The MURALE project: Image-based 3D modeling for archaeology

Luc Van Gool[1,3], Marc Pollefeys[1], Marc Proesmans[2], and Alexey Zalesny[3]
[1] ESAT-PSI, Kath. Univ. Leuven, firstname.lastname@esat.kuleuven.ac.be
[2] Eyetronics NV, Leuven and Los Angeles, proes@eyetronics.com
[3] IKT-BIWI, D-ELEK, ETH Zurich, lastname@vision.ee.ethz.ch

Abstract

Murale is a European IST project that will develop 3D capture and visualisation technology for archaeology. The project will put special emphasis on the usability on the site, by the archaeologists themselves. The paper describes techniques that are being developed by three of the Murale partners in particular. These comprise two methods to generate 3D models of objects, and approaches to deal with the textures of materials and terrain. Put together with the database and visualisation expertise brought in by the other partners, Murale will not only contribute to the enhanced visualisation of archaeological sites and finds, but also to a faster and more complete documentation of the progress of excavations. The ancient city of Sagalassos, one of the major excavation sites in the eastern part of the Mediterranean, will be used as the primary test site.

1 The MURALE project

This paper describes the planned contributions of MURALE, an IST (Information Society and Technology) project funded by the European Commission in order to advance the use of computer technology in archaeology.

The MURALE consortium consists of the following partners: Brunel university (UK), ETH Zurich (Switzerland), Eyetronics (Belgium), Imagination (Austria), the Technical University of Vienna (Austria), the University of Graz (Austria), and the University of Leuven (Belgium). The main areas of expertise of its researchers are archaeology, computer vision, and computer graphics. MURALE is about the development of technology, but from the start also wants to focus on its practical application by archaeologists on a test site. This site is the ancient city of Sagalassos, in what is now the southern part of Turkey.

The site at Sagalassos is one of the largest archaeological projects in the Mediterranean dealing with a Greco-Roman site over a period of more than a thousand years (4th century BC-7th century AD). One of the three greatest cities of ancient Pisidia, Sagalassos lies 7 km north of the village Aglasun in the province of Burdur, Turkey. Fig. 1 shows this location in more detail.

Figure 1: Sagalassos, the primary test site of the Murale project, lies in southern Turkey.

The ruins of the city lie on the southern flank of the Aglasun mountain ridge (a part of the Taurus-mountains) at a height between 1400 and 1650 metres. Fig. 2 shows the valley with Sagalassos against the mountain flank.

Figure 2: Overview of the Sagalassos site.

A team of the University of Leuven under the direction of professor Marc Waelkens has been excavating the whole area since 1990.

The issues of how to record and visualise the finds of excavations are as old as archaeology itself. With the advent of computer technology old procedures can be eased and even improved. At least as important will be the paradigm shifts that will follow. In keeping with the archaeological applications targeted in this paper, the concept of `typological rudiments' comes to mind. New technologies are often used at first as a pure substitute of older ones, to produce copies of the old tools, with improved characteristics. It is only after a while that the design of the tools themselves is adapted to fully capitalize on the added potential of the new technology. Several aspects of MURALE will allow archaeologists to solve old tasks with new means, while some others already add components that had been out of reach before.

These are the goals that MURALE plans to demonstrate at Sagalassos:

1. Provide tools that archaeologists themselves can use *in situ*. This is a crucial difference with some other efforts. All devices needed should be easy to bring to the site, they should work robustly under conditions that can be quite adverse (sun heat, dust, moisture), they should be easy to operate and carry around. Most excavation campaigns have little financial means and, hence, the devices should also be cheap. Phrasing it in information technology parlance, also in the era of computer wizards it should be the archaeologists who are firmly in control of content creation.

