

From Reassembly to Object Completion - A Complete Systems Pipeline

GEORGIOS PAPAIOANNOU, Athens University of Economics and Business
TOBIAS SCHRECK, Graz University of Technology
ANTHOUSIS ANDREADIS, Athens University of Economics and Business
PAVLOS MAVRIDIS and ROBERT GREGOR, Graz University of Technology
IVAN SIPIRAN, Pontificia Universidad Católica del Perú
KONSTANTINOS VARDIS, Athens University of Economics and Business

The problem of restoration of broken artefacts, where large parts could be missing, is of high importance in archaeology. The typical manual restoration can become a tedious and error-prone process, which also does not scale well. In recent years, many methods have been proposed for assisting the process, most of which target specialized object types or operate under very strict constraints. We propose a digital shape restoration pipeline consisting of proven, robust methods for automatic fragment reassembly and shape completion of generic three-dimensional objects of arbitrary type. In this pipeline, first we introduce a novel unified approach for handling the reassembly of objects from heavily damaged fragments by exploiting both fracture surfaces and salient features on the intact sides of fragments, when available. Second, we propose an object completion procedure based on generalized symmetries and a complementary part extraction process that is suitable for driving the fabrication of missing geometry. We demonstrate the effectiveness of our approach using real-world fractured objects and software implemented as part of the EU-funded PRESIOUS project, which is also available for download from the project site.

CCS Concepts: •Computing methodologies → Shape inference; Shape analysis; •Applied computing → Archaeology;

Additional Key Words and Phrases: Reassembly, object completion, symmetry detection, registration, segmentation, shape classification

ACM Reference Format:

Georgios Papaioannou Tobias Schreck, Anthousis Andreadis, Pavlos Mavridis, Robert Gregor, Konstantinos Vardis and Ivan Sipiran. 2016. From Reassembly to Object Completion - A Complete Systems Pipeline. *ACM J. Comput. Cult. Herit.* 10, 2, Article 8 (2017), 23 pages.
DOI: 0000001.0000001

1. INTRODUCTION

The physical restoration of cultural heritage (CH) objects from fragments found at excavation sites is a time consuming and difficult task, especially for large objects or large collections of fragments, given also the fact that the pieces may be dispersed in different physical collections. In order to provide cues

This work was supported by the EC FP7 STREP Project PRESIOUS, grant no. 600533. The examples shown as well as the software used for most of the pipeline can be found at the PRESIOUS project web site: <http://presious.eu>

Corresponding author's address: G. Papaioannou, 76 Patission Str., 10434, Athens, Greece; email: gepap@aubg.gr;
T. Schreck email: tobias.schreck@cgvtugraz.at; A. Andreadis email: anthousis@gmail.com; P. Mavridis email: pmavridis@gmail.com; R. Gregor email: gregor@tugraz.at; I. Sipiran email: isipiran@pucep.edu.pe; K. Vardis email: kvardis@hotmail.com

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2017 ACM. 1556-4673/2017/-ART8 \$15.00
DOI: 0000001.0000001

for the study and presentation of a restored object but also for strengthening its structural integrity, reassembled objects often undergo a repair process, where missing parts are speculated and physically created. The digital counterpart, *virtual object restoration*, has received significant research interest in the past years, mainly with regard to specialized object types, such as frescos and pottery. Moving to the virtual, computer-assisted domain provides numerous advantages, including the access to digital assets from remotely located physical finds and the ability to easily manipulate and explore 3D shapes, whose physical counterparts may be hard to handle. Above all, it benefits from the exploitation of robust and fast algorithms for computing and exhaustively testing hypotheses, even extreme ones, at a large problem scale.

In general, the computational restoration of structures is dramatically hindered by the deteriorated state of the parts involved. Fragments are typically exposed to weathering, erosion and physical stress, partially and sometimes unevenly damaging their surfaces. Fractured areas are often smoothed out, thus significantly reducing the effectiveness of local features in the tasks of shape comparison or registration. Furthermore, missing parts in the case of generic objects tend to result in room for more interpretations, greater ambiguity and sometimes make it impossible to determine the exact shape and number of the original objects from a collection of fragments.

In this paper we propose a shape restoration pipeline consisting of appropriate methods for automatic fragment reassembly and shape completion. The methodological approach followed and the implemented systems target generic three-dimensional shapes, such as broken architectural elements, ornaments, statues etc., but can also be used in constrained special cases such as flat or symmetrical objects. We demonstrate the effectiveness of our approach using real case studies of fractured objects, based on detailed precision evaluation experiments. Our key contributions are:

- A complete pipeline of working tools for performing a) reassembly of generic objects using pure geometric priors and b) object completion and missing part synthesis. For the first part, we employ robust registration algorithms to solve the 3D "puzzling" problem, under the presence of noise, data partiality and erosion. For the object completion, two data-driven approaches are used to synthesize missing parts based on symmetry, one suitable for point clouds or potentially open (partial) meshes and one better suited to objects with approximate or very partial planar symmetry.
- A fully implemented set of tools for performing the reassembly and symmetrical expansion of objects, available online.
- A thorough validation of the proposed methodology on real broken CH objects.

The remainder of this paper is structured as follows: Section 2 briefly summarizes most relevant work to our restoration system, Section 3 explains the steps of the pipeline, Sections 4 and 5 provide more details about the actual algorithms used, while indicative case studies and evaluation of results are given in Section 6.

2. BACKGROUND AND RELATED WORK

Before examining prior art in the virtual restoration domain or providing details about our method, it is informative to categorize the objects to be restored according to some broad geometric prior classification criteria: thickness, symmetry and presence of features on the intact surfaces (Figure 1). This taxonomy helps identify cases that can be efficiently handled by specialized algorithms (see specialized methods in Section 2) or exploit rich content, where available. For example, relatively thin surfaces can be treated as 2D embeddings in three-dimensional space, substituting the notion of a break surface with that of a break line in the reassembly stage. Symmetrical objects, such as vessels, can effectively



Fig. 1. A high-level taxonomy of broken artefacts with respect to geometric priors with some examples.

exploit the axial symmetry to reduce the optimization parameters for the pose estimation of the fragments (e.g. potsherds) and allow the symmetrical expansion of the resulting fragment clusters in order to complete them.

2.1 Reassembly

In computational archaeology, typically the *reassembly* or *puzzling* problem is described as the process for the identification of potentially fractured parts/regions of an object, the search for corresponding pieces within a fragment collection and finally the clustering and pose estimation of multiple parts that result in a virtual representation of (partially) reassembled objects. For non-flat, three-dimensional fragments, the process typically starts with the digitization (e.g. 3D scanning) of the physical fragments and continues with the pre-processing of the fragment geometry in order to extract usable geometric priors, such as break curves or surfaces, features and global descriptors. Subsequently, pair-wise combinations of the fragments are tested for alignment and a matching score is computed. This step usually begins with a global registration process that examines the solution search space for a good but rough alignment, which in turn initiates a local registration process in order to refine the solution. The complete set of pair-wise results subsequently drives the multi-part clustering and alignment, where complete objects are formed by finding the global position for each fragment. Existing methodologies can be split in three categories based on the dimensionality of the problem and the assumptions made for specific data collections.

The first category addresses the problem in two-dimensions, as for certain flat objects it is safe to make this assumption without affecting the quality of the solution. Two representative methods of this category, described in [Kong and Kimia 2001] and [da Gama Leitao and Stolfi 2002], use elastic curve matching, while others rely on textural information [Sağıroğlu and Erçil 2005] or a combination of both [Zhang and Li 2014]. Recent approaches also target cases with missing data. Koller and Levoy [2006] in order to address heavily eroded fresco fragments, use manually placed markings on 2D representations of the fragments. Huang et al. [2013b] use a similar idea for images, by extrapolating and matching salient feature curves, while Poleg and Peleg [2012] extrapolate whole image fragments and use the generated parts in order to perform the alignment with other pieces.

The second category includes methods that operate in a restricted 3D space by either reducing the search space parameters or by constraining the parts on existing template shapes. Typically these methods focus on fresco fragments and potsherds. The most representative work on fresco fragments is by [Brown et al. 2008], where the perimeter of each fragment (its *ribbon*) is uniformly sampled and registration operations are performed on a 2D plane using contour matching. In a similar fashion, Sanchez and Vidal [2012] discretize the fracture ribbon using orthographical projection on the Graphics Processing Unit (GPU) for a performance boost. Matching operations are performed over a shape descriptor using a hierarchical scheme.

In the case of potsherds, the problem is usually addressed by using estimates of the axis and profile curves with most notable the works of Kampel and Sablatnig [2004] and Willis and Cooper [2004], as well as the more recent approach by Son et al. [2013]. These estimation approaches are known to suffer in terms of robustness, when potsherds are very small or eroded. In the first case, the estimation of the axis may prove unstable, while in the latter, reliable correspondences cannot be easily found for smoothed out break curves. In order to address these issues, Cohen et al. [2013] perform the reconstruction using surface features (weighted curve moments) and template objects. Similar ideas are used for the reconstruction of skulls by Wei et al. [2011] and Li et al. [2011], where predefined template objects guide the assembly.

The final category involves methods that try to solve the reassembly problem in the general 3D case with free-form fragments. Several approaches of this category avoid the difficult global registration step by specifying constraints or enforcing particular matches [Parikh et al. 2007], [Mellado et al. 2010], [Palmas et al. 2013], but here we focus on more automated approaches. Papaioannou et al. [2001] were the first to address the problem by providing a complete reassembly pipeline. Under the assumption of nearly planar surfaces with adequate matching area, fragment matching is performed using GPU-accelerated distance queries. Huang et al. [2006] later presented another complete reassembly solution, that utilized features on the fractured regions in order to register adjacent fragments. Using multi-scale descriptors, feature groups are extracted and used in a forward search registration strategy. Results are refined using local registration both on the contact surface and the boundary edge of the fractured areas, an approach also adopted earlier in [Papaioannou and Karabassi 2003]. Altantsetseg et al. [2014] extract features on the fractured surfaces, which are subsequently grouped to generate curves. Fourier coefficients are computed for each curve and matching is performed using their total energy.

