



HAL
open science

SHREC'18 track: Retrieval of gray patterns depicted on 3D models

E. Moscoso Thompson, C Tortorici, N Werghi, S Berretti, Santiago Velasco-Forero, S Biasotti

► To cite this version:

E. Moscoso Thompson, C Tortorici, N Werghi, S Berretti, Santiago Velasco-Forero, et al.. SHREC'18 track: Retrieval of gray patterns depicted on 3D models. Eurographics Workshop on 3D Object Retrieval, 2018. <hal-01774421>

HAL Id: hal-01774421

<https://minesparis-psl.hal.science/hal-01774421v1>

Submitted on 30 Apr 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

SHREC'18 track: Retrieval of gray patterns depicted on 3D models

E. Moscoso Thompson^{1*}, C. Tortorici², N. Werghi², S. Berretti³, S. Velasco-Forero⁴, S. Biasotti^{1*}

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

² Electrical and Computer Engineering Department, Khalifa University

³ Department of Information Engineering, University of Firenze

⁴ PSL Research University - MINES ParisTech, CMM, Center for Mathematical Morphology

* Track organizers

Abstract

This paper presents the results of the SHREC'18 track: Retrieval of gray patterns depicted on 3D models. The task proposed in the contest challenges the possibility of retrieving surfaces with the same texture pattern of a given query model. This task, which can be seen as a simplified version of many real world applications, requires a characterization of the surfaces based on local features, rather than considering the surface size and/or bending. All runs submitted to this track are based on feature vectors. The retrieval performances of the runs submitted for evaluation reveal that texture pattern retrieval is a challenging issue. Indeed, a good balance between the size of the pattern and the dimension of the region around a vertex used to locally analyze the color evolution is crucial for pattern description.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation—Line and curve generation

1. Introduction

The challenge of this SHREC'18 track is to evaluate the performance of existing retrieval algorithms when a 3D surface is decorated with one or more gray patterns. The web page of the track is available at http://shrec.ge.imati.cnr.it/shrec18_color/. By a color pattern we mean a repeated decoration painted over the surface of a 3D model, e.g. circles, squares, stripes. This task represents a first attempt of addressing full-color patterns that are of large interest for practical applications, for example to retrieve artworks' patterns and artists' style in architectural and cultural heritage data. However, to create surface with simple and well-defined patterns, in this track we limit the colorimetric information to gray color variation. This definition includes black and white textures.

Patterns locally characterize the surface independently of the overall surface shape; for instance, the same floral decoration could characterize the border of a vase or the middle of a cup. Therefore, we consider a dataset classification which is independent of the overall surface. This aspect differentiates this track from previous SHREC contests for 3D texture retrieval and comparison [CBA*13, BCA*14, GFF*15] that focused on the joint comparison of geometry and texture, without specifically comparing the surface textures. The goal of this task is to find a robust and stable descriptor and a distance measure which can group the models on the basis of their patterns.

2. The dataset

The dataset consists of 300 surfaces characterized by different patterns. We created a set of 20 base models, selecting triangles meshes from the *cups and vases* classes of the SHREC'07 Watertight model contest [GBP07] and the *goblets* class of the COSEG [WAvK*12] datasets, see Figure 1. All models are oriented, connected and locally regular, triangle meshes with at most, two boundary components. The rationale behind the choice of these base models is the fact that these shapes do not possess a geometrically privileged point of view, many of them (e.g., cups) have an interior and exterior face and trivially projecting them in a plane is not possible without distorting the surface.

We also created 15 black and white textures: 10 textures (called *Single textures*) are characterized by a single pattern while 5 more textures (called *Double textures*) are obtained joining the 10 textures two by two, see Figure 2. The 15 textures have been applied to all the 20 base meshes using a semi-automatic procedure, which required a manual fix of the geometry of some models. Each pattern is applied so that it keeps the same scale over all the models. Then, the colorimetric information is stored in the vertices of the triangle mesh M . At the end of this process, every model is covered by a black and white pattern on at least the 30% of its surface, while the rest of the surface is only black or only white (see Figure 3d as an example). Assuming that the maximum value of the luminosity is 100, we uniformly alter the luminosity of each model with a fixed amount that is a randomly selected from the




Figure 1: The models used as base surfaces.

set $\{0, 8, 16, 24, 32, 40, 48, 56, 64, 72\}$. This results in a more or less faded version of the original model, see examples in Figure 3(b-c).