2. Generate 3D models of objects at different scales: from landscapes, over buildings and statuary, to small finds such as pot sherds. The acquisition should be fast and flexible enough not to interrupt the excavation longer than any of the traditional methods would. New technology should rather speed up work. The novel 3D technology should also introduce new possibilities. An example is matching parts for virtual restoration / anastylosis, possibly as a first step towards physical restoration. A second example is visualisation of the site's temporal evolution. Currently, the public all too often is given a single image of how the site has been or only some snapshots over time, which are difficult to relate to each other mentally. A third example is the 3D recording of newly excavated layers. Last but not least, there is the most obvious use of the 3D data, namely the realistic visualisation of the scene. This is useful for the public and the archaeologists alike. Archaeologists may e.g. use the terrain model to guess the positions of invisible infrastructure such as long gone roads or they may use the 3D city models to understand why fortifications like towers where built precisely at the spot where their ruins were found.

3. Provide a database, in which items can be stored efficiently and that supports efficient retrieval. Such databases can be made available over the internet. This turns them into powerful tools, as archaeologists want to compare their finds against those found elsewhere. In the case of Murale the emphasis will be put on potsherds, as Sagalassos was a major production site of pottery. Other excavations can use the elaborated typology and corresponding time scale to date layers in which Sagalassos ceramic ware is found. In order to compare pot shapes special search tools will be developed.

4. Integration of the 3D data and the database will support different functions. First, by having age information with the data, visualisation for different time periods is made possible. When the user selects a certain time period, the system itself can go out and look for buildings, objects, etc. from the specified period. Another way of linking the database with the 3D models is to provide annotations. If the user clicks on or points at parts of the scene, additional information can be popped up, e.g. as text or images, or text-to-speech can be used to further clarify what is shown.

In summary, Murale hopes not only to offer the public a more convincing and enticing impression of how Sagalassos developed over the centuries, but just as much to provide the archaeologists with tools that are effective and efficient in the field.

In this paper the emphasis is on work carried out by three of the partners: ETH Zurich, Eyetronics, and the University of Leuven. The results mainly pertain to 3D acquisition and visualisation:

1. 3D model extraction from video data
2. A 3D photo camera
3. Analysing and synthesising textures with 3D effects

These topics are described in sections 2, 3, and 4, resp. Although in the end, the Murale demonstrator will include 3D models of the terrain (landscape), of the existing ruins, of the statuary (sculptures and ornaments), and of the different finds such as pottery, only preliminary models can be shown at this point, as the project only started shortly before the time of writing.

2 Shape from video

The University of Leuven has developed a technique that allows producing 3D, textured models from multiple images as the only input. The images are supposed to be taken from different viewpoints as one walks around the scene of interest. Therefore, the only thing the archaeologists have to do is to take these images, either as a collection of subsequent stills, or as a video sequence. This can in general be done rather quickly and easily.

The only apparatus required is a normal camera and, of course, a computer to run the software that turns the images into 3D models. In contrast to what is required by most traditional photogrammetric and computer vision approaches, the motion and parameters of the camera need not be known. Motion parameters (rotation and translation) are typically called `extrinsic' camera parameters. Camera-specific parameters such as pixel dimensions and focal length are referred to as `intrinsic' parameters. The 3D modeling software extracts these parameters automatically, together with the 3D shape of the scene. As a result, also existing footage can be used to model scenes in 3D that no longer exist.

Much along the lines of work by Armstrong *et al.* [1], the method starts with the automatic tracking of image features over the different views. This is done in stages. First, a (Harris) corner detector is applied to yield a limited set of initial correspondences. These correspondences are found by correlating small neighbourhoods around corners in the same region in the next image. This only works if the motion between subsequent views is small. Therefore, methods are being developed to find correspondences between images that are taken from viewpoints that are farther apart. These are based on invariant neighbourhoods, image patches that automatically change their shape with viewpoint in order to systematically cover corresponding patches on objects' surfaces [13].

Once correspondences between the primary features have been established, the system can put in place geometric constraints that facilitate the search for further correspondences between points surrounded with a less salient image pattern than the corners. An example is the epipolar constraints, which serve to restrict the search for corresponding points to a line in the other image. The system tries to generate dense correspondences, i.e. to find a correspondence for almost all pixels [9]. Points, for which correspondences can be found, can in the final step be reconstructed in 3D.

