Tools for Virtual Reassembly of Fresco Fragments

by

Benedict Brown, Lara Laken, Philip Dutré, Luc Van Gool, Szymon Rusinkiewicz, Tim Weyrich

Tools for Virtual Reassembly of Fresco Fragments

Benedict Brown[a]*, Lara Laken[b], Philip Dutré[c], Luc Van Gool[c], Szymon Rusinkiewicz[d], Tim Weyrich[e]

[a]University of Pennsylvania, bjbrown@cis.upenn.edu
[b]Radboud University Nijmegen, l.laken@let.ru.nl
[c]University of Leuven
[d]Princeton University
[e]University College London

---

*Corresponding author. Work conducted while an FWO Post-Doctoral Fellow at the University of Leuven

Tools for Virtual Reassembly of Fresco Fragments

Benedict Brown, Lara Laken, Philip Dutré, Luc Van Gool, Szymon Rusinkiewicz, Tim Weyrich

Abstract

The process of reassembling fragmented wall paintings is currently prohibitively time consuming, limiting the amount of material that can be examined and reconstructed. Computer-assisted technologies hold the promise of helping humans in this task, making it possible to digitize detailed shape, color, and surface relief information for each fragment. The data can be used for documentation, visualization (both on- and off-site), virtual restoration, and to automatically propose matches between fragments. Our focus in this paper is on improving the workflow, tools, and visualizations, as they are used by archaeologists and conservators to scan fragments and find matches. In particular, we evaluate the system's performance and user experience in ongoing acquisition and matching work on material from a Roman excavation in Tongeren, Belgium. Compared to prior systems, we can acquire fragments approximately 10 times faster, and support a wider range of fragment sizes (from 1 cm to 20 cm in diameter).
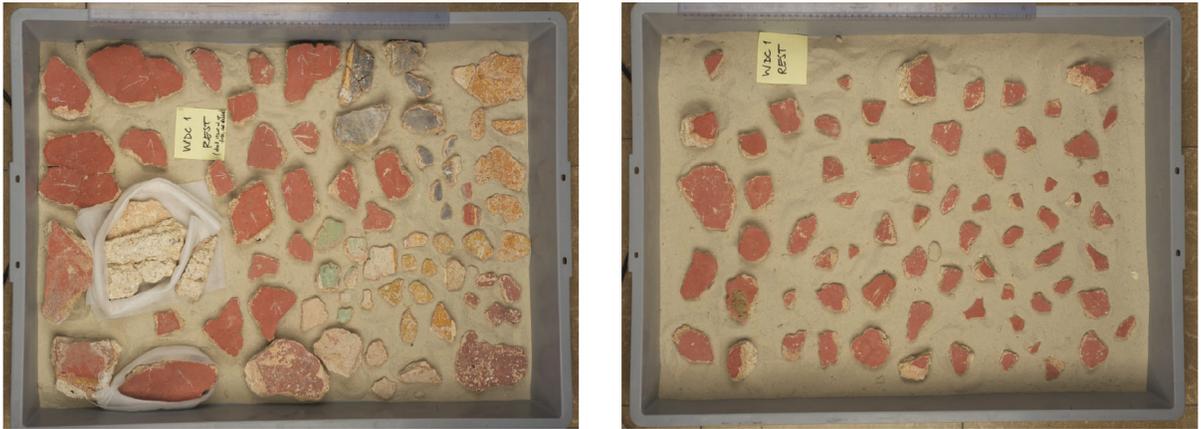
Figure 1: Two crates of fragments from the Tongeren Vrijthof excavation, containing 144 fragments. Even this small number of fragments is extremely difficult to assemble by hand (Copyright *Gallo-Romeins Museum Tongeren*).

Introduction

Although excavations sometimes reveal wall paintings more or less whole and *in situ*, very often they are fragmented and must be reassembled before being fully studied and displayed. The problem is essentially one of solving an enormous jigsaw puzzle of unknown design in which many pieces are the same color, many pieces may be missing, and there are few clear distinguishing features on the edges. Furthermore, fragments are generally fragile, and excessive handling is undesirable at best and highly damaging at worst. Because the assembly work is so time-consuming and labor-intensive, vast quantities of material remain in storage indefinitely, unexamined and unstudied.

The problem of reassembling "puzzles," including not only traditional jigsaw puzzles, but also fractured pottery, frescoes, manuscripts, and other types of objects, has been extensively studied in the computer science community. Recent surveys by Willis and Cooper [1] and by Kleber and Sablatnig [2] cover a range of techniques that have been developed. An associated challenge is acquiring computer representations of the thousands of fragments typically encountered at an excavation. Although there are many similarities in acquisition and matching across different types of material, the most successful approaches are specialized to take advantage of specific properties. For example, pottery assembly systems typically rely on detecting profile curves from the potsherds, whereas wall painting acquisition and assembly are both helped by assuming the painted surface is (nearly) flat.

