]. We computed the value of the minimum k_1 and maximum k_2 curvatures at each vertex adopting the implementation presented in [63]. The mean curvature K and the shape index SI are derived for each vertex, as $K = \frac{k_1+k_2}{2}$ and $SI = \frac{2}{\pi} \arctan\left(\frac{k_1+k_2}{k_1-k_2}\right)$, $k_2 \geq k_1$, respectively. In our experiments, we saw that the mean curvature better highlights the roughness of a model, thus we use it as the descriptor for this property. The corresponding descriptor is defined as the histogram of the distribution the mean curvature, in an interval of values that holds for the whole collection and is determined on the basis of the overall curvature variation over the collection.

We adopt the Earth Mover's distance [49] as the distance between two histograms.

4. *Compatibility in terms of the color distribution.* In the meshes we have considered, colorimetric information is associated to the vertices in terms of a RGB value. However, it is well known that the Euclidean distance in the RGB space does not correspond to the perceived distance between two colors. The CIELAB color space was designed to be perceptually uniform with respect to human color vision [64], meaning that the same amount of numerical change in its values corresponds to about the same amount of visually perceived change. Therefore, the Euclidean distance in this space is a good approximation of the perceived color distance. Among the descriptors adopted in the literature to code the colorimetric information for objects, we have considered the concatenated histogram of the three color channels (L, a and b) in the CIELab space [65, 18], see Figure 4 for its simplicity. For each color channel, we computed 100-bins histograms for all the models and we clamp them between the overall minimum bin-value and the overall maximum bin-value that are not empty. The color descriptor is composed by the concatenation of the three clamped histograms. Using the GRAVITATE use case as example the first 60 bins correspond to the interval [20,80] for the variation of the L-channel, the next 35 bins correspond to the interval [-15,20] of the a-channel and the last 65 bins correspond to the interval [-15,50] of the b-channel. The color descriptor for each model of this collection is a 160-bins

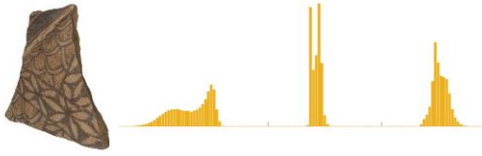


Fig. 4. A fragment and the corresponding concatenated color histogram with respect to the L, a and b channels.



Fig. 5. Left: The non-normalized D_2 signatures of two fragments; Right: the persistence space with respect to the distance from the center of mass and the average, geodesic distance for a vase model.

histogram. The single CIELAB channels contain a specific colorimetric information: the L channel represents the lightness, while the a and b channels respectively represents the green–red component, with green in the negative direction and red in the positive direction and the blue–yellow component, with blue in the negative direction and yellow in the positive direction. For this reason, we also considered the histogram of the single color channels as three separated descriptors, for a total of four color descriptors. Being the CIELab a non uniform space, thus we the Earth Mover’s distance as distance between two color histograms [49] to compensate the problem of non-uniformity.

5. *Compatibility in terms of the overall shape similarity*: Due to the large variability of shapes within the collections, the selection of a geometric signature yielding good results for all shapes is challenging. For instance, in our use cases, the collection is made of heads, feet, legs, busts, vases, broken arms and hands, and these artifacts can have different scale (e.g., small and big statues) or can have cracks and missing parts. The quite limited number of elements per class also prevents the use of learning techniques. We opted for a combined descriptor that mixes rough filters with scale invariant and non-rigid descriptions (called simply *Shape* descriptor). Namely, we adopted a combination of multiple global signatures, like compactness and hull packing [12], the spherical harmonics indexes [35], a non-normalized variation of the shape distribution signature with respect to random chords (this signature codes the probability distribution of the distance between two random points on a surface and is called D_2 in [37]) and the persistence spaces computed according to the average geodesic distance and the distance from the main axis of an object [66].

Figure 5(left) represents the non-normalized shape distribution D_2 [37] for a head and a foot model and Figure 5(right) shows the persistence space obtained with respect

to the distance from the center of mass and the average geodesic distance, evaluated according to [66]. All the distances adopted in this setting are metrics in the descriptors space.

As a second global shape descriptor, we considered the Spherical Harmonics indexes [35] alone, which is generally efficient for rigid matching and comparison of objects with spherical, or at least cylindrical, symmetries [17].

As concluding remarks for this section, it is worth noticing that the selections made for descriptors are not, and cannot be, guided by knowledge about the performance of specific descriptors on specific classes. Rather, the selection is made on the prospects opened by their combination in a search engine: there might be descriptors in the current literature that are better suited for a specific similarity task. However, the goal is to investigate how a user can navigate a dataset based on a given set of computed descriptors that capture different properties, like orthogonal axes in the similarity space.

4. Conceptual modelling of the similarity engine

The proposed comparison framework acts as a query-by-example search engine that combines a set of properties, following the ideas of faceted metadata for images [22, 57]. Starting from a model A (the query model) of a given collection of 3D models R , the list L of models in R that are considered similar to A is retrieved. The search is based on the activation of a number of properties (P_1, P_2, \dots) that the user selects before running the search. Each property corresponds to one of the shape signatures discussed in Section 3.2. The design of the search engine is compatible with prior classifications of the types of the objects collection. For instance, in case of archaeological datasets like the GRAVITATE use case the type of the query object (sherd/non-sherd) determines what properties are considered valid search criteria based on the nature of A (e.g.: if A is sherd-like, the filter related to properties that works on non-sherd-like object are not activated).

Each property P_i is interpreted as a *filter* (F_i), which removes from L the models that are dissimilar to A with respect to P_i . If the user judges the filter results too restrictive, it is possible to relax the filter severity incrementally adding more models to L and gradually enlarge the queue of models shown in the query window. Note that the list L could be empty.

From a more technical point of view, the single filters are distance matrices among models, one matrix for each property. When the user selects one or more criteria to compare a query model against the dataset, each criterion acts as a filter. The combination of more filters is done on the basis of the logic *and* operator. In practice, the model A is similar to the model B with respect to the two properties P_1 and P_2 if the two models A and B are similar with respect to both P_1 and P_2 . If L_1 is the list of the query results for the property P_1 and L_2 is the set of retrieved models for P_2 , the set L of the models that satisfy both P_1 and P_2 corresponds to the intersection of the two sets, i.e. $L = L_1 \cap L_2$. With these settings, the combined distance is a metric if all the distances for the single properties are metrics.

The query results that fulfill all the criteria are ranked according to a combined measure defined as the product of all the distances that are smaller than a threshold t : this threshold acts as a tuning parameter for the granularity of the search results, using the criteria described below.

A different threshold is automatically set for each shape signature, as follows. Given a distance matrix D , the k -th row Q_k stores the distance of the k -th model with respect to all the other elements in the dataset. Given a threshold t , the set V_k of models such that $Q_k(j) \leq t$ are considered valid query results for the model k with respect to the property stored in D .

The similarity between two models is defined as a score, represented in terms of a non-negative, real value that translates the distance between two signatures in a number. Depending on the shape signature adopted and the variety of the elements in a dataset, the image of the similarity distance spans different intervals of values, unless the similarity distance is normalized into interval $[0, 1]$. In our framework, we prefer to avoid normalizing the similarity distances and then combining the similarity scores derived from the methods described in Section 3.2. If a unique, global maximum is not known a priori, the normalization would depend on the dataset, and the similarity distance could be biased by the context. For this reason, we describe in the following the procedure to determine the initial thresholds for dataset exploration, adaptively tuning them to the dataset variations. For each matrix D , the threshold t is automatically predetermined as: $t = \text{average}_k\{t_k\}$, where $t_k = \text{argmin}\{\#\{V_k\} \approx 5\% \text{ of the elements in the model collection}\}$. With the symbol $\#$ we mean the set cardinality. Since t_k is an average value, it the size of L varies from model to model and, the list L could be empty.

By tuning the thresholds values, it is also possible to enlarge the search result set, offering more flexibility and interaction to the user. A value dt is determined as the average of the dt_k values that increase the number of elements of V_k of approximately 3%. The value dt is selected as many times the user decides to relax the threshold t . The thresholds values t and dt are automatically determined for every property and dataset.

