

Automatic 3D Object Fracturing for Evaluation of Partial Retrieval and Object Restoration Tasks - Benchmark and Application to 3D Cultural Heritage Data

Robert Gregor¹, Danny Bauer¹, Ivan Sipiran¹, Panagiotis Perakis², Tobias Schreck¹

¹University of Konstanz, Germany

²Norwegian University of Science and Technology, Trondheim, Norway



Figure 1: Illustration of our test data generation approach: Starting with an input CH object (left), an exchangeable cutter object is instanced multiple times (second from left). The intersection of the cutter surface and the CH object's volume results in new breaking edges that disconnect the object into fragments (middle). Subsequently, a randomized number of fragments are removed (right). In our approach, the number and position of cutter object instances as well as fragment removal are randomized within adjustable boundaries.

Abstract

Recently, 3D digitization and printing hardware have seen rapidly increasing adoption. High-quality digitization of real-world objects is becoming more and more efficient. In this context, growing amounts of data from the cultural heritage (CH) domain such as columns, tombstones or arches are being digitized and archived in 3D repositories. In many cases, these objects are not complete, but fragmented into several pieces and eroded over time. As manual restoration of fragmented objects is a tedious and error-prone process, recent work has addressed automatic reassembly and completion of fragmented 3D data sets. While a growing number of related techniques are being proposed by researchers, their evaluation currently is limited to smaller numbers of high-quality test fragment sets.

We address this gap by contributing a methodology to automatically generate 3D fragment data based on synthetic fracturing of 3D input objects. Our methodology allows generating large-scale fragment test data sets from existing CH object models, complementing manual benchmark generation based on scanning of fragmented real objects. Besides being scalable, our approach also has the advantage to come with ground truth information (i.e. the input objects), which is often not available when scans of real fragments are used. We apply our approach to the Hampson collection of digitized pottery objects, creating and making available a first, larger restoration test data set that comes with ground truth. Furthermore, we illustrate the usefulness of our test data for evaluation of a recent 3D restoration method based on symmetry analysis and also outline how the applicability of 3D retrieval techniques could be evaluated with respect to 3D restoration tasks. Finally, we discuss first results of an ongoing extension of our methodology to include object erosion processes by means of a physiochemical model simulating weathering effects.

1. Introduction

With recent advances in 3D data acquisition technology, increasing volumes of 3D object data become available. These not only comprise complete objects, but in many cases, also partial or fragmented objects. The latter is an important issue in digital archeology applications, where the problem is to reassemble or complete plausible shapes from sets of potentially incomplete and eroded fragments. Since the manual restoration process is tedious, error-prone and hardly less resource intensive than directly reassembling/restoring the real CH objects, automatic restoration of the digitized models is desirable. Recent advances in 3D similarity search are promising in supporting data-driven, scalable shape restoration. Typical tasks in object restoration include the reassembly of fragment sets as well as the completion and inpainting of candidate objects [GSP*14].

The development and evaluation of object restoration requires test data to evaluate inform the design of algorithms and the evaluation of precision and timing properties of the devised methods. However, currently available real test data is very limited. Creating additional data is expensive, due to the need to first fragment and scan physical objects. Furthermore, and especially in fragment data sets from the archeology domain, the ground truth (i.e. complete, original shapes) is often not available. This makes it difficult to quantitatively assess the performance of restoration methods. Computational approaches to fragmentation and erosion of 3D objects can provide a solution to overcome this experimental bottleneck and provide a larger variety of test objects. However, there are practical and conceptual problems involved in synthetic generation of fragment data from whole objects. First, the fragmentation process needs to be modeled. Ideally, this should include a plausible physical model of the fracturing process, including specification of structural and material properties of the objects, as well as time-dependent forces exerted on the objects. Furthermore, one may want to take into account long-term effects of erosion, which usually degrade the fragment shapes in several ways. Again, this will depend on a multitude of material and environmental factors, among others.