The number of vertices of the 300 models in this dataset ranges from 95K to 107K. Similarly to the outcome of many laser scan systems, the colorimetric information related to each model is stored as a RGB value associated to each vertex (no texture images are provided).

3. Ground truth

The challenge of this contest is to group the objects of the dataset according to the pattern impressed on them. The overall geometric shape of the object is not relevant; in practice the methods are asked to classify an object using the local, colorimetric properties. The runs submitted to the track have been evaluated according to the following classifications:

- The *Single Pattern Dataset* contains the 200 models of the dataset that are characterized by a single pattern (*Single textures*). Models in this dataset are grouped in 10 classes of 20 elements.
- The *Complete Dataset* contains also the 100 models characterized by *Double textures* that mix the 10 single textures. Models in this dataset are grouped in 15 classes, each class contains 20 elements.

For each setting, given a query model, its relevant models are the models with the same pattern(s) of the query. We highlight that, when considering the complete dataset, the models characterized by a double texture are not in the same class of the models that have its single texture components.

4. Participants

Six groups subscribed to the contest but only three of them sent their outcome. Two participants sent 3 runs and the other sent 1 run. Here we sum up the algorithms used in by each participant. Two baseline runs based on color histograms are also considered.

4.1. From Mesh to Image using Ordered Ring of Facets (ORF) for gray Texture Retrieval by Claudio Tortorici, Naoufel Werghi and Stefano Berretti (TWB*)

The algorithm extracts images holding the color texture information of a region of the surface. The images are extracted using the Ordered Ring Facet (ORF) [WTBdB15, TWB15], which allows to locally order the neighbors of a given central facet f_c . ORF generates concentric rings with different radius and azimuthal quantization; these information are used as polar coordinates. Using a polar to Cartesian transformation the 2D image is extracted (see Figure 4). In order to guarantee the images to cover all the color textures on the surface, the central facets, used for the image extraction, need to be "spread" on the overall surface. Therefore, starting from a random point of the surface, N points are selected, equally distant between each other in a geodesic sense. The algorithm, in fact, starting from a first random point, computes the geodesic distance against the rest of the surface points, then takes the most far and iterate the process. The point n needs to be the farthest one from the $(n - 1)$ previously selected points. Once the N figures are extracted, the LBP [OPH96] descriptor is computed over them. The comparison between two models is done as follows. A $N \times N$ matrix is generated, where N is the number of images per model and its entry (x, y) shows the distance between figure x of the first model, and figure y of the second, using Bhattacharyya distance (D_{bha}). In other words, for each image of the probe, the closest on the query model is taken and then all the N distances are summed.

$$R(i, j) = \sum_{x=1}^N \left(\min_{y=1}^N (D_{bha}(I_{i,x}, I_{j,y})) \right),$$

where $I_{i,x}$ and $I_{j,y}$ represents, respectively, the descriptor of the x^{th} and y^{th} image of the query i and model j .

Three runs are presented:

- *Run 1 - TWB1*: generating 5 images for each model;
- *Run 2 - TWB2*: generating 5 smaller images for each model;
- *Run 3 - TWB3*: generating 7 images for each model.

4.2. Histogram of Double Distance Transform by Santiago Velasco-Forero (V^*)

This method considers the texture on the mesh as binaries ones, i.e., only containing white (RGB color information in the vertices equal to $[255, 255, 255]$) and black (defined as not white color) vertices. A function that maps each interior vertex (exterior) on the mesh with the distance to the nearest vertex on the exterior (interior) set of vertices is defined. This function is inspired on the Distance maps [KKB96] (also known as Distance Transform for image processing community) for binary images.

In the following, a color shape is defined as $S = (V, S, F)$, where S is a closed triangles mesh in \mathbb{R}^3 and $F: S \rightarrow \mathbb{R}^3$ represents the mapping from each vertex in the mesh to a given color space. The descriptor is computed by these four steps:

1. For a given mesh, vertices are divided in two sets:

$$V_w = \{v \in V, F(v) = [255, 255, 255]\}, \quad V_b = V \setminus V_w,$$

which are respectively the set of white vertices and that of black vertices.




Figure 2: Textures (patterns) depicted on the models. Class 11 to 15 are Double Textures.

