

SHREC'18 track: Recognition of geometric patterns over 3D models

S. Biasotti^{1*}, E. Moscoso Thompson^{1*}, L. Barthe³, S. Berretti⁴, A. Giachetti⁵, T. Lejemble³, N. Mellado³, K. Moustakas⁶, Iason Manolas⁶, Dimitrios Dimou⁶, C. Tortorici⁷, S. Velasco-Forero⁸, N. Werghi⁷, M. Polig^{2*}, G. Sorrentino^{2*}, S. Hermon^{2*}

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

² The Cyprus Institute - Science and Technology in Archaeology Research Center (STARC)

³ CNRS, IRIT, Université de Toulouse, INPT, UPS, UT1C, UT2J, France

⁴ Dept. of Information Engineering, University of Firenze

⁵ Dept. of Computer Science, University of Verona

⁶ Electrical and Computer Engineering Department, University of Patras

⁷ Electrical and Computer Engineering Department, Khalifa University

⁸ PSL Research University - MINES ParisTech, CMM, Center for Mathematical Morphology

* Track organizers

Abstract

This track of the SHREC 2018 originally aimed at recognizing relief patterns over a set of triangle meshes from laser scan acquisitions of archaeological fragments. This track approaches a lively and very challenging problem that remains open after the end of the track. In this report we discuss the challenges to face to successfully address geometric pattern recognition over surfaces; how the existing techniques can go further in this direction, what is currently missing and what is necessary to be further developed.

Categories and Subject Descriptors (according to ACM CCS): I.3.6 [Computer Graphics]: Methodology and Techniques—

1. Introduction

The aim of this SHREC 2018 track was to evaluate the performance of automatic algorithms for recognizing geometric patterns over surfaces. By a geometric pattern we mean any repeated, relief variation over a surface embedded in the space. Three distinctive aspects characterize this contest with respect to previous SHREC'17 contest for retrieving relief patterns [BMTA*17]: (i) the surfaces may have none, one or more patterns and may also present features like eyes that are not in the list of features to be identified; (ii) the dataset is a real archaeological case study and not on an ideal dataset; (iii) instead of matching patterns, the task is to recognize if (and where) a query pattern is present over a complex surface. The algorithms were expected to recognize and differentiate between patterns and, possibly to be robust to different levels of abrasion in decorations of the material.

The track saw eleven registrations, thus demonstrating there is a strong interest in the problem proposed. Unfortunately, the outcome of the approaches considered to tackle the contest were unsatisfactory and the problem is still open. Besides the description of the problem and the benchmark (i.e., the dataset, the query patches and the ground truth) proposed in the contest, in this report we discuss the challenges to solve before being successfully able to perform pattern recognition over real object scans; we sketch some of the attempts done to address the problem and we outline those we foresee to be the most promising research avenues in this field.

1.1. Motivation

During archaeological excavations, Cultural Heritage artifacts are often found in various stages of incompleteness, with parts either totally missing or components of their overall structure being fragmented in various pieces. A tedious task of conservators is to re-compose such fragmented objects to their original shape. The reasons for such an investment are various: the more complete the shape of an object the better its typological classification, the amount of individual objects within a certain archaeological context are essential in deriving socio-cultural, economic or behavioral aspects of such contexts, a complete shape of an object is more indicative of its function in the past and finally a restored object has a higher aesthetic value than its fragmented parts and thus more appealing for its musealisation. Such a restoration process is normally performed on the physical fragments, and based on a visual assessment of the position of each fragment within the whole. In case of pottery objects, such as amphora, vessels of various kinds, cooking pots, etc., the restoration starts from the base of the object, in order to assess the vertical axis of the object, while other components are either joined together if they share common edges or are positioned in their spatial location based on their curvature (according to their visually assessed orientation), overall morphology or continuation of patterns along their external face (such as decorations, either painted, incised, etc.). A similar process is performed

on other types of fragmented objects, such as statues, jewels, frescoes, weapons, etc.