The above mentioned problems in creating fracture data sets are substantial and hard to solve all together. A thorough treatment would need to take into account complex physical, mechanical and chemical processes and would be governed by a vastly large computational complexity and parameter space that result out of the object's material properties and an even more broad and complex range of environmental conditions that can be expected to often change over time. Even if enough information of the environmental conditions can be plausibly compiled for a then also inherently specific case, doing so would clearly require expertise from various natural science domains. In this paper, we circumvent this issues and propose a simple but resource efficient baseline technique to coarsely approximate a fracturing process for

the generation of benchmark data sets for 3D object restoration. It is based on a volumetric operation involving a suitable (so-called) cutter object which is applied to an input object. Based on a random placement of this cutter object in conjunction with a size-based fragment selection, we create object fragments as the basis for benchmarking. Although this is a simple heuristic scheme which is not based on physical simulation, we show that we are able to obtain visually plausible data sets by this method. The applicability of the generated data is shown by visual inspection, and by exemplary application and evaluation of main processing steps in 3D restoration.

The remainder of this paper is structured as follows. In Section 2, we recall related work in 3D shape retrieval and benchmarking. In Section 3, we describe our methodology for creation of fracture benchmark based on volumetric object processing. Then, in Section 4 we apply our method and present a fracture benchmark data set, based on the well-known Hampson benchmark of whole objects from the Cultural Heritage domain. In Section 5 we show results of using our data for benchmarking of current shape restoration methods. In Section 6, we present a preliminary extension of our data generation process by including simulated weathering effects, and by discussing further extensions and applications of our benchmark. Finally, Section 7 concludes.

2. Related Work

2.1. Retrieval and Restoration Methods

Although there are a large number of methods for retrieval and restoration, we only focus on those closely related to the manipulation of fragmented objects in the cultural heritage domain. We divide the presentation in three main tasks: retrieval, reassembly and repair. In the context of cultural heritage, shape retrieval has proven to be effective because it provides tools for supporting tasks such as classification and restoration. For example, the categorization of several classes of pottery objects using view-based descriptors was proposed in [KPL*10] and [KC11]. The common pipeline of these two approaches considers the computation of several descriptors (2D Zernike moments, character-based encoding and profile features) and their combination for assessing the similarity between two CH objects. Similarly, but in the context of partial shape retrieval, Sfikas et al. [SPK*14] proposed to use panoramic views along with the Bag-of-Features approach to measure the dissimilarity between CH objects. Also, Savelonas et al. [SPS14a] applied a Fisher encoding on histograms of geometric features to tackle the problem of partial shape retrieval in pottery objects. On the other hand, Biasotti et al. [BCFS14] proposed to combine geometric and texture features for the categorization of homogeneous artifacts.

Another interesting problem in CH is object reassembly. Given a set of input fragments, the goal is to solve the 3D puzzle and reconstruct the complete object. An en-

compassing reassembly pipeline based on fracture surface identification for local and global matching of fragments has been introduced in [HFG*06a]. A variational approach was proposed by Litany et al. [LBB12] where the reconstruction was formulated as an optimization problem. Recently, Mavridis et al. [MAP15] proposed a registration-based algorithm to match pairwise fractured surfaces. Differently, Huang et al. [HFG*06b] took advantage of local features to propose an algorithm to find reliable correspondences in fractured surfaces. Likewise, Papaioannou et al [PKT02] defined a matching algorithm based on the computation of the geometric surface error of fractured regions. In a similar manner, Zheng et al. [ZHLW14] defined a geometric error between profile curves to assess the matching between broken objects. In contrast, a semi-automatic algorithm has also been described in [MRS10]. In this method, a user provides a first alignment which is further refined using a variation of the ICP algorithm.

A number of methods have been proposed to address the problem of CH object repair. Li et al. [LYW*11] specifically exploited the symmetric information of human skulls for repair. More generally, Sipiran et al. [SGS14] proposed a symmetry-preserving function which allows to detect approximated symmetry planes, even in very damaged objects. These planes are subsequently used to synthesize missing geometry. Also, Harary et al. [HTG14] devised an algorithm that take advantage of self-similarity to repair large holes in 3D meshes. On the other hand, Adán et al. [ASM12] described a semi-automatic procedure to repair complex CH objects guided by cone curvature features and manual alignment. Not only from the aforementioned methods, it is possible to identify a lack of a structured methodology to evaluate the methods in a consistent way. To enable a more informative comparison of the growing number of approaches in the field, we address the benchmarking and evaluation of CH-related restoration methods.