A schema of the components and behaviour of the search engine is shown in Figure 6. The schema also sketches the layout of the graphical user interface, which has been fully developed in the GRAVITATE project (Figure 6(Bottom)).

5. Datasets

The performance of the search framework has been evaluated in two domains using three different datasets. For the archaeological domain, we have used the GRAVITATE case studies and the Virtual Hampson collection. For a more generic search domain, we have selected the dataset proposed by the YCB benchmark [67], which consists of objects of daily life with different shapes, sizes, and textures.

5.1. The GRAVITATE case studies

The primary GRAVITATE [11] collection is composed by fragments of terracotta figurative statues discovered in Salamis, on the island of Cyprus, dating back to the seventh - early sixth

century BC [68]. Most of these statues are fractured, while most shards are faded and eroded: in the project, we worked on 241 digital models of Salamis statues fragments, acquired by laser scans.

Another set of models come from the Naukratis collection at the British Museum [69]. It is a collection of pottery vessels fragments from Naukratis, a Greek trading port on the Nile Delta, in Egypt, dated from the VII century BC to the VII century AD. The collection available for our studies is made of 72 digital fragments, acquired by photogrammetric scans. Figure 7 shows some examples of the fragments in the GRAVITATE use cases.

5.2. The Hampson collection

The second collection of archaeological artifacts is provided by the Virtual Hampson Museum ³ in the form of textured, triangle meshes. The dataset comprises 442 models, 395 of them are available for download, representing remains of Native Americans groups that were living in the northwestern portion of the central Mississippi valley from about 1450 to 1650 AD and that are referred to by archaeologists as the Nodena phase. An overview of this dataset is shown in Figure 8.

5.3. The YCB benchmark

This benchmark contains a set of 71 3D models generated from visual data that are commonly involved for human-robotic interaction. These models are available at <http://ycb-benchmarks.s3-website-us-east-1.amazonaws.com/>. The dataset contains everyday objects whose common trait is the fact they can be grasped with a robot, see Figure 9. Data were acquired with the scanning rig used to collect the BigBIRD dataset. Details on the dataset can be found in [70, 67].

6. Experiments on similarity reasoning

The considered collections are organized with respect to criteria that are different with respect to the similarity criteria that the archaeologists have identified as suitable for their findings. Indeed, these collections are organized through factors like the functionality of the models or the provenance and it is really complex (or impossible) to define a single descriptor that is able to assess such tricky classifications. Since we are not able to evaluate the search results with a rigid ground-truth, we judge their correctness based on their visual appealing, minding also the similarity criteria considered in each query search.

In the Sections (6.1 - 6.3), we discuss the outcome of some queries based on the ten shape signatures detailed in Section 3.2. Although the quality of the query results is obviously dependent on the descriptor(s) selected, we think that these examples highlight how the combination of multiple properties can tailor the engine to the users needs.

³VHM, <http://hampson.cast.uark.edu>

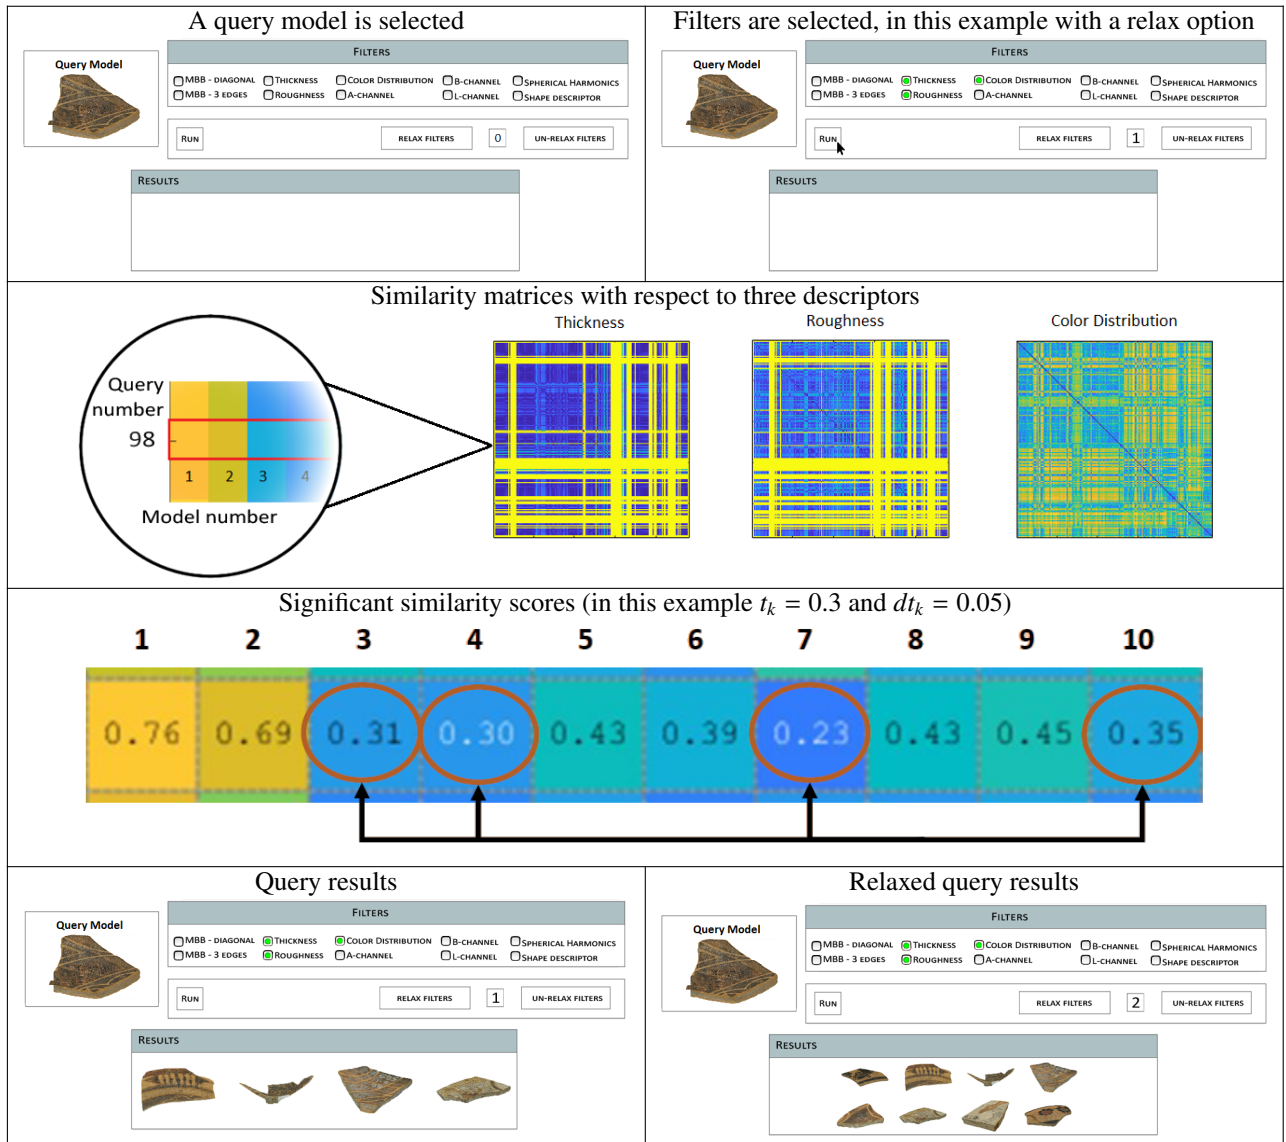


Fig. 6. An overview of the search engine behavior using three descriptors on a model of the Naukratis collection and the effect of a further relax.

As an exception, the experts provided a limited ground-truth for 46 of the 72 fragments in the Naukratis collection, identifying 8 groups of sherds that are likely to belong to the same object. This allows us to provide a quantitative analysis of the combination of two descriptors against the single ones, highlighting the advantages in using the search engine as a navigation tool instead of considering the single descriptors. This is discussed in Section 6.4.

