

Towards Automated 3D Reconstruction of Defective Cultural Heritage Objects

Robert Gregor¹, Ivan Sipiran¹, Georgios Papaioannou², Tobias Schreck¹, Anthousis Andreadis² and Pavlos Mavridis²

¹University of Konstanz

²Athens University of Economics and Business

Abstract

Due to recent improvements in 3D acquisition and shape processing technology, the digitization of Cultural Heritage (CH) artifacts is gaining increased application in context of archival and archaeological research. This increasing availability of acquisition technologies also implies a need for intelligent processing methods that can cope with imperfect object scans. Specifically for Cultural Heritage objects, besides imperfections given by the digitization process, also the original artifact objects may be imperfect due to deterioration or fragmentation processes. Currently, the reconstruction of previously digitized CH artifacts is mostly performed manually by expert users reassembling fragment parts and completing imperfect objects by modeling. However, more automatic methods for CH object repair and completion are needed to cope with increasingly large data becoming available.

In this conceptual paper, we first provide a brief survey of typical imperfections in CH artifact scan data and in turn motivate the need for respective repair methods. We survey and classify a selection of existing reconstruction methods with respect to their applicability for CH objects, and then discuss how these approaches can be extended and combined to address various types of physical defects that are encountered in CH artifacts by proposing a flexible repair workflow for 3D digitizations of CH objects. The workflow accommodates an automatic reassembly step which can deal with fragmented input data. It also includes the similarity-based retrieval of appropriate complementary object data which is used to repair local and global object defects. Finally, we discuss options for evaluation of the effectiveness of such a CH repair workflow.

1. Introduction

3D digitization of Cultural Heritage (CH) objects has recently gained significant ground, with applications ranging from preservation and presentation to analysis and re-use e.g., in the cultural industry sector. A range of high-to medium-precision acquisition techniques exist, recently complemented by widely available commodity-type methods and hardware. However, the acquired 3D data are frequently imperfect or incomplete with respect to the represented original objects. This may be due to many factors, including acquisition artifacts and inaccessible views, but also due to a problem specific to the CH domain, the deterioration and fragmentation of the original objects. Due to deterioration effects, local detail may have been eroded away over time and worse, objects are often fragmented or missing important parts of the shape. The repair of all of these effects typically requires manual effort of CH domain experts and is time-consuming and thus expensive.

Recent advances in 3D shape processing allow for intelligent, automated processing of defects that stem from physical defects of the previously digitized objects. Successful automatic repair and/or reconstruction approaches have been discussed for specific applications, exploiting domain knowledge and constraints which control the process. Examples include pottery data, exploiting rotational symmetry [KS03], or city models applying search for block schemes and supervised learning [LGZ*13]. However, such methods are specific to the addressed object class. More generic object repair may be possible if one assumes parametric templates of the object classes to be reconstructed [UF11]. Such approaches, however, inherently require definition of parametric repair templates and are increasingly computationally expensive with rising numbers of templates and template parameters.

Relying on recent advances in 3D Shape retrieval, we propose an approach that applies, extends and combines sev-

eral techniques to address various defect categories while maintaining a high level of automation. Among others, we propose two new data-driven techniques to synthesize missing parts. The first approximates symmetry planes in defective objects to infer missing parts by intra-object similarity. The second relies on a repository of template shapes and a heuristic for automated surface classification to deduce missing object parts by inter-object similarity. Whereas the first technique does not imply any need for external data, the second one can by design handle objects that don't have prevalent symmetry planes. Furthermore its applicability scales with the ever-growing availability of digitized 3D CH models. However, the second technique can also benefit from parametric templates without suffering in terms of scalability. The realization of our approach involves solving a range of 3D processing problems from retrieval of similar shapes based on partial shape information, registration and merging of shapes, repair of local mesh defects and transfer of shape detail.

Our contribution in this paper is to (a) characterize the problem of CH object repair based on a survey of typical shape defects occurring and a survey of related shape repair and reconstruction methods, (b) introduce and exemplify apply a comprehensive similarity-based CH object repair workflow, addressing the identified object imperfections while presenting first results. Then, (c) we propose an evaluation methodology for the effectiveness of the repair workflow.

2. Defects of CH Artifacts

The notion of a defect or flaw indicates a certain frailty or shortcoming of an object that prevents it from having certain qualities that are precondition for utilizing the object for a certain purpose. In the context of CH object digitization, we distinguish between defects that are inherent to the physical object, and defects that are introduced during acquisition, conversion or processing of its digital representation. The latter (also denoted as mesh defects, in case the representation is of type mesh) can be considered an ubiquitous problem in computer graphics and has been subject to extended research [ACK13]. However, there are few publications that specifically address identification or repair of the original physical defects of objects. Mesh defects such as singularities, holes or self-intersections can often be detected by rather simple and well-known algorithms. On the other hand, identifying physical defects of CH objects can in general be regarded as an ill-posed problem and is hard to do fully automatically. Typically, to repair such defects, background knowledge as e.g. the class of the object, its material, or environmental conditions that it was exposed to, needs to be brought in.