2.2. Retrieval Benchmarks

The number of datasets for evaluation of shape retrieval tasks is growing fast. They span from benchmarks for generic tasks (such as global retrieval, partial retrieval and non-rigid shape retrieval) to task-specific benchmarks such as molecule retrieval or face retrieval, among others. Indeed the Shape Retrieval Contest (SHREC) is an important resource to illustrate the large variety of benchmarks that currently exist. A recent compilation of 3D benchmarks including retrieval datasets can be found in [GZLW14]. In fact, as partial shape retrieval is closely related to the problem we want to tackle, it is worth to mention the recent comprehensive study of partial shape retrieval algorithms presented by Savelonas et al. [SPS14b]. Nevertheless, as the problem of using retrieval for analyzing fragment sets is relatively new, there is no benchmark that would accommodate this specific goal. On the other hand, from

the point of view of cultural heritage and to our knowledge, the Hampson dataset (http://hampson.cast.uark.edu/about_technology.htm) is the only one that provides scanned CH objects. However, this dataset is intended for classification tasks on mostly complete models. Therefore, and since we aim at producing a near-to-real CH benchmark for restoration, we decided to adopt the Hampson dataset to showcase the applicability of our approach.

3. Object Fracturing Methodology and Implementation

To generate data that can be used for benchmarking the effectiveness of 3D Retrieval and 3D Restoration methods on fractured CH Objects, we propose a method that produces artificial fragments from arbitrary, watertight input shapes. As outlined in section 1 and discussed in e.g. [CYFW14], simulating physically correct fracturing in high resolution has severe drawbacks. Usually, the simulation is computed on top of a tetrahedral object representation using a Finite Element Method (FEM) based approach. A large number of material and environmental parameters have to be provided to the model. This may be a difficult task to do even for experts, when the process needs to be applied to a larger number of shapes of different material, which in turn have been exposed to potentially different environmental conditions. Furthermore, FEM based fracturing simulation is already very expensive to compute for 3D objects with a medium level of detail that is clearly below high quality 3D CH models. According to [CYFW14], a simulation with models of reduced resolution yields significantly different characteristics than a simulation in higher resolutions.

However, for evaluation of retrieval and restoration methods, strictly correct physically fragments may not be needed, and heuristic methods may suffice. Many existing retrieval techniques do not rely on very detailed assumptions concerning the objects material or recurring, small-scale patterns in a shape that reflect certain physical effects under a given range of environmental conditions. In many cases, it can be expected that such techniques will have comparable effectiveness, regardless of the physical correctness as long as the overall characteristics of the shapes are similar enough. In the field of digital movie production, alternate methods have been proposed to efficiently produce larger numbers of fracturings. While their results are often considered as very plausible by the audience, the methods are not based on sound physical models. We adapt a state of the art method [AMM12] that usually relies on a 3D Modeling artist to provide so-called *cutter objects* and details about their instantiation to control the generation of new fragments. In our case, the cutter object resembles a sphere where small-scale details have been transferred from real digitized breaking edges of a certain material by semi-automatic mesh geometry cloning [TSS*11]. Figure 1 (second from left) illustrates two instances of a cutter object we modeled. Note that when relying on the aforementioned technique, alternative cutter objects based on other examples of real breaking




Figure 2: Comparison of real and synthetically generated defects within an object: A part was missing already in the input object (right). Our approach removed an additional part (center front). Also note original cracks (back right, front left) and synthetic cracks (back left, right) next to each other. Our method allows creating rich fragment test cases for retrieval and restoration applications from existing 3D CH Models.

edges can be generated easily. To compute the fracturing, we randomize a number of cutter object instances and their spatial arrangement within configurable boundaries and apply them to a given input object. Technically, after scale-normalization, the input and the cutter objects are converted to a hierarchical volumetric representation (Level Sets) before we rely on the implementation of the *Level Set Fracture Algorithm* as described in [AMM12]. To obtain disconnected sub-volumes within the Level Set, we generate new breaking edges by computing the intersection of the cutter object surface and the input object volume. The resulting Level Set is converted back to a mesh representation before its disconnected components – the newly obtained fragments of the input mesh – are split into separate manifold meshes. To reflect the fact that very small fragments usually cannot be recovered, or unambiguously associated to specific objects, we heuristically remove fragments that are smaller than a configurable threshold parameter. Next, all vertices of the remaining fragments are annotated according to their Hausdorff distance to the input object.