Feature-based approaches require surfaces with rich intrinsic geometric information. However, CH objects frequently suffer from weathering effects, which smooth out such features. Thuswaldner et al. [2009] in order to address this issue, combine the feature-based reassembly scheme of Huang et al. [2006] with the identification of planar surfaces and straight lines in order to align rectangular shaped fragments in the archaeological site of Ephesos. Zhang et al. [2015b] in a more recent approach, combine the use of template objects along with a feature based fractured region matching algorithm. The registration of the fragments with the template object is performed using the SHOT descriptor [Tombari et al. 2010]. For the fractured region matching, the authors extract and register feature curves from the fracture areas. Winkelbach and Wahl [2008] propose a feature-less approach for the pairwise registration problem that maximizes contact area, using a branch-and-bound search heuristic. In our work, in order to match fragments that share significant contact area, we use the feature-less approach of Mavridis et al. [2015b] that uses a progressive global-to-local optimization scheme utilizing an ℓ_p metric with low values of p in order to partially address noise and outliers. An extension of the later was presented in [Andreadis et al. 2015], where feature curves from the intact surfaces of the objects are also introduced in order to address the matching of adjacent fragments with a minimal or even absent matching surface.

2.2 Object Completion

When reassembly has taken place, parts of the object to be restored are often still missing. Also, due to possibly varying degrees of erosion affecting different reassembled parts, the level of preserved detail among clustered fragments may vary. In order to partially or fully compute a prediction of the missing geometry, several shape completion methods may be applied, distinguished by whether they rely on information intrinsic to the given geometry or external data.

2.2.1 Object Completion from Internal Reference. The use of internal references to accomplish a completion of damaged objects is based on the assumption that for man-made objects we can very often find enough information in the object itself to replicate the missing geometry. In a broad sense, the main concept involved is the search for self-similarities. This information can help to find internal relationships such as region replication and symmetries.

Symmetry is one of the main geometric priors used in the repair of CH objects. For example, a method to reconstruct ancient Chinese architecture was proposed in [Li et al. 2014], where global reflective symmetries are detected from a point cloud and used to reconstruct the large missing regions, while local symmetries are used to fill small holes. Likewise, Sipiran et al. [Sipiran et al. 2014] proposed a symmetry-invariant surface function based on the heat diffusion theory to reveal good approximations of symmetric points, even for damaged objects. The symmetric points can generate potential symmetry planes, which can be used to replicate missing geometry. More recently, Mavridis et al. [2015c] proposed an optimization formulation to generate missing geometry based on global symmetries in 3D objects. The method introduces a sparse optimization that enforces the partiality of overlap on replicated geometry.

A second variant of using internal references in object completion is the exploitation of local similarities. Harary et al. [2014a] proposed a method to complete the surface of a 3D object by looking for good correspondences in other regions of the object. The method uses a variation of the Heat Kernel Signature descriptors [Sun et al. 2009] to find correspondences. Subsequently, a careful merging method copies similar patches to produce a fair surface. Under the same approach, Harary et al. [2014b] used the HKS-based descriptors to complete regions, which were previously obtained through a combination of curve analysis and user interaction. Sharf et al. [2004] defined a volumetric approach to find local similarities between regions on the surface of an object. When a good set of correspondences is found, the method applies a combination of iterative rigid and non-rigid transformations in order to complete the missing surface.

2.2.2 Object Completion from External Reference. All of the object completion approaches mentioned so far assume that the missing parts can be generated by solely exploiting the existing geometry, since most man-made objects usually feature multiple occurrences of self-similarities. However, there are cases, where relying on self-similarity alone will not result in admissible objects, since distinguished, unique features and parts may be completely missing. As mentioned by Gregor et al. [2014] a restoration pipeline could thus employ strategies that rely on external reference. The approach outlined there queries a repository of external objects to find and transfer missing geometry to an incomplete cultural heritage object.

Essentially, finding suitable shapes in an external repository can be considered a variant of the partial 3D retrieval problem for which a substantial number of methods have been published so far. However, when incorporating such an approach in a restoration pipeline, several additional difficult problems arise. First, once a suitable external object has been retrieved, its scale needs to be adjusted to match the size of the incomplete object. Second, the parts of the external objects that should be transferred to the incomplete object need to be identified. Once suitable parts have been identified, their transfer to the incomplete object requires a non-rigid alignment to avoid implausible seams. Finally, if the level of detail of the external objects differs from the incomplete object, this could also lead to implausible results, unless local details are transferred from the intact surfaces of the incomplete object to the synthesized parts or vice versa.

Several interactive techniques have been proposed, which address related problems to a certain degree and are generally not specific to the CH domain. Mellado et al. [2016] proposed a semi-automatic registration approach based on the Growing Least Squares descriptor that rectifies the scale of the

source object in the process. However, since the method is intended for the combination of multi-modal or partial scans, it assumes uniform scaling and true similarity of the source and target object. In our case, none of these conditions hold, as the dimensioning and the structure of the repository objects may be significantly different. Lee et al. [2008] rely on a series of user-provided sketches that steers both the process of finding suitable external objects and the global scaling and alignment of the template to match the incomplete object in a plausible way. However the approach does not include a non-rigid registration and also does not compensate for different levels of local detail among incomplete and external objects, both of which impact the method's practical use for real-world collections of CH objects. A problem that is highly related to the latter shortcoming has been solved by Takayama et al. [2011] by proposing an interactive method that enables a user to select corresponding patches between several shapes in order to transfer small-scale local texture between them. Although several published methods address the problem of non-rigid registration between different 3D models [Li et al. 2008], [Zhang et al. 2015a], [Yang et al. 2015], their practical applicability within a completion system based on external references is uncertain.

3. PIPELINE OVERVIEW

The restoration pipeline consists of two broad workflows, reassembly and completion. Taking as input a number of digitized fragments, the reassembly stage produces clusters of properly aligned parts, which can be visualized and exported for further use. The resulting clusters are subsequently forwarded to the object completion stage, where symmetrical expansion of the geometry and fragment fusion can be performed. Also, using the symmetrically predicted shape of the whole object, the missing parts can be calculated and exported. The reassembly workflow is implemented within the Virtual Repair and Measurement Workspace (VRMW) reassembly system, while the object completion steps are performed via individual tools. More details about the individual methods and the software workflows involved are discussed in Sections 4 and 5 and an overview of the stages has also been presented as a poster at SIGGRAPH 2015 [Andreadis et al. 2015]. The software, which was developed as part of the PRESIOUS EU-funded project, is available for download from the project web site: <http://presious.eu>. The restoration pipeline stages and the data flow are illustrated in Fig. 2 and briefly discussed below.

3.1 Preprocessing

The meshes of the (potentially partially) scanned input parts are first pre-processed, in order to approximately classify surface regions as potentially fractured and intact. This is done primarily for robustness, to avoid subsequent pairwise registration operations resulting in trivial, yet undesirable matches

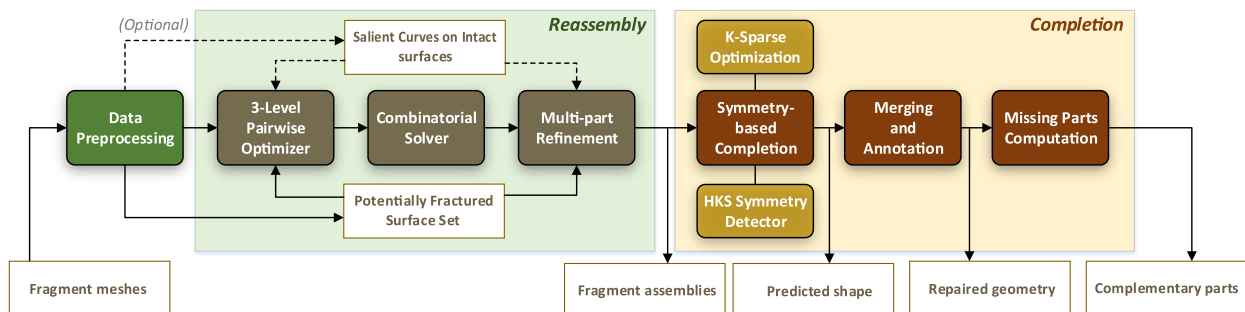


Fig. 2. The object restoration pipeline. Optional data flow paths are indicated with dashed lines.

between flat intact surfaces. Performing the surface registration only on the *potentially fractured surface set* also increases efficiency, as the exclusion of the intact facets from the geometric comparisons reduces both the number of redundant distance queries and the search space.

Feature Curves. After an attempt to automatically assemble fragments into clusters, certain fragment combinations may have a very small usable contact surface or none at all, especially if they have sustained heavy erosion or mechanical damage, or a significant part of their connecting structure is missing. In most cases, this leads to erroneous matches, that cannot be resolved with the comparison of contact surfaces only. For these fragments, we may perform an additional, user-guided processing step, where we extract and exploit more reliable constraints stemming from prominent linear or curved structures on the *intact* regions of the fragments. We detect such structures and extract and extrapolate curves that closely follow their trajectories. These *feature curves* take part in all fragment registration stages, when available, significantly improving the accuracy of alignment and the fidelity and extent of the final assemblies. This idea can be also used for the assembly of structures from individual (and potentially intact) elements belonging to a larger construction (e.g. entire stone blocks) that have inherently smooth, mostly featureless contact surfaces.

3.2 Reassembly

Pairwise optimization. Preprocessed fragments can be selected to participate in a reassembly task. The first part of reassembly is the generation of all combinations of potential matches between fragment pairs. Each identified fractured region on each fragment in the reassembly set is optimally aligned against fractured regions of all other pieces, using its surface geometry and a respective matching score is recorded. Optionally, we can trigger the inclusion of feature curves in the alignment process, as explained in the previous paragraph. The pairwise results are stored in a cache, so that they can be re-used in future experiments. This way, when introducing a new piece to a reassembly, the pairwise registration, which is one of the most costly operations of the reassembly pipeline, is incrementally performed only for the new fragment combinations formed.

Combinatorial solver and multi-part refinement. Having generated all combinations of fragment matches, the multi-part reassembly procedure is then initiated. Here, the set of matching pairs is explored and objects consisting of multiple fragments are generated, forming an acyclic graph of fragments and taking into account inter-part penetrations in the process. The pose of the individual fragments in the resulting clusters is finally refined using all geometric information available (surfaces and feature curves) to produce the final, reassembled objects.

Reassembly workflow. From the user's perspective, the above steps are completely abstracted; With the candidate fragments selected, the user triggers the reassembly task and after the multi-part refinement is complete, the resulting clusters are presented in the 3D workspace (Fig.3 - left). Manual exclusion or enforcement of combinations can take place between reassembly cycles and a new experiment can be initiated. Additionally, a reassembly report is produced with measurements (contact area, average distance) about the fidelity of the proposed clustering as shown in Fig.3 - right.