Finally, since the YCB dataset is not a Cultural Heritage related dataset, in Section 6.3 we discuss the generalization of the similarity criteria listed in Section 4 towards a more generic collection.

6.1. The GRAVITATE case studies

On the Salamis collection, we performed a preliminary distinction of the fragments into two broad categories: those shaped like a classical sherd and those representing still a volumetric component of the original object. This distinction was

validated by archaeological experts. For instance, with reference to Figure 7 (right), the heads, the figurines and the busts are examples of volumetric fragments, while the decorated fragments (e.g. top-right and bottom-right) or the fragments in Figure 7(Right) are examples of sherd-like fragments.

Figure 10(A-F) shows a set of query results for the GRAVITATE use cases. Rows 10(A-B) present the same query model with two different query options: in (A) only the Color Distribution filter is adopted, thus admitting models with different geometric shape. The combination of Color Distribution with Roughness (here, we are considering a sherd-like fragment of a statue and the overall smoothness of the fragment surface is relevant) is shown in 10(B) and provides a large set of fragments that are intuitively more similar (in terms of red-like color distribution and surface smoothness) to the query model. Row 10(C) presents the combination of Color Distribution, Roughness and Thickness at once. Again, the results are quite intuitive. The archaeologists classified the first six fragments (the

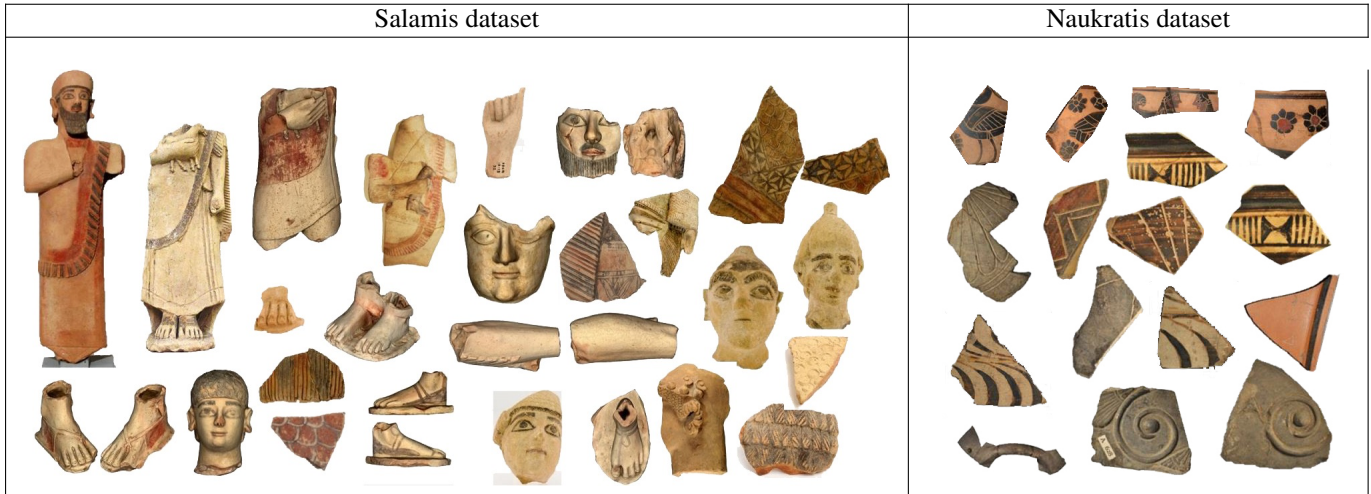


Fig. 7. Examples of fragments of Salamis terracotta figurines (right) and of Naukratis potteries (left).



Fig. 8. Examples of fragments from the VHM dataset.

query models and the first five results) as potential elements of the same vase. Nevertheless, also the other elements of the row 10(C) present quite homogeneous properties. Similarly, the example in row 10(D) highlights how using only Color Distribution and relaxation on the threshold we are able to select all the fragments in the dataset that potentially belong to the same group because they are made of the same material. One of the peculiar characteristics of this group is that all its elements are made by a gray stone, while the Shape, the Roughness and the Thickness of these fragments is largely variable. Finally, rows 10(E-F) look at the different capability of the global shape descriptors we have implemented in the search engine. The Shape property corresponds to a combination of global descriptors that are able to deal with both volumetric and topological shape distribution, thus mixing the head model with other heads but also hands. On the other hand, the Spherical Harmonics is a scale,

invariant, global representation where the overall shape of the object is represented with the coefficients on the spherical harmonics computed with respect to the center of mass: besides the heads in the dataset the query results are tricked by cylindrical objects, like the last model in the result set.

6.2. The Hampson collection

Figures 10(G-J) report some query results for the Hampson collection. The example 10(H) is based on the size of the model (described using the MBB(edges) descriptor) and shows us the possibility to navigate the dataset considering only fragments of a specific overall size. In this case, all the five models in the class are retrieved, without any false positive. The rows 10(G,I) combine other properties, such as Shape or Color Distribution, with the model size. The outcome of these queries highlights that the combination of properties targets the query results to

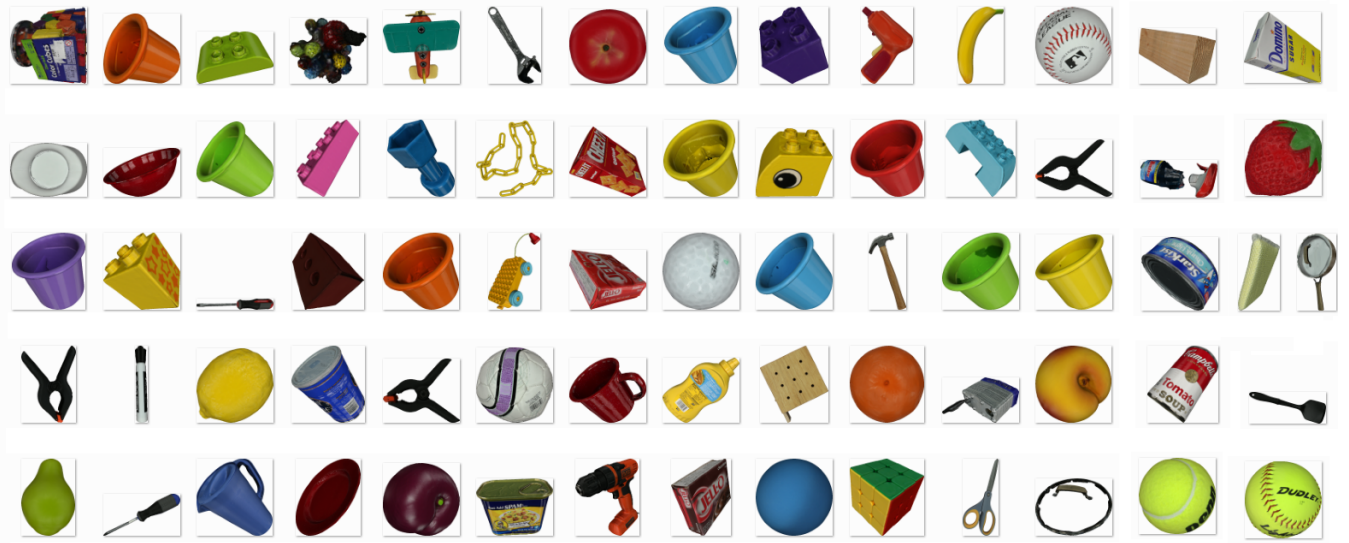


Fig. 9. Overview of the YCB dataset.

the query models. Finally, the row 10(J) shows an example of query results when combining Shape with Thickness when the query model is a bowl-like model.

Most of the examples on the Hampson collection highlight that a threshold relaxation is often applied. Indeed, when acting with multiple filters it is possible that the queue of the query results (i.e., the set of models satisfying all the properties selected) is empty. This is not surprising because the collection is not *complete* in the sense that there are not exemplars for all the possible property combinations. Moreover the thresholds are automatically set roughly evaluating the property variance on the whole collection and the query queue of some models can be very short. To overcome this effect, the *relax* option is crucial, allowing the user to complete the query search desired without being forced to ignore one of the similarity criteria.