Subsequently, a random subset of larger fragments is removed. This process can be controlled by configurable parameters for the maximum number of removed and the minimal number of retained fragments. During Fragment removal, an optional test can ensure that all remaining fragments share at least a complimentary breaking edge with another remaining fragment, which could e.g. prove useful for the evaluation of e.g. retrieval or symmetry-based restoration methods. If the test fails, a different fragment is selected randomly until either the minimum number of fragments to be retained or the maximum number of fragments to be removed is reached. To ensure shared, complementary breaking edges, we a priori construct a connectivity graph

based on spatially close vertices, annotated by high Hausdorff distance that belong to different fragments. We test for overall connectivity by counting the respective number of visited fragment nodes during breadth-first search when starting from each of the adjacent fragments. To coarsely approximate deformation of small-scale features (e.g. due to mechanical abrasion, or shot noise), a subsequent optional step can repeatedly introduce gaussian noise or smoothing. Finally, to synthesize cracks, the remaining fragments can optionally be merged back into a single manifold using Poisson Reconstruction [KBH06]. For testing reassembly methods, this step can be disabled. In both cases, adjacent fragments result in artifacts that are on a visual level very similar to digitization of real cracks (Fig. 2).

The implementation of our data generation method is provided under <http://fracture-benchmark.dbvis.de>. The main control flow is implemented in Groovy Script, where individual parameters can be configured. Most individual steps are delegated to various binaries. The Randomized Level Set Fracturing technique is implemented in a C++ based command line application that relies on OpenVDB's Level Set Fracturing algorithm. Several of the remaining mesh processing operations are delegated to Meshlab-Server and Blender. An additional C++ based command line tool may optionally be used to preprocess the fractured and unfractured objects by a sequence of Scale-Normalization, Poisson Reconstruction and Uniform Remeshing.

4. Fragment Benchmark Based on Hampson Collection

To demonstrate the effectiveness of our data generation approach, we selected the freely available Hampson Museum Dataset provided by the University of Arkansas. The collection contains an encompassing set of American Indian clay pottery. The collection currently contains 445 objects, however not all of them are watertight and contain interior object geometry. Hence we manually selected a subset of 84 watertight models from this collection. We applied our fracturing approach as described in Section 3 to create our *Hampson Fracture* data set. Specifically, we fixed the number of cutter objects to four instances. To exclude too small resulting fragments, we heuristically discarded fragments which contained less than one thousand vertices. From the resulting fragment set, we randomly omitted at most three fragments and retained at least one. No smoothing or Gaussian noise was applied, before the remaining fragments were merged by Poisson reconstruction. Figure 3 provides a comparison of textured original models and the results of our data generation.

In our benchmark, the fragment set represents the restoration/retrieval problem, while the underlying full object represents the ground truth. We provide the fractured data along with our implementation.




Figure 3: Original, digitized CH Objects from the Hampson Museum Dataset along with their artificially degraded counterparts. Fragment vertices are colored by Hausdorff distance to the input models. Note that only some of generated fragments have been randomly omitted whereas others are still included and cause cracks in the generated models (colored green to yellow, medium distance). Synthesized breaking edges are colored in blue (high distance).

5. Evaluation Assessment

In the following, we demonstrate how our benchmark can be used to evaluate shape retrieval and object restoration methods. The application also contributes a measure for quantifying the degree of shape completion, based on the ground truth information provided in the benchmark.

5.1. Shape Retrieval

In this exemplary experiment, we evaluate the performance of a given 3D retrieval technique. The goal is to retrieve the correct unfractured counterpart for each synthetically fractured model (see Section 4). While this task can be considered roughly similar to partial retrieval, fractured queries do not contain holes and gaps but additional breaking edges where larger parts are missing. Using the generated data in this way, the obtained measures can provide a first indication of the invariance of a retrieval technique against defects encountered in fractured CH Objects.