Recently, Brown *et al.* [3] introduced a system for computer-aided reassembly of wall paintings from the Akrotiri excavations in Thera (Santorini), Greece. The idea is to efficiently capture a virtual 3D model of each fragment along with high-resolution color and texture

information of the front surface. The computer tries to fit pairs of fragments based on the shape information on the fragments' sides (equivalent to physically trying all possible configurations to see if fragments "lock" together) as well as on color, plaster surface texture, and surface roughness information [4]. The ultimate goal is a system that can be used by archaeologists and conservators to document and help assemble frescoes without close participation of computer scientists.

In this paper we describe our experience using an improved version of this system to help assemble a Roman fresco from Tongeren, Belgium. As can be expected, the fragments are fragile, so the excessive handling necessary for traditional reconstruction methods is undesirable. Moreover, the necessary financial resources and coterie of archaeologists to do so are unavailable. We have therefore embarked on a campaign to scan *every* available fragment of this fresco and see how much the system can increase productivity. At the same time we are improving the tools to find more matches and eliminate the need for trained computer scientists.

Our system provides several benefits for the reconstruction of wall paintings:

- It takes approximately three minutes to completely scan a single fragment, including both 3D shape data and 2D color and surface texture. This is approximately four times faster than the system used in Brown *et al.* [3], yet it obtains more detailed shape information and supports larger fragments (up to 20 cm in diameter as opposed to only 10 cm);
- Acquisition is simple and does not require touching fragments to each other, unlike manual matching. It is therefore easily performed by interns with minimal training;
- Initially verifying computer-proposed matches requires no access to the physical material, and can be performed offsite or potentially even over the Internet by anyone;
- The acquisition process yields a database of high-quality imagery of every fragment that can be used for further reproduction tasks, eliminating the need to independently photograph fragments for catalogs, internal documentation, and publications;
- The computer can propose many correct matches that are otherwise very difficult to find, greatly reducing the labor involved in reconstruction.

The productivity gains of our system therefore come not only from increasing the number of matches that can be found in a given amount of time, but also from reducing the need to separately find

| Benedict Brown, Lara Laken, Philip Dutré, Luc Van Gool, Szymon Rusinkiewicz, Tim Weyrich

and photograph fragments for reproduction and from allowing a wide range of people to assist in the reassembly process.

In the next section we describe the fresco that we have investigated. Following that, we present our improvements to the scanning and matching procedures in more detail and the matches we found with it. We close with a discussion of the system's current usefulness and the required improvements.

The Tongeren Vrijthof Fresco

The fragments we have scanned were found at the Vrijthof square in Tongeren, in a rubble layer filled with plaster fragments next to the fourth century city wall, just outside the Church of Our Lady (*Onze-Lieve-Vrouwe-basiliek*) (Figure 4) [5]. They belong to one out of more than ten distinct paintings with different mortar compositions [6]. Most were probably originally located in the house that was excavated beneath the church, a luxurious urban dwelling in the center of the second-century *municipium Tungrorum* [7]. Fragments of the same paintings were discovered underneath its floors dating to the second half of the second century AD and were probably applied in the previous building phases, between the Batavian revolt in 69 – 70 AD and the enormous fire that destroyed the city just after the middle of the second century. Decoration aside, the fragments of this painting are easily discerned from the other frescoes by the pink mortar layer directly under the lime coat (Figure 2). The painted system could be largely reconstructed: it consists of a stone imitation in the lower zone (dado) and a black main zone with red panels (Figure 5).

The red panels are bordered by thin vegetable stems in yellow shades and are flanked at the wall's ends by vertical black strips containing vines with yellow and red flowers. A green band frames the main zone. It ran along the corner of the wall, preserved at the right side, and probably also along the joint with the ceiling, as can be recognized by the slightly curved surface. On the black background between the red panels was a series of green and pink vegetable stems comparable to the yellow ones bordering the panels; they seem to have formed an *aedicula* or some sort of a pilaster, but this is

Figure 2. Side view showing the two layers of mortar that identify this painting (Copyright *Gallo-Romeins Museum Tongeren*).

difficult to see because these fragments are covered with a white incrustation layer that is difficult to remove without demolishing the painted surface. One fragment showing part of a lion's body and mane or human hair (Figure 3) indicates that fantastic animals like sphinxes adorned the panels, or possibly the top of the pilaster. There was only a small distance between these motifs and the upper green band along the joint with the ceiling, as can be seen on another fragment.

The dado has an *opus sectile* imitation of marble slabs in an orange frame between red side panels. Its central tondo is painted in orange with flecks of pink, surrounded by a pearl string of white ovals and beads. Two pink-flecked gray triangles point toward the central circle; the remaining surfaces above and below contain greenish yellow marbling. A reddish purple plinth with white splashes runs underneath. On most fragments a flat side indicates where the bottom of the painting touched the ground. The lower surface of this plinth is covered with concretions of pinkish mortar; we do not know the reason, but this may have been caused by splattering when an *opus signinum* floor was laid against the wall.