6.3. YCB Benchmark

Aiming at a generalization of the criteria identified for archaeological context, we have revisited the model properties in order to identify the search axes in the conceptual dimensions of the similarity space that could be useful in generic contexts. We do not argue, obviously, that these are the only and the most relevant ones and we point the reader to specific studies in this direction, such as [71].

1. *Overall size*: in a generic context, this property is still relevant, and it reflects the need of considering the volume occupied by the model, not only in terms of absolute volume but also in terms of proportion among the three dimensions of the spatial embedding. item *Local shape properties*: this concepts involves properties related to local surface properties, like roughness and thickness.
2. *Color*: similarly to the CH use case, with colorimetric similarity we mean similarity by global texture behavior.
3. *Overall shape*: this criterion corresponds to the classical global matching measures. It admits similarity by global shape distribution but also with respect to articulations.

We test these criteria over the YCB collection that presents a large variety of colors, shapes (spherical-like, elongated, kitchen utensils, boxes with sharp edges, cans, etc.). The examples in Figure 11(A-B) shows that the Color Distribution property is far more effective in a dataset with such a larger variety of colors than the GRAVITATE use case and the Hampson dataset, which are populated of artifacts that are mostly brown (terracotta) with colors ranging from yellow to red, at most. Figure 11(C-E) instead shows searches done using the Spherical Harmonics as global shape signature. The results are satisfactory, as the closest models have similar global shape to that of the query model. The addition of the Thickness to the query search 11(E) highlights how this property keeps models with a shape distribution similar to the query cup model. Similarly, Figure 11(F) provides an example of query results when the combined shape signature is considered. Still, some false positives pop up in the results, but this is due to the nature of the descriptors and the fact that the false positives are actually very similar to the query models in terms of elongated shape distribution. Finally, the examples in 11(G-H) show examples of queries under overall size and the combination of color and overall shape for some utensil-like models.

6.4. Discussions and quantitative evaluation

The aim of this search engine is to provide a tool for the navigation of 3D model collections where multiple similarity axes can be identified. Depending on the user needs and tasks, the similarity notion can be different and vary in time; for instance, a professional user could be interested in details like manufacturing aspects and specific decorations, while a more generic user could search for all statues with black stripes. The data collections we are targeting generally are not fully classified and a well specified ground-truth is not available. Therefore, a generic, meticulous analysis of the performance of the shape descriptors combinations proposed by the search engine against the single methods is impossible.