As an exemplary retrieval techniques, we use the DSR descriptor [Vra05], as well as Heat Kernel Signatures (HKS) [SOG09] and Scale Invariant HKS [BK10], both in combination with a soft-quantized Bag of Features with dictionary size of 256, trained on the unfractured objects. We retrieve the input models using different distance functions such as L_p where $p \in \{0.1, 0.3, 0.6, 0.9, 1, 2\}$ as well as statistical distance measures, such as χ^2 and Kullback-Leibler, Jeffrey and Jensen-Shannon Divergence. We measure the success rate as the fraction of fractured objects, for which the correct unfractured object is ranked within the first n results, where $n \in \{1, 5, 10, 25\}$. While we have to note that the used approaches are not in particular designed for partial retrieval, we can observe interesting results (Fig. 4). Given existing experiments for global retrieval on these methods, it was not

to be expected, that DSR outperforms HKS, which in turn is not clearly outperformed by SIHKS. Furthermore, it is interesting to see that for our fractured objects, the DSR descriptor is very sensitive to the choice of distance function. This contrasts our results of the same descriptors and distance functions when measuring R-Precision with other benchmarks, such as the well known Princeton Shape Benchmark [SMKF04].

Note that this is an exemplary experiment. A comparison with Partial Retrieval methods is currently work in progress. Further refinement of the setup for retrieval evaluation could e.g. include an object classification, a higher total number of fractured objects, multiple fracturings per input object and the introduction of small scale deformations, such as noise or erosion simulation (sec 6.2).




Figure 4: Fraction of successfully retrieved unfractured objects. Due to space limitations, we show the results for L_1 and L_2 and the respective best performing other distance measure tested.

5.2. Object Restoration

In this experiment we evaluate the ability of a symmetry-based completion algorithm. Note however, that the following evaluation methodology can also be applied to other

CH object restoration methods. We selected the method proposed by Sipiran et al. [SGS14]. Unlike the original evaluation in [SGS14], here we provide a more systematic evaluation which can lead us to objectively test the robustness of completion algorithms. Each object in the benchmark is associated to a set of fragments which were obtained by our fragmentation tool. For this test, we randomly discarded one fragment. In the evaluation, we compare the results to the ground truth provided by the unfractured input object.

More formally, let O_i be the i -th object in the benchmark. Let $F_i = \{f_i^j\}$ be a set of fragments associated to O_i . Let $f_i^* \in F_i$ be the discarded fragment. After applying the algorithm to the set of fragments $F_i \setminus \{f_i^*\}$, we obtain a completed object C_i . Our criterion to evaluate the robustness of completion algorithms with respect to the completion is the ratio of volumes

$$E_{\text{completion}}(O_i) = \frac{\text{vol}(O_i \cap C_i)}{\text{vol}(O_i \cup C_i)} \quad (1)$$

It is worth to note that $E_{\text{completion}}$ is a conservative measure for the robustness. This is because we penalize the congruent geometry ($\text{vol}(O_i \cap C_i)$) with the divergent geometry ($\text{vol}(O_i \cup C_i)$). A clear example where this penalization is important is when the original object is completely covered by the completed object, but in addition the completed object introduces more geometry than needed. In that case, the ratio will decrease depending on the amount of divergent geometry. Also note that a perfect completion will give a $E_{\text{completion}} = 1$. Note however that some Hampson models originally have (comparatively small) missing parts and cracks, which introduces a source of error in our exemplary, quantitative evaluation.




Figure 5: Difference between effectiveness and partiality of the fragmented objects. Each bar represents the magnitude of how much net true geometry was completed. Positive bars illustrate successful completions, whereas negative bars represent inappropriate completions.

It is also possible to compute the degree of missing geometry for each object which will be useful as a reference for comparison:

$$D(O_i) = \frac{\text{vol}(F_i \setminus \{f_i^*\} \cap O_i)}{\text{vol}(F_i \setminus \{f_i^*\} \cup O_i)} = \frac{\text{vol}(F_i \setminus \{f_i^*\})}{\text{vol}(O_i)} \quad (2)$$

Note that $D(O_i)$ is the volume ratio of the input fragments

compared to the original object. This quantity gives us an idea about how much information was removed by discarding one fragment.

To compute the volume of a 3D shape, we implemented the algorithm proposed in [ZC01]. Briefly, the algorithm accumulates the signed volume of the tetrahedron formed by each triangle in the shape and the origin. In our experiments, we executed the completion algorithm in 76 sets of fragments (one per object in the benchmark of fractured objects). On average, we obtained $E_{\text{completion}} = 0.8267$. Nevertheless, this average alone is not a good indicator for the effectiveness. A more meaningful comparison is the gain of effectiveness; that is the difference $E_{\text{completion}}(O_i) - D(O_i)$ for each object. This difference can reveal whether completion methods really complete some geometry or not. Fig. 5 shows the gain of effectiveness for each query object. The completion algorithm obtained a positive gain in 45 out of 76 objects. This fact reveals that in average the method is able to infer and complete missing geometry. After inspecting some random objects with negative gain, we discovered a reason for the negative difference: When the missing fragment is present across the symmetry planes, the method is not able to reproduce some parts of the completion. Finally, an example of result obtained in our evaluation can be seen in Fig. 6.