We can classify important types of defects that we encounter in CH objects:

Small-scale decay results in a high number of small



Figure 1: Defective objects from Nidaros Cathedral, Trondheim, Norway: Fragments of a doorway arch (top left), reassembled embrasure with missing parts and cracks (top right), reassembled column base with large cracks (bottom left), reassembled tombstone with hole and cracks (bottom right).

and irregular changes to the objects geometry and texture. Examples of such processes are chemical weathering and mechanical abrasion. As these defects are normally densely collocated, they often affect larger, connected areas of the object surface. Although the spectrum of small-scale features of an object can be severely affected, the effect on the overall shape of an object is usually small.

Missing of small fragments encompasses chipped off slivers, fissures and larger scratches which reduce the volume of the object. These effects can be caused by physical weathering or other mechanical stress. In comparison to the previous defect category, these defects are usually less in number but cause stronger local effects on the surface geometry. Thus they can be more easily distinguished and the geometrical information for their repair is more likely to be derived from its surroundings. However, their exact geometric appearance is still dependent on the objects material and environmental conditions.

Sediments on the surface increase the volume of the object and can have a varying intensity which especially affects the texture as well as small-scale geometric features of the surface. The affected area may be large, depending on the root cause. Some parts of sediments may be manually removed in preparation of 3D acquisition.

Missing of Large Parts leads to significant changes of the overall shape and can arise due to mechanical stress or as a long-term effect of physical and chemical weathering. Missing large parts also correspond to breaking edges that introduce small-scale local features which are not related to the original shape. If broken-off large parts were recovered and digitized, they can be used for reassembly.

For parts that can not be recovered, all explicit information on the geometry and texture is lost. Due to their varying size, it is in many cases difficult to infer their geometry directly from the remaining recovered parts. Under such circumstances, missing larger parts could look alike to missing small scale fragments or residues. For example, a mug missing its handle could resemble a vase. A fragment of an embrasure could be misclassified as a decoration of a pillar. Often, even CH domain experts have difficulties in predicting the class or completion of a CH artifact which is missing larger parts. Then, domain specific knowledge and assumptions may be used to create hypotheses on the original appearance and meaning of the object.

To summarize, when dealing with CH artifacts in practice a range of imperfections will be encountered. Figure 1 illustrates some of the defects occurring in real CH object data. To obtain a digitization that is closer to the original geometry of the object, these need to be dealt with.

3. Taxonomy of 3D Shape Repair Methods

In this section, we present a brief survey of 3D shape repair methods that can be applied to digitization of defective CH objects. Based on an extensive literature study, we can distinguish a) general methods which can address mesh structure defects [ACK13], and b) methods which rely on shape similarity to repair objects. While general mesh repair methods are useful for automatically removing smaller defects or global mesh inconsistencies, similarity-based methods are needed to repair or replace larger shape parts. We sub-divide the similarity-based approaches into those that address Modeling (or Assembling), Reassembly, Inpainting, and Symmetry-based methods. While the former two are mainly based on exploiting similarity to externally provided shapes, the latter two focus on exploiting similarity within the shape itself. Fig. 2 shows our taxonomy. We next survey some recent work in each of these categories, before giving a comparative summarization.

3.1. Modeling / Assembly

In approaches such as the classic “Modeling by Example” [FKS*04] and its extensions, the user interactively models the reconstruction by reusing parts retrieved from an external shape repository. Repository queries are performed by user-provided meta-data keywords or by content-based similarity to the object within user-denoted query regions. Subsequently, the user selects a shape from the retrieved candidate list for geometry transfer. In [LF08] the original query




Figure 2: A basic taxonomy of 3D Object repair methods. We distinguish similarity-driven repair methods that exploit intra or inter-shape similarity to synthesize a plausible overall shape.

approach is extended to use 2D sketches that can be drawn on top of the object. Furthermore, the shape repository is extended to contain the segmentation of objects as well as semantic annotations of objects and their parts. These are used to increase retrieval accuracy while reducing online user input. E.g. based on annotations of already assembled parts, the system can propose relevant parts even before the user initiates the next query. Both approaches can be regarded as highly versatile, yet they require external example data from which to obtain information for the repair process. To address this, in [KHS10] an approach is described that can be used to automate the creation of such repository data by machine-learning-based segmentation and part-labeling. Although further extensions as [CKGK11, KCKK12] aim to automate the generation of repository data even more, modeling inherently requires a high number of online user interventions.

3.2. Reassembly

While also primarily relying on inter-object similarity, reassembly approaches differ from modeling in that they operate on a set of fragment objects that are expected to be parts of the reconstruction. Reassembly can be driven by two entirely different processes: Contact surface (multi-part) registration and part-to-whole matching, which, in many cases, can also be combined.