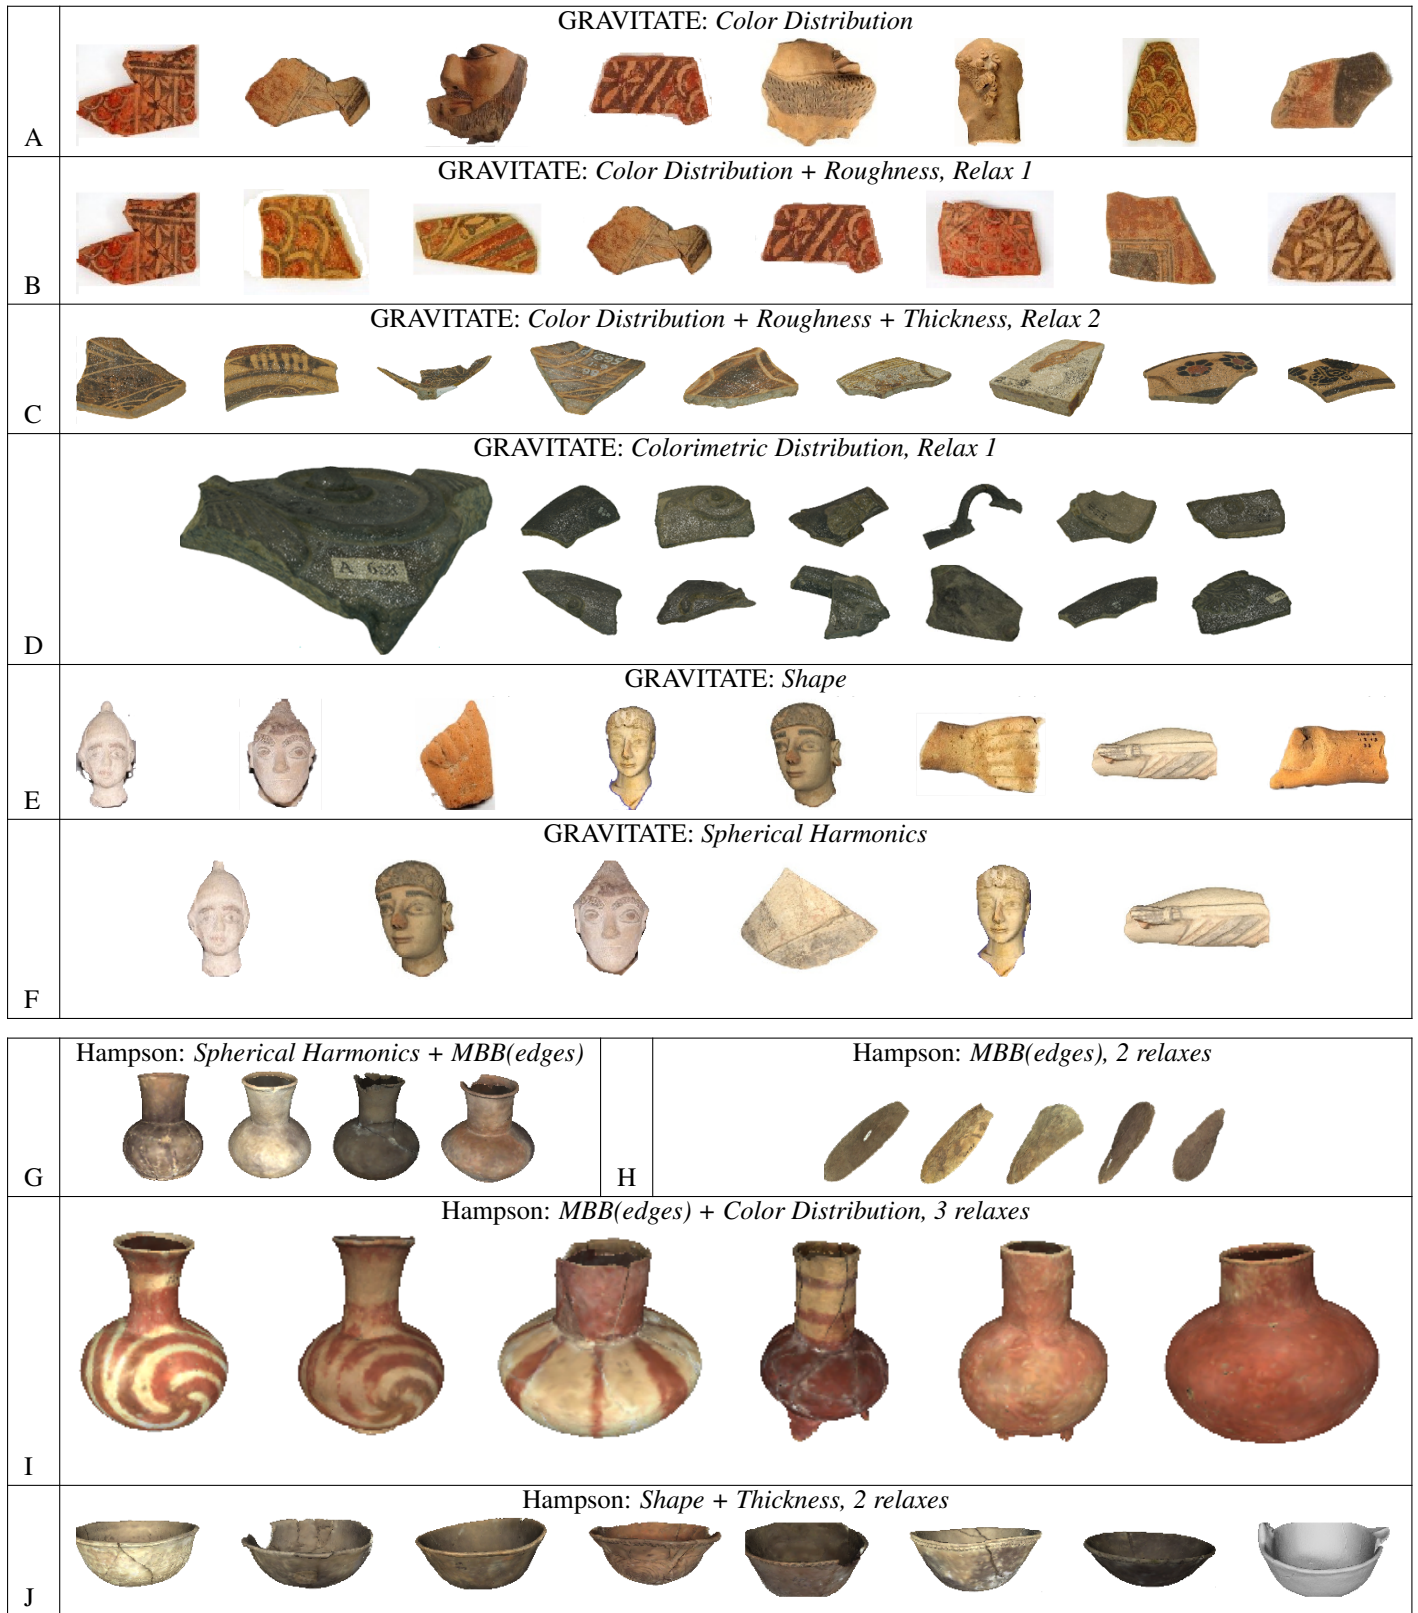


Fig. 10. Examples of query search on models of the GRAVITATE use case and Hampson collection. The first model in each cell is the query model.

As a possible case study, we analyze the search engine performance over the groups of fragments in the Naukratis collection that are likely to belong to the same object. The intuition behind this test is that fragments that belong to the same object

share the same colorimetric distribution and, if they are contiguous, they have compatible thickness. Therefore, we analyze the ability to group each fragment with the other elements of the object it belongs to for the single shape descriptors (color

5
6
7
8

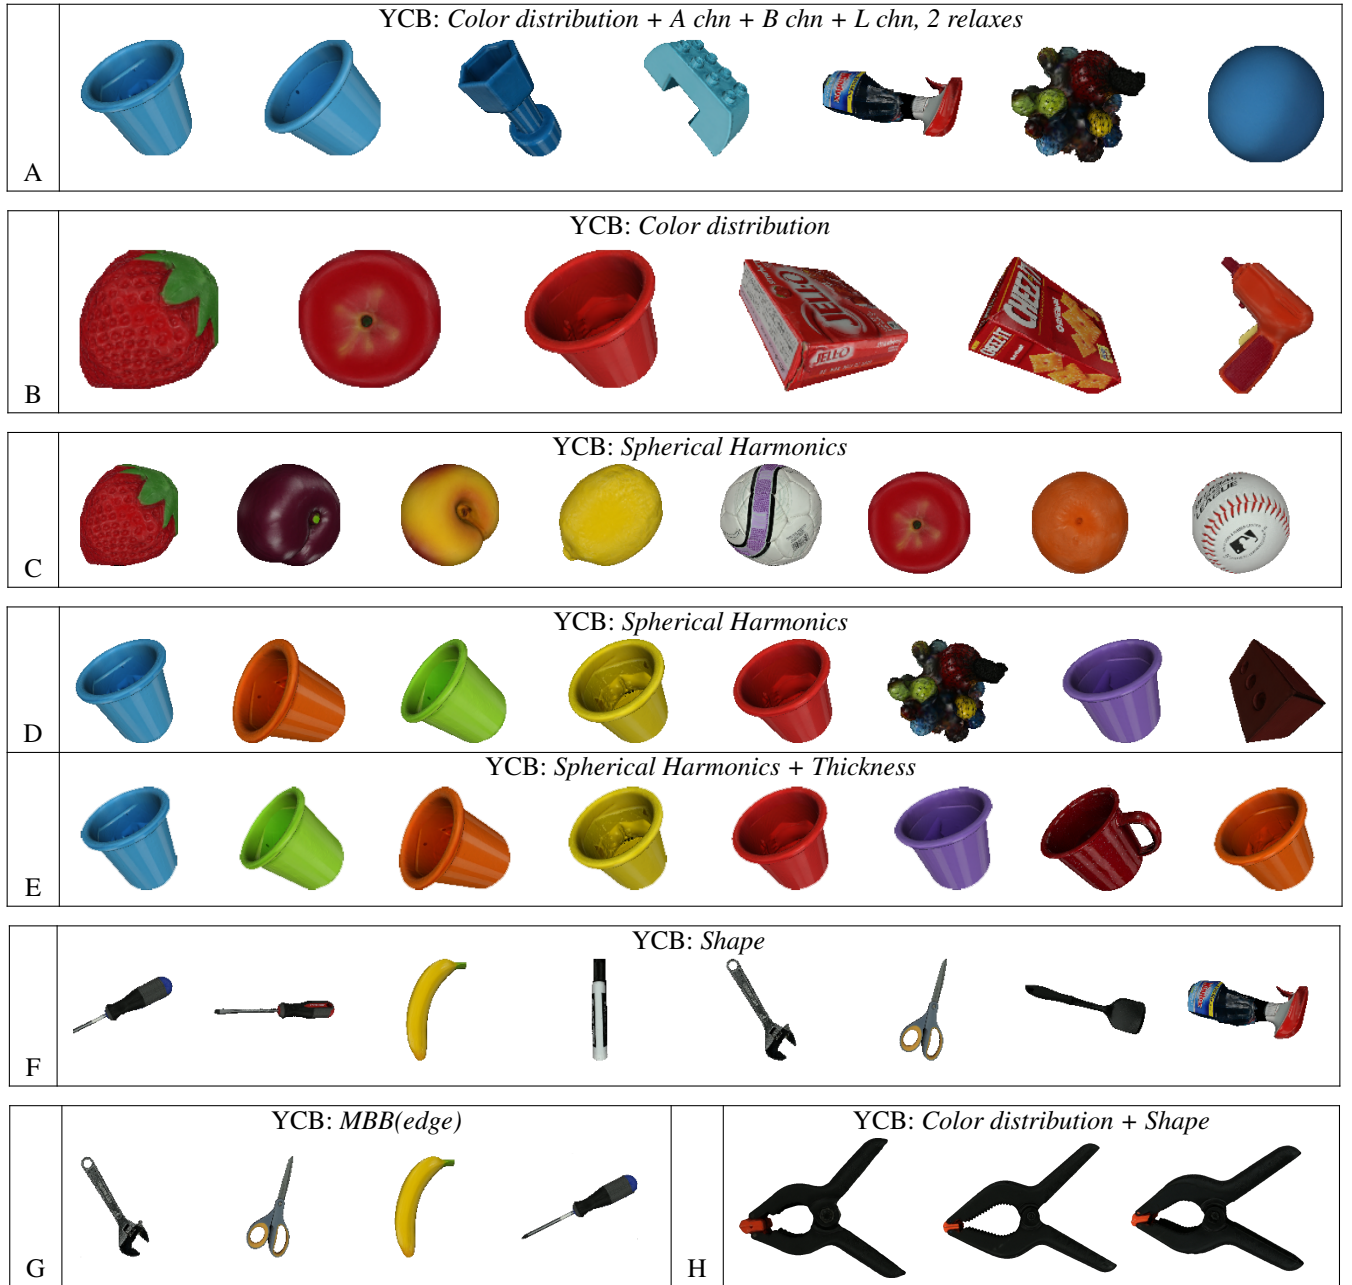


Fig. 11. Examples of query search on models of the YCB dataset. The first model in each cell is the query model.

distribution and thickness) and their combination. In turn, we use each model as a query against the whole Naukratis collection. As performance measures we consider the percentage of true positives and the number of false positives that appear in each query result. The main intuition is that the similarity combination adopted in this search engine diminishes the number of false positives because the query results of the combinations much satisfy multiple criteria at the same time.