Since archaeological excavations often yield thousands of artifacts, the restoration work occurs once these objects are already stored in deposits. Moreover, it may occur that fragmented belonging to a same object are stored separately and thus the restorer must rely on drawings or digital acquisitions of the fragments in order to propose possible joining of fragments, often supported by similarities of decorative patterns along their external face or the recognition of common patterns and decoration style. In case fragments are stored in different locations, under separate administrations, a physical restoration will be almost an impossible task. Therefore, it is imperious to advance digital approaches that facilitate the finding and retrieval of fragments that possibly relate to each other and furthermore develop algorithms that support the restoration process of such fragments. Thus, a restoration based on quantitative methods rather than visual observation may optimize the entire restoration process, making it faster and more accurate. Moreover, having 3D digital replicas that can be manipulated within a 3D space enable not only their re-joining, but allows further analyses, such as investigation of morphological characteristics of such patterns, how they were applied, etc..

1.2. Problem statement

The goal of this track is to identify if and where a given geometric pattern (represented as a query patch fully characterized by a single pattern) is located over the surface of a 3D model. Geometric patterns are repeated, local relief variations on the surface; for instance a chiseled decoration over an archaeological artifact. We admit that more than one geometric pattern might be present on the surface of 3D models. Figure 1 illustrates with colors how the recognition of two patterns is expected to run on a set of four fragments. The relevance or non-relevance of the models for the given queries and where the patterns are located were established and annotated a priori with the help of archaeologists.

1.3. Report organization

The remainder of the paper is organized as follows. First, in Section 2 we describe the dataset and the query patches proposed in this contest and the ground truth. Then, we discuss the intrinsic difficulties of recognizing patterns over mesh surfaces (Section 3), the obstacles the participants found in this very contest (Section 4) and, finally, we sketch some of the methods considered in this contest and how we think it would be possible to overcome the current limitations (Section 5). Conclusions and final remarks end the paper.

2. Dataset and ground truth

All the 3D models comes from fragments of archaeological artifacts adopted as use cases in the EU H2020 project GRAVITATE [GRA]. The models were organized into two sets: the Query set (QS) and the Model set (MS).

- *QS*: it contains 8 triangle meshes representing a single pattern. All the patches are surfaces with one boundary, almost

flat from a global point of view. The triangulations representing the patches have no fixed number of vertices. These patches were tailored from fragments of the dataset set. Overall, we have 6 possible patterns. The pattern classes are: *eyebrows*, *oblique fringe*, *fringe*, *long incisions*, *spirals*, *stamped circlets*. In particular, there are two different patches for the *stamped circlets* and *oblique fringe* classes. Figure 2 illustrates the definition of the patches in QS.

- *MS*: it contains 30 triangle meshes, representing 30 different archaeological fragments. Of these, 25 models are characterized by at least one geometric pattern, while the other have no patterns on them. From a geometrical point of view, all the models in MS represent a single, closed surface. All triangle meshes are watertight and do not contain self intersections or degenerate triangles. All triangle meshes are provided at the highest resolution available and there is no a fixed number of vertices, also called vertex resolution (reported in Figure 3). The number of vertices of the meshes in MS ranges from 150356 to 6800671. Besides full resolution models, simplified versions with 50K and 100K vertices of the models in MS were available; these simplified meshes were generated with the tool [MPS17]. All the models in MS are reported in Figure 3, with a detailed informations on which pattern is chiseled on each one of them (if any).

The initial response to the contest shows that this is a very lively problem, as 11 groups positively react to the contest call. Despite this, at the results delivering deadline no one had a complete output to show (only a membership matrix was submitted, with quite poor performances). Therefore, in the following we are discussing the reasons of this lack of methods for 3D pattern recognition and how to direct future research efforts to solve such an open problem.

3. Geometric pattern recognition challenges

Recognizing geometric patterns over surfaces is more complex than simply matching two surfaces. The straightforward extension to 3D models of the techniques adopted for image object detection is not possible because 3D models have peculiar characteristics that require ad-hoc techniques. For instance, the use of meshes instead of grids prevent the adoption of methods that take advantage of the regular, grid structure of the images. In the following, we report a list of issues that in our opinion deal with geometric pattern recognition.