The limited set of corner correspondences also yields the necessary data to perform a fully automated calibration of the camera and hence to determine the camera projection matrices for its different, subsequent positions. Once these matrices are available, the 3D model of the observed scene can be produced. Earlier versions of the system required that intrinsic camera parameters like the focal length remained fixed during image acquisition if one was to arrive at metric structure, i.e. to obtain a 3D model that is correct up to a scale. But if one has limited *a priori* knowledge about some intrinsic parameters, like a known pixel aspect ratio or the fact that rows and columns in the images are orthogonal, then others like focal length can be allowed to change [7, 11, 12]. Sagalassos will become an interesting testing ground to validate and test the robustness and quality of such self-calibrating 3D modeling process.

Fig. 3 shows a part of a bathhouse that has been modeled in 3D using this shape-from-video technique.

Figure 3: Top: three views of the Roman bathhouse at Sagalassos. Bottom: three views of the model constructed from 6 images of which 3 are the ones in the top row.

It shows 3 of 6 images that were taken of the Roman bathhouse, and then used for the creation of its 3D model. The bottom row shows 3 views of this model. For the moment only few of such models have been produced, but it is the goal of the project to record in 3D several of the ruins in this way. The same method will be applied to model the Sagalassos landscape. Several images have already been taken along the top of a hill overlooking the excavation site. This already yielded a global model of the valley. In all cases the intrinsic and extrinsic parameters of the camera were unknown.

Under Murale, another type of 3D modeling will be used, that is of particular importance to archaeology. Excavations are carried out layer by layer. Although drawings and photographs are made to document the course of action, a full 3D representation would be preferable. Hence, one of the goals of Murale is exactly this: to build 3D models of the different layers, which can subsequently be virtually scrolled through, with layers being dynamically added or removed, with visualisation from any viewpoint. Again, the archaeologists only have to take images, as is already current practice.

Figure 4: The shape-from-video approach allows the archaeologists to reconstruct the stratigraphy found at the excavations from the same photographs they already use to document their finds.

Fig. 4 shows two views of the same excavation, one before and one after an additional layer had been removed.

Another application that Murale will investigate is the use of 3D acquisition technology for the support of virtual or real restoration and anastylosis, i.e. to use the 3D shapes of building blocks, sherds, and pieces in general to see how they can fit together. If the building or the artefact to which the pieces belong is of high scientific or artistic value a real restoration can then follow.
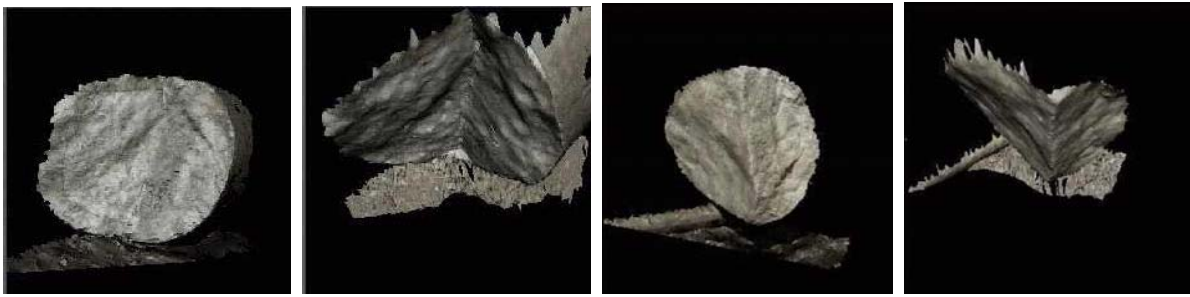
Figure 5: 3D models of the fractured surface of broken pillars, which tumbled down during the earthquake that destroyed Sagalassos in the 7th century.

Fig. 5 shows 3D models of the fractured surfaces of a series of pillars that tumbled down and broke as a result of the 7th century earthquake. These models will be used to test how the pieces could be puzzled together. Rather than having to physically move around the pieces - each weigh several tons - the computer can first look for good matches. Only when successfully can the pieces then be brought together, without trial and error at that stage.