The dimensions of the painted system can be deduced from the right side of the wall and the dado, since the first *opus sectile* panel is almost completely preserved. Its central point in the middle of the tondo makes it easy to duplicate its height and width and also gives us a clue about the width of the panels in the main zone. Their heights can then be estimated based on an average width-to-height aspect ratio of 2:3. The full height of the walls is estimated at about 2.85 m, while their widths depend on how many red panels they had: this wall may have had one, two or three red panels. The stems or pilaster decoration in the black fields between the red panels were
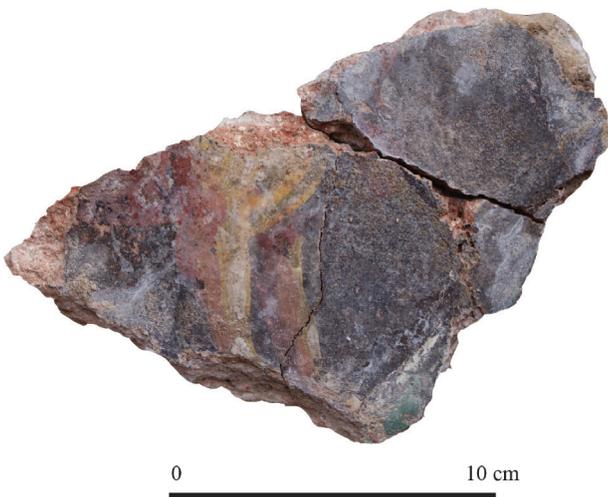

Figure 3. Fragment of a lion or sphynx (*Copyright Gallo-Romeins Museum Tongeren*).

0   10 cm

Benedict Brown, Lara Laken, Philip Dutré, Luc Van Gool, Szymon Rusinkiewicz, Tim Weyrich
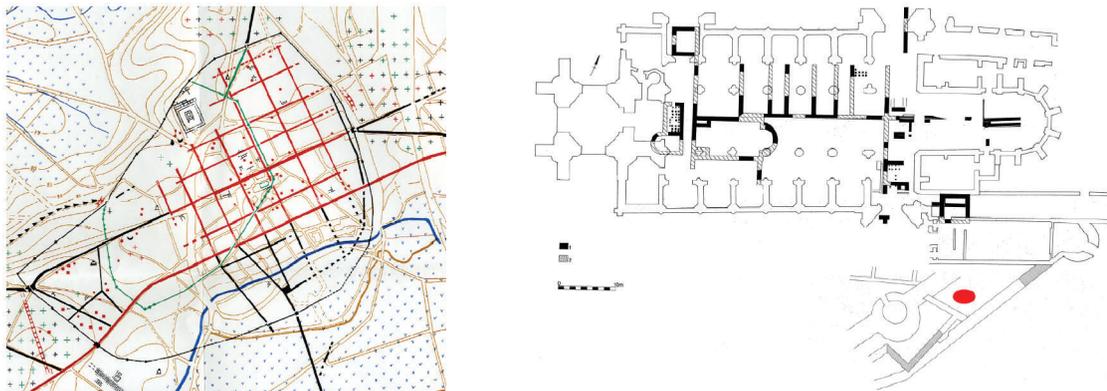
Figure 4. Left: Map of Tongeren in Roman and modern times. In green the fourth century city wall and the Early Medieval church (*Drawing Agentschap Onroerend Erfgoed [VIOE], after [8], figs. 17 and 34*). Right: Plan of the Roman domus below the Church of Our Lady and the remains of the fourth century city wall excavated on the Vrijthof square. The plaster fragments were found at the red dot, just outside the city walls (Copyright VIOE).
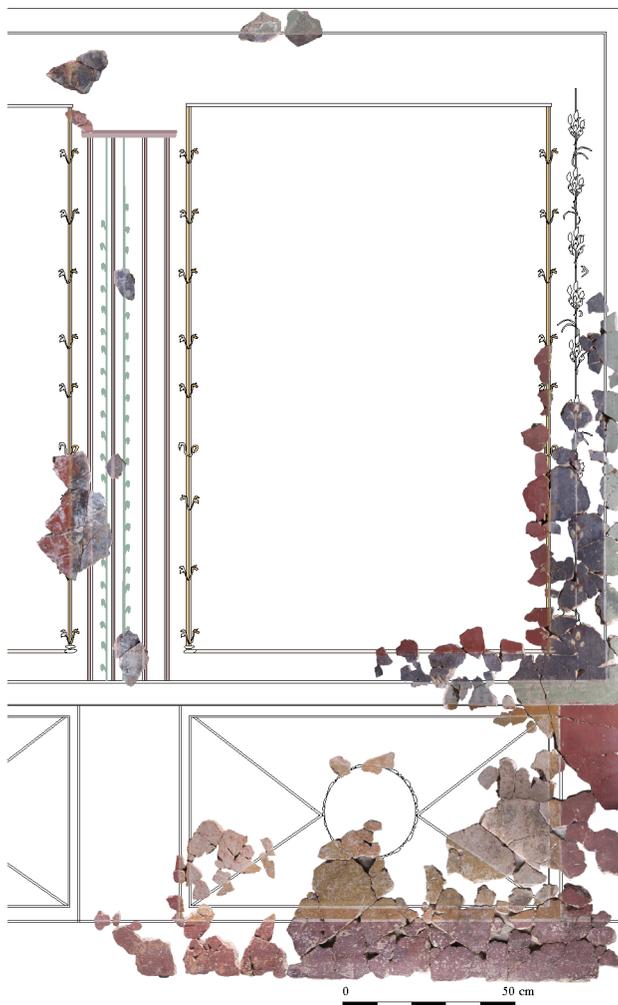

Figure 5. Partial reconstruction of Vrijthof Wall Decoration 1, which we are trying to further complete by finding additional matches (*Copyright Gallo-Romeins Museum Tongeren*).

supposedly painted in symmetrical pairs, flanking either a central panel, or two other red panels and a richly decorated black field in the middle.