2. Then, the k -nearest neighbor graph (k -NNG) of \mathbb{V}_w on \mathbb{V}_b and the k -NNG \mathbb{V}_b on \mathbb{V}_w are computed.
3. The following quantities are computed:

$$NormD_{k,\mathbb{V}_w}(\mathbf{v}) = Dist_1(\mathbf{v}, \mathbb{V}_w) / Dist_k(\mathbf{v}, \mathbb{V}_w),$$

$$NormD_{k,\mathbb{V}_b}(\mathbf{v}) = Dist_1(\mathbf{v}, \mathbb{V}_b) / Dist_k(\mathbf{v}, \mathbb{V}_b);$$

respectively $\forall \mathbf{v} \in \mathbb{V}_w$ and $\forall \mathbf{v} \in \mathbb{V}_b$

4. The proposed feature with parameter (k,b) is the concatenation of two histograms computed on the normalized distance, i.e.:

$$Feat(\mathcal{S}, k, b) = \left[H\left(\frac{NormD_{k,\mathbb{V}_w}(x)}{|\mathbb{V}_b|, b}\right), H\left(\frac{NormD_{k,\mathbb{V}_b}(\mathbf{v})}{|\mathbb{V}_w|, b}\right) \right],$$

where $H(\cdot, b)$ defines the histogram with equal-width b bins. Thus, the proposed descriptor produces a vector of dimension $2b$, with values between $(0, 1]$.

The three runs submitted to this contest (respectively $V1$, $V2$ and $V3$) differ in how the distances between models are defined:

- $D_1(\mathcal{S}_i, \mathcal{S}_j) = \|Feat(\mathcal{S}_i, 5, b = 10) - Feat(\mathcal{S}_j, 5, b = 10)\|$,
- $D_2(\mathcal{S}_i, \mathcal{S}_j) = \|Feat(\mathcal{S}_i, 9, b = 10) - Feat(\mathcal{S}_j, 9, b = 10)\|$,
- $D_3(\mathcal{S}_i, \mathcal{S}_j) = \|Feat(\mathcal{S}_i, 13, b = 10) - Feat(\mathcal{S}_j, 10, b = 10)\|$,

where $\|\cdot\|$ is the Euclidean norm.

4.3. Edge Local Binary Pattern by Elia Moscoso Thompson and Silvia Biasotti (MT)

This is an extension of the Local Binary Pattern [OPH96] to 3D surfaces, called *edgeLBP*, [MTB18]. For each vertex v of a triangulation mesh T , a ring of radius R is defined as the intersection of the mesh edges with a sphere of radius R centered in v . For each ring, the piecewise linear curve C that represents the intersection of the mesh with the sphere is detected. Due to the unpredictable structure of the triangulation (the vertex distribution, the mesh connectivity, etc., are not a-priori known and uniform along T), each curve C is oversampled or sub-sampled with P points, s_1, \dots, s_P , see Figure 5. Each ring is then represented by the samples s_1, \dots, s_P , which are points in \mathbb{R}^3 . Concentric rings (N_r) for each vertex are used in order to achieve a multi-ring description. The radius R of the largest sphere, the number of rings N_r and the number of samples P are the parameters of the method. Additional details on the definition and computation of the edgeLBP descriptor can be found in [MTB18]. Finally, the local property used to code the pattern evolution (for instance, the gray-scale value for images) is a scalar function $h : V \rightarrow \mathbb{R}$, where V is set of vertices of T . In this track, $h(v)$ is the value of L -channel associated to v . For each ring, the edgeLBP is computed over the samples s_1, \dots, s_P . For each model, the final descriptor is the normalized histogram of the LBP values




Figure 3: Models in the dataset of this contest, with different fades. The models in (a-b) are fully covered by a single pattern, in (c) the model presents a faded double pattern while the surface of the model in (d) is partially covered with the pattern decoration and the rest is uniformly black.




Figure 4: Examples of 5 figures extracted from 4 different surfaces.

of every ring for all the vertices of model. The distances among the descriptors is the Bhattacharyya distance [DD09].

The run submitted to the contest adopts the following settings: $R = 0.04$, $P = 15$ and $N_r = 5$. In particular, R is defined looking at 5 random models of the dataset and defined as the bigger ring that covers a meaningful portion of the pattern around each vertex.