- *Type of representation*: while for images the grid structure is unique, predictable and regular, this is no longer true for boundary representations of 3D objects, like mesh tessellations, point clouds, and so on. Moreover, two representations of the same object are not unique and can be really different from each other (number of vertexes, vertex distribution, etc.). Also, model acquisition can be affected by noise and/or errors (see Session 4). Not having an ideal, exact reference template both for patterns and models increases the difficulty of the task.
- *Pattern definition and size*: the concept of pattern as meant in this contest is not formally defined. Here we distinguish patterns from local features: patterns contain a repeated configuration of some surface property over the surface while by features we mean local changes on the surfaces without a repetition rule. With reference to the model of a statue head like that in



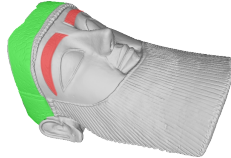



Query Patches	Dataset Models			
	M1	M2	M3	M4
$Q_1(CR)$ 				
$Q_2(EB)$ 	$V_1 = \{23, 32, \dots, 2322\}$	$V_2 = \{3, 33, \dots, 177\}$ $V_3 = \{77, 83, \dots, 1815\}$	-empty lists-	$V_4 = \{1, 186, \dots, 2987\}$
Q_1 $\{M_1 : \emptyset\} \mid \{M_2 : V_2\} \mid \{M_3 : \emptyset\} \mid \{M_4 : V_4\}$		Q_2 $\{M_1 : V_1\} \mid \{M_2 : \emptyset\} \mid \{M_3 : V_3\} \mid \{M_4 : \emptyset\}$		

Figure 1: An example of expected results for a dataset of 4 models and 2 query patches. In each model M , the vertexes expected to be recognized are colored according to the pattern they depict. Each list V_i contains the vertices highlighted with the same color. The lists in bottom boxes are the expected results for the two query patches Q_1 and Q_2 .









Q_1 (v.r. = 88950) 	Q_2 (v.r. = 20250) 	Q_3 (v.r. = 147810) 	Q_4 (v.r. = 204740) 
CR	EB	OF	FR
Q_5 (v.r. = 118000) 	Q_6 (v.r. = 31640) 	Q_7 (v.r. = 124650) 	Q_8 (v.r. = 180800) 
OF	SP	CR	LI

Figure 2: Summary table for QD. v.r stands for vertex resolution.

Figure 3(12), the stamped circles on the helmet delimit a region with a geometric pattern, while the eyes, the mouth, the ears, etc., represent features chiseled on the model. This definition of pattern, despite being intuitive, is not formal. In particular, the size of a pattern is a crucial point to be identified because it cannot be easily generalized to every model and type of pattern and

it is necessary to distinguish what is shape and what is decoration. Methods in the literature for local feature characterization that tackle the problem of considering local regions instead of a point-wise characterization (e.g. [SPS16]) strongly depend on the size of the area considered. Moreover, the same model could present patterns of different size. This means that the pattern size









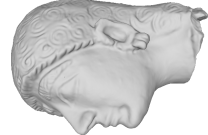


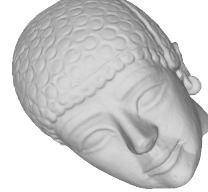
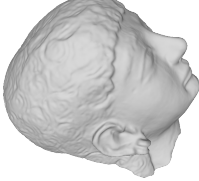





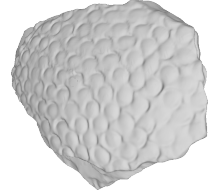
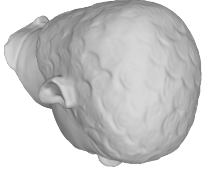

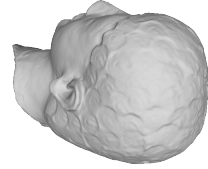
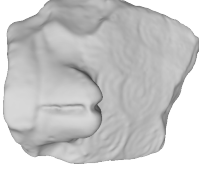


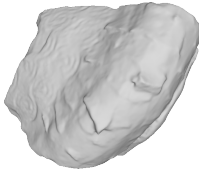




Pattern Legend: eyebrows (EB), oblique fringe (OF), fringe (FR), long incisions (LI), spirals (SP), stamped circlets (CR)				
TOTAL: 4 eyebrows, 5 oblique fringe (OF), 3 fringe, 5 long incisions, 6 spirals, 7 stamped circlets				
1 (v.r. = 410477)  LI	2 (v.r. = 1757300)  -	3 (v.r. = 1408798)  -	4 (v.r. = 4335854)  OF	5 (v.r. = 446529)  OF
6 (v.r. = 414944)  CR	7 (v.r. = 385918)  FR	8 (v.r. = 303620)  CR	9 (v.r. = 492226)  SC EB	10 (v.r. = 946896)  OF
11 (v.r. = 152359)  CR	12 (v.r. = 2373398)  CR	13 (v.r. = 536386)  SP	14 (v.r. = 1063910)  LI	15 (v.r. = 2097095)  -
16 (v.r. = 5312259)  LI CR EB	17 (v.r. = 491542)  LI	18 (v.r. = 1418735)  -	19 (v.r. = 897796)  CR	20 (v.r. = 172935)  SP
21 (v.r. = 1320780)  -	22 (v.r. = 1668277)  SP	23 (v.r. = 150356)  SP	24 (v.r. = 6800671)  EB LI CR	25 (v.r. = 500528)  EB
26 (v.r. = 622012)  SP	27 (v.r. = 771663)  LI	28 (v.r. = 1507431)  OF	29 (v.r. = 1048793)  FR	30 (v.r. = 2467033)  FR