As the method produces the list of intrinsic and extrinsic camera parameters one could also add virtual objects to the video sequences that were used as input. The University of Leuven has just started to explore such `augmented reality'. Such techniques can be used to give the general public an idea of how the city looked liked at different times, by dynamically superimposing 3D CAD reconstructions of buildings, streets, walls, etc on videos of the site in its current state. Fig. 6 shows a preliminary example.
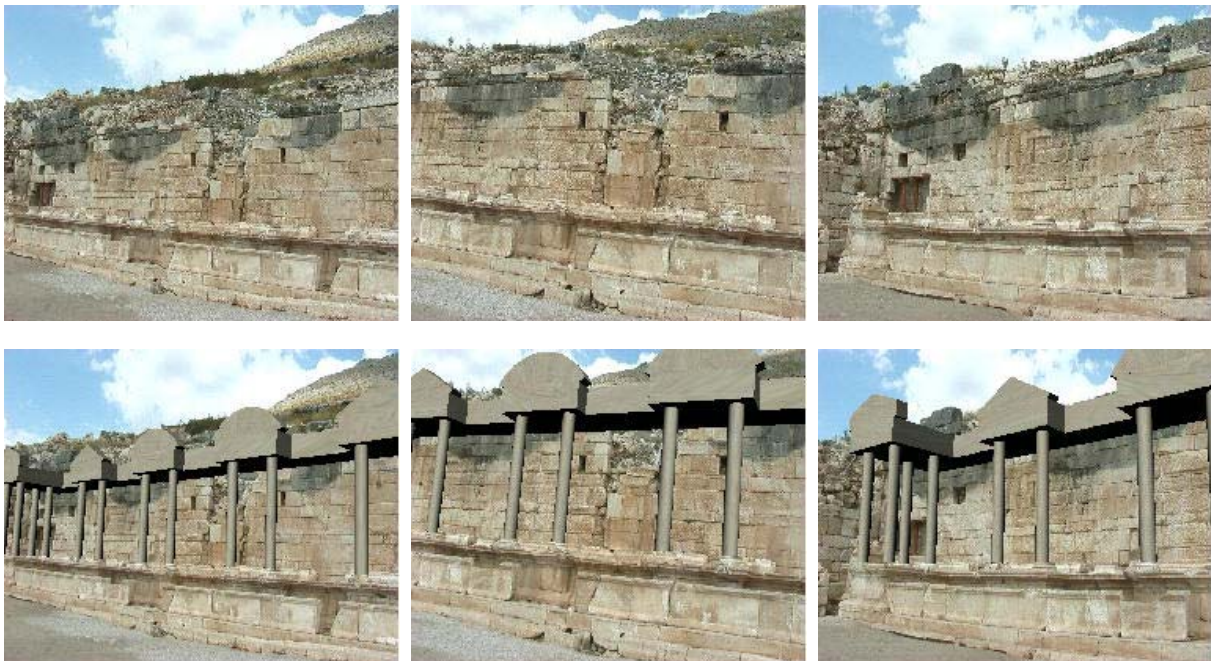

Figure 6: Top row: original frames of a video taken at Sagalassos. Bottom row: with part of the original structures superimposed.

3 A 3D photo camera

The `passive' technique outlined in the previous section cannot deal with untextured parts of a scene. Its geometric precision is also limited. These are major problems with objects such as statues or ornaments, which typically consist of untextured stone, but the shape of which should be extracted with high precision. For the 3D shape extraction of such objects, Eyetronics brings in its `active' ShapeSnatcher system [4].