Figure 5: (a): The black lines represent the piecewise, linear curves that correspond to the intersection between the model M and a set of spheres centered in a vertex of M . (b): The black dots highlight the intersection points between the spheres and the edges of the triangulation of M . (c): The blue-red dots represent the uniformly-spaced samples s_1, \dots, s_P over the curves.

The unit of measure of R is the same of the triangulation. $P = 15$ corresponds to the ideal subdivision of a planar circle into sectors of 24° while $N_r = 5$ is quite a standard choice for the LBP.

4.4. Baseline Method: Color distribution on a model in CieLAB space color(HBI & HnBI)

The colorimetric information of each model M is stored in a histogram h_M of 16 bins. The RGB values of a model are converted into CieLAB coordinates (a - channel, b - channel, L - channel) and the L - channel values are stored in a histogram that codes the probability of a vertex of being in one of the 16 intervals. That is, the n - th bin value l is the probability that a random vertex of the model has l as its L - channel value.

This algorithm is run with with two different settings:

- *HBI*: only the values of the L - channel between 10 and 100 are stored in the histogram,
- *HnBI*: all the L - channel values are stored in the histogram.

Thanks to their simplicity in computing and reproducing them, these runs are considered as baseline benchmarks for this contest.

5. Evaluation measures and results

The retrieval performance of each run has been evaluated in both the Single Pattern and Complete datasets. The evaluation process has been based on the following evaluation measures: Average precision-recall curves, Nearest Neighbor (NN), First Tier (FT), Second Tier (ST), e -measure, mean Average Precision (mAP) and normalized Discounted Cumulated Gain (nDCG).

Average precision-recall curves. Precision is the fraction of retrieved items that are relevant to the query. Recall is the fraction of the items relevant to the query that are successfully retrieved. Plotting the two quantities in the reference frame recall vs. precision, we get a curve: the larger the area below such a curve, the better the performance under examination. In particular, the precision-recall curve of an ideal retrieval system would result in a constant curve equal to 1. For each query, we have a precision-recall curve. By




Figure 6: The Precision-Recall plots for all the runs reported in this contest.

taking the average on all the queries, we get the average precision-recall curve. The *mean Average Precision (mAP)* corresponds to the area between the horizontal axis and the average precision-recall curve, the maximum mAP value is equal to 1. Precision-recall plots for all the runs on both datasets are reported in Figure 6.

Nearest Neighbor (NN), *First tier (FT)*, *Second tier (ST)* and *e-Measure (e)*. These measures focus on the fraction of models of the query's class also appearing within the top k retrievals [SMKF04]. The final score, always ranging in $[0, 1]$, is an average over all the models in the database.

Normalized Discounted Cumulated Gain (nDCG). It is based on the assumption that relevant items are more useful if appearing earlier

Single Pattern Dataset						
Run	NN	FT	ST	e	nDCG	mAP
HBI	0.52	0.141	0.23	0.097	0.501	0.233
HnBI	0.595	0.52	0.773	0.425	0.698	0.564
TWB1	0.71	0.454	0.632	0.338	0.679	0.516
TWB2	0.615	0.46	0.675	0.362	0.681	0.523
TWB3	0.755	0.502	0.688	0.376	0.711	0.577
V1	0.82	0.45	0.725	0.386	0.743	0.571
V2	0.82	0.51	0.731	0.39	0.753	0.593
V3	0.815	0.474	0.71	0.376	0.733	0.557
MT	0.905	0.699	0.863	0.551	0.865	0.749

Table 1: Summary evaluation over the single pattern dataset.

Complete Dataset						
Run	NN	FT	ST	e	nDCG	mAP
HBI	0.467	0.115	0.16	0.139	0.548	0.184
HnBI	0.557	0.437	0.646	0.509	0.75	0.481
TWB1	0.593	0.367	0.514	0.419	0.758	0.408
TWB2	0.45	0.379	0.553	0.439	0.75	0.424
TWB3	0.593	0.417	0.564	0.455	0.795	0.46
V1	0.77	0.412	0.584	0.474	0.798	0.472
V2	0.79	0.433	0.594	0.481	0.808	0.493
V3	0.78	0.394	0.569	0.463	0.786	0.462
MT	0.87	0.646	0.805	0.594	0.888	0.697

Table 2: Summary evaluation over the complete dataset.

in a search engine result list. The nDCG is based on the graded relevance of a result with respect to the query and then normalized with respect to the ideal outcome of that query [GBP07].