3.3 Completion

Object Completion. After the reassembly of fragments, there often remain missing parts in the object, since not all original fragments may have been available to the reassembly stage, or since parts may have eroded away. In the absence of information about the complete original shape, one may rely on self-similarity, or similarity to existing complete objects to generate plausible completions. We support two approaches for objects exploiting the idea of self-similarity. The first approach computes

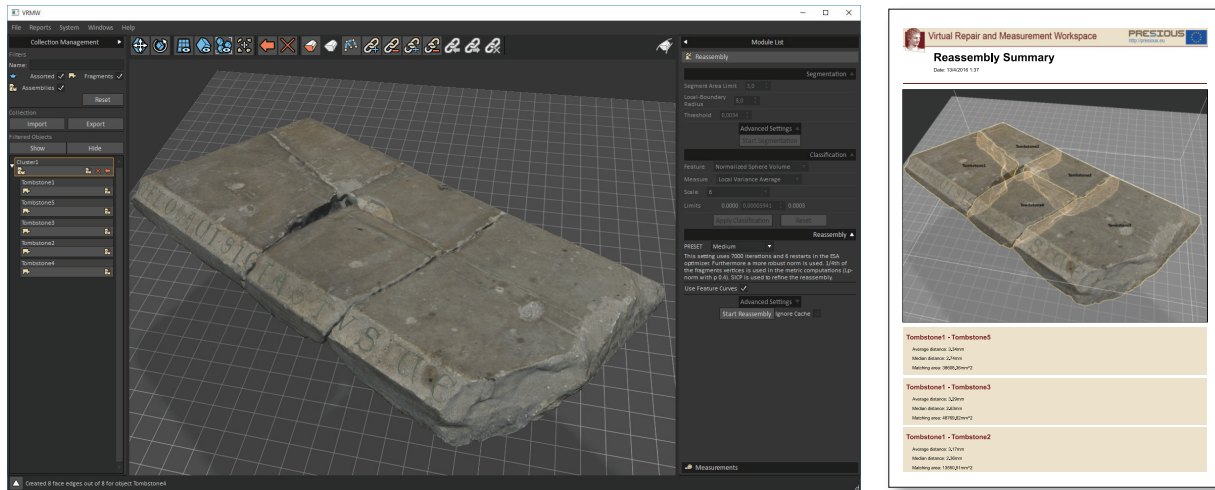


Fig. 3. Reassembly example using the VRMW software. Left: the workspace and reassembly options. Right: first page from the generated reassembly report.

candidate symmetry planes based on local shape features. A voting scheme decides for a best ranked symmetry plane and completes the object by symmetrical geometry replication. Our second approach is applicable in case of general symmetry or when we cannot rely on presence of local features. This approach, which is not restricted to planar symmetries, computes a partial self-registration of the incomplete object with itself, hence repeating local object structures to eventually complete the partial shape. For details please refer to Section 5.1.

Merging and Annotation. The reassembled fragments along with the geometry of the completion step are then merged to form a *finished* object, i.e. the geometry that represents the shape of the original, unbroken object, as closely as possible. The merging step may include filtering operations to smooth out small discontinuities, which remain after the registration steps, such as smaller cracks between reassembled and completed parts of the shape (see Section 5.2). As it is important to retain information about the nature and origin of the geometry, we automatically markup (annotate) the parts of the restored shape as belonging to a) the fragment set, b) detected fracture surfaces, and c) the completed (generated) shape.

Missing Parts Extraction. After successful completion, we may extract the parts missing from the assembly by subtracting the assembled parts from the symmetrical expansion, using appropriate operators (see Section 5.3). The extracted models can be used for physical restoration of the shapes, e.g. by 3D-printing them.

We note that the goal of the completion stage is to propose plausible complete shapes, keeping in mind that the original artifact may have looked differently. Hence, completion is an optional step, useful among others for interpretation-based visualization or physical restoration purposes (see below). The annotation step described above provides relevant information to inform domain-specific analysis on the reassembled and restored shapes. We also note that, while our completion methods work automatically, this stage is intended to be a supervised procedure, since user interaction may be required to refine the completion process or reject implausible completion candidates.

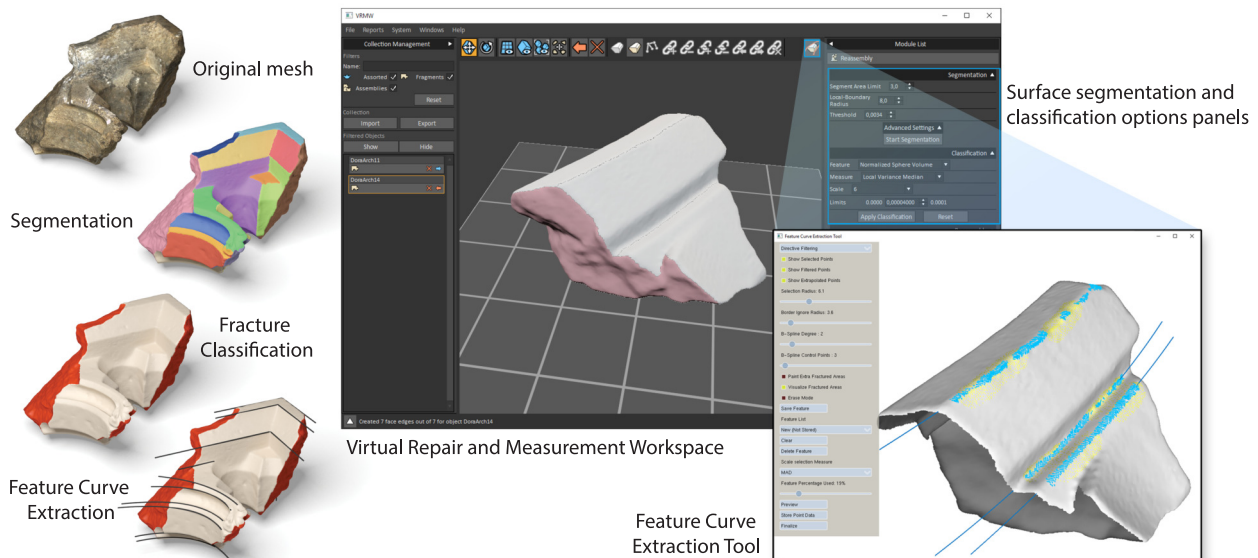


Fig. 4. Fragment preprocessing steps (left) and corresponding workflow in the VRMW software. The two compulsory pre-processing steps (segmentation and classification - top inset) are optionally followed by a feature curve extraction procedure using the feature curve extraction tool (bottom right), to increase the geometric priors available during the registration stages.

4. REASSEMBLY

In this section we provide more details about the three main tasks in the reassembly pipeline, i.e. the fragment preprocessing, the pairwise registration and compatibility testing as well as the multi-part registration and clustering of fragments.

4.1 Fragment Pre-processing

The fragment pre-processing is illustrated in Fig. 4 - bottom). We initially extract distinct facets of the fragmented objects using either a straightforward region-growing segmentation algorithm (default operator) or hierarchical agglomerative clustering ([Andreadis et al. 2014]). In both cases, the deviation of surface normals in local neighborhoods is used as the stopping criterion. The segmentation is re-evaluated in a cleanup step, in order to eliminate small segments, which are assimilated by larger, neighboring regions. The resulting segments are classified as intact or fractured by using a threshold-based approach that relies on measurements of local surface descriptors. The VRMW user can choose between two surface descriptors, *local bending energy* [Huang et al. 2006] and *sphere volume integral invariant* [Pottmann et al. 2009], computed over three different scales. For each descriptor-scale pair, any of the following statistical properties can be used: segment average, median or variance and segment local variance of average and median. The local variance measurements are computed as the variance of the metric in an Euclidean neighborhood around each surface point.

The segmentation of the input fragments and the clustering are performed within the VRMW software (Fig. 4 - middle top). The complete segmentation functionality, along with all options available in the VRMW system, can also be found as a Meshlab plugin, available through the official repository and the PRESIOUS project web site.

As mentioned in the pipeline overview, fragments with heavily damaged or very smooth fracture surfaces provide little or no usable contact support for surface-based registration. If problematic reg-

istrations appear in the reassembly solution, the user can generate a number of feature curves corresponding to salient structures on the intact surfaces of a fragment to further constrain the registration process, by requiring the continuity and registration of these features between adjacent parts [Andreadis et al. 2015].

The feature curves are generated using the Feature Curve Extraction Tool, launched from within the VRMW main application. With the graphical user interface of the tool the user can create groups of features, add or remove mesh points to them using a brush tool directly on the 3D model, set curve estimation parameters and constraints and export the data to VRMW or other auxiliary visualization applications. To define a feature curve, the user creates strokes on the mesh that include or cross part of the desired salient feature. Using mean curvature at multiple scales as a local descriptor, a subset of the selected points, where the local minima and maxima of the descriptor appear, are marked as belonging to a salient feature. Subsequently, we perform a thinning process using the skeletonization method of Huang et al. [Huang et al. 2013a] and the resulting point set is approximated using a parametric curve (2nd degree B-spline) via least-squares fitting.

Points near the curve trajectories with similar descriptor values are assimilated and adjoining curves are iteratively merged and re-evaluated. The resulting feature curves are densely sampled and presented to the user, along with their extrapolated parts beyond the extents of the fragment that will be used in the matching step. Since we are only interested in features that span across multiple fragments, we keep feature curves with at least one end near a fractured facet and discard the rest. Furthermore, each feature curve is associated with the fractured facets it is adjacent to or crosses. During the pairwise matching stage, only feature curves associated with the particular facets under examination are considered.

4.2 Pairwise Registration

The core of our reassembly system is a hybrid surface- and curve-based pairwise global registration approach, which is performed for all pairs of facets identified as potentially fractured that belong to distinct fragments. Our approach combines the feature-less rigid geometric registration of Mavridis et al. [2015b] for the fractured surface of the facets with the alignment of the extracted feature curve point sets of the intact object's surface, as proposed by Andreadis et al. [2015]. While several feature-based approaches have been proposed for the fractured faces, the extraction of reliable features is challenging, especially under the presence of heavy erosion, partial geometry and noise. On the contrary, several works such as [Aiger et al. 2008], [Mellado et al. 2014] and [Mavridis et al. 2015b], demonstrated that automatic global rigid registration can be performed without relying on features.