Figure 3: Summary table for MD. v.r stands for vertex resolution.

cannot be uniquely defined and needs to be tuned according to the type of query given in input, calling for adaptive approaches.

- *Local-global nature of the characterization*: Similarly, while we are used to think that the characterization of the object belongs to the local or global description, this task seems to fall in both fields, without belonging to either of those entirely. Most methods that attempt to perform such a task extend or specialize techniques initially born for local or global description. This new mentality, which is really hard to grasp, is still unexplored and could be the key to define effective pattern recognition techniques.
- *Surface embedding*: the same pattern can be chiseled on surfaces with different manifold embedding, for instance a helmet and a cuirass. Differently from images that can be considered flat everywhere, surfaces may present a very general bending in the space. This fact also implies that planar projections may introduce significant distortions of the geometric pattern over the surface. Unless the problem clearly states that only a limited number of configurations is admissible, it is not possible to specialize the pattern recognition problem as a surface fitting or registration problem. In general, an effective description of the pattern must be able to distinguish between the variation of the surface decorations from the overall variation of the surface they lay in.

4. Challenges of the dataset

The choice of a real archaeological case study and not on an ideal dataset where the 3D models are laser scans acquisition of fragments of statues that are degraded and partially abraded makes the contest particularly challenging.

- *Noise and data acquisition*: being the patterns considered in this contest defined by small geometric variations on the surface, the presence of noise can significantly alter the nature of the pattern. In general, it is really hard to remove noise from a mesh without also affecting the geometric variations that define a pattern. A possible solution could be the definition of noise invariant descriptors, a task which is almost impossible without knowing the nature of the noise and the pattern. Similarly, data acquisition conditions (the resolution of instruments, various sources of imprecision) or incompleteness (presence of occlusions, misalignments, surface specularities) challenge the efficiency of the descriptions and increase the difficulty of the pattern recognition task. Therefore, the understanding and modeling of all the sources of uncertainty and incompleteness is the starting point to improve the existing description methods, the results they produce and their quality.
- *Data size*: some of the models in the dataset are modeled at high definition (more than 6 million vertices), thus dealing with them is computationally demanding. The reason for providing the participants with the highest resolution available as they are stored in the STARC dataset [STA] was twofold: first, to limit the approximation error due to the algorithm used for the data simplification; second: to provide the best quality of the data currently available. However, the computational complexity is a bottleneck aspect of most methods and calls for the definition of approaches able to take advantage from parallel architectures.
- *Need of a training set*: the track participants agree that an initial

training set on which to set the parameters could significantly help. However, in the Cultural Heritage domain, artifacts are often unique and the creation of training sets can be very difficult or limited to specific models. While now this dataset can be used as ground truth for feature testings, it would be almost mandatory to have bigger datasets with a sufficiently high variety of pattern embeddings.

5. Possible research directions

Several methods have been considered to tackle the track challenge, ranging from statistical, multi-scale vertex characterization [MGB*12] to divide-and-conquer techniques aimed at isolating the sub-regions with a uniform pattern and then applying retrieval techniques to the single components.

As mentioned in the previous sections, noise, data incompleteness, variability of the patterns size/scale, variation of the surface embedding, necessity of training, computational complexity are all crucial aspects to be considered when dealing with geometric pattern recognition.

In the following we discuss some of the research directions that in our opinion it is worth on further investigating.