`Active' systems bypass the problem of insufficient texture by projecting a pattern onto the scene. The 3D shape is extracted by analysing the displacements/deformations of the pattern when observed from a different direction (see Jarvis [8] and Besl [2] for an overview). Typically, such methods have relied on the projection of single points or lines

and on scanning the scene to gradually build a 3D description point by point or line by line. Eyetronics' ShapeSnatcher contains no moving parts as it directly projects and analyses a complete line grid instead. The components needed are a normal slide projector, camera, and computer. The grid, provided on a slide, is projected onto the scene, which is observed from a different viewing angle by the camera. The software then allows extracting the 3D shape of the complete patch visible to both. The angle between the viewing and projection directions can be kept pretty low, around 7 to 10 degrees, such that as small a part of the surface as possible goes lost to occlusions. With the ShapeSnatcher the 3D positions of up to 360,000 points can be extracted from a single image. Thus far, approaches that projected several lines or other patterns simultaneously had heavily depended on the inclusion of a code into the projected structure, which kept the resolution of the extracted shapes low [3, 14, 10].

Fig. 7 shows the set-up and a detail of an image from which 3D information can be extracted.

Figure 7: Top: The active system only consists of a normal slide projector and camera, and a computer. The camera takes an image from a direction that is slightly different from the direction of projection. Bottom: A regular square pattern is projected on the scene, as seen in this detailed view. In this case, the grid covers the complete face. 3D co-ordinates are calculated for all the line intersections, resulting in the simultaneous measurement for thousands of points.

The result is not a mere 3D-point cloud, but a connected surface description (the grid in 3D). In order to also extract the surface texture, the lines of the grid are filtered out. Obviously, an alternative for static objects is to take another image without the grid. This is supported by the software and often done in an archaeological setting, where objects are static and one would like to have the best texture information possible. Although not very salient, the reduction in texture quality when obtained by filtering out the lines, is an issue.

Fig. 8 shows the partial 3D reconstruction of a statue of Dionysos, the god of fertility and wine, that has been excavated at Sagalassos.
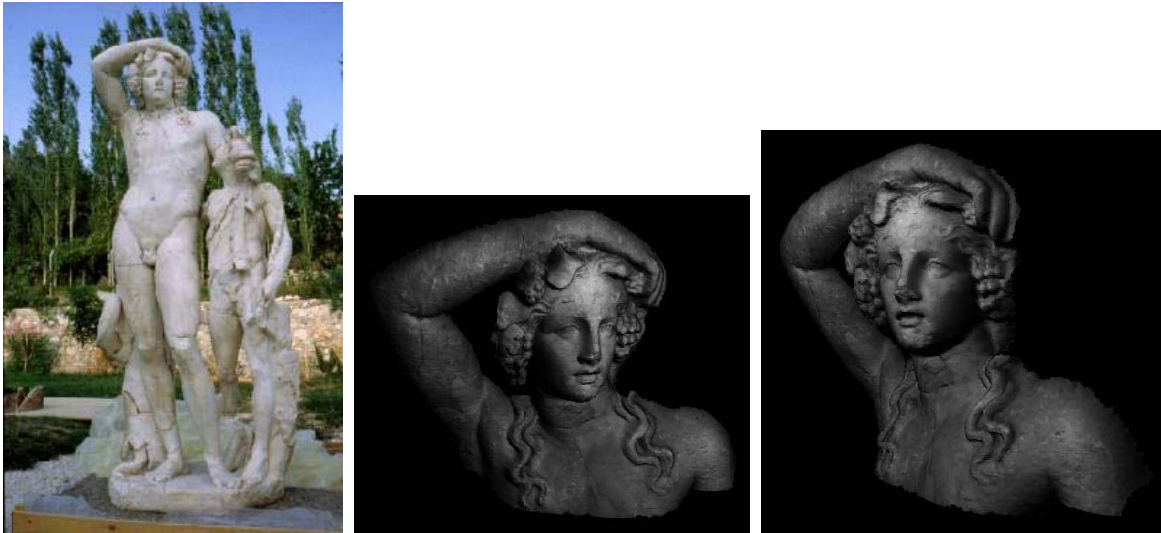

Figure 8: Left: Statue of Dionysos, Middle and Right: Two views of the extracted 3D model.

Such objects obviously cannot be modeled from a single image. Several need to be combined in order to build a complete model. Eyetronics also provide their ShapeMatcher software, which supports the knitting of several, partial 3D patches. Given the fact that such (several tons) heavy objects need to be modeled, it is important that the acquisition system can be brought to the pieces and not vice versa. It would for instance be difficult to put such a statue into the working volume of a laser scanner, if it would fit in there at all and if the scanner would be able to withstand the pressure.