Results on the Single Pattern dataset are reported in Table 1, while those on the Complete Dataset are in Table 2.

Confusion matrix. To each run we associate also a confusion matrix CM , that is, a square matrix whose order is equal to the number of classes in the dataset. For a 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 . Thus, the classification matrix CM of an ideal classification system should be a diagonal matrix, such that the element $CM(i, i)$ equals the number of items belonging to the class i . Confusion matrices are reported in Figure 7 for the Single Pattern Dataset, while those of the Complete dataset are in Figure 8.

Tier images visualize the matches of the NN, FT and ST. The value of the element $T(i, j)$ 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), where $|C|$ is the number of elements of the class C . When the models of a class are grouped along each axis, the optimal tier image clusters the black/red square pixels on the diagonal. The best tier image per participant is reported in Figure




Figure 7: Confusion matrices over the single pattern dataset.




Figure 8: Confusion matrices over the complete dataset.

9 and Figure 10, for the Single pattern dataset and the Complete dataset respectively.

6. Discussions and conclusions

The recall-precision curves in Figure 6 and the statistics in Table 1 and 2 indicate that the pattern retrieval task is challenging on this dataset. Even if the proposed dataset is not a real world dataset (with a huge variety of shapes and color patterns) it presents two important challenges that the pattern retrieval on 3D surfaces problem must face: (i) the impossibility of encoding the colorimetric information of a mesh in a single image and (ii) the presence of the same pattern over a number of different shape embeddings.

The dataset presents several elements of complexity: (i) single and double patterns; (ii) degradation of the color pattern from black to light gray; (iii) different coverage of the pattern over the surface (from 30% to 100%). From the precision recall curves, we see that 7 of the 9 runs behave qualitatively equally, aside HBl (which performances are really low) and edgeLBP (which overcome the other methods). Moreover, the degradation of the precision recall performance in terms of curve trend is approximately equivalent for all runs, with the exception of the HnBl run that is slightly more constant.

Interestingly, the performance of the methods is quite stable over the single pattern and the complete datasets, with a quite limited degradation on the complete one (the largest decrease for the NN value is approximately 15% for the TWB* runs). Besides the NN evaluation, in which there is an average variance of values, the FT,

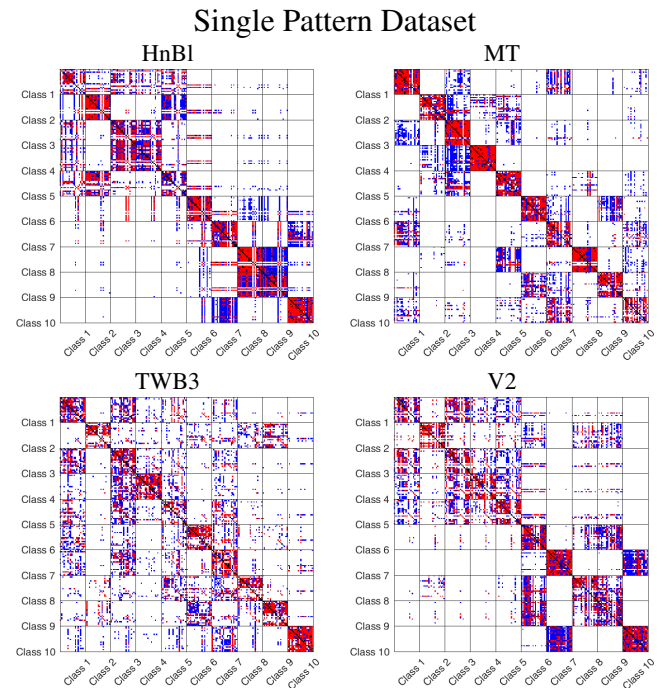


Figure 9: The best tier image per participant on the single pattern dataset. NN are marked in black, FT in red and ST in blue.

Complete Dataset




Figure 10: The best tier image per participant on the complete dataset. NN are marked in black, FT in red and ST in blue.

ST, e and nDCG measures seem to be similar for most methods, polarizing in a range of values for each measure ([0.45,0.50] for FT, [0.65,0.75] for ST and so on). The outliers are HBI, whose performances are really low, as expected for an histogram-based method without any local analysis of the gray distribution, and edgeLBP, whose performances overcome the other runs of approximately 10%.