Figure 6: Left: fragment set. Right: result of the completion algorithm proposed in [SGS14]

6. Discussion and Extension of Our Approach

6.1. Plausibility

While our current fracturing approach is not based on a physically founded simulation process, we observe by visual inspection that it can produce plausible results. It can be computed efficiently and only requires setting few intuitive parameters. Still, certain details of the produced geometry deviate from what one might expect of real fracturing. E.g., in the case of clay objects, it is often observed that the breaking edges exhibit a *flocking* effect where a large number of small, exposed parts break away (flock out) which in turn can have a significant impact on the overall characteristics of the fracture surface. If e.g., the task is to benchmark methods for segmenting and annotating fracture surfaces, depending on the studied material, inclusion of flocking effects should be done. A solution to this could be to apply an appropriate small scale morphological operator as a post-processing step after the initial fracturing process has taken place.

Furthermore, our approach does not reflect other small scale object defects introduced by e.g., mechanical wear and abrasion of real fragmentation processes. Extending our approach in this direction could benefit the research of a new class of shape retrieval techniques which could be invariant with respect to such abrasion effects. We also note that in our approach we use a roughly spherical cutter object, which reflects global properties of the resulting fragments surface. By modeling cutter objects with different global shapes, one may vary the types of obtained fracture shapes, adapting it to characteristics of specific problems to develop restoration methods for.

Finally, our current implementation discards existing texture coordinates and does not synthesize new 2D textures, which otherwise could be useful for testing multi modal restoration methods.

6.2. Inclusion of Erosion Effects




Figure 7: Two digitized and manually reassembled fragment sets made of limestone (top: column base, bottom: embrasure) before (left) and after (right) a simulation of 150 years of chemical weathering in polluted atmosphere.

In addition to mechanical influence, objects also deteriorate over time. Erosion due to chemical or meteorological factors can especially affect the characteristics of small-scale features on the object. We are currently extending our approach by an erosion simulation approach that assumes homogeneous material as well as spatially and temporally uniform exposure to environmental effects. For each vertex of an input object, a shift is computed depending on the object's assumed material properties and environmental parameters.

At this stage, we experiment with material parameter estimations for the chemical weathering of marble and limestone due to interaction with carbon dioxide (CO_2) in unpolluted atmosphere regions, as well as sulfur and nitrogen dioxide (SO_2 and NO_2) in polluted atmosphere regions (e.g. close to industrial exhaust). Our chemical models and parameter estimations are based on [YJCG94, YCJ*96, GB99]. In particular the following effects are simulated: *Dry deposition in polluted environments* reflects the creation of crust, usually due to chemical material transported in polluted environments. The crust mostly consists of gypsum and results of reactions of the stone material with SO_2 and NO_2 . *Recession by Acid Rain* reflects the loss of material mainly due to reactions of water with the stone material and SO_2 and NO_2 . To a lesser extent, recession also takes place in unpolluted atmosphere due to reactions caused by the presence of CO_2 . In both cases, Ca^{2+} ions are washed out of the objects.

First results are provided in Figure 7. We argue that considering erosion processes is an important problem in the definition of benchmarks for object restoration tasks. One intriguing research challenge we see in this context, is to work on shape retrieval methods which provide invariance with respect to fragmentation and erosion processes.

7. Conclusion

We presented a methodology to automatically generate 3D fracture test data for shape restoration and retrieval tasks. The method is a first step to fill a gap in experimental data provision, which is typically expensive to obtain from real objects, and often lacks ground truth. Our approach is based on a simple volume operation based on an appropriately modeled cutter object. Although it is not based on physical fracturing processes, it shows plausible results and encompasses only few, intuitive parameters. We create a benchmark data set based on the Hampson 3D repository. We demonstrated how our benchmark can be applied to evaluate existing object restoration and retrieval tasks. We also discussed extension possibilities for our generation method including simulation of erosion processes.