In multi-part registration, the fragments are first automatically segmented to detect potentially fractured surfaces. Subsequently, all combinations of marked segments are subjected to a rigid registration procedure, where both the estimated relative pose and the matching error (residual) are computed. This (potentially offline) processing results in a large number of possible fragment pair configurations, which are finally resolved by a combinatorial optimization stage to form clusters of aligned fragments. Two prominent automatic approaches are reported in this area; Papaioannou et al. [PKT01] perform a rigid registration by optimizing the

local residual gradient and contact area between the fracture surfaces. Residuals are measured via the projection of the fractured segment pair on a common plane using hardware rasterization and stochastic pose optimization is used for the fast alignment of the surfaces. Results are further refined by performing contour matching on fracture boundary segments [PK03]. Multi-part reassembly is achieved by a genetic algorithm that globally optimizes the pairwise links. Huang et al. [HFG*06] extract multi-scale features on the fracture surfaces and perform a forward search for an optimal alignment of local descriptor tuples. Combinatorial multi-part reassembly is handled via graph optimization, by incrementally clustering and merging best-matching fragments and testing for interpenetration. Both methods are generic, i.e. not restricted to cultural heritage objects or particular types of shapes and require no online user input.

In the part-to-whole matching, fragments are compared against an objective goal shape and their pose is primarily guided by their best-fit registration on the latter. A prominent example is the method by Wei et al. [WYLL11], where the reassembly is specifically geared towards the reassembly of human skulls models. In contrast to the multi-part registrations above, the approach also operates on fragments that do not share breaking edges. Local shape descriptors are extracted from the fragments and a pose is estimated by feature registration and distance minimization. Subsequently, the alignment of the fragments is globally optimized by shifting them to eliminate fragment intersection. Finally, missing patches between the fragments are synthesized using non-rigidly deformed patches of the skull template. This approach is reported to work without any need for online user input.

3.3. Inpainting

Very similar to 2D image inpainting, 3D Shape Inpainting synthesize local object details in smaller defective areas based on the features found in the surrounding or remote but similar intact areas. Kawai et al. [KSY08, KSY09] proposed a method that does not require any user input after candidate target areas to reuse are identified from the shape, and thereby improves over previous approaches such as [BSK05, BGK06]. In the method, first the boundary of a whole where an object needs to be repaired, is detected. All vertices on the boundary are connected to a new vertex that is placed in their center of gravity. For determining features that are to be matched and reproduced by the hole-filling, the geodesic vicinity around the hole is chosen as a target area. In the next step, points are iteratively added and deleted from the solution to match the point density and distribution of the target region. In a final step, a cost-function that compares the principal curvature of vertices on target and source area is minimized by shifting the inserted vertices.

Similarly, Harary et al. [HTG14] proposed to use self-similarity to inpaint holes in surfaces. When a hole is iden-

tified in the surface, a rough initial triangulation is produced. Subsequently, the algorithm computes Heat Kernel Signatures [SOG09] for every vertex in the mesh and a new descriptor for small regions is proposed by averaging the HKS's of vertex in the region. Then, for each vertex in the initial region, the nearest neighbor patch in the HKS space is used as hypothesis to inpaint. Finally, a blending step is performed to combine all the patches for the points in the initialized region. In contrast to the Inpainting method from Kawai et al., this approach is not limited to source areas that are located in direct vicinity of the target hole.

3.4. Symmetry-based completion

Many man-made objects are roughly mirror, point or radially symmetric (*global symmetry*). Even more man made objects can be decomposed into segments that are in some way symmetrical on their own (*partial symmetry*). Almost all man made objects can be regarded as partially symmetric if instead of euclidean (*extrinsic*) distances, geodesic (*intrinsic*) distances of features on the segment's surface are considered. There are numerous publications [MPWC12] on the detection of global, partial as well as extrinsic and intrinsic symmetry and their application to areas such as shape segmentation, recognition, compression, missing parts detection and reconstruction. However, rather few methods are actually robust to cope with larger missing parts and the often imperfect (i.e. inexact) symmetries found on digitized CH objects.