Given the point clouds of two potentially fractured facets \mathcal{X} and \mathcal{X}' , belonging to two distinct fragments with surface geometry \mathcal{G} and \mathcal{G}' , we search for the rigid transformation M , that simultaneously minimizes both the distance between them (\mathcal{F}_{surf} term) and the distance between the feature curves associated with each one (\mathcal{F}_{curve} term):

$$\arg \min_M (c \cdot \mathcal{F}_{surf} + (1 - c) \cdot \mathcal{F}_{curve}). \quad (1)$$

where c is a balancing constant (typically set to 0.5).

In particular, the \mathcal{F}_{surf} term is calculated based on the distance of points on the fractured surface \mathcal{X} of the first fragment from the surface \mathcal{G}' of the second and vice versa. The individually measured Euclidean distances are concatenated in a distance vector whose ℓ_p metric is computed. We perform the same calculation for the other fractured surface and the sum of the two ℓ_p metrics comprise the \mathcal{F}_{surf} term. The use of an ℓ_p metric with $0 < p < 1$ de-emphasises the contribution of outliers and non-contacting parts of the fragments, significantly increasing the robustness of the registration.




Fig. 5. Pairwise registration. Left: Potentially fracture surface point sets ($\mathcal{X}, \mathcal{X}'$), surface feature curves (S, S') and extrapolated curve samples (E, E') – see text. Middle top: example of truncated (signed) distance field. Middle bottom: Registration results without and with penetration testing and penalization. To better illustrate the mutual penetrations, the fragment in front is rendered with transparency. Right: Reassembled object using both surface- and curve-based metrics.

For the \mathcal{F}_{curve} term, we must minimize the distance between the feature curve points on the intact surface of one fragment and the extrapolated points on the curve of the other and vice versa, as illustrated in Fig. 5 - left. In this figure, we denote as \mathcal{E} and \mathcal{E}' the extrapolated curve point sets and as S and S' the surface feature-curve point sets associated with fractured facets \mathcal{X} and \mathcal{X}' respectively. Again, the (two-way) curve distances are combined with the ℓ_p metric to form \mathcal{F}_{curve} .

The extrapolated feature points \mathcal{E} and \mathcal{E}' are essentially a hypothesis we make for the shape of the extents of the features S and S' on the intact surfaces and the farther a point on the extrapolated curves is from the fractured surface, the larger the uncertainty about its true location becomes. In order to properly model this, when using these points in the registration process, we apply an exponential falloff weight to the measured distances the farther the extrapolated samples are from their associated fracture surface. These weights are determined during feature curve extraction.

Eq. 1 combines the potentially contradicting terms \mathcal{F}_{surf} and \mathcal{F}_{curve} in a single process, with the additional requirement for a global optimum, leading to a non-convex optimization problem. To solve it, we rely on the three-level coarse-to-fine search strategy of Mavridis et al. [2015b], [2015a] that is both efficient and robust. The method first coarsely aligns the two fragments based on a given pair of regions (facets) identified as fractured on each part. For facets with incompatible surface area, an alignment procedure based on the Random Sample Consensus (RANSAC) strategy is performed. Similarly sized facets are aligned using the centroids and average facet normals. Subsequently, a reduced-search-space global search is performed using a Simulated Annealing variant followed by a local refinement of the alignment employing one of the Iterative Closest Point methods available in VRMW, such as Sparse ICP [Bouaziz et al. 2013].

A very important factor for the convergence efficiency of the above optimization method is the use of fast data structures for the calculation of the distance function. Due to their small number, curve points are directly indexed using k-d trees and their distance is accurately calculated. Point-to-surface queries for the iteration-heavy levels of the global search employ a discretized, truncated approximation of the target surface distance field (Fig. 5 - middle top). The refinement step, consisting of only a few

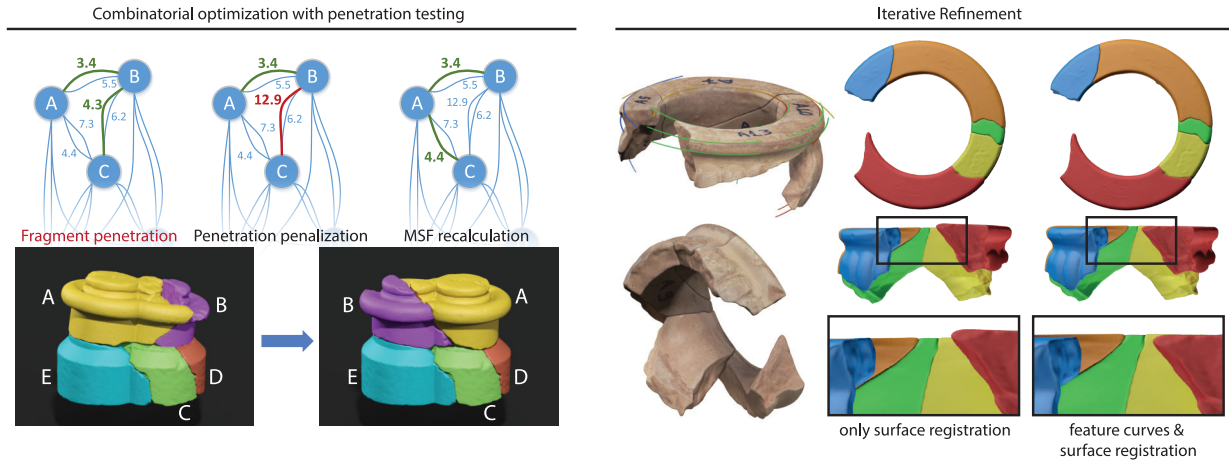


Fig. 6. Multipart registration stages. Left: Example of penetration-based re-weighting of pairwise combinations during the combinatorial optimization. Right: Iterative solution refinement employing both contact surfaces and feature curves. The refinement step required 5 seconds to complete.

iterations, switches to accurate distance queries. Distance fields are automatically generated for the input fragments by the VRMW and persistently stored for future use.

The calculated distance field is also used for the prevention of inter-penetrating parts. Especially for closed meshes, the distance field can register negative offsets from the fragment's surface, which during point-to-surface distance calculations can be penalized, repelling the contacting surfaces as a result (Fig. 5 - middle bottom). For curve-constrained reassembly, penetration is also detected and prevented by the constrained offsetting of one of the fragments along the extrapolated curve trajectories.

4.3 Multipart Reassembly

Given the set of all pairwise registration results and corresponding matching scores, a multi-part optimization stage is performed to produce one or more clusters of fragments. Multi-part reassembly consists of two steps, the combinatorial optimization of pairwise matches and a final geometric refinement of the fragments' pose in the cluster(s).

We represent the solution space as a graph, where each node is a fragment and pairwise matches are the edges between nodes. The discovery of prevailing combinations and therefore solutions to the 3D puzzling problem is directly mapped to finding the Minimum Spanning Forest of the graph, which is computed using Kruskal's algorithm. Since pairwise matches are only locally evaluated up to this point, this process might lead to overlap of partial clusters. In order to eliminate these occurrences, we employ penetration tests during the merging operations of sub-forests. If a penetration is found, the edge weight is readjusted by a fixed value and the algorithm continues (see Fig. 6 - left). If a dead end is detected, i.e. all other valid edges have been tested and the penetration still occurs, we backtrack by re-adjusting all the weights to their original values and modifying the weight of the previously accepted edge. For a reassembly set with N fragments and an average number of facets F per fragment, the overall complexity of the multipart registration process is $O(N^3 \log N)$, since we may have $O(N)$ restarts of Kruskal's algorithm with complexity $O(FN^2 \log(FN^2)) = O(N^2 \log N)$.

A common issue of the multi-part reassembly is the propagation of small misalignments down the chain of connected parts. In order to improve the registration of the reassembled object, we refine the position of the fragments with an iterative multi-part local registration step that uses Sparse ICP and

takes into account both feature curves and surface points (Fig. 6 - right). We avoid having to optimize an arbitrarily large number of parameters, which correspond to the transformation parameters of all fragments; instead of simultaneously optimizing all fragment transformations in a single optimization problem (bundle adjustment), we perform sequential multi-part alignment iteratively, until stabilization. In each iteration, every fragment is repositioned against all other fragments, which are treated as a single part. For each fragment, we minimize the distance of all of its points in contact with all other fragments and all of its associated extrapolated points on the feature curves against all other feature curves, using the same metrics as in the pairwise alignment. Points on the current fragment are considered to be in contact with other fragments if they are closer than a threshold to other surfaces, regardless of whether they belong to a fracture or not. Simultaneously optimizing contact surfaces and feature curves allowed us to accurately rectify distortions even when the contact of fragments could not be cross-validated with other neighboring pieces, as is the case in the example of Fig. 6 - right.

5. OBJECT COMPLETION

The assembled pieces can drive an object completion stage, as explained in the pipeline overview. The predicted shape, which in our work relies on generalized symmetries, can be used for the estimation of the missing geometry either for visualization or for fabrication purposes.

5.1 Symmetrical Expansion of Clusters

The first step in our completion pipeline is based on the observation that many cultural heritage objects exhibit properties of self-similarity and symmetry. Therefore, the missing parts of a globally symmetric partial object can be predicted by using information from the existing parts. However, the biggest challenge in this process is the detection of partial symmetries in the incomplete dataset.

Symmetries can be generally classified as planar (reflectional), rotational and translational. To address the problem of missing parts completion in CH objects, we have integrated in our system the highly-accurate planar symmetry detection approach by Sipiran et al. [2014]. This method is based on the detection of salient features in the input dataset, using a modified version of the *Heat Kernel Signature* [Sun et al. 2009]. Because this approach can be less accurate for smooth featureless surfaces and to handle the remaining cases of rotational and translational symmetries, we have also integrated in our system the registration-based featureless symmetry detection approach of Mavridis et al. [2015c]. This approach does not require any preprocessing and is more general, but it can be less accurate when the detected symmetries depend on small-scale geometric features.

In the remainder of this section we briefly describe these two approaches. For further details, the interested reader is referred to the relevant publications.

5.1.1 Symmetry Analysis based on Heat Diffusion. The heat diffusion theory on surfaces is a key concept to understand 3D structures in multiple scales. The process of heat transfer on a 3D object can be formulated as a partial differential equation that can be efficiently solved through the spectral analysis of an appropriate differential operator defined on the surface of the object. The heat kernel (the solution of the equation) can be computed from the eigenvalues and eigenvectors of the differential operator and it contains multi-scale information about the properties of the 3D object, including the information about symmetries.