Future work includes applying the benchmark to compare and improve the effectiveness of object restoration and retrieval methods. The fracturing and erosion process should be extended to include physical and mechanical factors, which could provide additional realistic properties to the fragment data. One idea is to guide the fracturing along structural properties, such as e.g., medial axes or weak areas of structural integrity. A challenge will be to strike a balance between the realism of the process, its computational complexity and usability regarding user parameters choice. Finally, the results of the fracture and erosion generation should be quantitatively compared against real-world examples, found e.g., in digital archeology applications.

Acknowledgements

We thank the anonymous reviewers for their constructive comments. This work was supported by the European Commission in context of the FP7 STREP Project PRESIOUS, grant no. 600533, <http://presious.eu/>.

References

- [AMM12] ALDEN M., MELICH G., MUSETH K.: Efficient and Seamless Volumetric Fracturing. *SIGGRAPH Talk* (2012). 3, 4
- [ASM12] ADÁN A., SALAMANCA S., MERCHÁN P.: A hybrid human-computer approach for recovering incomplete cultural heritage pieces. *Computers & Graphics* 36, 1 (2012), 1–15. Cultural Heritage. 3
- [BCFS14] BIASOTTI S., CERRI A., FALCIDIENO B., SPAGNUOLO M.: Similarity Assessment for the Analysis of 3D Artefacts. Klein R., Santos P., (Eds.), Eurographics Association, pp. 155–164. 2
- [BK10] BRONSTEIN M. M., KOKKINOS I.: Scale-invariant heat kernel signatures for non-rigid shape recognition. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (2010), pp. 1704–1711. 5
- [CYFW14] CHEN Z., YAO M., FENG R., WANG H.: Physics-inspired adaptive fracture refinement. *ACM Transactions on Graphics* 33 (2014), 1–7. 3
- [GB99] GAURI K. L., BANDYOPADHYAY J. K.: Carbonate Stone, Chemical Behaviour, Durability and Conservation. *John Wiley & Sons, Inc.*, 1999. (1999). 7
- [GSP*14] GREGOR R., SIPIRAN I., PAPAIOANNOU G., SCHRECK T., ANDREADIS A., MAVRIDIS P.: Towards automated 3D reconstruction of defective cultural heritage objects. In *Proc. EG Workshop on Graphics and Cultural Heritage* (2014), Eurographics, pp. 135–144. 2
- [GZLW14] GUO Y., ZHANG J., LU M., WAN J.: Benchmark Datasets for 3D Computer Vision. *IEEE Industrial Electronics and Applications (ICIEA)* (2014). 3
- [HFG*06a] HUANG Q.-X., FLÖRY S., GELFAND N., HOFER M., POTTMANN H.: Reassembling fractured objects by geometric matching. *ACM Trans. Graphics* 25, 3 (2006), 569–578. 3
- [HFG*06b] HUANG Q.-X., FLÖRY S., GELFAND N., HOFER M., POTTMANN H.: Reassembling Fractured Objects by Geometric Matching. *ACM Trans. Graph.* 25, 3 (July 2006), 569–578. 3
- [HTG14] HARARY G., TAL A., GRINSPUN E.: Context-based coherent surface completion. *ACM Trans. Graph.* 33, 1 (Feb. 2014), 5:1–5:12. 3
- [KBH06] KAZHDAN M., BOLITHO M., HOPPE H.: Poisson surface reconstruction. *Symposium on Geometry Processing* (2006). 4
- [KC11] KOUTSOUDIS A., CHAMZAS C.: 3D pottery shape matching using depth map images. *Journal of Cultural Heritage* 12, 2 (2011), 128–133. 2
- [KPL*10] KOUTSOUDIS A., PAVLIDIS G., LIAMI V., TSIFAKIS D., CHAMZAS C.: 3D pottery content-based retrieval based on pose normalisation and segmentation. *Journal of Cultural Heritage* 11, 3 (2010), 329–338. 2
