

Edge-based LBP description of surfaces with colorimetric patterns

E. Moscoso Thompson¹ and S. Biasotti¹

¹ Istituto di Matematica Applicata e Tecnologie Informatiche ‘E. Magenes’ - CNR

Abstract

In this paper we target the problem of the retrieval of colour patterns over surfaces. We generalize to surface tessellations the well known Local Binary Pattern (LBP) descriptor for images. The key concept of the LBP is to code the variability of the colour values around each pixel. In the case of a surface tessellation we adopt rings around vertices that are obtained with a sphere-mesh intersection driven by the edges of the mesh; for this reason, we name our method edgeLBP. Experimental results are provided to show how this description performs well for pattern retrieval, also when patterns come from degraded and corrupted archaeological fragments.

Categories and Subject Descriptors (according to ACM CCS): : [Computer Graphics [I.3.6]]: Methodology and Techniques— Information storage and retrieval [H.3.3]: Information search and Retrieval

1. Introduction

Thanks to advances in the modeling techniques and to the availability of cheaper yet effective 3D acquisition devices, we see a remarkable increase of the amount of 3D data available. Many sensors are able to acquire not only the 3D shape but also its *texture*; this is the case, for instance, of the Microsoft Kinect device. The creation of an increasing number of 3D models has opened new opportunities to study the past, by giving access to plenty of representations of artifacts close to their original form. At the same time, Cultural Heritage owns a growing mass of non-interpreted 3D data, which call for innovative solutions for the analysis of data. In this context, local descriptors, feature recognition and similarity measures become indexes to the informative content of 3D models, and are essential to categorize objects and to recognize a style, e.g. to attribute objects to a given society or to a given author. A typical problem the archaeologists face when dealing with collections of fragments is to determine their compatibility. Compatibility is generally determined by multiple factors: geometric correspondence, same material and, possibly, if there are not evidently matching fragments, continuity consideration on the fragment skin (colour, texture) [Pe16].

Within the large scenario of Cultural Heritage, we focus on the analysis and description of color patterns. The idea is to recognize the same decoration, for instance a repeated lotus leaf, independently of the support (e. g., the surface bending) on which it is depicted. Therefore, this work will contribute to the definition of a compatibility measure among artifacts based on skin decorations. To approach this problem, we consider a novel extension of the Local Binary Pattern description to surface tessellations based on the evolution of the color over concentric circles around a vertex.

To determine these circles we adopt a sphere - edge intersection strategy and for this reason we name our approach edgeLBP. As application of the edgeLBP description, we propose the retrieval and classification of color patterns over surfaces.

The remainder of the paper is organized as follows. Section 2 briefly reviews the literature on the retrieval of textured images and surfaces. Section 3 introduces the elements of our method, i.e. the edgeLBP operator and how we store it in a descriptor. Section 4 presents and analyses the retrieval and classification performances of the method over two datasets, while conclusive remarks end the paper, Section 5.

2. State of art

A typical strategy to detect textures on images is to consider local patches that describe the behavior of the texture around a group of pixels. Examples of these descriptions are the Local Binary Patterns (LBP) [OPH96, OPM02], the Scale Invariant Feature Transform (SIFT) [Low04] and the Histogram of Oriented Gradients (HOG) [DT05]. The generalization of these descriptions to (even textured) surfaces has been explored in several works, such as the PANORAMA views of the 3D objects [PPTP10], the mesh-HOG [ZBH12] and the meshLBP [WTBB16, WTBB15]. In general, the methods for matching textured 3D shapes adopt a combination of geometric and colorimetric descriptors. Possible choices of the colorimetric descriptors are: feature-vectors, where the color is treated as a general property of the shape, [Suz01], or its subparts in [GG16]; local or global views of the objects [WCL*08, PZC13]; point-to-point correspondences among sets of feature points (e.g., the CSHOT descriptor [TSDS11]); the evolution of the sub-level sets according to the persistent homology settings [BCGS13].




Figure 1: In (a) the ring of the pixel i is shown; while (b) and (c) show two examples of concentric rings.

These methods mainly address the shape matching problem without focusing on the surface details and local colorimetric variations. On the contrary, when looking for patterns, locality and scale are the two key aspects. A detailed evaluation and comparison of methods for 3D texture retrieval and comparison can be found in [Be16] and several SHREC contests [Ce13, Be14, Ge15]. However, all these contests focused on the joint comparison of geometry and texture, without considering the comparison of the purely colorimetric information that characterizes the surface decorations.

At the best of our knowledge, the Mesh Local Binary Pattern (meshLBP) approach [WTBB16, WTBB15, WBB15] is the unique approach that explicitly addresses pattern analysis over surfaces. The meshLBP extends the LBP [OPH96] to triangle meshes. The main idea behind the meshLBP is that triangles play the role of pixels and the 8-neighbor connectivity in an image is ideally substituted by a 6-neighbor connectivity around triangles. Rings on the mesh are computed using a uniform, region growing, triangle-based expansion. From the practical point of view, the meshLBP encodes a pattern efficiently, providing a compact representation of it.