Although a projector and camera already better fit the description of a portable system, Eyetronics will carry this aspect a step further. The projector is being replaced by a flash. This makes it easy to walk around a piece and to photograph it from whatever angle that is required, after which the pieces can be knitted together into a single object description. Eyetronics has already such a device, but competition by the sun poses extremely hard conditions on such 3D-photo camera.

A special set-up will be produced for potsherds. The outline has to be delineated precisely and the thickness of these parts has to be measured with good precision. In this case, speed is of particular importance. Literally millions of sherds are excavated each year in Sagalassos. Although it is obviously not the goal to extract the 3D shapes of all of these, there still is an appreciable number which are big enough to infer at least a substantial part of their silhouette, work that will be done by the Technical University of Vienna. It is current practice that experienced archaeologists in the team draw the silhouettes by hand. A disadvantage is that experiments have shown that the drawings of different experts may also differ quite substantially. Although some trials with laser based methods have been carried out in the past, the speed of such automated scanning in the end proved much lower than that of manual sketching, and the trials were aborted for that reason. A one-shot system like the ShapeSnatcher may overturn this problem. Fig. 9 shows preliminary results of acquired 3D shapes for one of the many sherds found on the site.

Figure 9: Left: A sherd found at Sagalassos. Middle: 3D reconstruction, seen from the `outside', with the external and internal surfaces already aligned. Right: View of the 3D reconstruction seen from the `inside'.

4 Image-based texture synthesis

Only a rather rough model of the landscape (terrain model) will be built. Its resolution will not match that of the 3D models of the ruins, and this is the case for both the geometry and the texture. Visualisation of virtual Sagalassos will require that one can freely navigate through the scene. As one moves from building to building and crosses the terrain in between, a noticeable and disturbing difference in visual quality between the buildings and the terrain would appear. ETH Zurich develops tools to map realistic texture on the terrain, and to model the textures of other surfaces such as those of the building materials used on the site.

Precise modeling of the landscape geometry and texture all over the tens of square kilometres spanned by the site would cost an enormous amount of time and memory space. Also, such precise modeling is not really required for general visualisation purposes. Then it typically suffices to cover the landscape with a texture that looks detailed and realistic, but that does not necessarily correspond to the real texture on that particular part of the site. There is no need to precisely capture every bush or natural stone. Thus, as a compromise we model the terrain texture on the basis of selected example images of real Sagalassos texture. The model is then covered with similar textures of the right type. Texture synthesis is based on texture models learnt automatically from the example images. Such example image is shown in fig. 10.

Figure 10: Image showing terrain texture at the Sagalassos site.

The resulting texture models are very compact and can be used to generate arbitrarily large patches of texture that look very similar to the example textures.

Textures are synthesised as to mimic the pairwise statistics of the example texture. This means that the joint probabilities of different colours at pairs of pixels with a fixed relative position are approximated as closely as possible. Just including all pairwise interactions in the model is not a viable approach and a good selection needs to be made [6]. ETH have opted for an approach that makes a selection as to keep this set minimal but that on the other hand brings the statistics of the synthesised textures very close to that of the example textures [15]. Interaction type selection follows an iterative approach, where pairwise interactions are added one by one to the texture model, the synthetic texture is each time updated accordingly, and the statistical difference between the example texture and the synthesised texture is analysed to decide on which further addition to make. The set of pairwise interactions selected for the model (from which textures are synthesised) is called the *neigbhourhood system*. The complete texture model consists of this

neighbourhood system and the *statistical parameter set*. The latter contains the joint probabilities for the selected relative pixel positions, also called *cliques*.