1. *Noise*. Typical ways to address noised data are local, patch fitting with parametric surfaces or mesh denoising prior to descriptor extraction [Tau95, HP04], even if these methods are unable to differentiate between geometric features and noise. To overcome these limitations, the iterative denoising scheme in [ALMF18] automatically estimates noise level and geometric feature subspace size (statistical characteristics). Other data-driven approaches [WLT16] for mesh denoising could be employed in cases where data exhibiting similar types of noise exist. The definition however, of an integrated shape representation and description that would be able to describe the local surface geometry without being affected by noise, even for specific types of noise, is still a challenge.
2. *Data incompleteness*. Building local geometric pattern descriptors that are robust to missing-data is a challenging issue. Approaches following the principles of texture inpainting, for instance extracting the shape description on the basis of the surface normal distributions, could be the basis for potential solutions for this problem.
3. *Multi-scale characterization*. Extending single point characterization approaches to a regional level is not straightforward. Volumetric or multi-ring characterizations of the model are worth to be considered [MPS*04, GMGP05]. To extend point-wise descriptors at a more global level, a top-down analysis procedure could be an appropriate strategy. High-level primitives could be computed from local descriptors so that the retrieval problem is more simple and intuitive. For instance, this approach has been used for symmetry detection using lines [ASC11] and one can imagine to compare models and queries based on such a line feature.
4. *Reducing the problem to image pattern recognition*. A possible solution is to adopt a local parametrization/projection of the pattern on a flat surface and evaluate the texture descriptors over that projection. This strategy should provide a method not depending on the quality of point cloud/meshes and the possibility

of exploiting the wide amount of texture analysis methods in the literature. Apart from the computational complexity and the possible pattern distortion introduced with this procedure, that might be high depending on the parametrization choice and the density of the patch sampling, the biggest issue of the method is related to the metric used to compare the patch descriptors.

5. *Definition of proper training sets.* A texture descriptor is typically a vector with a huge dimensionality and in order to derive a way to localize patterns in model we should be able to characterize the pattern versus the other potential "background" patterns that can be found in the models of interest. This cannot be done using just *positive* examples, but requires also the *negative* ones.
6. *Learning.* In image processing, recognition tasks, e.g. finding and classifying objects inside images, are usually accomplished working on large databases of annotated images [EEVG*15] with localization and labeling of objects from which algorithms can learn characterizations of both objects and background. The solution in this case can rely on convolutional neural networks, with methods not directly exploitable in mesh processing, even if there are recent works applying CNN-like approaches to the mesh domain [MBBV15, BBL*17]. For instance, deep learning techniques have recently given remarkable results in shape segmentation and classification [QYSG17]. Although this contest does not contain a large dataset, it could be interesting to see how deep neural networks can be trained on these only 8 queries, still made of hundreds of thousands of vertices. Alternatively, one of the large set of existing learning algorithms such as bag-of-words classification [Lav12] or discriminant analysis [CDF*15] could give an initial partitioning together with representatives descriptors of the queries that would facilitate further matching procedure.

6. Conclusions and future insights

The goal of this track was to recognize a given relief pattern over a set of triangle meshes obtained from laser scans. Unfortunately, none of the methods tested on this benchmark gave satisfactory results. The attempted approaches are various but the difficulty of automatically estimating a unique pattern size, the dimension of the data and the subsequent computational complexity, the need of a data representation which is unique and incorporate data uncertainty seem to be the mostly common problematics faced by the participants.

Some future paths for the research on this topic are proposed. A mainstream research trend in Computer Vision is the adoption of learning techniques for supervised classification. However, in the Cultural Heritage scenario envisaged in this contest only a few shapes of a certain class or type exist, and often, they do not correspond to existing shapes. Therefore, in this specific contest we see the room for developing both learning and direct approaches. On one hand, learned techniques require extensive training data in addition to practical configuration expertise and computational resources [BBL*17] but can be potentially applied to thousands mesh elements. On the other hand, direct approaches yield the adoption of ad-hoc features and algorithms in scenarios where the computational resources are limited or where there is a lack of substantial training data.

Finally, an appropriate definition of a mid-level region characterization would balance between local scale (vertex and very close neighborhood) and global scale (the whole shape) descriptions.

Acknowledgements The work is developed within the research program of the "H2020" European project "GRAVITATE", contract n. 665155, (2015-2018). This research was (partially) funded by the CIMI Labex (ANR-11-LABX-0040)