3. The edgeLBP

We extend the LBP to surfaces using rings defined on the basis of a sphere-mesh intersection. In Section 3.1 we briefly sum up the definition of the LBP definition. Our extension to surface tessellations is described in Section 3.2, while Section 3.3 details the edgeLBP descriptor and the distance adopted to compare two descriptors.

3.1. Local Binary Pattern for gray-scale images

The *Local binary pattern (LBP)* and its variants prove to be a good solution for the classification of patterns in images [LFG*17]. Given a gray-scale image I , the LBP describes the pattern in I coding the local variation of the gray-scale values (encoded with a function $h : I \rightarrow [0, 255]$) around each pixel of I . More extensively, for each pixel $i \in I$, a *ring* of pixels around i (called $ring_i$) is considered (see Figure 1) and a 8-digit binary array str_i defined as follow:

$$str_i(j) = \begin{cases} 1 & \text{if } h(i) < h(i_j) \\ 0 & \text{otherwise} \end{cases}$$

where i_j is the j -th pixel of the ring around i , sorted clockwise and starting from the top-left pixel. The LBP operator of a pixel i is defined by:

$$LBP(i) = \sum_j str_i(j) \alpha(j),$$




Figure 2: (a): in blue, two rings defined on the basis of triangles; (b): the ring around the vertex v is defined by mesh vertices (red dots).

where α is a weight function. Throughout this paper we consider $\alpha_1(j) = 1, \forall j$. Notice that in this case, the $LBP(i)$ value is independent of the ordering of $ring_i$. Finally, the LBP descriptor of the pattern in I is defined as the histogram of the values $LBP(i)$.

The LBP operator was extended to multiple rings around each pixel in I , see Figure 1(b-c). The descriptor of the LBP multi-ring is the concatenation of the histograms of the LBP values of each single ring, e.g., an array or a matrix.

3.2. Definition and implementation of the edgeLBP operator

We extend the multi-ring LBP operator to deal with surface tessellations through a sphere-mesh intersection technique, called the *edge Local Binary Pattern (edgeLBP)*. By a surface tessellation, we mean a polygonal mesh $T = (V, E, F)$, which is a collection of vertices V , edges E and faces F defining the surface of an object. In our settings, we assume that the faces of the tessellation are convex polygons; examples of admissible surface representations are triangle and quad meshes, [BLP*13].

We assume that the surface property can be stored as a scalar function h defined on the vertices of the tessellations, formally, $h : V \rightarrow \mathbb{R}$. In our settings, we consider two choices for the function h : (i) the *L-channel* from the CIE Lab color space [AKK00, HP11]; (ii) the gray-scale value defined as $0.21R + 0.72G + 0.07B$ (R , G and B are the channels of the RGB color space).

The concept of ring is crucial for the LBP operator: while a pixel grid has the same connectivity everywhere, surface tessellations can be widely *irregular*, thus the *ring* definition over them is not obvious. By irregular we mean that the vertices can be non uniformly distributed over the surface and the faces of the tessellation may have different area, shape and number of edges. Figure 2 depicts two possible *ring* definitions exclusively made of mesh elements (triangles in Figure 2(a) and vertices in Figure 2(b), resp.): in both cases, the irregularity of the mesh elements strongly influences these of rings.

We define the *ring* of a vertex $v \in V$ as the intersection of the surface tessellation with a sphere of radius R centered in v . Such an intersection is represented by the set of points $\mathcal{R} = \{p_1, p_2, \dots, p_k\}$ that approximate the intersection between the sphere and the surface. Figure 3 shows a number of concentric rings over a triangle mesh. To determine a *ring* around a vertex v , we follow a mesh expansion approach driven by the Euclidean distance from the vertex v , as summarized in the following steps:




Figure 3: (a): in black, multiple closed curves defined by the set of points $P_i \in \mathcal{R}$; (b): the black dots correspond to the elements p_i of the three central curves in (a).

1. All the edges that are incident in v are added to a list L .
2. Starting from an edge $e = (v, v_1) \in L$, the intersection between e and the sphere centered in v with radius R is evaluated. If there actually is an intersection, it is stored as a new point p_i , otherwise, if e completely falls inside the sphere, we add to L all the edges that are incident to v_1 . The edge e is removed from L and labeled as visited. The value $h(p_i)$ on p_i is given by the linear interpolation of the values that h assumes in v and v_1 .
3. The step 2 is repeated $\forall e \in L$, until the list is empty.

To achieve a multi-ring representation, for any vertex $v \in V$ we consider N_r rings, $\{ring_1^v, \dots, ring_{N_r}^v\}$. Let S_l^v be the surface portion of T that contains v and has the $ring_l^v$ as its boundary, $l = 1 \dots N_r - 1$, then the relation $S_l^v \subset S_{l+1}^v$ holds for each l . When extending the edgeLBP evaluation to multiple rings, the algorithm takes advantage of the nested nature of the rings and extracts S_l^v with respect to increasing values of the radius R .