More details of the performance of the methods over the different classes are shown in the confusion matrices and tier images, see Figures 7, 8, 9 and 10. Confusion matrices help us understand which are the hardest patterns to correctly deal with. For the baseline methods, the confusion matrices show that there are multiple pattern classes that are tricky for them. This is especially true for the HBI run on the Complete dataset. For the methods who submitted more runs, the confusion matrices show a quite stable behavior on the classes of the benchmark, despite the runs different settings. On all runs, the tier images confirm the trends revealed by the confusion matrices.

In conclusion, we observe that all runs submitted to this track are based on feature vectors. In the case of the baseline methods (simple histograms of the L-channel values) the description does not encode any information of gray distribution in the vertex neighbor. All other methods propose strategies for coding the colorimetric information evolution in a region around the vertices. From the retrieval results it seems that methods that works directly on the mesh are independent of the tessellation (MT and V*) achieve better performances.

Acknowledgments

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

References

- [BCA*14] BIASOTTI S., CERRI A., ABDELRAHMAN M., AONO M., HAMZA A. B., EL-MELEGY M., FARAG A., GARRO V., GIACHETTI A., GIORGI D., GODIL A., LI C., LIU Y.-J., MARTONO H. Y., SANADA C., TATSUMA A., VELASCO-FORERO S., XU C.-X.: Retrieval and Classification on Textured 3D Models. In *Eurographics Workshop on 3D Object Retrieval* (2014), Bustos B., Tabia H., Vandeborre J.-P., Veltkamp R., (Eds.), The Eurographics Association. 1
- [CBA*13] CERRI A., BIASOTTI S., ABDELRAHMAN M., ANGULO J., BERGER K., CHEVALLIER L., EL-MELEGY M., FARAG A., LEFEBVRE F., GIACHETTI A., GUERMOUD H., LIU Y.-J., VELASCO-FORERO S., VIGOUROUX J., XU C.-X., ZHANG J.-B.: SHREC'13 Track: Retrieval on Textured 3D Models. In *Proceedings of the 6th Eurographics Conference on 3D Object Retrieval* (Girona, Spain, 2013), EG 3DOR'13, Eurographics Association, pp. 73–80. 1
- [DD09] DEZA M. M., DEZA E.: *Encyclopedia of Distances*. Springer Berlin Heidelberg, 2009. 4
- [GBP07] GIORGI D., BIASOTTI S., PARABOSCHI L.: *Watertight Models Track*. Research Report 09, IMATI, Genova, 2007. 1, 5
- [GFF*15] GIACHETTI A., FARINA F., FORNASE F., TATSUMA A., SANADA C., AONO M., BIASOTTI S., CERRI A., CHOI S.: Retrieval of Non-rigid (textured) Shapes Using Low Quality 3D Models. In *Eurographics Workshop on 3D Object Retrieval* (2015), Pratikakis I., Spagnuolo M., Theoharis T., Gool L. V., Veltkamp R., (Eds.), The Eurographics Association. 1
- [KKB96] KIMMEL R., KIRYATI N., BRUCKSTEIN A. M.: Sub-pixel distance maps and weighted distance transforms. *Journal of Mathematical Imaging and Vision* 6, 2 (Jun 1996), 223–233. 2
- [MTB18] MOSCOSO THOMPSON E., BIASOTTI S.: Edge-based LBP description of surfaces with colorimetric patterns. In *Eurographics Workshop on 3D Object Retrieval* (2018), Telea A., Theoharis T., (Eds.), The Eurographics Association. 3
- [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. 2, 3
- [SMKF04] SHILANE P., MIN P., KAZHDAN M., FUNKHOUSER T.: The Princeton shape benchmark. In *Shape Modeling International* (June 2004). 5
- [TWB15] TORTORICI C., WERGI N., BERRETTI S.: Boosting 3D LBP-based face recognition by fusing shape and texture descriptors on the mesh. In *2015 IEEE International Conference on Image Processing (ICIP)* (Sept 2015), pp. 2670–2674. 2
- [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 (Nov. 2012), 165:1–165:10. 1
- [WTBdB15] WERGI N., TORTORICI C., BERRETTI S., DEL BIMBO A.: Local binary patterns on triangular meshes: Concept and applications. *Computer Vision and Image Understanding* 139 (2015), 161 – 177. 2