By exploiting the heat diffusion theory, we proposed a multi-scale function \mathcal{H} that is invariant to reflectional symmetries (Fig. 7(b)). That is, if two points x and y on the surface are symmetric correspondences, then it follows that $\mathcal{H}(x) \approx \mathcal{H}(y)$. The challenge in the detection of symmetries in broken or damaged objects is to consider global information (symmetries) and local information (invariance to




Fig. 7. Stages of Symmetry analysis with Heat Diffusion. (a) Original input shape. (b) Surface function $\mathcal{H}(\mathbf{x})$ invariant to symmetry and missing parts. (c) Local maximum vertices are detected as potential symmetric points. (d) The candidate plane with the highest votes after the symmetry validation method.

missing geometry) at the same time. However, with our multi-scale analysis we can deal with varied degrees of damaged objects.

Symmetric correspondences. The first step of our algorithm is to compute the function \mathcal{H} on the surface of the input 3D object. We only preserve the points with local maximal values of the function, which correspond to outstanding regions in the 3D object. We noted that this subset of points always contains symmetric correspondences (see Fig. 7(c)). Therefore, every possible pair of points is a potential indication of the presence of reflectional symmetry.

Symmetry validation. For every pair of detected points, we generate a hypothesis of symmetry that requires validation. Our algorithm computes the bisector plane for a given pair of points and accumulates evidence of the presence of symmetry for that plane. We randomly select pairs of points in both sides of the symmetry plane and evaluate several geometric properties: distance to the plane, normal coherence and local feature similarity. We associate random point pairs to corresponding weights which are finally accumulated as votes for the planes. Planes with the highest votes are considered as symmetry planes (see Fig. 7(d)).

Refinement. The symmetries found in the previous step are only approximate and cannot be directly used to mirror the input shape. We need to apply a registration algorithm in order to refine the small misalignment between the input and the reflected shapes. We used the well-known Iterative Closest Point algorithm for this final refinement step.

All of the above steps involved in the symmetrical object expansion using heat diffusion are performed using our Symmetry-based Completion Tool, which can be found at the PRESIOUS project web site.

5.1.2 k-Sparse Optimization Approach. This partial symmetry detection algorithm is based on the simple observation that if we use a registration algorithm to align a shape with itself, then by recording all possible solutions we are essentially detecting the symmetries and self-similarities of an object. The completion of a partial object is performed in three stages: first, the underlying symmetries of the input are detected, then these symmetries are used to iteratively replicate parts of the object and finally accumulated registration errors are relaxed. These steps are detailed below.

Symmetry detection. As previously noted, rotational and translational symmetries are detected by registering the input object with itself. For reflectional symmetries, we register the object with its reflection (along a random plane). If we apply the symmetry transformation M to the set of existing points \mathcal{X} , something that we will refer as *symmetric expansion* or *replication*, we will get a new set of

points \mathcal{X}' . The key observation of the k -Sparse Optimization algorithm [Mavridis et al. 2015c] is that, if we detect and use an optimal symmetry transform that results in a total overlap between \mathcal{X} and \mathcal{X}' , then we do not get any new points after the symmetrical expansion, since both \mathcal{X} and \mathcal{X}' represent the same points. In order to generate new surface points, a *suboptimal* transformation is required, that guarantees that at least k -percent of the points (or surface) of the objects is *non-overlapping*. In the original k -sparse publication, the authors have modified the well-known Super4PCS algorithm to reject solutions that do not satisfy the above criteria, however similar modifications can be performed in other registration approaches, too.

Iterative Replication. A single symmetrical expansion might not be enough to complete all the missing parts of an object, as shown in Figure 8. To this end, the algorithm should run iteratively, until the shape of the object has converged and new non-overlapping points are not generated or until a pre-defined maximum number of iterations is reached. To make this process more efficient, the same self-similarity transformation is reused in all iterations, thus reducing the cost of additional iterations and replications.

Global Error Relaxation. Since the method iteratively applies the same alignment transformation, registration errors are accumulated and propagated in every iteration. To overcome this problem, we apply a global error relaxation approach, similar to Multiview ICP (*bundle adjustment*).

5.2 Merging of Reassembled and Completed Objects

The output of the reassembly and the subsequent symmetrical expansion is still a set of independent parts, possibly with overlapping and duplicate geometry. The next step in our framework is to generate a unique, watertight model for the completed object. The current method to accomplish this goal is based on a Poisson reconstruction method [Kazhdan and Hoppe 2013] that removes the internal fracture surfaces and guarantees a watertight, manifold model as output (Figure 9 - left).

5.3 Complementary Shape Extraction

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 be obtained by computing the difference of the completed object 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. This final step requires the elimination of the fractured surfaces classified using the algorithm described in Section 4.1.



Fig. 8. Iterative completion of the embrasure shape using rotational symmetry. The k -sparse completion algorithm requires three iterations to converge. In the fourth iteration, new geometry is not generated and the algorithm terminates.




Fig. 9. The symmetrical expansion, merging and complementary geometry calculation process.

Nevertheless, a simple boolean difference between the volumetric representations of two shapes is not suitable for the determination of accurate missing parts, since it depends on the resolution of the volumetric representation. Instead, we devised a method based on point cloud processing to get the as-accurate-as possible missing geometry, which is explained below and illustrated in Figure 9. The inputs are the original shape S and the symmetric counterpart S' . Our method is divided in two stages, one for extracting the external (intact) surfaces of the missing geometry and one for generating its fracture surfaces. To compute the intact surface of the missing part, we proceed in the following way:

Geometry filtering by distance. For each vertex \mathbf{v}' in S' , locate the nearest neighbor \mathbf{v} in S . If $\text{dist}(\mathbf{v}', \mathbf{v}) < \rho$ (where ρ is an adequately small threshold), then \mathbf{v}' is marked for removal. Conversely, if $\text{dist}(\mathbf{v}', \mathbf{v}) > 10 \times \rho$ (points are far enough) and the normals of \mathbf{v} and \mathbf{v}' are similar, then point \mathbf{v}' is also marked. Finally, marked vertices and their corresponding triangles are removed from S' .

Geometry filtering by fracture support. The previous filtering could break the connectivity of S' , so it is mandatory to check which connected components should remain. After identifying the connected components, for each one we record the percentage of nearest points on the original shape belonging to fractured facets and only keep connected components with more than 50% of vertex proximity to a fracture. This effectively eliminates patches nearly overlapping the intact surfaces of the fragments.

To compute the fractured surface of the missing part, we proceed in the following way:

Approximate missing part computation. The first step is to compute an approximation of the missing part by applying a volumetric boolean difference $B = S' - S$. It is worth mentioning that in this step resolution is not important, as we only need a very rough approximation to guide the search for the accurate fracture surface. Note that the boolean operation guarantees that B and S are disjoint, and therefore the only shared surface between both is the fracture surface we are looking for.

Extraction of the fractured surface. We perform the extraction applying a nearest neighbor search in point clouds. For each vertex in S , we only keep the vertices with a nearest neighbor in B . Note that this is feasible given the observation in the previous step. Both the intact and fractured surfaces of the missing part are merged to obtain the final complementary geometry. Some examples of our complementary shape extraction algorithm can be observed in Fig. 10.



Fig. 10. Examples of complementary shape extraction. Top: the reassembled objects. Middle: extracted complementary shapes after subtracting the merged reassembled parts from their symmetrical expansion. Bottom: the isolated complementary parts.

6. CASE STUDIES AND EVALUATION

In order to evaluate our pipeline, we performed several experiments with real cultural heritage data, a subset of which is presented throughout our figures. In this section we comment on our results at a high level. A very detailed evaluation of individual algorithms and measurements can be found in the PRESIOUS project 1st Evaluation Report, Sections 3.1 and 3.2 [Theoharis et al. 2015].

6.1 Case Studies

Most of the objects presented are from the archaeological site of the Nidaros Cathedral in Trondheim, Norway but there are also parts scanned at the archaeological site of Elefsis, Greece. Some of the examples in this paper also include contemporary clay objects, which were involved in the quantitative evaluation of our methods. The inclusion of non-CH objects in our test data helped us also cover a wider range of shape classes according to the categorization presented in Section 2.

For the reassembly task, since for actual fragmented archaeological finds the results could not be compared against a reference geometry, we evaluate the quality of the resulting experiments mostly by visual inspection. In order to perform a quantitative evaluation of our method, we performed experiments with non-archaeological models, which we also scanned before fracturing in order to have real

reference models. We have also performed individual tests for pairwise alignment performance evaluation under different degrees of fragment degradation and multi-part reassembly tests with single-object parts and mixed puzzles.

The above limitations also apply to the case of the object completion task, since the real completed object is not available for comparative study. Therefore, we only use CH objects in the evaluation to check the suitability and robustness of our algorithms. Again, this evaluation was mostly performed with visual inspection. Between the criteria used to evaluate robustness we consider the generation of manifold meshes, the introduction of surface defects and the smooth transitions among real and synthetic geometry. Additionally, we created a benchmark of synthetically fractured CH objects using the Hampson dataset as reference [Gregor et al. 2015]. This benchmark contains the complete object and fragments that can be used to test how well a completion algorithm performs. For this evaluation we only performed experiments on the method by [Sipiran et al. 2014].

6.2 Reassembly Evaluation

The correctness of the clustering and alignment of fragments for real archaeological datasets relies on visual inspection, as there are no ground-truth data for this purpose. Nevertheless, in most cases, inspection of the resulting 3D models can easily indicate failures and deviations, especially when a slight misalignment propagates a potential error down an entire chain of connected parts.

For the evaluation of our reassembly pipeline, we performed all assemblies using the VRMW reassembly system. After visual inspection, through the collection management system, sets of fragments were selected and submitted for surface-based registration. The fragment sets may constitute mixed puzzles, i.e. collections of fragments belonging to different original artefacts. Since pairings are persistently cached and therefore their computational cost is amortized among many experiments, we typically used the more time-consuming, high-precision settings for the reassembly engine, i.e. many iterations for the global registration plus SparseICP refinement and multi-part alignment refinement. With values of p close to 0.2 for the ℓ_p metric, fragments with even a small albeit usable contact area were successfully and precisely matched (for example, see columns 2-4 in Fig. 11).