In general, the sphere-surface intersection can produce multiple, closed curves that bound either a multiple connected or a disconnected portion of the surface, as detailed in [MPS*04]. Using a region growing approach, we dynamically consider only the S_l^v components. Therefore, S_l^v is always a connected region that contains v ; however, it can become multiply connected. If all the N_r components of S_l^v are simply connected and all the N_r rings do not intersect the surface boundary (if any), the v is considered an *admissible* vertex for the edgeLBP, otherwise it is *non-admissible*.

3.2.1. Ring re-ordering and sampling

Each *ring* is represented as the piecewise, linear curve C determined by the segments (p_i, p_{i+1}) , $p_i \in \mathcal{R}$. Then, the curve C is oriented *counter-clockwise* with respect to the vector in v normal to T . We select As the starting point for ordering C , we select the point \tilde{p} such that:

$$\tilde{p} = \operatorname{argmax}_{p_i \in \mathcal{R}} h(p_i).$$

In case of symmetries around a vertex, multiple choices of the starting point are possible: we select the candidate point that is the farthest from the other elements of \mathcal{R} . The stability of the starting point of a *ring* is confirmed in numerous experiments we performed on meshes of different resolution, where by mesh resolution we mean the number of vertices of the mesh. Figure 4 shows the vector field generated by the difference between \tilde{p} and v ($\tilde{p} - v$) all over the mesh. The orientation of the field indicates the position




Figure 4: Arrows represent the starting point of the rings in meshes representing the same surface but sampled with a different number of vertices (40K, 16K and 8K vertices, resp.).

of \tilde{p} . The pictures show a detail of the field over a mesh with 40K vertices and two mesh sub-samplings with 16K and 8K vertices: the overall orientation of the field (and therefore the choice of \tilde{p}) is robust to different mesh samplings. In case of multiple rings, \tilde{p} is selected only on the biggest ring $ring_{N_r}$; for each concentric ring, the starting point is the point p_i , which is the closest one to \tilde{p} .

Generally the number of elements $p_i \in \mathcal{R}$ varies from one ring to another, because of the increasing radius of the sphere and the irregularity of the tessellation, see Figure 3(b). To have the same number of elements on every ring, we sample C with P points, where P is a fixed number, called the *spatial resolution*. The results of this sampling is S , a set of equidistant samples of C , s_j with $j = 1, \dots, P$. In details, the equidistant re-sampling is performed as follows:

- we set the expected distance δr between two successive points in S as $\delta r = \frac{2\pi R}{P}$;
- we set $s_0 = \tilde{p}$ and extract the points s_j on C such that
$$|s_{j-1} - s_j| \approx \delta r, \quad j = \{1, \dots, P\}.$$

The value $h(s_j)$ is linearly approximated from the values the function h assumes on the extrema of the corresponding segment in C .

3.2.2. Choice of the ring radii

With the edgeLBP we are interested to code local variations on the surface, therefore the radius R should be kept small with respect to the overall dimension of the surface. This implies that the choice of the radius R is crucial for the type (and the size) of the patterns we are going to identify; indeed it must be not too large to avoid to mix global and local surface information and not too small to become insignificant. In practice, the multiply connected regions appear in case of topological noise, like small handles and mesh self-intersections; in our experiments over thousands of tessellations we never met meaningful admissibility problems.

We opt for a uniform distribution of the ring radii values. Denoting R_{max} the maximum radius and N_r the number of rings, the value of the ring radii will be $\frac{R_{max}}{N_r}, 2\frac{R_{max}}{N_r}, \dots, R_{max}$.

3.3. Similarity assessment

Once the function h is evaluated over the sample sets of the rings around v , the edgeLBP value on v straightforwardly follows from the classic LBP definition, see Section 3.1.

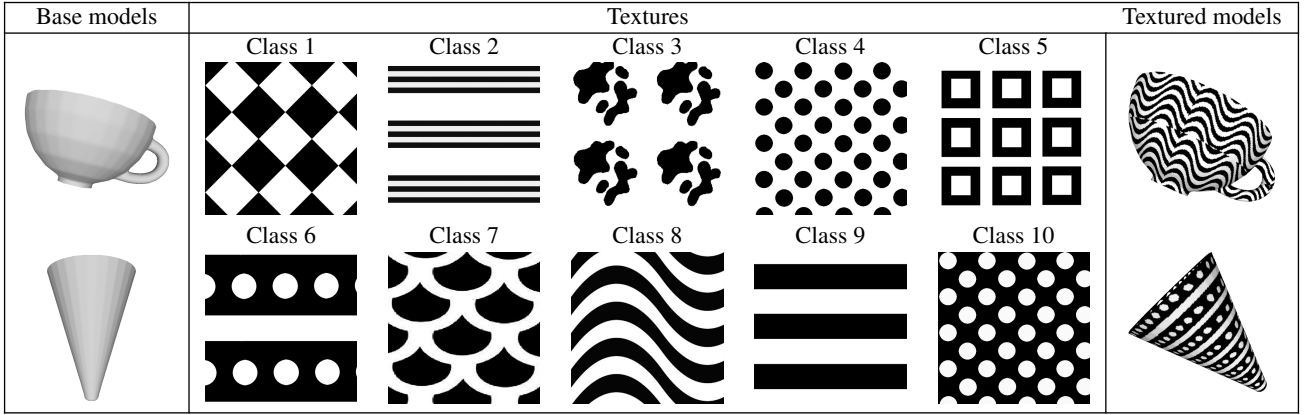


Figure 5: Left: Two of the original models. Center: The ten patterns imprinted on the models of the CPP dataset. Right: Two examples of the textured models of the CCP dataset.