In practice, we consider 46 fragments of the Naukratis dataset, grouped in 8 different classes (a sample model per class is reported in Figure 12). Notice, that the classes are numbered according to the archaeological catalogue; thus, there is not continuity in the class labels. In summary, for every frag-

ment, we analyze its query results with respect to the following runs:

- *Search 1* - filter: Color Distribution;
- *Search 2* - filter: Thickness;
- *Search 3* - filter: Combination of Color Distribution and Thickness.

In all the runs, the we equally set the relax filter to two. We overview the performance of the three filters, *Search 1*, *Search 2* and *Search 3* in Figure 13. In the ideal case, we want the combined query (*Search 3*) to hold the same true positive percentage as in *Search 1* and 2 but with a smaller number of false



Fig. 12. The list of classes used for a quantitative evaluation on the Naukratis dataset.

archaeological fragments is that they do not have a unique interpretation and multiple similarity criteria can be applied depending on the archaeological interests. The classical partial matching problem interpreted as a part-in-whole problem is not enough to address similarity among fragments because the original model is generally missing and the archaeologists only own small fragments that cannot be completely reassembled. In this sense, it is necessary to develop forms of multi-signatures that specifically target the different fragment parts; for instance, the external part for decorations and internal one for recognition the print of the technical turning.

Many problems are still open and need further efforts to be solved. For instance, the CH artifacts present quite complex features such as decorations, style elements and patterns (either colorimetric and geometric) that require the development of specific descriptors [72, 73]. In this direction, we are currently working on localizing the descriptors to sub-parts only and on the integration of further descriptors into the search engine. The quantitative analysis of the combination mechanism performed in Section 6.4 highlights that, to be effective, the single descriptors must have high true positive rates rather than a small number of false positives: this could be an advice for the development of future descriptors that aim at modeling a single similarity aspect. Same goes for skin continuity, which, in terms of compatible overall skin curvature, is crucial for re-assembling artifacts. Currently, it is addressed for models with symmetries, such as potteries, for instance considering the the partial axial symmetry of the surface [74].

As mentioned in Section 3.2, in the CH domain and in small datasets like the YCB there is a lack of substantial training data that in our opinion partially prevent the adoption of learning techniques. Indeed, learning techniques are based on supervised information, requiring extensive training data in addition to practical configuration expertise and computational resources [75]. On the contrary, in our use cases only a few shapes of a certain class or type exist, and they do not correspond to existing, complete shapes. Using methods like those addressed in this paper, a domain expert might be better able to determine the type of properties needed (occurrence of feature points, curvature statistics, repeated patterns, etc.) and which may be interpreted to some extent.

Acknowledgments

Acknowledgments will be properly added in the final version of the paper.

References

- [1] Velkamp, RC, Hagendoorn, M. State-of-the-Art in shape matching. In: Lew, M, editor. Principles of Visual Information Retrieval. Springer-Verlag; 2001, p. 87–119.
- [2] Bustos, B, Keim, DA, Saupe, D, Schreck, T, Vranić, DV. Feature-based similarity search in 3D object databases. *ACM Comput Surv* 2005;37(4):345–387.
- [3] Tangelder, JWH, Velkamp, RC. A survey of content based 3D shape retrieval methods. *Multimedia Tools Appl* 2008;39(3):441–471.
- [4] Santini, S, Jain, R. Similarity measures. *IEEE Trans Pattern Anal Mach Intell* 1999;21(9):871–883.

positive results. Each element of the diagrams corresponds to a model and the height of the columns measures the performance of search engine when that model is adopted as the query. Models that belong to the same class have same color. In the diagrams on the left, for each model, we report the percentage of the elements of the same class of the query model correctly retrieved (true positives). The value 1 means that all the elements of the class are shown in the list of the query results, 0.5 means that on a half of the elements are correctly retrieved when that fragment is adopted as the query. In the diagrams on the right, we report the number of false positives that are present in each query. Theoretically, this number could equal the dataset size less the elements of that query group. We observe that some of the classes (the first, third, fifth, sixth and seventh classes) for *Search 1* and *2* have comparable percentages of true positives and numbers of false positives. For *Search 3*, we observe that while the percentage of true positives is slightly lower overall, the number of false positives is far away lower than in the other searches. Overall, these results show how the combination proposed in this search engine works and highlight that it can manage queries based on multiple descriptors diminishing the number of false positives in the query results.

7. Conclusions

Working with real use cases poses a lot of problems, which span from the need of keeping results consistent across different resolutions to the management of damaged models.

We have shown that the search engine conceptual model can be generalized across different domains. However, each domain has its own needs and the same object can be seen from different points of view. Therefore, future efforts will be devoted to a specialization of descriptors and measures to the different domains, especially in the CH domain, since that is where this search engine is born. For instance, a peculiar aspect of the

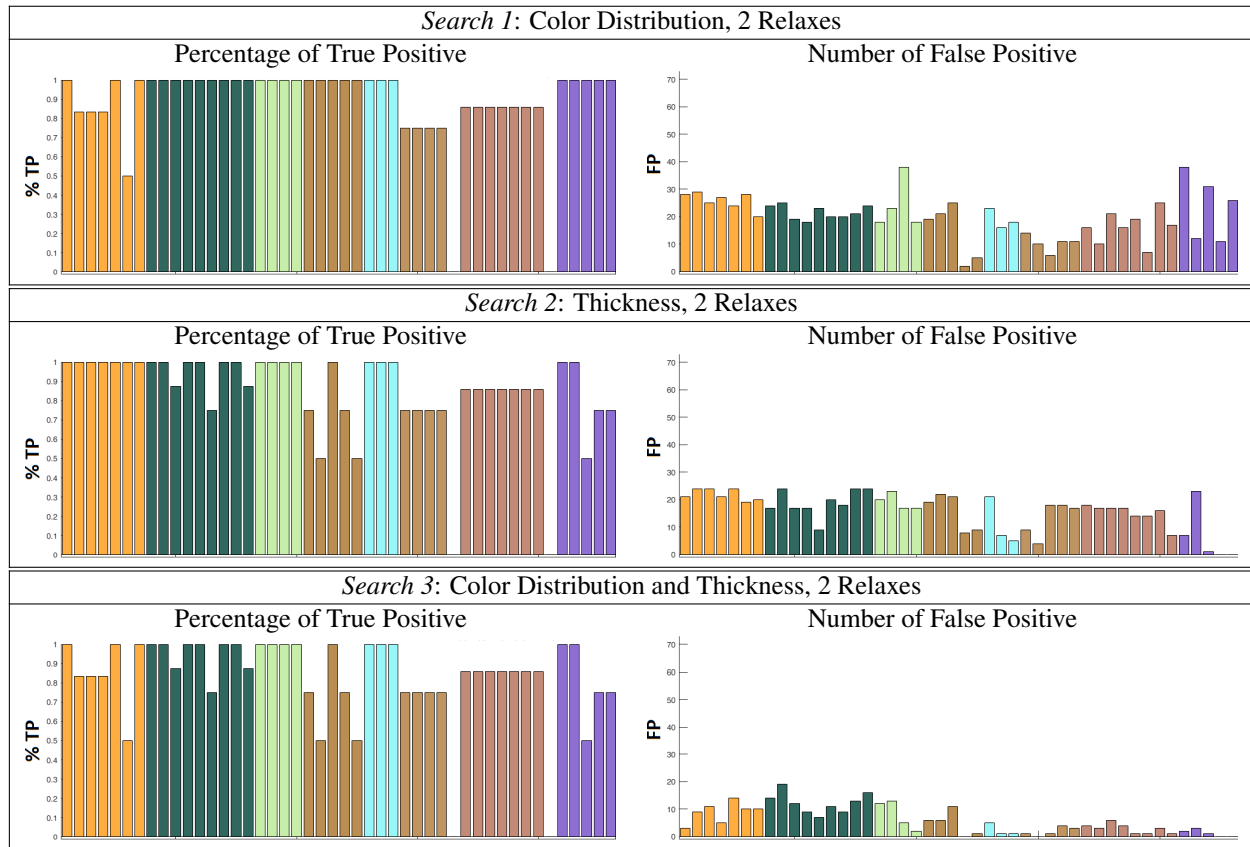


Fig. 13. Qualitative test on the Naukratis collection. Each column represent the a query search based on a Naukratis model that has been labelled by the experts.