A very smooth transition of the cracked region into an intact one can lead to poor identification of potentially fractured areas. As mentioned in the pipeline overview, erosion can also significantly smooth fractured surfaces or remove large parts that result in incompatible fragments. Both problems can lead to erroneous registration and inadmissible clustering solutions. Our system offers two ways to address such issues: the use of feature curves and the manual white- or black-listing of combinations. As explained in Section 4.1, the curve extraction is a per-fragment pre-processing operation that requires user input. Therefore enabling the curve-based assembly is intended for experiments where contact-based registration is inadequate. Similar to the segmentation and clustering, once performed, the generated data accompany a fragment in all future operations. The feature curves were involved in the correct identification of matches between fragments with particularly reduced or trivial contact area, such as the curved ending of the stone architectural element shown in 5 (right) and the stone sarcophagus of Fig. 11 (first column). Most importantly, feature curves, when enabled, significantly improved the alignment of parts both in the pairwise registration stage and the multi-part refinement (see example of Fig. 6 - right). Obviously the feature curves are not applicable for surfaces devoid of any sharp features. In these cases, only the surface-based criterion should be enabled.

For large puzzles (in our experiments the largest collection consisted of 64 pieces), it is likely that a number of mismatched pairs appear in the multipart reassembly, especially when the collection contains small fragments with uncharacteristic broken surfaces. In such cases, the manual disassociation of the erroneously matched pieces is necessary via the black-listing mechanism provided.



Fig. 11. Additional reassembly results for fragment sets of varying partiality and surface degradation. Feature curves are exploited in the first and last example, while for the rest we only use the contact surfaces.

In order to be able to quantify the performance of our assembly system, we performed a series of experiments with non-archaeological clay objects, which were scanned before fracturing. Using the scanned intact model as reference, we first precisely aligned the scanned fragments with it, thus creating a *reference assembly*. Then, we performed an automated reassembly, and after locking a single fragment to the respective part in the reference assembly, compared the deviation of the two assemblies. All experiments confirmed that the resulting surface distance error was below the error margin of the scanning process, i.e. less than 1.5mm in our case, for objects with a typical diameter of more than 20cm. Detailed measurements can be found in the online appendix and [Theoharis et al. 2015].

6.3 Object Completion Evaluation

To evaluate the object completion algorithm, we implemented a set of tools to check its robustness. The first test inspected the manifoldness of the result by counting the number of non-manifold edges in the output. None of the generated parts from the tested CH objects had non-manifold edges. The test for the introduction of defects counted the number of edges in boundaries, the number of isolated vertices, the number of degenerate triangles and the number of points with inter-penetration. Again, none of these defects manifested in our results.

To quantify the performance of our completion algorithm, we needed a ground-truth case for both the completion itself and the complementary shape extraction. For this reason, we implemented a fracture generation tool that produces synthetic fragments. Our tool was based on random Boolean operations applied using a "cutter" object. When applied to a pre-selected set from the Hampson dataset, the

fracture generation tool produced a set of fragments for every shape in the collection. Finally, we selected a random subset of fragments for each shape (leaving out a single fragment for each object) and applied our completion algorithm. The result was compared against the non-fractured object. Moreover, since the missing fragment was already identified, we could easily compare it with the complementary geometry generated by our algorithm.

The criterion to compare the results to the ground-truth was the volumetric divergence defined as the ratio between the volume of the intersection of the two shapes divided by the volume of the union of the two shapes. A divergence of one means a perfect matching. Our algorithm obtained a divergence of 0.8267 on average using the proposed benchmark. More details on the evaluation and results can be seen in [Theoharis et al. 2015] and [Gregor et al. 2015] and the synthetically generated benchmark data can be downloaded from <http://fracture-benchmark.dbvis.de>.

7. DISCUSSION AND CONCLUSIONS

In this paper we have proposed, implemented and demonstrated a computational pipeline for assisting experts in the restoration and study of CH objects, consisting of thoroughly-tested, state-of-the-art algorithms. The reassembly stages are integrated in a single application environment and complemented by distance and contact area measurement facilities and user-friendly reporting. Considering the versatility and utility of automated shape completion, it is worth noting that despite the fact that unarguably, synthesized parts are of little value to the CH expert in terms of interpretation and study, the geometry of a symmetrically or otherwise repaired object is very useful in restoration tasks. Most notably, it provides a reference model for computationally extracting complementary parts for visualization and fabrication in order to assist the physical restoration process, especially when large parts of an object are missing. Second, as shown in [Andreadis et al. 2015], the expanded mesh can provide feedback for the reassembly, thus enabling the plausible and relatively accurate positioning of isolated parts.

Currently our object completion pipeline relies on generalized symmetries, which when applicable, can provide a plausible complete shape that is also usable for complementary geometry extraction. However completion via retrieval is a promising future research subject, one that we have so far only scratched the surface of, since it heavily relies on interpretation and hidden semantics at a higher level than the geometry of the objects.

Although the implemented software is at a prototype level, it is the first time a cohesive digital restoration system is made widely and publicly available. We believe that through wide deployment, experimentation and feedback, the community can be inspired to advance and adopt technological tools in this direction.

REFERENCES

- Drol Aiger, Niloy J. Mitra, and Daniel Cohen-Or. 2008. 4-points Congruent Sets for Robust Surface Registration. *ACM Transactions on Graphics* 27, 3 (2008), #85, 1–10.
 - Enkhbayar Altantsetseg, Katsutsugu Matsuyama, and Kouichi Konno. 2014. Pairwise matching of 3D fragments using fast fourier transform. *The Visual Computer* 30, 6-8 (2014), 929–938.
 - Anthousis Andreadis, Róbert Gregor, Ivan Sipiran, Pavlos Mavridis, Papaioannou Georgios, and Tobias Schreck. 2015. Fractured 3D Object Restoration and Completion. In *ACM SIGGRAPH 2015 Posters (SIGGRAPH '15)*. ACM, New York, NY, USA, Article 74, 1 pages. DOI : <http://dx.doi.org/10.1145/2787626.2792633>
 - Anthousis Andreadis, Pavlos Mavridis, and Georgios Papaioannou. 2014. Facet Extraction and Classification for the Reassembly of Fractured 3D Objects. In *Eurographics 2014 - Posters*, Mathias Paulin and Carsten Dachsbacher (Eds.). The Eurographics Association. DOI : <http://dx.doi.org/10.2312/egp.20141060>
 - Anthousis Andreadis, Georgios Papaioannou, and Pavlos Mavridis. 2015. Generalized Digital Reassembly using Geometric Registration. In *Digital Heritage*, Vol. 2. IEEE, 549–556.
- ACM Journal on Computing and Cultural Heritage, Vol. 10, No. 2, Article 8, Publication date: 2017.

- Sofien Bouaziz, Andrea Tagliasacchi, and Mark Pauly. 2013. Sparse Iterative Closest Point. *Computer Graphics Forum (Symposium on Geometry Processing)* 32, 5 (2013), 1–11.
- Benedict J. Brown, Corey Toler-Franklin, Diego Nehab, Michael Burns, David Dobkin, Andreas Vlachopoulos, Christos Doumas, Szymon Rusinkiewicz, and Tim Weyrich. 2008. A System for High-volume Acquisition and Matching of Fresco Fragments: Reassembling Thera Wall Paintings. *ACM Trans. Graph.* 27, 3 (2008), 84:1–84:9.
- Fernand Cohen, Zexi Liu, and Taslidere Ezgi. 2013. Virtual reconstruction of archeological vessels using expert priors and intrinsic differential geometry information. *Computers & Graphics* 37, 12 (2013), 41 – 53.
- Helena Cristina da Gama Leitao and Jorge Stolfi. 2002. A multiscale method for the reassembly of two-dimensional fragmented objects. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 24, 9 (Sep 2002), 1239–1251.
- Robert Gregor, Danny Bauer, Ivan Sipiran, Panagiotis Perakis, and Tobias Schreck. 2015. Automatic 3D Object Fracturing for Evaluation of Partial Retrieval and Object Restoration Tasks: Benchmark and Application to 3D Cultural Heritage Data. In *Proc. Eurographics Workshop on 3D Object Retrieval (3DOR)*. Eurographics Association, 7–14.
- Robert Gregor, Ivan Sipiran, Georgios Papaioannou, Tobias Schreck, Anthousis Andreadis, and Pavlos Mavridis. 2014. Towards Automated 3D Reconstruction of Defective Cultural Heritage Objects. *Proc. Eurographics Workshop on Graphics and Cultural Heritage* (2014), 135–144.
- Gur Harary, Ayellet Tal, and Eitan Grinspun. 2014a. Context-based Coherent Surface Completion. *ACM Trans. Graph.* 33, 1, Article 5 (Feb. 2014), 5:1–5:12 pages.
- Gur Harary, Ayellet Tal, and Eitan Grinspun. 2014b. Feature-Preserving Surface Completion Using Four Points. *Comput. Graph. Forum* 33, 5 (Aug. 2014), 45–54.
- Hui Huang, Shihao Wu, Daniel Cohen-Or, Minglun Gong, Hao Zhang, Guiqing Li, and Baoquan Chen. 2013a. L1-medial Skeleton of Point Cloud. *ACM Trans. Graph.* 32, 4, Article 65 (July 2013), 8 pages.
- Hui Huang, Kangxue Yin, Minglun Gong, Dani Lischinski, Daniel Cohen-Or, Uri Ascher, and Baoquan Chen. 2013b. Mind the Gap: Tele-Registration for Structure-Driven Image Completion. *ACM Trans. Graph. (Proc. SIGGRAPH ASIA)* 32 (2013), 174:1–174:10. Issue 6.
- Qi-Xing Huang, Simon Flöry, Natasha Gelfand, Michael Hofer, and Helmut Pottmann. 2006. Reassembling Fractured Objects by Geometric Matching. *ACM Trans. Graph.* 25, 3 (July 2006), 569–578.
- Martin Kampel and Robert Sablatnig. 2004. On 3D mosaicing of rotationally symmetric ceramic fragments. In *Proc. 17th International Conference on Pattern Recognition (ICPR)*, Vol. 2. IEEE, 265–268.
- Michael Kazhdan and Hugues Hoppe. 2013. Screened Poisson Surface Reconstruction. *ACM Trans. Graph.* 32, 3 (July 2013), 29:1–29:13.
- David Koller and Marc Levoy. 2006. Computer-aided reconstruction and new matches in the forma urbis romae. *Bullettino Della Commissione Archeologica Comunale di Roma* 2 (2006).
- Weixin Kong and B.B. Kimia. 2001. On solving 2D and 3D puzzles using curve matching. In *Proc. IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, Vol. 2. II–583–II–590 vol.2.
- Jeehyung Lee and Thomas Funkhouser. 2008. Sketch-Based Search and Composition of 3D Models. In *EUROGRAPHICS Workshop on Sketch-Based Interfaces and Modeling*.
- Er Li, Xiaopeng Zhang, and Yanyun Chen. 2014. Symmetry Based Chinese Ancient Architecture Reconstruction from Incomplete Point Cloud. In *Proc. 5th International Conference on Digital Home (ICDH)*. 157–161.
- Hao Li, Robert W. Sumner, and Mark Pauly. 2008. Global Correspondence Optimization for Non-Rigid Registration of Depth Scans. *Computer Graphics Forum (Proc. SGP'08)* 27, 5 (July 2008).
- Xin Li, Zhao Yin, Li Wei, Shenghua Wan, Wei Yu, and Maoqing Li. 2011. Symmetry and template guided completion of damaged skulls. *Computers & Graphics* 35, 4 (2011), 885 – 893. Semantic 3D Media and Content.
- Pavlos Mavridis, Anthousis Andreadis, and Georgios Papaioannou. 2015a. Efficient Sparse {ICP}. *Computer Aided Geometric Design* 35 – 36 (2015), 16 – 26. DOI: <http://dx.doi.org/10.1016/j.cagd.2015.03.022> Geometric Modeling and Processing 2015.
- Pavlos Mavridis, Anthousis Andreadis, and Georgios Papaioannou. 2015b. Fractured Object Reassembly via Robust Surface Registration. In *Eurographics 2015 - Short Papers*. The Eurographics Association. DOI: <http://dx.doi.org/10.2312/egsh.20151005>
- Pavlos Mavridis, Ivan Sipiran, Anthousis Andreadis, and Georgios Papaioannou. 2015c. Object Completion using k-Sparse Optimization. *Computer Graphics Forum* 34, 7 (2015), 13–21. DOI: <http://dx.doi.org/10.1111/cgf.12741>
- Nicolas Mellado, Dror Aiger, and Niloy J. Mitra. 2014. Super 4PCS Fast Global Pointcloud Registration via Smart Indexing. *Computer Graphics Forum* 33, 5 (2014), 205–215.
- Nicolas Mellado, Matteo Dellepiane, and Roberto Scopigno. 2016. Relative scale estimation and 3D registration of multi-modal geometry using Growing Least Squares. (2016). <http://vcg.isti.cnr.it/Publications/2016/MDS16>