Related to the skull reassembly method of [WYLL11] (see also Section 3.2), in [LYW*11] missing skull fragments are synthesized according to their counterparts across the global, extrinsic symmetry plane of the skull. First, a uniform sampling of the previously aligned fragments is obtained. In a second step, the sampled vertices are paired according to their local principal curvature and shape diameter function. Reportedly, this combination is robust against missing large parts. Then the transformation between normals and direction of minimum and maximum curvature of matched pairs are computed. Each of these transformations corresponds to a pair-specific ideal symmetry plane. To determine the global plane, these transformations are clustered using a mean-shift algorithm. As the reassembly proposed in [WYLL11], this approach does not require any online user input. The approach assumes the existence of exactly one global extrinsic symmetry plane and has thus a limited versatility. Besides the hard-coded assumption about symmetry characteristics, there is actually no need for external data. In Jiang et al. [JXCZ13] an approach is published that uses intrinsic symmetry for completion of defective models. Symmetry detection is based on an a priori skeleton extraction. Once the skeleton is computed, the algorithm finds symmetric correspondences between the bones. Regions around the skeleton can be completed by copying them from the corresponding symmetric bones. While this method is suitable for repairing medium sized defects, it would not

be suitable if large parts are missing since the skeleton would change dramatically.

3.5. Comparison of Similarity-Based Repair Methods

As the discussed similarity-oriented repair methods have been developed for different use cases and applications, it is hard to directly compare them. Yet, we can compare them according to qualitative criteria regarding their expected applicability for CH object repair. Specifically, we consider the following comparison criteria along which we can discuss the identified methods:

large parts repair estimates the effectiveness of a method for reconstruction of missing large parts.

small parts repair estimates the effectiveness of a method for reconstruction of areas that are affected by small scale defects such as e.g. chipped of fragments, fissures, scratches or small-scale decay.

input actions classifies the number of user input actions that are required during the repair of an object. e.g. one or more input actions are required per input object, fragment or defect

input complexity classifies the methods according to their most complex input. Low complexity resembles e.g. the choice of a solution out of a small list of candidates, whereas manual alignment or CAD-like geometry input would be classified as high complexity.

external data summarizes a methods overall need for external data, such as e.g. Shape Repositories or other object specific parameters. External data, when provided offline, can often reduce the amount or complexity of user inputs during the repair process, however this external data usually cannot be computed automatically.

versatility limits enumerates the most limiting aspects of a method according to their application to generic defective CH objects.

Table 1 summarizes the discussed methods and comparative criteria. Note that these are our own assessments based on a literature analysis but more substantiation and also, empirical comparison of the methods to each other is a work yet to be done.

4. A Workflow for Automated Reconstruction

When examining the comparison of the repair methods above, one can easily observe, that none of the methods is simultaneously highly versatile and also well suited for repairing both small-scale defects and missing large parts. This leads to the conclusion that, to obtain a more versatile approach for automated reconstruction, several methods have to be combined. However, special care has to be taken that their sequence does not degrade the quality of results of the individual methods and that expensive intermediate results are reused where possible. Figure 3 summarizes our proposed combination and sequence of mesh processing and repair methods. While every step could be executed without

the need for any online user input, several steps require external data that has to be provided a priori. For the workflow to produce results, not all completion steps have to contribute repairs in every case. In fact, for many objects, at least one of Reassembly, Symmetry- or Template-based Completion cannot always be applied successfully. If necessary, robustness and solution quality could optionally be improved by relying on additional user input. E.g. every step could potentially be executed multiple times with slightly different parameter sets to produce a list of solution candidates from which the user could choose interactively. Alternatively, the user could modify parameters such as thresholds directly on the fly[§], while the corresponding processing results were visualized for feedback.

Our contribution is the conceptual combination and sequence of the methods within the overall workflow. In principle, for each processing step it is in many places possible to identify alternate methods that implement similar functionality. However, to show the feasibility of the workflow, we sketch approaches for each step and also show first, preliminary results in the following sections.

4.1. Surface Classification




Figure 4: Results of automated classification of column-base fragments into exterior surface (blue) and contact surface (red) using the approach from [AMP14], classification confidence is encoded by saturation

The first step in the workflow is an automated surface classification. Its results are used to improve the accuracy of subsequent processing steps. While it does not alter the geometry in any way, the step aims to detect breaking edges based on curvature and vertex normals. In the case of CH Objects made of e.g. marble or limestone, breaking edges can be identified by continuous occurrence of certain high-frequency, small-scale features. We propose to use the approach presented in [AMP14]. Preliminary classification results for fragments of a column base are shown in Figure 4.

[§] e.g. by dragging simple sliders

repair type / method	large parts	small parts	input actions	input complexity	external data	versatility limits
Modeling Funkhouser et al. [FKS*04]	+	-	several per defect	3d query boxes, ...)	fully detailed models	amount of external data & input
Modeling Lee et al. [LF08]	+	-	several per defect	2d sketches, ...	fully detailed models	amount of external data & input
Reassembly Papaioannou et al. [PK03]	+	-	none	n/a	none	shared breaking edges
Reassembly Wei et al. [WYLL11]	+	+	none	n/a	none [†]	skull fragments only
Inpainting Kawai et al. [KSY09]	-	+	one per defect	defect selection	none	source area adjacent to target, fading quality for larger targets
Inpainting Harary et al. [HTG14]	-	+	one per defect	defect selection	none	fading quality for larger targets
Symmetry [‡] Li et al. [LYW*11]	+	+	none	n/a	none	dominant global symmetry plane
Symmetry Jiang et al. [JXCZ13]	o	+	none	n/a	none	not robust to larger missing parts

Table 1: Comparison of related repair approaches according to their applicability for reconstruction of CH Objects**Figure 3:** Overview of our workflow for automated repair of defective CH objects

The approach first segments the shape according to differences of adjacent vertex normals. In a second step, adjacent segments are merged. Finally, merged segments are classified based on a semi-supervised machine learning approach. While the surface classification does not rely on online user input, it can be adapted to different materials by appropriate training.