More than 450 fragments belonging to this decoration were found, of which about two thirds could be manually reassembled into bigger and smaller groups of matching fragments. How much time this required is difficult to estimate since it was assembled and studied simultaneously with several other decorations; however the effort for all decorations was restricted to only a few months spread over a period of more than two years. The groups of joining fragments have been photographed and integrated in a reconstruction drawing of the painted system, and the fragments have been physically integrated in a restored panel of one of the walls, ready for display in the Gallo-Roman Museum of Tongeren [9]. The remaining fragments are mainly monochromatic, or integration in the restored panel without knowing their exact location would not add anything to the reconstruction of the painting, so they were left out.

The manually assembled fragments are therefore largely those which are "easy" to match manually based on their decoration, while the unassembled ones are largely "hard-to-match" fragments. The latter consist mainly of red fragments, on which little manual matching effort has been spent. More effort has been spent on the imitation marble, so there are fewer of these fragments remaining, and we would expect to find fewer new matches among them.

Scanning Procedure

The scanning and matching procedures have been designed specifically for small, flat objects. Just as a person relies on the properties of the material—identifying the flat, front surface of a fresco, or the profile curve of a potsherd for example—so too can the computer. In fact, this is essential to achieve robust, automated operation. Thus, while parts of our system are very general, many others assume the object is flat. Although this is not always true of plaster fragments, it covers the vast majority of cases. Other systems have been proposed for other types of material, such as pottery [10], that exploit specific properties of their domains.

Data acquisition proceeds in several phases. Each fragment must be individually numbered, either before or as it is scanned. No specific numbering system is imposed, so existing identifiers can be used, or a method that is convenient to the particular fresco or institution may be adopted. The fragment shape is then acquired with the 3D scanner, and its color and surface texture with the 2D (flatbed) scanner. For both procedures, the operator must enter the fragment's ID number and click a scan button. 3D scanning requires two sequences of scans

 Benedict Brown, Lara Laken, Philip Dutré, Luc Van Gool, Szymon Rusinkiewicz, Tim Weyrich

Figure 6. A fragment being scanned with the in-hand 3D scanner. We fully scan the front and back in approximately two minutes.
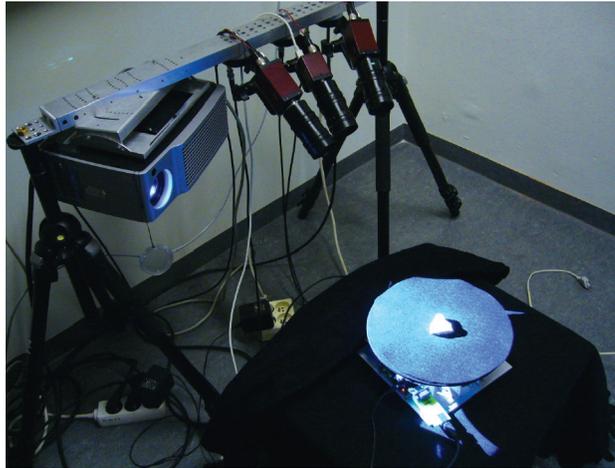
(front and back), while 2D scanning requires four scans of the front to acquire surface texture and one of the back. The computer automatically assembles all this data together into a complete model of the fragment which the user verifies and corrects as necessary, again using a specially designed program. These steps are described in more detail in Brown *et al.* [3].

A major drawback of the previous system was the commercial 3D scanner it used, which while inexpensive ($3000) is also slow and somewhat unreliable. On average, it took 10 minutes to fully scan a fragment in 3D. It was also difficult to scan fragments that were less than 1 cm thick, less than 2 cm in diameter, or more than 10 cm in diameter. Furthermore it was difficult to integrate into a custom fragment scanning application because we did not control its interface to the computer. Of these, the speed problem can be mitigated by operating several scanners simultaneously, but this increases the costs and the likelihood of entering the wrong ID number for a fragment. Higher speed commercial 3D scanners typically cost upwards of $25000, especially if they come with the technical support necessary to integrate them into a custom scanning application.

For the Vrijthof fresco, we have adopted the in-hand scanner design of Weise et al. [11], which uses three cameras and a projector to obtain 3D scans in real time (Figure 6). This scanner is fast enough that the user can simply hold an object in the hand and rotate it in front of the scanner. Nevertheless, we continue to place fragments on a motorized turntable as in the previous system because it is convenient and robust, and allows us to obtain higher quality data. The resulting system is approximately five times faster than its predecessor, and we have been able to scan both thin, tiny fragments, and large, thick fragments (20 cm in diameter and several centimeters thick) without any adjustment to the scanner settings.