- Nicolas Mellado, Patrick Reuter, and Christophe Schlick. 2010. Semi-automatic Geometry-driven Reassembly of Fractured Archeological Objects. In *Proc. 11th International Conference on Virtual Reality, Archaeology and Cultural Heritage (VAST'10)*. Eurographics Association, 33–38.
- Gregorio Palmas, Nico Pietroni, Paolo Cignoni, and Roberto Scopigno. 2013. A computer-assisted constraint-based system for assembling fragmented objects. In *Proc. of Digital Heritage 2013 International Congress*, Vol. 1. IEEE, 529–536.
- Georgios Papaioannou and Evaggelia-Aggeliki Karabassi. 2003. On the automatic assemblage of arbitrary broken solid artefacts. *Image and Vision Computing* 21, 5 (2003), 401–412. DOI: [http://dx.doi.org/10.1016/S0262-8856\(03\)00008-8](http://dx.doi.org/10.1016/S0262-8856(03)00008-8)
- Georgios Papaioannou, Evaggelia-Aggeliki Karabassi, and Theoharis Theoharis. 2001. Virtual Archaeologist: Assembling the Past. *IEEE Computer Graphics and Applications* 21 (2001), 53–59.
- Devi Parikh, Rahul Sukthankar, Tsuhan Chen, and Mei Chen. 2007. Feature-based Part Retrieval for Interactive 3D Reassembly. In *Proc. 8th IEEE Workshop on Applications of Computer Vision (WACV '07)*. IEEE Computer Society, Washington, DC, USA, 14–.
- Yair Poleg and Shmuel Peleg. 2012. Alignment and mosaicing of non-overlapping images. In *Computational Photography (ICCP), 2012 IEEE International Conference on*. 1–8.
- Helmut Pottmann, Johannes Wallner, Qi-Xing Huang, and Yong-Liang Yang. 2009. Integral Invariants for Robust Geometry Processing. *Comput. Aided Geom. Des.* 26, 1 (2009), 37–60.
- Belenguer C. Sanchez and Eduardo Vendrell Vidal. 2012. Archaeological fragment characterization and 3D reconstruction based on projective gpu depth maps. In *Virtual Systems and Multimedia (VSMM)*. IEEE, 275–282.
- Mahmut Şamil. Sağıroğlu and Aytul Erçil. 2005. A Texture Based Approach to Reconstruction of Archaeological Finds. In *Pro. 6th International Conference on Virtual Reality, Archaeology and Intelligent Cultural Heritage (VAST'05)*. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 137–142.
- Andrei Sharf, Marc Alexa, and Daniel Cohen-Or. 2004. Context-based Surface Completion. *ACM Trans. Graph.* 23, 3 (Aug. 2004), 878–887.
- Ivan Sipiran, Robert Gregor, and Tobias Schreck. 2014. Approximate Symmetry Detection in Partial 3D Meshes. *Computer Graphics Forum (proc. Pacific Graphics)* 33 (2014), 131–140.
- Kilho Son, Eduardo B Almeida, and David B Cooper. 2013. Axially Symmetric 3D Pots Configuration System Using Axis of Symmetry and Break Curve. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 257–264.
- Jian Sun, Maks Ovsjanikov, and Leonidas Guibas. 2009. A Concise and Provably Informative Multi-scale Signature Based on Heat Diffusion. In *Proc. Eurographics Symposium on Geometry Processing (SGP '09)*. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 1383–1392. <http://dl.acm.org/citation.cfm?id=1735603.1735621>
- Kenshi Takayama, Ryan Schmidt, Karan Singh, Takeo Igarashi, Tamy Boubekeur, and Olga Sorkine. 2011. GeoBrush: Interactive Mesh Geometry Cloning. *Computer Graphics Forum* (2011). DOI: <http://dx.doi.org/10.1111/j.1467-8659.2011.01883.x>
- Theoharis Theoharis, Christian Schellewald, Panagiotis Perakis, Anthousis Andreadis, Georgios Papaioannou, Pavlos Mavridis, Michalis Savelonas, Konstantinos Sfikas, Ioannis Pratikakis, Fotis Arnaoutoglou, Ivan Sipiran, Robert Gregor, Tobias Schreck, and Dirk Rieke-Zapp. 2015. *PRESIOUS Project - First Evaluation Report*. Technical Report D5.4. PRESIOUS EU Project, 600533. http://presious.eu/file_downloads/PRESIOUS-D5.4-v2.4.pdf
- Barbara Thuswaldner, Simon Flory, Robert Kalasek, Michael Hofer, Qi-Xing Huang, and Hilke Thürr. 2009. Digital anastylis of the Octagon in Ephesos. *Journal on Computing and Cultural Heritage* 2, Article 1 (July 2009), 27 pages. Issue 1.
- Federico Tombari, Samuele Salti, and Luigi Di Stefano. 2010. Unique Signatures of Histograms for Local Surface Description. In *Proc. 11th European Conference on Computer Vision Conference on Computer Vision: Part III (ECCV'10)*. Springer-Verlag, Berlin, Heidelberg, 356–369.
- Li Wei, Wei Yu, Maoqing Li, and Xin Li. 2011. Skull Assembly and Completion Using Template-Based Surface Matching. In *2011 International Conference on 3D Imaging, Modeling, Processing, Visualization and Transmission*. 413–420.
- Andrew R. Willis and David B. Cooper. 2004. Bayesian assembly of 3D axially symmetric shapes from fragments. In *Proc. IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, Vol. 1. I-82–I-89 Vol.1.
- Simon Winkelbach and FriedrichM. Wahl. 2008. Pairwise Matching of 3D Fragments Using Cluster Trees. *International Journal of Computer Vision* 78, 1 (2008), 1–13.
- Jingyu Yang, Ke Li, Kun Li, and Yu-Kun Lai. 2015. Sparse Non-rigid Registration of 3D Shapes. In *Proc. Eurographics Symposium on Geometry Processing (SGP '15)*. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 89–99. DOI: <http://dx.doi.org/10.1111/cgf.12699>
- Kang Zhang and Xin Li. 2014. A graph-based optimization algorithm for fragmented image reassembly. *Graphical Models* 76, 5 (2014), 484–495. Geometric Modeling and Processing 2014.
- ACM Journal on Computing and Cultural Heritage, Vol. 10, No. 2, Article 8, Publication date: 2017.

Kang Zhang, Wuyi Yu, Mary Manhein, Warren Waggenspack, and Xin Li. 2015b. 3D Fragment Reassembly Using Integrated Template Guidance and Fracture-Region Matching. In *Proc. IEEE International Conference on Computer Vision (ICCV)*. 2138–2146.

Ran Zhang, Xuejin Chen, Takaaki Shiratori, Xin Tong, and Ligang Liu. 2015a. An efficient volumetric method for non-rigid registration. *Graphical Models* 79 (2015), 1 – 11. DOI: <http://dx.doi.org/10.1016/j.gmod.2015.01.003>

Received June 2016; revised — 2016; accepted — 2016

Online Appendix to: From Reassembly to Object Completion - A Complete Systems Pipeline

GEORGIOS PAPAIOANNOU, Athens University of Economics and Business

TOBIAS SCHRECK, Graz University of Technology

ANTHOUSIS ANDREADIS, Athens University of Economics and Business

PAVLOS MAVRIDIS and ROBERT GREGOR, Graz University of Technology

IVAN SIPIRAN, Pontificia Universidad Católica del Perú

KONSTANTINOS VARDIS, Athens University of Economics and Business

A. EXTENDED REASSEMBLY RESULTS EVALUATION

In this appendix we provide detailed measurements for selected reassembly cases used in our evaluation process. To validate our alignment method in a more concise and measurable manner, as explained in the paper, we performed experiments with non-archaeological models as well, which we scanned before fracturing. Using the scanned original as a reference model, we generated *reference reassemblies* and measured their RMS Hausdorff distance from our resulting assemblies, with the procedure explained in Section 6.2.

Three test categories of increasing reassembly difficulty are provided: a) single object, b) multiple objects in the same reassembly set (mixed puzzle) and c) reassembly of heavily damaged fragments. The first two categories use only the contact surfaces for the registration, while the third also exploits the feature curves.