4.2. Fragment Reassembly

In case the input consists of several disconnected parts, the Fragment Reassembly step targets their reconstruction. Any multi-part registration technique that prevents fragment intersection, such as the ones presented in section 3.2, can be used. For our repair pipeline, we are developing a custom rigid registration method that penalizes penetrations at runtime, so that no post-adjustment of the pose is required. The multi-part reassembly is expressed as a node clustering problem and is addressed by finding a minimum spanning tree, or forest, when the input fragments correspond to multiple reassembled objects. Preliminary results are shown in Fig. 5.

4.3. Symmetry-based Completion

The third step attempts to leverage symmetry to synthesize missing parts. Our current implementation is similar in spirit

**Figure 5:** Fragment reassembly results. The multi-part registration results in a minimum spanning tree of the best matching links and their respective relative transformations.

to [MGP06] in that a set of local features are used to detect symmetry planes.

First, a simplified model is computed from the input by applying a Poisson Reconstruction [KBH06] that combines aligned fragments (from the previous step), fixes minor mesh defects and reduces the number of vertices. Second, Heat Kernel Signatures [SOG09] are computed on the simplified model. Then, for a given set of random samples, our algorithm computes candidate planes using correspondences in the feature space. Subsequently the planes are organized as a 4D point collection. Finally, the final set of candidate planes




Figure 6: First results of surface completion by detection of multiple global extrinsic symmetry planes: input shape (top left), highlighted by color the synthesized parts according to a vertical plane (top right), diagonal plane (bottom left) and horizontal plane (bottom right)

is detected by clustering the 4D points using the DBSCAN algorithm.

After initial symmetry planes have been established, surface patches are mirrored and aligned to existing parts by performing an ICP [RL01] based registration (Figure 6). After matching, plane detection, mirroring and registration have been performed on the simplified model, the computed transformations are also applied to the initial high-resolution model. In case any parts have been mirrored in this step, the simplified model and corresponding HKS have to be re-computed again before they are passed to the next step. Our implementation does not yet rely on the previously established surface classification. However it is expected that an exclusion of breaking edges from symmetry detection and registration will increase accuracy and robustness of the approach[¶]. A revised technique to the approach outlined here will be published in [SGS14] in greater level of detail along with further preliminary results.

4.4. Template-based Completion

The fourth step in the workflow attempts to retrieve and align a template from an repository to the input object. While Symmetry-based Completion might often be the most effective way of synthesizing missing parts, there are cases where the input object is not symmetric, does not provide highly distinctive local features or is too severely damaged

[¶] This could e.g. be implemented by a weighting scheme to feature distances that incorporates their distance to nearby breaking edges, a similar scheme is used in section 4.4

for symmetry detection. To address this, the Template-based Completion attempts to retrieve and align matching template shapes from a repository to understand the overall shape. (Figure 7). The readily available HKS (see section 4.3) are




Figure 8: Populating the Template Repository: automatically generated parameter configurations for a single, parametric CSG description of a column base.

reused to compute Scale Invariant Heat Kernel Signatures (SI-HKS) [BK10]. SI-HKS have been reported to deliver state of the art retrieval performance when used in conjunction with a spatially sensitive Bag of Features (SS-BoF) approach as described in [BBGO11] while at the same time being insensitive to non-rigid, local scalings. The template repository is populated by a set of primitive shapes and by automatically generated instances of parametric surface descriptions (see Figure 8). Although templates are not required to contain local details, the repository can still be populated by arbitrary watertight templates. To increase retrieval performance, all templates are automatically re-meshed off-line by a Poisson Reconstruction to roughly match the expected resolution and sampling density of the simplified models

The SS-BoF from [BBGO11] is extended by a weighting scheme and an additional composite distance function to exploit the previously established surface classification (section 4.1). The weighting scheme is applied to the computation of a second global descriptor of the query object. Identical to the first descriptor, it resembles a matrix that encodes the occurrence frequency of nearby combinations of feature vectors (i.e. *geometric words*) from the vocabulary^{||}. Every geometric word that is found close to^{**} a breaking edge has its weight reduced accordingly. After appropriate normalization, the resulting global descriptor is more insensitive to areas around breaking edges.