- [LBB12] LITANY O., BRONSTEIN A., BRONSTEIN M.: Putting the pieces together: Regularized multi-part shape matching. In *Computer Vision - ECCV 2012. Workshops and Demonstrations*, Fusiello A., Murino V., Cucchiara R., (Eds.), vol. 7583 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 2012, pp. 1–11. 3
- [LYW*11] LI X., YIN Z., WEI L., WAN S., YU W., LI M.: Symmetry and template guided completion of damaged skulls. *Computers & Graphics* 35, 4 (2011), 885–893. Semantic 3D Media and Content. 3
- [MAP15] MAVRIDIS P., ANDREADIS A., PAPAIOANNOU G.: Fractured object reassembly via robust surface registration. In *Eurographics Conference (short paper)* (2015). To appear. 3
- [MRS10] MELLADO N., REUTER P., SCHLICK C.: Semi-automatic geometry-driven reassembly of fractured archeological objects. In *Proceedings of the 11th International Conference on Virtual Reality, Archaeology and Cultural Heritage* (Aire-la-Ville, Switzerland, Switzerland, 2010), VAST'10, Eurographics Association, pp. 33–38. 3
- [PKT02] PAPAIOANNOU G., KARABASSI E.-A., THEOHARIS T.: Reconstruction of three-dimensional objects through matching of their parts. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 24, 1 (Jan 2002), 114–124. 3
- [SGS14] SIPIRAN I., GREGOR R., SCHRECK T.: Approximate symmetry detection in partial 3D meshes. *Computer Graphics Forum* 33, 7 (2014), 131–140. 3, 6
- [SMKF04] SHILANE P., MIN P., KAZHDAN M. M., FUNKHOUSER T. A.: The princeton shape benchmark. In *2004 International Conference on Shape Modeling and Applications (SMI)* (2004), pp. 167–178. 5
- [SOG09] SUN J., OVSIANIKOV M., GUIBAS L.: A Concise and Provably Informative Multi-Scale Signature Based on Heat Diffusion. *Computer Graphics Forum* 28, 5 (July 2009), 1383–1392. 5
- [SPK*14] SFIKAS K., PRATIKAKIS I., KOUTSOUDIS A., SAVELONAS M. A., THEOHARIS T.: Partial matching of 3D cultural heritage objects using panoramic views. *Multimedia Tools and Applications, Springer* (2014), 1–15. 2
- [SPS14a] SAVELONAS M. A., PRATIKAKIS I., SFIKAS K.: Fisher encoding of adaptive fast persistent feature histograms for partial retrieval of 3D pottery objects. In *Eurographics Workshop on 3D Object Retrieval* (2014). 2
- [SPS14b] SAVELONAS M. A., PRATIKAKIS I., SFIKAS K.: An overview of partial 3D object retrieval methodologies. *Multimedia Tools and Applications, Springer* (2014), 1–26. 3
- [TSS*11] TAKAYAMA K., SCHMIDT R., SINGH K., IGARASHI T., BOUBEKEUR T., SORKINE O.: GeoBrush: Interactive mesh geometry cloning. *Computer Graphics Forum* 30, 2 (2011), 613–622. 3
- [Vra05] VRANIC D. V.: DESIRE: a composite 3d-shape descriptor. In *Proceedings of the 2005 IEEE International Conference on Multimedia and Expo, ICME 2005, July 6-9, 2005, Amsterdam, The Netherlands* (2005), pp. 962–965. 5
- [YJC*96] YERRAPRAGADA S., CHIRRA S., JAYNES J., LI S., BANDYOPADHYAY J., GAURI K.: Weathering rates of marble in laboratory and outdoor conditions. *Journal of Environmental Engineering*, 122(9):856–863 (1996). 7
- [YJCG94] YERRAPRAGADA S., JAYNES J., CHIRRA S., GAURI K.: Rate of weathering of marble due to dry deposition of ambient sulfur and nitrogen dioxides. *Analytical Chemistry*, 66(5):655–659 (1994). 7
- [ZC01] ZHANG C., CHEN T.: Efficient feature extraction for 2D/3D objects in mesh representation. In *Image Processing, 2001. Proceedings. 2001 International Conference on* (2001), vol. 3, pp. 935–938 vol.3. 6
- [ZHLW14] ZHENG S. Y., HUANG R. Y., LI J., WANG Z.: Reassembling 3D Thin Fragments of Unknown Geometry in Cultural Heritage. *ISPRS Annals of Photogrammetry, Remote Sensing and Spatial Information Sciences* (May 2014), 393–399. 3