Given the surface tessellation T , its *edgeLBP descriptor* is labeled D_T . The entry $D_T(n, m)$ is defined as the histogram that counts how many vertices have an edgeLBP value equal to m on the $ring_n$. Since in the experiments we are mostly interested in the distribution of the edgeLBP values, we adopt $\frac{D_T}{n_v}$ as the edgeLBP descriptor, where n_v is the number of the admissible vertices. Through this normalization of T we achieve robustness to the number of vertices of the surface representation.

We define the dissimilarity between two tessellations A and B as the distance between their corresponding edgeLBP descriptors D_A and D_B . Since the edgeLBP can be thought as a matrix, any feature vector distance is suitable to evaluate the similarity between two edgeLBP descriptors. We analysed the Euclidean distance between matrices, the *Earth Mover's Distance* as defined in [RTG00] and the *Bhattacharyya distance*. The Bhattacharyya distance d_{Bha} between two distributions ϕ and ψ of a scalar random variable X has the following definition:

$$d_{Bha}(\phi, \psi) = \sqrt{1 - BC(\phi, \psi)}, \quad BC(\phi, \psi) = \sum_{x \in X} \sqrt{\phi(x)\psi(x)},$$

where BC is called the *Bhattacharyya coefficient*. Then, for a set of surface tessellations, the dissimilarity values are stored in a *distance matrix* $DM(i, j) = d(D_i, D_j)$, where d is the distance between the descriptors of the tessellation i and j . Diagonal values of $Dist(i, i)$ are zero.

4. Experimental results

In this Section we introduce the datasets and the evaluation measures adopted to analyse the retrieval performance of the edgeLBP. We present the edgeLBP performances and discuss its robustness to different tessellations of the same surface.

4.1. Dataset

To evaluate the edgeLBP ability of effectively discriminating pattern variations, we used two datasets:

- the *Cups, Pots and Pans* dataset (or *CPP* for short) is created

from triangle meshes in the SHREC'07 Watertight model contest [GBP07] and the COSEG [WAvK*12] datasets (see Figure 5(Left)). The original meshes do not have any texture or colorimetric information. From 20 base models and 10 black and white textures representing a pattern (see Figure 5(Center)) we derived 200 models, applying each texture to every model with a semi-automatic algorithm. The proper RGB value was added to the mesh vertices discarding any other colorimetric information (see Figure 5(Right)). At the end of this process, each model is covered by one of the 10 patterns for at least the 30% of its surfaces while the rest of the surface is only black or only white. The number of vertices of the 200 models ranges from 95K to 107K.

- the *Artifacts* dataset is derived from the laser scans of CH artifacts stored in the STARC repository [SH07] and selected as test-beds in the Gravitare EU project [GRA]. The colorimetric information comes as a RGB value associated to each mesh vertex. Differently from the CPP dataset, this second dataset contains full-color information, with a predominance of red, yellow and brown nuances. From these fragments we identified 10 classes of different patterns (see Figure 6); then, for each type of pattern, we tailored 4 representative patches coming from different fragments, for a total of 40 patches. Every patch is made of approximately 40K vertices.

The edgeLBP algorithm is used to perform colorimetric pattern retrieval on the CPP and Artifact datasets, separately.

4.2. Evaluation measures

The evaluation tests have been performed using a number of classical information retrieval measures, namely the Nearest Neighbor, First Tier, Second Tier, Discounted Cumulative Gain, e-measure, Precision-Recall plot, confusion matrices and tier images.

Nearest Neighbor, First Tier, Second Tier These measures aim at checking the fraction of models in the query's class also appearing within the top k retrievals. In detail, for a class with $|C|$ members, $k = 1$ for the Nearest Neighbor (NN), $k = |C| - 1$ for the first tier



Figure 6: Representatives of the 10 classes considered in the Artifacts dataset.

(FT), and $k = 2(|C| - 1)$ for the second tier (ST). Note that all these values range from 0 to 1.

Discounted cumulative gain The Discounted Cumulative Gain (DCG) is an enhanced variation of the Cumulative Gain, which is the sum of the graded relevance values of all results in the list of retrieved objects of a given query. The definition of DCG adopted in this paper can be found in [JK02].