References

- [ALMF18] ARVANITIS G., LALOS A., MOUSTAKAS K., FAKOTAKIS N.: Feature preserving mesh denoising based on graph spectral processing. *IEEE Transactions on Visualization and Computer Graphics PP*, 99 (2018), 1–1. 5
- [ASC11] AUBRY M., SCHLICKWEI U., CREMERS D.: The wave kernel signature: A quantum mechanical approach to shape analysis. In *2011 IEEE International Conference on Computer Vision Workshops (ICCV Workshops)* (Nov 2011), pp. 1626–1633. 5
- [BBL*17] BRONSTEIN M. M., BRUNA J., LECUN Y., SZLAM A., VANDERGHEYNST P.: Geometric deep learning: Going beyond Euclidean data. *IEEE Signal Processing Magazine* 34, 4 (July 2017), 18–42. 6
- [BMTA*17] BIASOTTI S., MOSCOSO THOMPSON E., AONO M., HAMZA A. B., BUSTOS B., DONG S., DU B., FEHRI A., LI H., LIMBERGER F. A., MASOUMI M., REZAEI M., SIPIRAN I., SUN L., TATSUMA A., FORERO S. V., WILSON R. C., WU Y., ZHANG J., ZHAO T., FORNASE F., GIACHETTI A.: Retrieval of Surfaces with Similar Relief Patterns. In *Eurographics Workshop on 3D Object Retrieval* (2017), Pratikakis I., Dupont F., Ovsjanikov M., (Eds.), The Eurographics Association. 1
- [CDF*15] CAPUTO M., DENKER K., FRANZ M., LAUBE P., UMLAUF G.: Support Vector Machines for Classification of Geometric Primitives in Point Clouds. In *Curves and Surfaces, 8th International Conference, Paris 2014* (2015), Boissonnat J.-D., Cohen A., Gibaru O., Gout C., Lyche T., Mazure M.-L., Schumaker L., (Eds.), Springer, pp. 80–95. 6
- [EEVG*15] EVERINGHAM M., ESLAMI S. M. A., VAN GOOL L., WILLIAMS C. K. I., WINN J., ZISSERMAN A.: The pascal visual object classes challenge: A retrospective. *International Journal of Computer Vision* 111, 1 (Jan 2015), 98–136. 6
- [GMGP05] GELFAND N., MITRA N. J., GUIBAS L. J., POTTMANN H.: Robust global registration. In *Symposium on Geometry Processing* (2005), pp. 197–206. 5
- [GRA] GRAVITATE: Discovering relationships between artefacts using 3D and semantic data. EU H2020 REFLECTIVE project. 2
- [HP04] HILDEBRANDT K., POLTHIER K.: Anisotropic Filtering of Non-Linear Surface Features. *Computer Graphics Forum* (2004). 5
- [Lav12] LAVOUE G.: Combination of bag-of-words descriptors for robust partial shape retrieval. *The Visual Computer* 28, 9 (2012), 931–942. 6
- [MBBV15] MASCI J., BOSCAINI D., BRONSTEIN M. M., VANDERGHEYNST P.: Geodesic convolutional neural networks on riemannian manifolds. In *2015 IEEE International Conference on Computer Vision Workshop (ICCVW)* (Dec 2015), pp. 832–840. 6
- [MGB*12] MELLADO N., GUENNEBAUD G., BARLA P., REUTER P., SCHLICK C.: Growing least squares for the analysis of manifolds in scale-space. *Computer Graphics Forum* 31, 5 (2012), 1691–1701. 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 (Jan 2004), 227–248. 5

- [MPS17] MORTARA M., PIZZI C., SPAGNUOLO M.: Streamlining the Preparation of Scanned 3D Artifacts to Support Digital Analysis and Processing: the GRAVITATE Case Study. In *Eurographics Workshop on Graphics and Cultural Heritage* (2017), Schreck T., Weyrich T., Sablatnig R., Stular B., (Eds.), The Eurographics Association. 2
- [QYSG17] QI C. R., YI L., SU H., GUIBAS L. J.: Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *Advances in Neural Information Processing Systems* (2017), pp. 5105–5114. 6
- [SPS16] SAVELONAS M. A., PRATIKAKIS I., SFIKAS K.: Fisher encoding of differential fast point feature histograms for partial 3D object retrieval. *Pattern Recognition* 55 (2016), 114 – 124. 3
- [STA] STARC repository. <http://public.cyi.ac.cy/starcRepo/>. 5
- [Tau95] TAUBIN G.: A signal processing approach to fair surface design. In *Proceedings of the 22Nd Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 1995), SIGGRAPH '95, ACM, pp. 351–358. 5
- [WLT16] WANG P.-S., LIU Y., TONG X.: Mesh denoising via cascaded normal regression. *ACM Trans. Graph.* 35, 6 (Nov. 2016), 232:1–232:12. 5