- [5] Biasotti, S, Cerri, A, Bronstein, A, Bronstein, M. Recent trends, applications, and perspectives in 3D shape similarity assessment. *Computer Graphics Forum* 2016;35(6):87–119.
- [6] Shilane, P, Min, P, Kazhdan, M, Funkhouser, T. The Princeton shape benchmark. In: *Shape Modeling International*. IEEE Computer Society; 2004, p. 167–178.
- [7] Biasotti, S, Falcidieno, B, Giorgi, D, Spagnuolo, M. The Hitchhiker's guide to the galaxy of mathematical tools for shape analysis. In: *ACM SIGGRAPH 2012 Courses*. SIGGRAPH '12; New York, NY, USA: ACM; 2012, p. 17:1–17:33.
- [8] Biasotti, S, Falcidieno, B, Giorgi, D, Spagnuolo, M. 3D objects exploration: Guidelines for future research. In: *9th Eurographics Workshop on 3D Object Retrieval, 3DOR 2016, Lisbon, Portugal, May 8, 2016*. 2016, p. 9–12.
- [9] Jetter, HC, Gerken, J, Zöllner, M, Reiterer, H, Milic-Frayling, N. Materializing the query with facet-streams – a hybrid surface for collaborative search on tabletops. In: *The ACM CHI Conference on Human Factors in Computing Systems (CHI 2011), May 7–12, 2011, Vancouver, Canada*. ACM; 2011, p. 3013–3022.
- [10] Biasotti, S, Moscoso Thompson, E, Spagnuolo, M. Experimental Similarity Assessment for a Collection of Fragmented Artifacts. In: *Telea, A, Theoharis, T, Veltkamp, R, editors. Eurographics Workshop on 3D Object Retrieval*. The Eurographics Association; 2018, p. 103–110.
- [11] GRAVITATE: Discovering relationships between artefacts using 3D and semantic data. 2015–2018. EU H2020 REFLECTIVE project.
- [12] Corney, J, Rea, H, Clark, D, Pritchard, J, Breaks, M, Macleod, R. Coarse filters for shape matching. *IEEE Comput Graph Appl* 2002;22(3):65–74.
- [13] Kim, VG, Li, W, Mitra, NJ, DiVerdi, S, Funkhouser, T. Exploring collections of 3D models using fuzzy correspondences. *ACM Trans Graph* 2012;31(4):54:1–54:11.
- [14] Liu, Z, Bu, S, Zhou, K, Gao, S, Han, J, Wu, J. A survey on partial retrieval of 3D shapes. *J Comput Sci Technol* 2013;28(5):836–851.
- [15] Cardone, A, Gupta, SK, Karnik, M. A survey of shape similarity assessment algorithms for product design and manufacturing applications. *Journal of Computing and Information Science in Engineering* 2003;3:109 – 118.
- [16] Iyer, N, Jayanti, S, Lou, K, Kalyanaraman, Y, Ramani, K. Three-dimensional shape searching: state-of-the-art review and future trends. *Computer-Aided Design* 2005;37(5):509 – 530. *Geometric Modeling and Processing* 2004.
- [17] Schulz, A, Shamir, A, Baran, I, Levin, DIW, Sithi-Amorn, P, Matusik, W. Retrieval on parametric shape collections. *ACM Trans Graph* 2017;36(4).
- [18] Biasotti, S, Cerri, A, Falcidieno, B, Spagnuolo, M. 3D artifacts similarity based on the concurrent evaluation of heterogeneous properties. *J Comput Cult Herit* 2015;8(4):19:1–19:19.
- [19] Smeulders, AWM, Worring, M, Santini, S, Gupta, A, Jain, R. Content-based image retrieval at the end of the early years. *IEEE Trans Pattern Anal Mach Intell* 2000;22(12):1349–1380.
- [20] Leifman, G, Meir, R, Tal, A. Semantic-oriented 3D shape retrieval using relevance feedback. *The Visual Computer* 2005;21(8-10):865–875.
- [21] Giorgi, D, Frosini, P, Spagnuolo, M, Falcidieno, B. 3D relevance feedback via multilevel relevance judgements. *The Visual Computer* 2010;26(10):1321–1338.
- [22] Hearst, MA. Natural search user interfaces. *Commun ACM* 2011;54(11):60–67.
- [23] van Kaick, O, Zhang, H, Hamarneh, G, Cohen-Or, D. A survey on shape correspondence. *Comput Graph Forum* 2011;30(6):1681–1707.
- [24] Biasotti, S, Falcidieno, B, Giorgi, D, Spagnuolo, M. Mathematical tools for shape analysis and description. *Synthesis Lectures on Computer Graphics and Animation* 2014;6(2):1–138.
- [25] Biasotti, S, Cerri, A, Giorgi, D, Spagnuolo, M. PHOG: Photometric and Geometric Functions for Textured Shape Retrieval. *Comput Graph Forum* 2013;32(5):13–22.
- [26] Laga, H, Mortara, M, Spagnuolo, M. Geometry and context for semantic