Precision-Recall, mAP and e-measure The *Precision* and *Recall* are common measures for retrieval evaluation. Recall is the ratio of the number of relevant records retrieved to the total number of relevant records, while precision is the ratio of the number of relevant records retrieved to the size of the return vector [Sal65]. Precision and recall always range from 0 to 1. Often, precision and recall are plot as a curve in the reference frame recall vs. precision [BYRN99]: the larger the area below such a curve, the better the performance under examination. As an additional index, we consider the mean Average Precision (mAP), which is the portion of area under a precision-recall curve. Finally, we consider the *e-measure* e [Rij79], which is a quality measure of the first models retrieved for every query. The *e-measure* depends on the *Precision* and *Recall* values by the relation: $e = \frac{2}{Precision^{-1} + Recall^{-1}}$.

Confusion matrices and Tier images Each classification performance can be associated with a confusion matrix CM , that is, a square matrix whose dimension is equal to the number of classes in the dataset. For the row i in CM , the element $CM(i, i)$ gives the number of items which have been correctly classified as elements of the class i ; similarly, elements $CM(i, j)$, with $j \neq i$, count the items which have been misclassified, resulting as elements of the class j rather than elements of the class i . Similarly, the tier image TI visualizes the matches of the NN, FT and ST. The value of the element $TI(i, j)$ is: *black* if j is the NN of i , *red* if j is among the $(|C| - 1)$ top matches (FT) and *blue* if j is among the $2(|C| - 1)$ top matches (ST). For an ideal classification matrix, CM becomes the diagonal matrix while the TI clusters the black/red square pixels on the diagonal.

4.3. Results

In this Section we discuss the retrieval and classification performance of the edgeLBP. For simplicity, we report only the results obtained with the Bhattacharyya distance because in our experiments it performs better than the other distances considered.

We performed multiple runs with different settings, changing the number (N_r) of rings and the number of samples (P) on them, together with different R associated to the N_r -th ring (called R_{max}). The value of R is based on the size of the patterns in it: we randomly picked 3 models of that dataset and choose one or more R_{max} values that were properly scaled for the dataset. The parameters N_r and P are initially set with what we consider the default settings: $P = 15$, $N_r = 5$. Similarly we consider $h = L - channel$ of the CIELAB color space as the default setting of the function h . Different choices of h , P and N_r are discussed for the Artifacts dataset.

CPP dataset We tested the edgeLBP on this dataset using the default settings and adapting the R_{max} to the size of the wanted pattern ($R_{max} = 0.04mm$), in what in this paper is called *Run1*. As baseline methods to compare against the edgeLBP descriptor we consider two variations of the color histograms. *Hist1* outputs descriptors based on a 16-bin histogram normalized on his minimal and maximal L values. *Hist2* is similar, but no normalization is applied to the values of L . In addition, we also consider the meshLBP descriptor as implemented in the Matlab toolbox [mes].

Figure 7(Top) reports the numerical evaluation measures. Figure 7(Middle) compares the recall vs precision curves of all the methods. Figure 7(Bottom) reports the confusion matrix and the tier image of edgeLBP and the meshLBP runs. The classification and retrieval results obtained over this dataset are very promising and highlight how the edgeLBP encoding captures the pattern distribution over the surface. The edgeLBP overcome simple histogram-based descriptions that, in practice, measure the percentage of color distribution without any control around vertices and also the meshLBP description that bases the ring definition on mesh elements. The positive edgeLBP performance is confirmed in the recent SHREC'18 track for gray color patterns [MTW*18].

Artifacts dataset

This dataset is challenging because of the quality of the original fragments, as their colorimetric patterns are degraded and damaged. Table 1 reports the NN, FT and ST evaluations for different parameter settings of the edgeLBP. Confusion matrices for the two best radius values are reported in Figure 8, along with the relative Tier Images. The number of models in this dataset is too small to consider meaningful the other evaluation measures.

The edgeLBP achieves good retrieval and classification results for most classes. We observed, as expected, that the correctness of the classification is mainly driven by the size of R , rather than P and N_r . As a final note, we tested our algorithm using gray scale values as h function: the results obtained with it were pretty much the same as those obtained with $h = L$. We think that this is due to which information both L of CielAB color space and the gray scale encodes.

	NN	FT	ST	e	mAP	nDCG
edgeLBP	0.985	0.801	0.97	0.66	0.859	0.94
meshLBP	0.94	0.615	0.805	0.54	0.691	0.87
Hist1	0.3	0.301	0.415	0.27	0.354	0.58
Hist2	0.61	0.522	0.774	0.51	0.57	0.76