A.1 Reassembly of a Single Object

In the experiments of this section we expect a single island of fragments as a result of the multi-part reassembly. The robustness of our registration process is evaluated here as its capability to a) discriminate correct pairs and b) properly group the fragments without leaving isolated islands in the solution. For the qualitative assessment of the single-object reassembly, when no ground truth original object is available (i.e. for actual archaeological datasets), visual inspection of the resulting models confirms the correct shape and helps detect severe skewing of the alignment. We also measure and report the maximum penetration depth detected. For test sets with ground truth data, we additionally measure the RMS Hausdorff distance against the reference reassembly. Figure 12 demonstrates some results using single-object collections and in particular, two archaeological data sets and two test sets with ground truth data.

A.2 Reassembly of Mixed Objects

Here we evaluate the performance of our algorithms in the case of mixed puzzles. Fragments from multiple objects are used and the expected result is a set of multiple clusters of reassembled objects, as well as potentially isolated fragments. In order to achieve that, during the graph construction stage we filter out pairwise matches with a matching error higher than a threshold value (the *split threshold* reported in Figure 13), which signify a bad or impossible registration (e.g. due to fragment penetration).

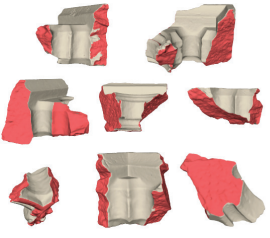
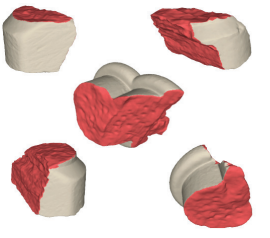
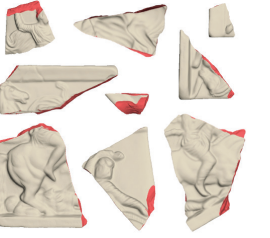
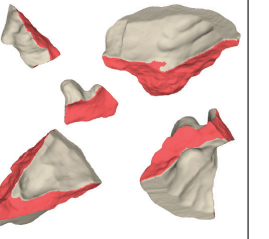

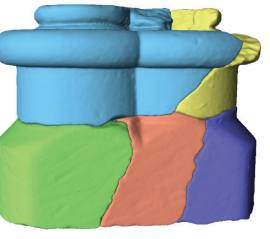

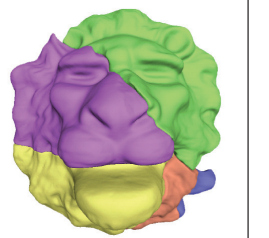
	Dora Embrasure	Dora Column Base	Metope	Lion Head
Fragments				
Reassembled object				
	Penetration: 0.46mm Pairwise tests: 526 Hausdorff RMS: - Time: 591s	Penetration: 0.41mm Pairwise tests: 112 Hausdorff RMS: - Time: 213s	Penetration: 0.20mm Pairwise tests: 233 Hausdorff RMS: 1.14mm Time: 250s	Penetration: 0.36mm Pairwise tests: 38 Hausdorff RMS: 0.55mm Time: 89s

Fig. 12. Object reassembly of a single object. We report the total computation time, the number of pairwise tests and the maximum computed penetration. RMS Hausdorff distance from the ground truth is also provided, where available. For the results shown, no final multi-part refinement was necessary.

The criteria and respective measurements for the mixed puzzle problems presented are identical to the single-object experiments, with the exception of the cluster split threshold, a parameter that can be adjusted from the VRMW interface.

In Figure 13 the first two examples our methodology manages to produce the two desired clusters. This is attributed mainly to the good discrimination achieved by the matching error of the fragments. In the example of the third column, we show a failure scenario. With the advantage of hindsight, we now that the fragments should form 4 clusters, but a total of 8 were generated. The cluster at the bottom-left of the example consists of two fragments that failed to form a cluster using the default parameters. In the case of the top-left cluster our algorithm generated 4 sub-clusters. Since some of the fragments only share a trivial contact surface or are heavily damaged and flaked, only after introducing the feature curves was it possible to discover substantial support to form a single object.

As collections of fragments become large, the pairwise registration cost becomes the predominant factor. For each new fragment processed and introduced to a collection, the pairwise registration score between all of its fractured facets and the corresponding fractured facets of the fragments already in the set must be evaluated and stored. For the largest collection of fragments from the same origin that we had in our possession, the Nidaros Cathedral dataset, we measured the incremental update time when introducing the last piece in a set already containing 63 fragments. The 64th fragment had 6 facets identified as potentially fractured and formed 2268 pairs with the other fragments in the collection, resulting in a total pairwise registration time of 19 minutes.

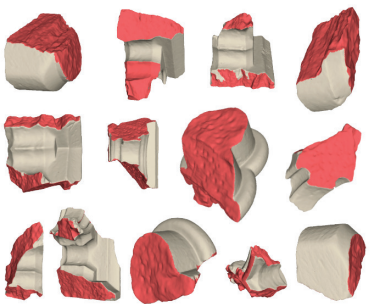


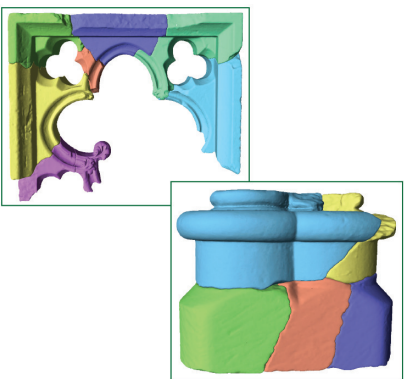
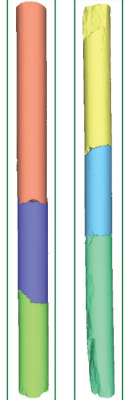
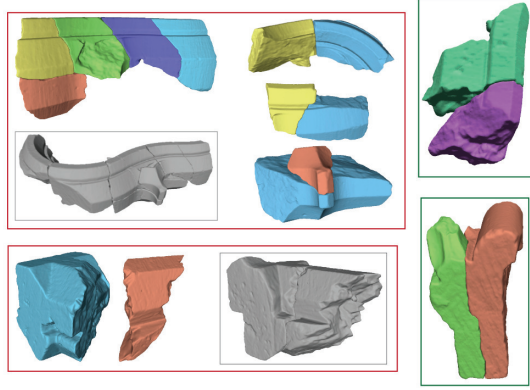
	Embrasure and Column Base	Dora Columns	Dora Arch Set
Fragments			
Reassembled object			
	Split threshold: 2.0 Pairwise tests: 1233 Clusters: 2/2 Time: 1420s	Split threshold: 2.0 Pairwise tests: 149 Clusters: 2/2 Time: 151s	Split threshold: 1.5 Pairwise tests: 2620 Clusters: 8/4 Time: 2625s

Fig. 13. Object reassembly with mixed puzzles. We report the total time, the number of pairwise tests and generated/expected clusters. Correctly reassembled clusters are highlighted with a green outline, while red indicates partial or erroneous ones. In the second case, we visualize the expected cluster in grey. For every test case, we also provide the error threshold that was used for splitting the clusters. For the results shown, no final multi-part refinement was necessary.

A.3 Reassembly of Heavily Damaged Fragments

This is the final case of our evaluation for the reassembly of objects, where we experiment with fragments with large parts of the fractured areas missing. In such cases, contact-based matching approaches fail as the fractured surfaces do not contain enough information for the alignment process. In order to solve these cases we enable the feature curves. In Figure 14 we present complete object reassemblies that were obtained using registration with both the contact surface and the feature curves criteria enabled. Due to the fact that these results were obtained in a semi-automatic way, as user evaluation of some intermediate alignment results and blacklisting of certain invalid pairs was required, we do not present timings.

It is noteworthy that by exploiting the feature curves constraint, we managed to discover some new combinations of fragments in the Dora Arch set of fragments from the Nidaros Cathedral, which were nearly impossible to detect by visual inspection of the physical fragments or their virtual counterparts,

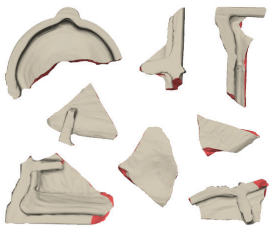
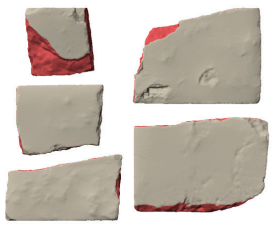
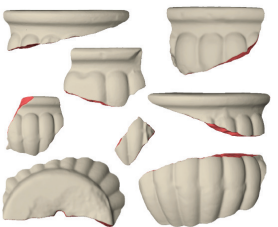
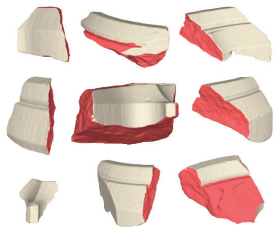
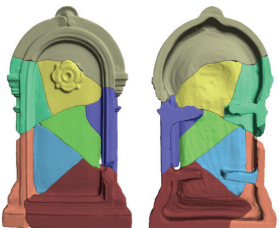
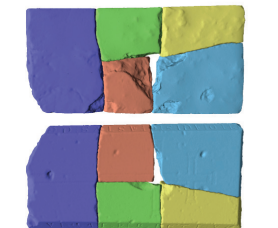
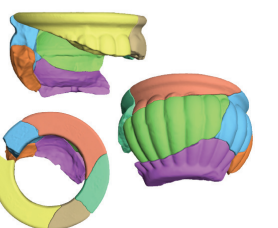
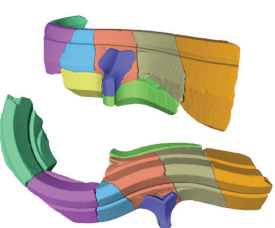
	Fountain	Tombstone	Ornate Pot	Dora Arch
Fragments				
Reassembled object				
	Penetration: 0.01mm Pairwise tests: 273 Hausdorff RMS: 1.16mm	Penetration: 0.81mm Pairwise tests: 157 Hausdorff RMS: -	Penetration: 0.46mm Pairwise tests: 422 Hausdorff RMS: -	Penetration: 0.68mm Pairwise tests: 772 Hausdorff RMS: -

Fig. 14. Reassembly results that were obtained using both the contact surfaces and the f-curves. We report the number of pairwise tests, the maximum detected penetration of the assembled parts and the Hausdorff distance RMS from the ground truth, where available.

alone. The joined parts in the proposed reassembly, that employs feature curves, contact surfaces and white-/black-listing is presented in the last column of Figure 14.