A sketch of the algorithm is as follows:

step 1: Collect the complete 2nd-order statistics for the example texture, i.e. the statistics of all pairwise interactions. (After this step the example texture is no longer needed) As a matter of fact, the current implementation doesn't start from all pairwise interactions, as it focuses on interactions between positions within a maximal distance.

step 2: Generate an image filled with independent noise and with values uniformly distributed in the range of the example texture. This noise image serves as the initial synthesised texture, to be refined in subsequent steps.

step 3: Collect the full pairwise statistics for the current synthesised image.

step 4: For each type of pairwise interaction, compare the statistics of the example texture and the synthesised texture and calculate their `distance'. For the statistics the intensity difference distribution (normalised histograms) were used and the distance was simply Euclidean. In fact, the colour distributions of the images were added also, where `singletons' played the role of an additional interaction. The current implementation uses image quantization with 32 gray levels.

step 5: Select the interaction type with the maximal distance (cf. step 4). If this distance is less than some threshold go to step 8 – the end of the algorithm. Otherwise add the interaction type to the current (initially empty) neighborhood system and all its statistical characteristics to the current (initially empty) texture parameter set.

step 6: Synthesize a new texture using the updated neighourhood system and texture parameter set

step 7: Go to step 3.

step 8: End of the algorithm.

After the 8-step analysis algorithm we have the final neighborhood system of the texture and its statistical parameter set. This model is very small compared to the complete 2nd-order statistics extracted in step 1. Typically only 10 to 40 pairwise interactions are included and the model amounts from a few hundreds to a few thousands bytes. Nevertheless, these models have proven effective for the synthesis of realistically looking textures of a wide variation. Fig. 11 shows a collage of textures (left) and the corresponding synthetic versions (right).
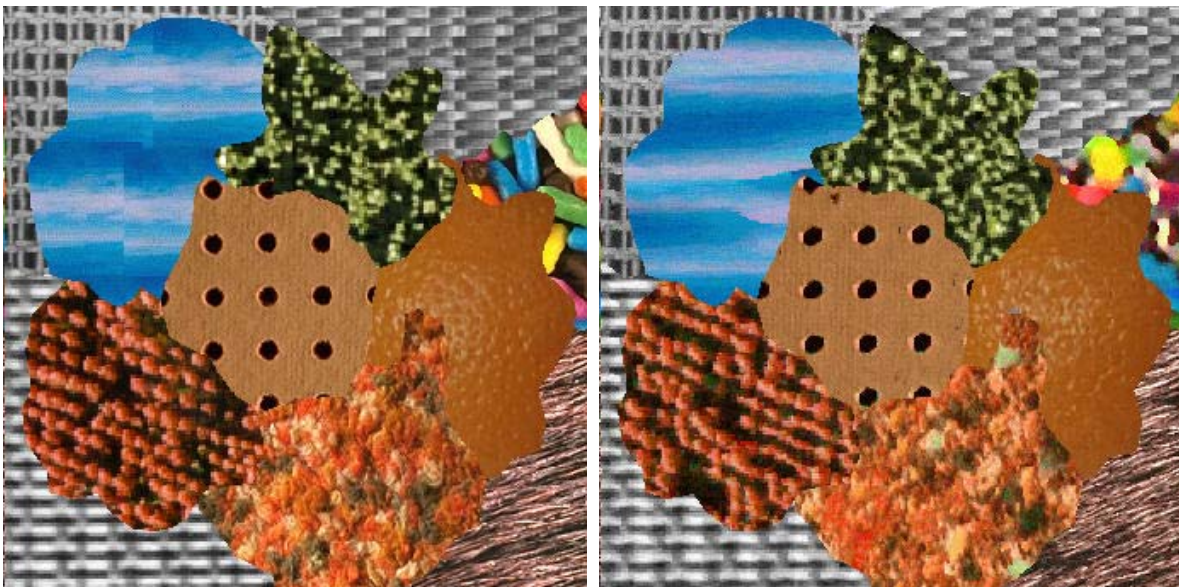

Figure 11: Left: collage of original textures. Right: collage of corresponding, synthetic textures.