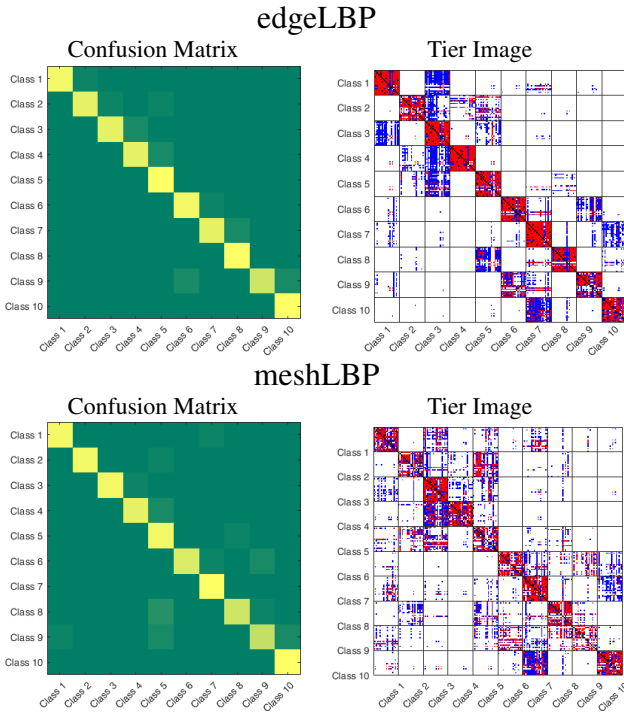



Figure 7: Performance evaluation on the CCP dataset. Top: the NN, FT, ST, e-measure, mAP and nDCG evaluation measures. Middle: the Precision-Recall curves. Bottom: the confusion matrix and tier image of the edgeLBP and the meshLBP runs.

Table 1: The NN, FT and ST scores for some runs of the edgeLBP on the Artifacts dataset. The * in the fourth row means that in these settings we adopt h^3 instead of h (here, h corresponds to the L-channel). R is expressed in mm.

Parameter Settings	NN	FT	ST
$P : 15, N_r : 5, R_{max} : 0,2$	0.775	0.789	1
$P : 15, N_r : 5, R_{max} : 0,3$	0.75	0.811	0.989
$P : 15, N_r : 5, R_{max} : 0,5$	0.75	0.711	0.889
$P : 15, N_r : 5, R_{max} : 0,7^*$	0.725	0.667	0.756
$P : 12, N_r : 7, R_{max} : 0,5$	0.75	0.789	0.9
$P : 12, N_r : 7, R_{max} : 0,2$	0.775	0.856	0.978
$P : 18, N_r : 5, R_{max} : 0,7$	0.7	0.667	0.744

Table 2: Evaluation measures of the performances on the CPP dataset resampled with 40K vertices.

	NN	FT	ST	e	mAP	nDCG
edgeLBP	0.95	0.688	0.857	0.59	0.761	0.9
meshLBP	0.77	0.517	0.703	0.47	0.58	0.79

4.4. Robustness over different surface tessellations

The strength of the edgeLBP is its ring definition, which is robust to different surface tessellations: in this Section we experimentally discuss this robustness. To this aim we re-sample the triangles meshes with a decreasing number of vertices. The triangle mesh re-sampling with x vertices is done with the MeshLAB tool [CCC*08] that approximates the original mesh preserving its geometry as much as possible with the given number of vertices (for instance, $x = 40K$ vertices). This process generally modifies the mesh connectivity and the area of the triangles, discards the smallest details and keeps the overall shape, unless the number of vertices drastically diminishes and the new vertices are too few to preserve it.

First, we re-sampled the meshes in the CCP dataset with 40K vertices. On this dataset, we compare the outcome of the edgeLBP with the default settings with the meshLBP, see Table 2. If compared with the performances on the original CCP dataset in Figure 7, the edgeLBP degrades less than the meshLBP, demonstrating of being more robust to mesh degradation and re-sampling.

Second, we selected 3 patches from the Artifacts dataset and sub-sampled them with 32K, 24K, 16K and 8K vertices (see Figure 9). These four meshes are compared against the original patch (that has 40K vertices). These four distance values provide an estimate of the error the descriptors do when working with the simplified meshes.

We performed two runs for both the edgeLBP and meshLBP:

- *Run1:* $P = 12, N_r = 7$. These settings are the setting used by the meshLBP as default. Both meshLBP and edgeLBP are run with these settings. For the edgeLBP we set $R_{max} = 0.5mm$.
- *Run2:* $P = 15, N_r = 5$. These settings are those that we consider default for the edgeLBP. Both the algorithms are run with these settings. As in *run1*, we set $R_{max} = 0.5mm$.

Figure 10 represents the distance between the original model and its




Figure 8: Confusion matrices and tier images for two of the best runs of the edgeLBP on the Artifacts dataset. In the tier images, the black dots represent the NN element, the red dots correspond to points in FT while blue ones are the ST.



Figure 9: The degradation of one of the model used to test the robustness of the descriptor of both edgeLBP and meshLBP. The number on each image is the respective vertex resolution.

four approximations with respect to both edgeLBP and meshLBP, for all the three original meshes. Since the scale of the distances adopted by the meshLBP and edgeLBP is different, we normalize them with respect the range of the distance values among these patches. From Figure 10, we can see that in both runs the edgeLBP produces more stable descriptors, as the errors are lower than those of the meshLBP (except in one case, the model 1 in *Run2*). In our opinion the nature of the ring definition of the two methods is crucial being both methods based on the LBP concept. Indeed, the meshLBP creates rings of different size when the vertex density decreases becoming quite sparse when the number of vertices of the mesh is significantly reduced. This is not the case of the edgeLBP, as the radius of each ring is always the same (R), for each mesh.