Because we built this scanner ourselves, we have all the freedom necessary to optimize its setup and operation for scanning fresco fragments and can integrate it tightly into our scanning application. On the other hand, research prototypes are never as polished as commercial products, and this scanner takes considerable expertise to set up and calibrate correctly. Once configured, however, no special expertise is required to operate it.

Matching

There are many ways to match fragments. Perhaps the most obvious are to assume fragments found near each other are likely to match and to look for matching color or decorative motifs. However fragments may be dispersed over a wide area, and fragments from multiple walls may be collected in one location. Furthermore, large, solid-color areas are common. Subtler cues such as mortar type and plaster surface texture can help narrow the search, but frequently an exhaustive search among dozens or hundreds of similar fragments to find pairs whose *edges* fit together tightly is necessary. Doing so is not only time-consuming, but necessarily poses a risk to the original material.

All of the cues listed in the previous paragraph can be simulated by the computer using the scanned data. For the Vrijthof fresco, we continue to rely on finding pairs of edges that lock together, although incorporating other cues—especially surface color—would be very beneficial. The computer can very rapidly test every possible configuration of every pair of fragments, reporting only those possibilities that are both inherently plausible and among the best possible configurations for the specific pair. The user then examines these possibilities, and picks out the plausible candidates to physically test on the original material.

At the current stage the computer does not find *every* match, nor are its proposals particularly accurate. For the Vrijthof fresco, the computer proposed 6103 matches, of which we have confirmed 18 are correct. This may not sound helpful, but with appropriate visualizations it is possible to examine all proposals in only a few hours.
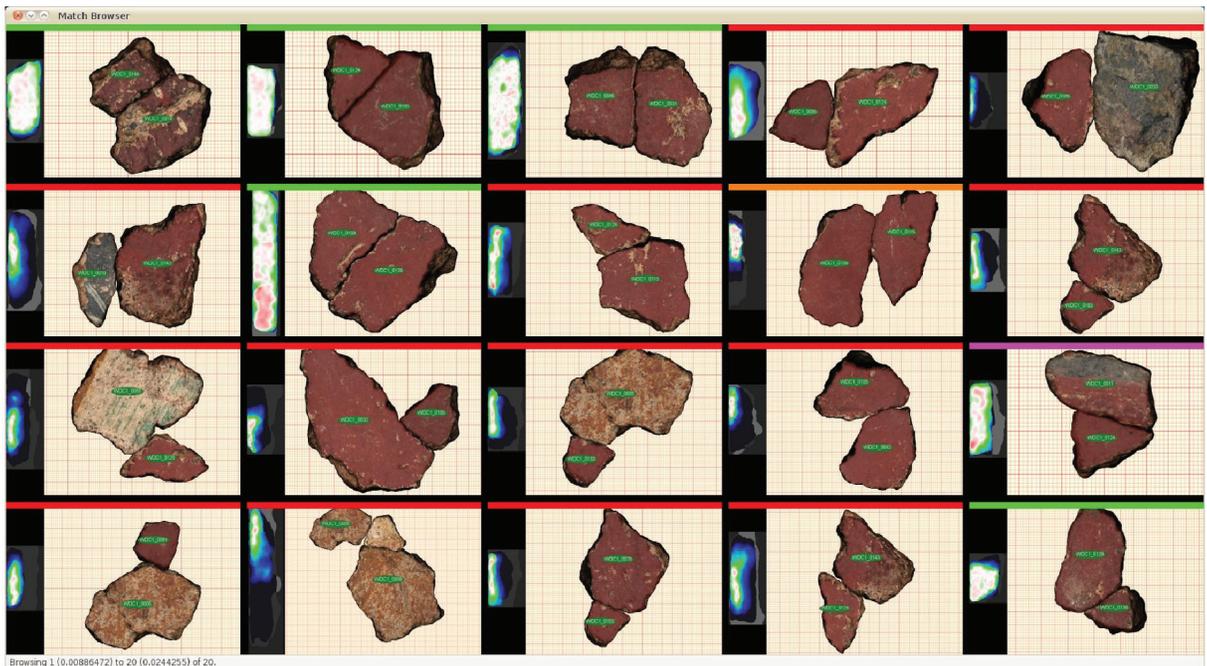
To sift through the proposals, we have developed a program that displays thumbnails of 20 proposed matches at a time (Figure 7). Each match is accompanied by a cross section of how the fragments fit together, color-coded so that white means they touch exactly, red means they intersect slightly, and green and blue indicate small and large gaps. Correct matches usually have a cross-section that is mostly white with wisps of red (some slight intersection results from the way the computer calculates the configurations).

| Benedict Brown, Lara Laken, Philip Dutré, Luc Van Gool, Szymon Rusinkiewicz, Tim Weyrich

Within this program, the user can mark a proposed match as correct, possible, impossible, or unknown (the default state). On the assumption that most proposals are obviously wrong, it is also possible to mark *all* proposals on view as wrong with a single keystroke. For the most part, it therefore takes only a second or two to evaluate 20 proposals. An additional important feature is that the program automatically calculates which configurations are mutually incompatible, or conflicting (because they would result in overlapping fragments). When a proposal is marked as correct, all incompatible proposals are automatically marked as conflicting. The user can also view the thumbnails of all proposals that conflict with a selected one.