^{||} the vocabulary is extracted by k-means clustering of the SI-HKS of all templates in the repository, the SI-HKS of the query object are quantized to their nearest neighbor in the vocabulary

^{**} due to the nature of the HKS and its SI-HKS transform, a diffusion distance is expected to lead to better results than simple euclidean distance




Figure 7: Template-based Completion, Merging and Inpainting of a reassembled Column-Base

With the previous approach, it is possible to determine a similarity score between an input shape and target shapes in the repository. However, scaling and alignment (to find the overall shape) need to be applied as an additional step. In this direction, we are also developing a joint formulation to deal with similarity assessment and scaling/alignment in one step while also being robust against missing of large parts. The idea is to formulate the similarity evaluation as an optimization problem in some descriptor space such that small changes in the geometry induce small changes in the features. We follow the methodology of [OLGM11] to formulate a non-linear optimization problem that is able to find the best transformation that fits a target object to the query as much as possible. First results of this technique, optimizing similarity according to the global shape distribution descriptor parametrized by non-isometric scaling operations can be seen in Fig.9.




Figure 9: First results of our alignment approach: After initial pose estimation, a template shape (red) is scaled non-isometrically until optimal similarity to reassembled fragments of a tombstone (blue) is reached.

4.5. Merging & Annotation

The previous three processing steps aligned fragments and synthesized or retrieved additional shapes. In contrast to the simplified model, in which the reassembled fragments and the mirrored parts (section 4.3) were already merged in low resolution, the aligned high-resolution fragments and generated parts are spatially aligned but still completely disconnected, separate shapes. In the merging step, all aligned parts are combined into a single watertight mesh. Merging could be implemented by a Poisson Reconstruction (a higher octree depth should be used to preserve small scale features) or by directly computing a boolean union of the volumetric

representations^{††}. Subsequently, the resulting vertices have to be annotated by their distance to the various encountered surface types that were used for the merge (i.e. surface areas classified as defective, exterior, stemming from a template shape or having been synthesized based on symmetry).

4.6. Automated Inpainting

In this step, we aim at transferring geometry and additional information such as textures to the synthesized data. Our idea is to use self-similarity between HKS signatures to determine similar local regions which will be used as basis for the transfer. The use of heat diffusion has been successfully applied for inpainting of holes in [HTG14]. Although our problem is a bit different due to the nature of the defects we want to recover, we believe that self-similarity is still a useful tool to this stage. The result of the inpainting step is regarded as the completed object.

4.7. Missing Parts Computation

The final step in our workflow can be used to export the predicted missing parts of the object. These could in turn be supplied to a 3D printer to produce suitable parts for a physical reconstruction of the CH object. Given the previously computed result in the workflow, missing parts can basically be obtained by computing the difference of the completed object (i.e. result of the inpainting step) and the aligned fragment set (i.e. result of the reassembly stage). To address limitations of current 3D printers, a subsequent post-processing step is required to discard parts with too small diameter.

5. Evaluation Methodology

The repair workflow described in Section 4 is a complex setup. We are currently experimenting with different implementation alternatives for the various repair workflow steps, considering different example data. While our system is not fully operational yet, we devised a set of criteria by which we can compare alternative instantiations and identify effective combinations. We next outline the devised evaluation scheme.

Options for evaluation can be categorized into more user-

^{††} e.g. using a readily available, efficient volumetric data structure as described in [Mus13]

and application-centric aspects and technical, quantitative measures. While many of the criteria can be used for comparing an entire workflow implementation, all of them can be used to compare and evaluate alternative implementations of individual processing steps. Prior to evaluating them, a set of defective CH objects for a certain application scenario has to be defined.

5.1. General Criteria

Plausibility of Results assesses the overall practical benefit of a workflow implementation for reconstruction of a defined set of defective CH objects. As plausibility can hardly be defined in a formal way, this criterion should be obtained by with the aid of CH domain experts. Automatically reconstructed objects, digitizations of mostly intact CH objects or manual reconstructions could be rated for their plausibility by experts in double blind tests. The rating differences across all three categories provide an indication on the applicability of the tested implementation for the provided object types and their defect characteristics.

Versatility Limits assess the aspects of a workflow that limit its applicability to generic CH objects on a qualitative level. It roughly resembles the criteria used for comparing related work in section 3.

External Data Requirements assess the amount and properties of required external data on a qualitative level. High external data requirements can heavily impact practical use if a high number of different object types is targeted. In contrast to the criterion applied in section 3, the data requirements can be obtained in a more precise way if a fixed set of input objects is used.

User Input Requirements assess amount and complexity of required user input. If a fixed set of input objects are used, the criteria could be measured in a quantitative way by e.g. counting clicks and keystrokes. In context of a user study, a set of reconstruction tasks can be defined where the time required for completion is tracked^{††}.