5. Discussions and conclusive remarks

We defined an extension of the LBP on surfaces, whose strength is the robustness to the surface tessellation. In this paper we used this technique to successfully retrieve and classify colorimetric patterns on mesh surfaces. The edgeLBP also performed the best to the SHREC'18 track on retrieval of colorimetric patterns [MTW*18]. Besides synthetic datasets, we tested our algorithm on samples coming from a challenging dataset made of corrupted and degraded artifacts of the EU GRAVITATE project test beds [GRA], achieving promising results. Further extensions are planned and possible.




Figure 10: The plots represent the distance of the four simplified meshes from the original ones, with respect to the meshLBP and the edgeLBP descriptors. The labels in the horizontal axis highlight to the number of vertices of the mesh.

For instance, it is possible to adopt this approach for the description of geometric patterns, encoding the geometric variations with scalar properties of the mesh, like mean curvature or shape index. Moreover, we think that for full color patterns better results could be achieved using all the colorimetric information, for instance the L, a, and b channels of the CIELab space. In this direction, we are currently working on the extension of the edgeLBP to multidimensional properties.

Finally, we think it is worth investigating the automatic recognition and localization of multiple patterns on surfaces. Current experiments are performed on surfaces fully characterized by a single pattern at a time and the similarity distance is defined on the global fragment skin. Next plans include the combination of the shape description step with segmentation techniques and the aggregation of parts made of vertices with similar local descriptions.

Acknowledgments

The work is developed within the research program of the “H2020” European project “GRAVITATE”, contract n. 665155, (2015-2018).