This conflict analysis helps resolve the common case of fragments with short, straight edges that appear visually to match many things. It is precisely the plausible, yet incorrect, proposals that take the most time to evaluate, and viewing the list of all conflicting proposals has proven very effective at finding the correct match among them or determining that all are incorrect.

Efficiently evaluating matches with the match browser requires experience both with the fresco matching problem and the visualizations. Therefore, while it is easy to use the program, a new user can be expected to test many more proposals on the original fragments than an experienced user. With access to the original material for physically testing matches, we expect many users to become proficient after only a few sessions. For a student or intern, we

Figure 7. The match browsing application displays thumbnails of proposed matches along with a cross section of the matching area. The user marks correct matches (green bar), possible matches (orange bar), and incorrect matches (red bar). When the user confirms a match, all conflicting matches are marked with a purple bar. For instance, the 13th proposal above is impossible because it occupies the same edge of fragment WDC1_0124 as the second match, which has been confirmed correct.

believe this represents a reasonable learning curve. However, one could imagine asking the members of the public to evaluate match proposals via an on-line game or paid mechanism such as Amazon Mechanical Turk. In such a case, we would expect the proposals deemed probable by the general public to be overbroad; a final review by an experienced user would easily solve this problem.

Results and Discussion

As part of our ongoing project to reconstruct the extant material from all Vrijthof wall paintings, we have scanned two crates of fragments, comprising 144 fragments. The scanning and processing was performed by a computer scientist, and required approximately two days. For the next phase, we replaced an outdated flatbed scanner with a newer, faster one and slightly improved the speed of the 3D scanner. The remaining 28 crates of fragments were being scanned by an undergraduate archaeology student and a computer scientist in about one month.

The two crates of unassembled fragments from Wall Decoration System 1 (Figure 1) discussed in this paper have been examined to a limited extent, and obviously matching fragments have either been glued together (in which case we treat them as a single fragment for our purposes) or marked with chalk or charcoal (in which case we hope the computer proposes the match also). The fragments are predominantly red (about 100), containing one previously known (but unglued) match, and little time has been spent searching these fragments manually. The remaining fragments are purple, orange, gray, yellow, green, and black; a few contain stripes. They contain at least five known matches.

The matching algorithm proposed 6103 matches. Using our match browsing interface, a computer scientist was able to review the first 4000 proposals in approximately two hours, testing likely proposals on the actual fragments and confirming 18 correct matches. Of those, only four matches were previously known. On the other hand, only one unknown match came from a group of fragments that had been extensively analyzed, and the computer did not find two of the previously known matches.

Although 18 matches out of 4000 proposals may seem like few, the fundamental question is the efficiency of this process compared to manually searching for matches. An archaeologist therefore reviewed each of the found matches to subjectively determine whether it would have been "easy" or "hard" to find by hand. Some specific examples are discussed below.

Our conclusion is that the system finds matches considerably more efficiently than a person when the decoration does not provide clear

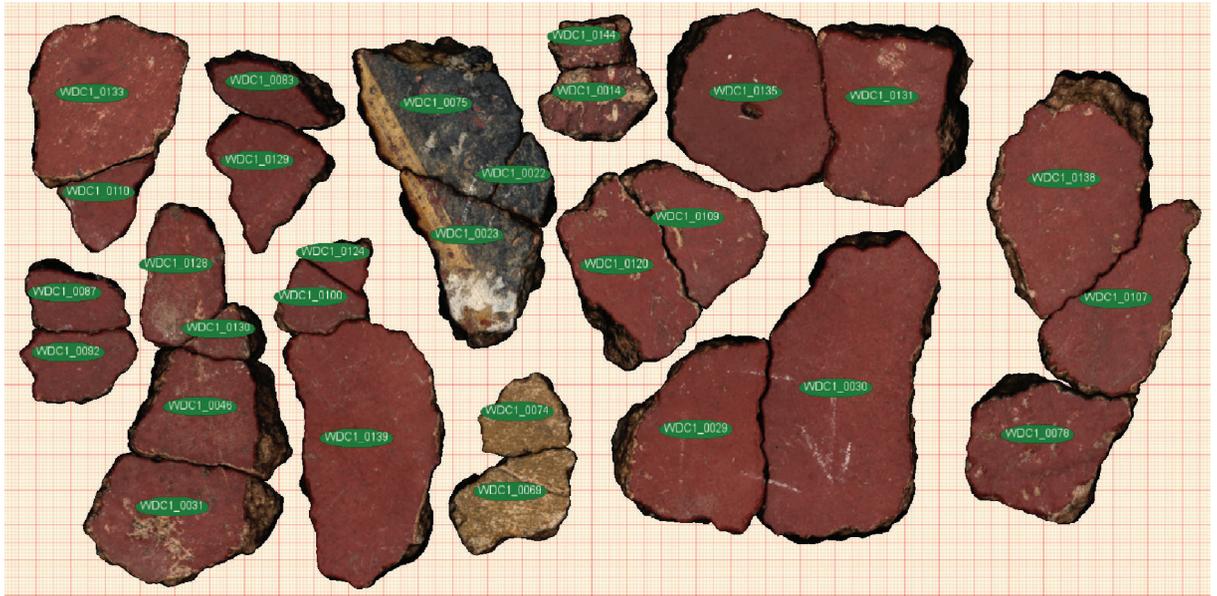Benedict Brown, Lara Laken, Philip Dutré, Luc Van Gool, Szymon Rusinkiewicz, Tim Weyrich