- correspondences and functionality recognition in man-made 3D shapes. *ACM Trans Graph* 2013;32(5):150:1–150:16.
- [27] Spagnuolo, M. Shape 4.0: 3D shape modeling and processing using semantics. *Computer Graphics and Applications, IEEE* 2016;36(1):92–96.
- [28] Chen, DY, Tian, XP, Shen, YT, Ouhyoung, M. On visual similarity based 3D model retrieval. *Comput Graph Forum* 2003;22(3):223–232.
- [29] Kleiman, Y, Fish, N, Lanir, J, Cohen-Or, D. Dynamic maps for exploring and browsing shapes. *Computer Graphics Forum* 2013;32(5):187–196.
- [30] Gao, L, Cao, Y, Lai, Y, Huang, H, Kobbelt, L, Hu, S. Active Exploration of Large 3D Model Repositories. *IEEE Trans Vis Comput Graph* 2015;21(12):1390–1402.
- [31] Xu, K, Zhang, H, Cohen-Or, D, Chen, B. Fit and diverse: Set evolution for inspiring 3D shape galleries. *ACM Trans Graph* 2012;31(4):57:1–57:10.
- [32] Huang, SS, Shamir, A, Shen, CH, Zhang, H, Sheffer, A, Hu, SM, et al. Qualitative organization of collections of shapes via quartet analysis. *ACM Trans Graph* 2013;32(4):71:1–71:10.
- [33] Ling, H, Jacobs, DW. Shape classification using the inner-distance. *IEEE Trans Pattern Anal Mach Intell* 2007;29(2):286–299.
- [34] Shapira, L, Shamir, A, Cohen-Or, D. Consistent mesh partitioning and skeletonisation using the shape diameter function. *Vis Comput* 2008;24(4):249–259.
- [35] Kazhdan, M, Funkhouser, T, Rusinkiewicz, S. Rotation invariant spherical harmonic representation of 3D shape descriptors. In: *EG/ACM SIGGRAPH Symposium on Geometry Processing*. Eurographics Association; 2003, p. 156–164.
- [36] Sun, J, Ovsjanikov, M, Guibas, L. A concise and provably informative multi-scale signature based on heat diffusion. *Comput Graph Forum* 2009;28(5):1383–1392.
- [37] Osada, R, Funkhouser, T, Chazelle, B, Dobkin, D. Shape distributions. *ACM Trans Graph* 2002;21(4):807–832.
- [38] Ankerst, M, Kastenmüller, G, Kriegel, HP, Seidl, T. Nearest neighbor classification in 3D protein databases. In: *7th Int. Conference on Intelligent Systems for Molecular Biology*. AAAI Press; 1999, p. 34–43.
- [39] Körtgen, M, Park, GJ, Novotni, M, Klein, R. 3d shape matching with 3d shape contexts. In: *The 7th Central European Seminar on Computer Graphics*. 2003, p. 2.
- [40] Sidi, O, van Kaick, O, Kleiman, Y, Zhang, H, Cohen-Or, D. Unsupervised co-segmentation of a set of shapes via descriptor-space spectral clustering. *ACM Trans Graph* 2011;30(6):126:1–126:10.
- [41] Talton, JO, Gibson, D, Yang, L, Hanrahan, P, Koltun, V. Exploratory modeling with collaborative design spaces. *ACM Trans Graph* 2009;28(5):167:1–167:10.
- [42] Kim, VG, Li, W, Mitra, NJ, Chaudhuri, S, DiVerdi, S, Funkhouser, T. Learning Part-based Templates from Large Collections of 3D Shapes. *ACM Trans Graph* 2013;32(4).
- [43] Averkiou, M, Kim, VG, Zheng, Y, Mitra, NJ. Shapessynth: Parameterizing model collections for coupled shape exploration and synthesis. *Comput Graph Forum* 2014;33(2):125–134.
- [44] Zheng, Y, Liu, H, Dorsey, J, Mitra, NJ. Ergonomics-inspired reshaping and exploration of collections of models. *IEEE Transactions on Visualization and Computer Graphics* 2015;.
- [45] Fish, N, Averkiou, M, van Kaick, O, Sorkine-Hornung, O, Cohen-Or, D, Mitra, NJ. Meta-representation of shape families. *ACM Trans Graph* 2014;33(4):34:1–34:11.
- [46] Rustamov, RM, Ovsjanikov, M, Azencot, O, Ben-Chen, M, Chazal, F, Guibas, L. Map-based exploration of intrinsic shape differences and variability. *ACM Trans Graph* 2013;32(4):72:1–72:12.
- [47] Mahalanobis, PC. On the generalised distance in statistics. In: *Proceedings National Institute of Science, India*. 1; 1936, p. 49–55.
- [48] Deza, MM, Deza, E. *Encyclopedia of Distances*. Springer Berlin Heidelberg; 2009.
- [49] Rubner, Y, Tomasi, C, Guibas, LJ. The Earth Mover's Distance as a metric for image retrieval. *International Journal of Computer Vision* 2000;40(2):99–121.
- [50] Ovsjanikov, M, Li, W, Guibas, L, Mitra, NJ. Exploration of continuous variability in collections of 3D shapes. *ACM Transactions on Graphics* 2011;30(4).
- [51] Koutsoudis, A, Pavlidis, G, Liami, V, Tsiafakis, D, Chamzas, C. 3D pottery content-based retrieval based on pose normalisation and segmentation. *Journal of Cultural Heritage* 2010;11(3):329 – 338.
- [52] Koutsoudis, A, Chamzas, C. 3D pottery shape matching using depth map images. *Journal of Cultural Heritage* 2011;12(2):128 – 133.
- [53] Ovsjanikov, M, Ben-Chen, M, Solomon, J, Butscher, A, Guibas, L. Functional maps: A flexible representation of maps between shapes. *ACM Trans Graph* 2012;31(4):30:1–30:11.
- [54] Smeets, D, Keustermans, J, Vandermeulen, D, Suetens, P. meshSIFT: Local surface features for 3D face recognition under expression variations and partial data. *Comput Vis Image Und* 2013;117(2):158 – 169.
- [55] Biasotti, S, Cerri, A, Aono, M, Hamza, AB, Garro, V, Giachetti, A, et al. Retrieval and classification methods for textured 3D models: a comparative study. *The Visual Computer* 2016;32(2):217–241.
- [56] Garro, V, Giachetti, A. Scale space graph representation and kernel matching for non rigid and textured 3D shape retrieval. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 2016;38(6):1258–1271.
- [57] Yee, KP, Swearingen, K, Li, K, Hearst, M. Faceted metadata for image search and browsing. In: *SIGCHI Conference on Human Factors in Computing Systems*. CHI '03; New York, NY, USA: ACM; 2003, p. 401–408.
- [58] Stamatiopoulos, M, Anagnostopoulos, C. The Thickness Profile method: A new digital 3D approach for reassembling unpainted archaeological ceramic pottery. 2016.
- [59] Angelo, LD, Stefano, PD, Pane, C. Automatic dimensional characterisation of pottery. *Journal of Cultural Heritage* 2017;26:118 – 128.
- [60] Gadelmawla, E, Koura, M, Maksoud, T, Elewa, I, Soliman, H. Roughness parameters. *Journal of Materials Processing Technology* 2002;123(1):133 – 145.
- [61] Phillips, F, Todd, JT. Perception of local three-dimensional shape. *Journal of experimental psychology Human perception and performance* 1996;22 4:930–44.
- [62] Koenderink, JJ, van Doorn, AJ. Surface shape and curvature scales. *Image and Vision Computing* 1992;10(8):557 – 564.
- [63] Cohen-Steiner, D, Morvan, JM. Restricted Delaunay triangulations and normal cycle. In: *19th Annual Symposium on Computational Geometry*. New York, NY, USA: ACM; 2003, p. 312–321.
- [64] Albuz, E, Kocalar, ED, Khokhar, AA. Quantized cielab* space and encoded spatial structure for scalable indexing of large color image archives. In: *Acoustics, Speech, and Signal Processing, ICASSP '00; vol. 6*. 2000, p. 1995–1998.
- [65] Suzuki, M. A Web-based retrieval system for 3D polygonal models. In: *IFSA World Congress and 20th NAFIPS International Conference*. Joint 9th; vol. 4. 2001, p. 2271–2276.
- [66] Biasotti, S, Cerri, A, Frosini, P, Giorgi, D. A new algorithm for computing the 2-dimensional matching distance between size functions. *Pattern Recogn Lett* 2011;32(14):1735–1746.
- [67] Calli, B, Singh, A, Bruce, J, Walsman, A, Konolige, K, Srinivasa, S, et al. Yale-CMU-Berkeley dataset for robotic manipulation research. *Int J Rob Res* 2017;36(3):261–268.
- [68] Karageorghis, V, Karageorghis, J, Foundation, AL. The coroplastic art of ancient Cyprus. Nicosia : A.G. Leventis Foundation; 1991. At head of title: A.G. Leventis Foundation.
- [69] British Museum - Naukratis: An introduction. 2017. URL: www.britishmuseum.org/research/projects/all_current_projects/naukratis_the_greeks_in_egypt/the_site_and_its_history.aspx.
- [70] Calli, B, Walsman, A, Singh, A, Srinivasa, S, Abbeel, P, Dollar, AM. Benchmarking in manipulation research: Using the Yale-CMU-Berkeley object and model set. *IEEE Robotics Automation Magazine* 2015;22(3):36–52.
- [71] Gardenfors, P. Conceptual spaces as a framework for knowledge representation. *Mind and Matter* 2004;2(2):9–27.
- [72] Biasotti, S, Moscoso Thompson, E, Barthe, L, Berretti, S, Giachetti, A, Lejembre, T, et al. Recognition of Geometric Patterns Over 3D Models. In: Telea, A, Theoharis, T, Veltkamp, R, editors. *Eurographics Workshop on 3D Object Retrieval*. The Eurographics Association; 2018, p. 71–77.
- [73] Moscoso Thompson, E, Biasotti, S. Description and retrieval of geometric patterns on surface meshes using an edge-based LBP approach. *Pattern Recognition* 2018;82:1 – 15.
- [74] Sipiran, I. Analysis of partial axial symmetry on 3d surfaces and its application in the restoration of cultural heritage objects. In: *2017 IEEE International Conference on Computer Vision Workshops (ICCVW)*. 2017,

- 1 p. 2925–2933.
- 2 [75] Bronstein, MM, Bruna, J, LeCun, Y, Szlam, A, Vandergheynst, P.
- 3 Geometric deep learning: Going beyond Euclidean data. IEEE Signal
- 4 Processing Magazine 2017;34(4):18–42.