References

- [AKK00] ALBUZ E., KOCALAR E. D., KHOKHAR A. A.: Quantized cielab* space and encoded spatial structure for scalable indexing of large color image archives. In *ICASSP '00* (2000), vol. 6, pp. 1995–1998. 2
- [BCGS13] BIASOTTI S., CERRI A., GIORGI D., SPAGNUOLO M.: Phog: Photometric and geometric functions for textured shape retrieval. *Computer Graphics Forum* 32, 5 (2013), 1383–1392. 1
- [Be14] BIASOTTI S., ET AL.: Retrieval and Classification on Textured 3D Models. In *3DOR'14* (2014), Bustos B., Tabia H., Vandeborje J.-P., Veltkamp R., (Eds.), The Eurographics Association. 2
- [Be16] BIASOTTI S., ET AL.: Retrieval and classification methods for textured 3D models: a comparative study. *The Visual Computer* 32, 2 (2016), 217–241. 2
- [BLP*13] BOMMES D., LÉVY B., PIETRONI N., PUPPO E., SILVA C., TARINI M., ZORIN D.: Quad-Mesh generation and processing: A survey. *Computer Graphics Forum* 32, 6 (2013), 51–76. 2
- [BYRN99] BAEZA-YATES R. A., RIBEIRO-NETO B.: *Modern Information Retrieval*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1999. 5
- [CCC*08] CIGNONI P., CALLIERI M., CORSINI M., DELLEPIANE M., GANOVELLI F., RANZUGLIA G.: MeshLab: an Open-Source Mesh Processing Tool. In *EGIT'08* (2008), Scarano V., Chiara R. D., Erra U., (Eds.), The Eurographics Association. 6
- [Ce13] CERRI A., ET AL.: SHREC'13 Track: Retrieval on Textured 3D Models. In *3DOR'13* (Girona, Spain, 2013), EG 3DOR'13, Eurographics Association, pp. 73–80. 2
- [DT05] DALAL N., TRIGGS B.: Histograms of oriented gradients for human detection. In *CVPR* (2005), vol. 1, pp. 886–893. 1
- [GBP07] GIORGI D., BIASOTTI S., PARABOSCHI L.: *Watertight Models Track*. Research Report 09, IMATI, Genova, 2007. 4
- [Ge15] GIACHETTI A., ET AL.: Retrieval of Non-rigid (textured) Shapes Using Low Quality 3D Models. In *3DOR'15* (2015), Pratikakis I., Spagnuolo M., Theoharis T., Gool L. V., Veltkamp R., (Eds.), The Eurographics Association. 2
- [GG16] GARRO V., GIACHETTI A.: Scale space graph representation and kernel matching for non rigid and textured 3D shape retrieval. *IEEE Trans. Pattern Anal. Mach. Intell.* 38, 6 (2016), 1258–1271. 1
- [GRA] GRAVITATE: Discovering relationships between artefacts using 3D and semantic data. EU H2020 REFLECTIVE project. 4, 7
- [HP11] HUNT R. W. G., POINTER M. R.: *Measuring Colour, Fourth Edition*. Wiley, 2011. 2
- [JK02] JÄRVELIN K., KEKÄLÄINEN J.: Cumulated gain-based evaluation of IR techniques. *ACM Trans. Inf. Syst.* 20 (2002), 422–446. 5
- [LFG*17] LIU L., FIEGUTH P., GUO Y., WANG X., PIETIKÄINEN M.: Local binary features for texture classification: Taxonomy and experimental study. *Pattern Recognition* 62 (2017), 135 – 160. 2
- [Low04] LOWE D. G.: Distinctive image features from scale-invariant keypoints. *Int. J. Comput. Vision* 60, 2 (2004), 91–110. 1
- [mes] meshLBP. <https://it.mathworks.com/matlabcentral/fileexchange/48875-mesh-lbp>. 5
- [MPS*04] MORTARA M., PATANÉ G., SPAGNUOLO M., FALCIDIENO B., ROSSIGNAC J.: Blowing bubbles for multi-scale analysis and decomposition of triangle meshes. *Algorithmica* 38, 1 (2004), 227–248. 3
- [MTW*18] MOSCOSO THOMPSON E., TORTORICI C., WERGHİ N., BERRETTI S., FORERO S. V., BIASOTTI S.: SHREC'18 track: Retrieval of gray patterns depicted on 3D models. In *3DOR'18* (2018), Telea A., Theoharis T., (Eds.), The Eurographics Association. 5, 7
- [OPH96] OJALA T., PIETIKÄINEN M., HARWOOD D.: A comparative study of texture measures with classification based on featured distributions. *Pattern Recognition* 29, 1 (1996), 51–59. 1, 2
- [OPM02] OJALA T., PIETIKÄINEN M., MAENPAA T.: Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *IEEE T. Pattern Anal. Mach. Intell.* 24, 7 (2002), 971–987. 1
- [Pe16] PHILLIPS S. C., ET AL.: GRAVITATE: Geometric and Semantic Matching for Cultural Heritage Artefacts. In *GCH'16* (2016), Catalano C. E., Luca L. D., (Eds.), The Eurographics Association. 1
- [PPTP10] PAPADAKIS P., PRATIKAKIS I., THEOHARIS T., PERANTONIS S.: Panorama: A 3D shape descriptor based on panoramic views for unsupervised 3D object retrieval. *Int. J. Comput. Vis.* 89, 2 (2010), 177–192. 1
- [PZC13] PASQUALOTTO G., ZANUTTIGH P., CORTELAZZO G. M.: Combining color and shape descriptors for 3D model retrieval. *Signal Process-Image* 28, 6 (2013), 608 – 623. 1
- [Rij79] RIJSBERGEN C. J. V.: *Information Retrieval*, 2nd ed. Butterworth-Heinemann, Newton, MA, USA, 1979. 5
- [RTG00] RUBNER Y., TOMASI C., GUIBAS L. J.: The Earth Mover's distance as a metric for image retrieval. *Int. J. Comput. Vis.* 40, 2 (Nov. 2000), 99–121. 4
- [Sal65] SALTON G.: The evaluation of automatic retrieval procedures—selected test results using the smart system. *American Documentation* 16, 3 (1965), 209–222. 5
- [SH07] STARCK J., HILTON A.: Correspondence labelling for wide-timeframe free-form surface matching. In *ICCV* (2007), pp. 1–8. 4
- [Suz01] SUZUKI M.: A Web-based retrieval system for 3D polygonal models. In *NAFIPS* (2001), vol. 4, pp. 2271–2276. 1
- [TSDS11] TOMBARI F., SALTİ S., DI STEFANO L.: A combined texture-shape descriptor for enhanced 3D feature matching. In *ICIP* (2011), pp. 809–812. 1
- [WAvK*12] WANG Y., ASAFI S., VAN KAICK O., ZHANG H., COHEN-OR D., CHEN B.: Active co-analysis of a set of shapes. *ACM Trans. Graph.* 31, 6 (2012), 165:1–165:10. 4
- [WBB15] WERGHİ N., BERRETTI S., BIMBO A. D.: The mesh-LBP: A framework for extracting local binary patterns from discrete manifolds. *IEEE Trans. Image Processing* 24, 1 (2015), 220–235. 2
- [WCL*08] WU C., CLIPP B., LI X., FRAHM J.-M., POLLEFEYS M.: 3D model matching with Viewpoint-Invariant Patches (VIP). In *CVPR* (2008), pp. 1–8. 1
- [WTBB15] WERGHİ N., TORTORICI C., BERRETTI S., BIMBO A. D.: Local binary patterns on triangular meshes: Concept and applications. *Comput Vis Image Underst* 139 (2015), 161–177. 1, 2
- [WTBB16] WERGHİ N., TORTORICI C., BERRETTI S., BIMBO A. D.: Boosting 3D LBP-based face recognition by fusing shape and texture descriptors on the mesh. *IEEE Trans. Inf. Forensics Security* 11, 5 (2016), 964–979. 1, 2
- [ZBH12] ZAHARESCU A., BOYER E., HORAUD R.: Keypoints and local descriptors of scalar functions on 2D manifolds. *Int. J. Comput. Vision* 100, 1 (2012), 78–98. 1