Figure 8: Using our scanning system, we were able to find the 18 matches among 29 fragments shown here. Of these matches, only the three striped fragments, and match 29 – 30 (marked with chalk) were previously known. The thickest red lines in the background millimeter grid represent 1 cm.

cues, although we do not pretend it will eliminate all manual matching. For example, among the 100 red fragments we scanned, it would be unrealistic to manually find 18 "easy" matches in two days, simply because there are so many possibilities to try. On top of that, many of the proposed matches were deemed hard or essentially impossible to find by hand because the fragments' shapes do not provide any indication that they might match. Adding in the benefits of a high quality fragment database, the broad range of people that can perform this work, and reduced impact on the source material, we believe the system in its current form already delivers a substantial benefit.
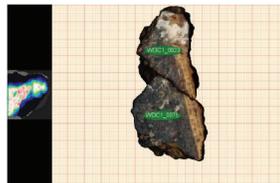

23 – 75 This is one of the four matching pairs that had been found but not glued together before scanning. The stripe on the front surface makes this an "easy" match to find, even among a large set of fragments. The computer probably didn't provide any matching benefit here, even if acquisition is useful unto itself.
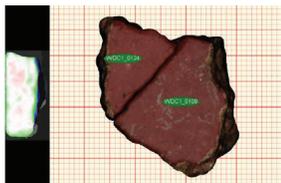
100 – 124 Although these fragments are small, the cross section shows the match is clearly correct (mostly white interspersed with light red). Evaluation of the fragments shows the fracture is very straight with no erosion, indicative of a "recent" break. Due to the clean break, an archaeologist would likely find this match without computer assistance, although with difficulty due to the number of red fragments.
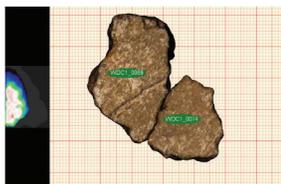
69 – 74 This pair comes from a set of approximately twenty, originally yellow fragments on which considerable manual assembly effort has been spent without finding anything. We therefore have objective confirmation that this is a "hard" match. The short, featureless matching edge is typical of many fragment matches, and is a good example of why we think the computer is a powerful tool to aid in reassembly.
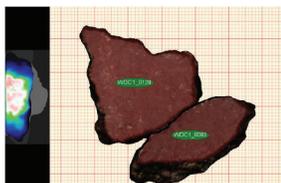
83 – 129 The matching edge on this pair has only a slight curve, and does not follow any natural corners or features on the fragment edges. As a result, this is a "very hard" match, that would be nearly impossible to find by hand.

78 – 107 Even knowing these fragments match, an experienced archaeologist was unable to fit them together without referring to the thumbnail. This match would likely only be found as part of a larger assemblage of fragments or if they were excavated (and documented) in the correct configuration.

The partially reconstructed state of this fresco is quite common, making it important to not only assemble the remaining fragments, but also to automatically attempt to match them to the previously assembled sections. Since we obviously cannot scan reassembled, mounted frescoes, our current system is unable to propose such

| Benedict Brown, Lara Laken, Philip Dutré, Luc Van Gool, Szymon Rusinkiewicz, Tim Weyrich
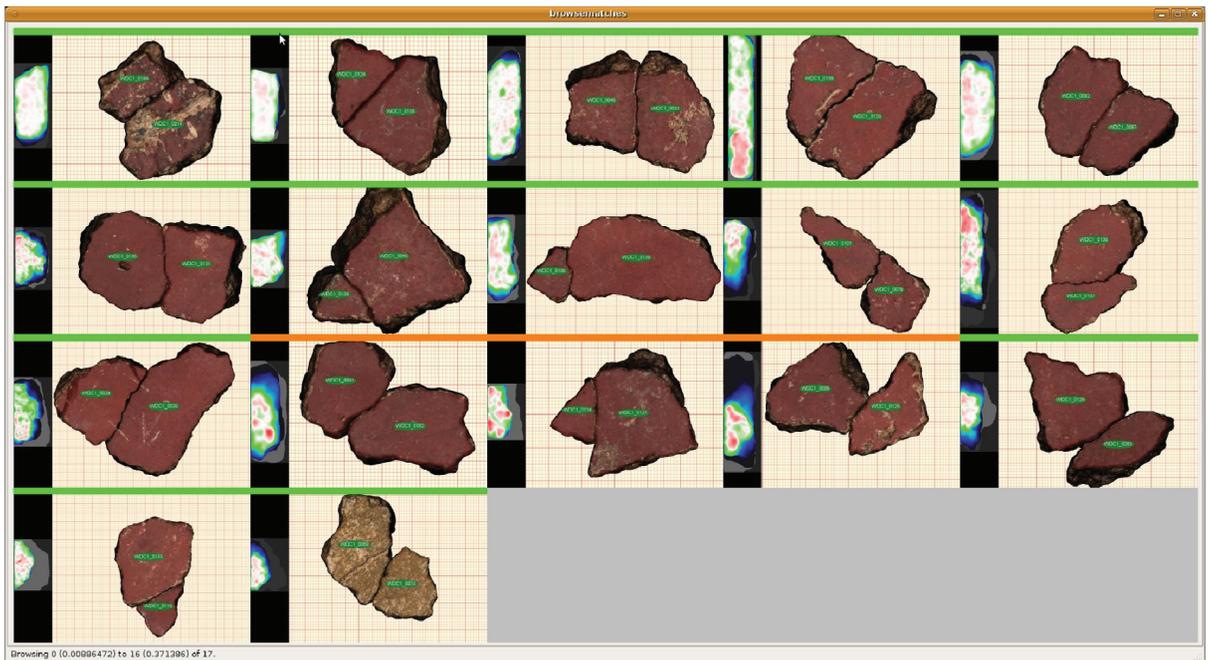
Figure 9. A computer scientist was able to identify and confirm these 13 matches in about two hours, as well as identify three potential matches (those with an orange stripe) that required expert verification. All three potential matches proved incorrect, but five more correct matches were found with additional searching. The total time to acquire data and examine match proposals was two days.