5.2. Technical Measures

General Effectiveness is the most basic measure for the fraction of a fixed set of input data that can be correctly addressed by a workflow implementation or an individual processing step. Correctness includes formal aspects of the result such as manifoldness of output meshes. In context of a user study, a threshold on plausibility ratings could also be taken into account. For individual processing steps, ground truth can be defined (e.g. predefined symmetry planes, alignments for assembled parts, relevant template shapes)

that has to be met by a method within certain boundaries.

Efficiency assesses the consumption of resources such as memory and processing time for a fixed set of input objects. Special care has to be taken to exclude time spent waiting for user input (if required).

Robustness refines the general effectiveness measures by extending the input set by objects with artificial defects. E.g. additional fragments could be omitted or unrelated fragments could be added for reassembly, defects of mesh structure can be inserted, noise could be added or small scale features could be obfuscated by smoothing.

Retrieval Performance: Surface Classification, Reassembly as well as Symmetry-based and Template-based Completion all rely on extraction and matching of local features. While well known measures such as Precision, Recall and Mean Average Precision could readily be applied to the retrieval of complementary parts and templates, they could be adapted for the Retrieval of corresponding vertices, if suitable surface annotations are provided as ground truth.

Geometric Accuracy measures the undesired change of geometric features during processing. For workflow implementations, this could e.g. be achieved by simply aggregating the distance of vertices on intact input patches (provided as ground truth) to their nearest neighbors on the output shape. For individual processing steps such as e.g. the Reassembly, the volume of fragment intersections could also be used.

Scalability measures could be obtained experimentally for individual steps. E.g. by extending the input set with up-sampled meshes or using larger fragment sets by adding large unrelated fragments while tracking resource consumption (see *efficiency* criterion).

6. Conclusions and Future Work

After categorizing defect types that are encountered in CH objects and establishing a taxonomy and comparison of related similarity driven repair methods, we proposed a workflow that combines and extends several methods while simultaneously reducing the need for online user input. We provided promising first results and recommendations how to evaluate our workflow implementation as well as individual steps. In Future Work, a benchmark that encompasses a range of defective CH objects needs to be created. Furthermore, a user study with CH domain experts could verify the usefulness of our approach.

Besides evaluation, several extensions to individual steps could lead to further improvements. Most prominently, the support for an automated segmentation step that divides the reassembled model into more primitive parts before they are processed by the Template-based Completion could reduce

^{††} also note the tracking of processing time for the efficiency criteria

the required number of template shapes in case of complex non-symmetric objects. External data needs and quality of results could also be improved by support for more generic non-rigid alignment. Symmetry-based completion could benefit from support of radial or point-based symmetry. Also a completion based on intrinsic symmetry could be useful if performed after the template shape retrieval or directly integrated into the Inpainting step.

Acknowledgments

We would like to 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^{§§}.