matches. However we can obtain information on the coloration, decoration, and outlines of reconstructed sections from photographs and by incorporating matching techniques based on this information into our system, we expect to be able to support matching to previously assembled frescoes as well.

The overwhelming majority of matches that our system proposes are implausible. While this is mitigated by the ability to rapidly scan large numbers of proposals in our match browser, it is of course preferable to propose mainly correct—or at least plausible—matches. As the overall number of fragments increases, the number of proposals quickly explodes, making improved relevance of match proposals essential. We are working to improve the quality of our matching algorithms by improving the ranking of proposals based on edge shape, incorporating additional cues [4], and searching for *clusters* of matching fragments as opposed to only pairs. For example, the match proposals for the Vrijthof fresco would be much better if we also considered the fragment color.

Not all fragments in the Vrijthof fresco include the front surface. In some cases where the upper and lower layers of mortar have separated, the *front* of one fragment may match the *back* of another. Similarly two fragments may have matching edges, but front surfaces will not line up because one fragment is missing its upper mortar layer. Our current matching approach cannot handle these cases. In other cases, there is mud or debris attached to the front surface, confusing

our algorithm. These cases cover only a small number of fragments in our current fresco, and do not seem to be a major hindrance. However, new algorithms may be necessary to handle frescoes with many such fragments.

Acknowledgments

References

[1] A. Willis and D. Cooper. "Computational reconstruction of ancient artifacts," *IEEE Signal Processing Magazine* 25:4, June 2008, 65–83.

[2] F. Kleber and R. Sablatnig. "A survey of techniques for document and archaeology artefact reconstruction," *10th International Conference on Document Analysis and Recognition*, July 2009, 1061 – 1065.

[3] B. Brown, C. Toler-Franklin, D. Nehab, M. Burns, D. Dobkin, A. Vlachopoulos, C. Doumas, S. Rusinkiewicz, and T. Weyrich. "A system for high-volume acquisition and matching of fresco fragments: Reassembling theran wall paintings," *ACM Transactions on Graphics* (Proc. SIGGRAPH) 27:3, 2008.

[4] C. Toler-Franklin, B. Brown, T. Weyrich, T. Funkhouser, and S. Rusinkiewicz, "Multi-feature matching of fresco fragments," *ACM Transactions on Graphics* (Proc. SIGGRAPH Asia) 29:6, 2010.

[5] A. Vanderhoeven and G. Vynckier. "Archeologische kroniek van Tongeren: De opgraving aan het Vrijthof," *Tongerse Annalen* 9:4, 1995, 59–61.

[6] L. Laken, "Wall-paintings in Atuatuca Tungrorum: Preliminary report on the plaster fragments from the Vrijthof in Tongeren (Belgium)," in *Atti del X congresso dell'AIPMA (Association Internationale pour la Peinture Murale Antique), Napoli 17–21 settembre 2007,* a cura di Irene Bragantini, vol. II, Napoli 2010, 865–869.

[7] A. Arts, V. Vander Ginst, A. Vanderhoeven, and G. Vynckier. "Het archeologisch onderzoek in de O.L.V.-basiliek van Tongeren. De Romeinse en vroegmiddeleeuwse sporen," *Monumenten, Landschappen en Archeologie* 28:1, 2009, 8–31.

[8] W. Vanvinckenroye, *Tongeren, Romeinse stad.* Tielt, 1985.

[9] R. Nunes Pedroso, "Tongres 'Vrij[t]hof', Musée Archéologique de Tongres: Rapport de restauration," (Centre d'Etude des Peintures Murales Romaines, rapport interne), Soissons, 2010.

[10] A. Karasik and U. Smilansky. "3D scanning technology as a standard archaeological tool for pottery analysis: practice and theory," *Journal of Archaeological Science* 32:6, 2008, 841–853.

[11] T. Weise, B. Leibe, and L. Van Gool, "Fast 3D scanning with automatic motion compensation," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2007.