References

- [ACK13] ATTENE M., CAMPEN M., KOBELT L.: Polygon mesh repairing. *ACM Computing Surveys* 45, 2 (2013), 1–33. 2, 3
- [AMP14] ANDREADIS A., MAVRIDIS P., PAPAIOANNOU G.: Facet Extraction and Classification for the Reassembly of Fractured 3D Objects. In *EG 2014 - Posters* (2014), Paulin M., Dachsbacher C., (Eds.), Eurographics Association, pp. 1–2. 5
- [BBGO11] BRONSTEIN A., BRONSTEIN M., GUIBAS L., OVSIJANIKOV M.: Shape google: Geometric words and expressions for invariant shape retrieval. *ACM Trans. on Graphics* 30, 1 (2011), 1–20. 7
- [BGK06] BENDELS G., GUTHE M., KLEIN R.: Free-form modelling for surface inpainting. In *Proc. of the 4th Int. Conf. on Computer Graphics, Virtual Reality, Visualisation and Interaction in Africa (Afrigraph 2006)* (2006), pp. 49–58. 4
- [BK10] BRONSTEIN M., KOKKINOS I.: Scale-invariant heat kernel signatures for non-rigid shape recognition. In *Proc. of Conf. Computer Vision and Pattern Recognition* (2010), pp. 1704–1711. 7
- [BSK05] BENDELS G. H., SCHNABEL R., KLEIN R.: *Fragment-based Surface Inpainting*. Tech. rep., 2005. 4
- [CKGK11] CHAUDHURI S., KALOGERAKIS E., GUIBAS L., KOLTUN V.: Probabilistic reasoning for assembly-based 3D modeling. *ACM SIGGRAPH 2011* 30, 4 (2011), 1. 3
- [FKS*04] FUNKHOUSER T., KAZHDAN M., SHILANE P., MIN P., KIEFER W., TAL A., RUSINKIEWICZ S., DOBKIN D.: Modeling by example. *ACM Trans. on Graphics* 23, 3 (2004), 652–663. 3, 6
- [HFG*06] HUANG Q.-X., FLÖRY S., GELFAND N., HOFER M., POTTMANN H.: Reassembling fractured objects by geometric matching. In *ACM SIGGRAPH* (2006), SIGGRAPH '06, ACM, pp. 569–578. 4
- [HTG14] HARARY G., TAL A., GRINSUN E.: Context-based coherent surface completion. *ACM Trans. on Graphics* 33, 1 (2014). 4, 6, 8
- [JXCZ13] JIANG W., XU K., CHENG Z.-Q., ZHANG H.: Skeleton-based intrinsic symmetry detection on point clouds. *Graphical Models* 75, 4 (July 2013), 177–188. 4, 6
- [KBH06] KAZHDAN M., BOLITHO M., HOPPE H.: Poisson surface reconstruction. In *Proc. of Symposium on Geometry Processing* (2006), pp. 61–70. 6
- [KCKK12] KALOGERAKIS E., CHAUDHURI S., KOLLER D., KOLTUN V.: A probabilistic model for component-based shape synthesis. *ACM Trans. on Graphics* 31, 4 (July 2012), 1–11. 3
- [KHS10] KALOGERAKIS E., HERTZMANN A., SINGH K.: Learning 3D mesh segmentation and labeling. *ACM Trans. on Graphics* 29, 3 (2010). 3
- [KS03] KAMPEL M., SABLATNIG R.: An automated pottery archival and reconstruction system. *Journal of Visualization and Computer Animation* 14, 3 (2003), 111–120. 1
- [KSY08] KAWAI N., SATO T., YOKOYA N.: Surface completion by minimizing energy based on similarity of shape. In *IEEE Int. Conf. on Image Processing* (2008), pp. 1532–1535. 4
- [KSY09] KAWAI N., SATO T., YOKOYA N.: Efficient surface completion using principal curvature and its evaluation. In *Proc. Int. Conf. Image Processing* (2009), pp. 521–524. 4, 6
- [LF08] LEE J., FUNKHOUSER T.: Sketch-based search and composition of 3D models. In *Proc. Fifth Eurographics Conf. on Sketch-Based Interfaces and Modeling* (2008). 3, 6
- [LGZ*13] LIN H., GAO J., ZHOU Y., LU G., YE M., ZHANG C., LIU L., YANG R.: Semantic decomposition and reconstruction of residential scenes from lidar data. *ACM Trans. Graph.* 32, 4 (July 2013), 66:1–66:10. 1
- [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. 4, 6
- [MGP06] MITRA N., GUIBAS L., PAULY M.: Partial and approximate symmetry detection for 3D geometry. *ACM Trans. on Graphics* 25, 3 (2006), 560–568. 6
- [MPWC12] MITRA N. J., PAULY M., WAND M., CEYLAN D.: Symmetry in 3D Geometry: Extraction and Applications. *Computer Graphics Forum* 32, 6 (2012), 1–23. 4
- [Mus13] MUSETH K.: VDB: High-resolution sparse volumes with dynamic topology. *ACM Trans. on Graphics* 32, 3 (2013). 8
- [OLGM11] OVSIJANIKOV M., LI W., GUIBAS L., MITRA N. N. J.: Exploration of continuous variability in collections of 3D shapes. *ACM Trans. on Graphics* 34, 4 (2011), 1. 8
- [PK03] PAPAIOANNOU G., KARABASSI E.-A.: On the automatic assemblage of arbitrary broken solid artefacts. *Image and Vision Computing* 21, 5 (2003), 401–412. 4, 6
- [PKT01] PAPAIOANNOU G., KARABASSI E.-A., THEOHARIS T.: Virtual archaeologist: Assembling the past. *Computer Graphics and Applications, IEEE* 21, 2 (2001), 53–59. 3
- [RL01] RUSINKIEWICZ S., LEVOY M.: Efficient variants of the ICP algorithm. *Proc. 3rd Int. Conf. on 3-D Digital Imaging and Modeling* (2001), 145–152. 7
- [SGS14] SIPIRAN I., GREGOR R., SCHRECK T.: Approximate symmetry detection in partial 3d meshes. *Computer Graphics Forum* (to appear in 2014). 7
- [SOG09] SUN J., OVSIJANIKOV M., GUIBAS L.: A Concise and Provably Informative Multi-Scale Signature Based on Heat Diffusion. *Computer Graphics Forum* 28, 5 (2009), 1383–1392. 4, 6
- [UF11] ULLRICH T., FELLNER D. W.: Generative object definition and semantic recognition. In *3DOR* (2011), pp. 1–8. 1
- [WYLL11] WEI L., YU W., LI M., LI X.: Skull Assembly and Completion Using Template-Based Surface Matching. *Proc. Int. Conf. on 3D Imaging, Modeling, Processing, Visualization and Transmission*, 1 (May 2011), 413–420. 4, 6

§§ <http://presious.eu>