This texture synthesis approach can handle quite broad classes of textures. Nevertheless, it has problems with capturing complex semantic orderings or texels with specific shapes. The method has mainly been used for coloured textures, as is also required for the Sagalassos virtual site. In the case of colour images pairwise interactions are added that combine

intensities of different bands. The shortest 4-neighborhood system and the vertical interband interactions were always preselected because experiments showed that they are important for the vast majority of the texture classes. Fig. 12 shows a synthesised textured for the example image in fig. 10.

Figure 12: A synthesised texture based on the example image of fig. 10.

Fig. 13 shows part of the site with the original terrain model texture (left) and with synthesised texture mapped onto the landscape (right).

Figure 13: View of the old bathhouse and surrounding landscape at Sagalassos. Left: view with the original landscape texture. As this is a view strongly zooms in onto this model, the texture is of insufficient quality. Right: the landscape texture has been replaced by synthetic texture.

It does not suffice to synthesise textures. Where different textures meet clear seams would appear. Hence, a texture knitting tool was developed. Consider fig. 14.

Figure 14: Example of texture knitting. Left: image comprising two types of Sagalassos texture, with rocks on top and grass below. Right: knitted rock and grass textures.

The image on the left shows a composition of rock and grass texture images, both taken at Sagalassos. The right image shows the result obtained with the texture knitting algorithm. Knitting is based on learning a texture model from the zone around the border between the two textures. Then, new texture is generated in a zone around the border, not necessarily the same, based on the border zone texture model. In the case of a single texture, the seams between separately generated patches can be removed by simply applying that texture's model near the border.

Future work on texture synthesis will include the simulation of 3D effects that appear once the virtual camera moves around through the scene. Texture typically finds its origin in fine-grain, 3D surface geometry. Effects such as self-occlusions cannot be modeled satisfactorily if a fixed texture is mapped onto surfaces. The texture analysis and synthesis method has already been extended so that models and hence also textures for different viewpoints can be generated. An example is shown in fig. 15.
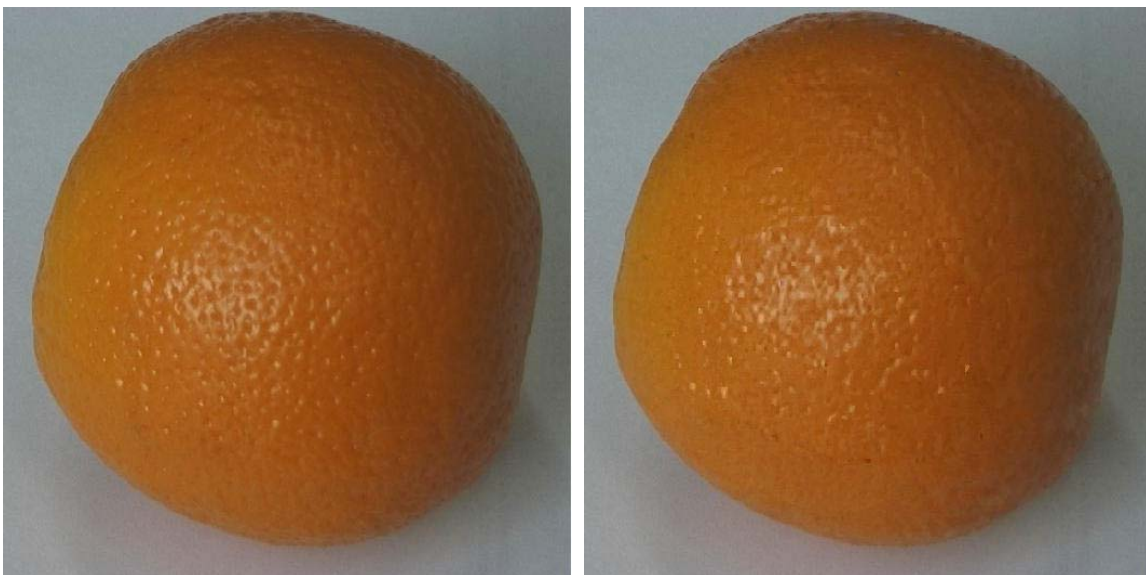

Figure 15: Left: a real orange (not found in Sagalassos;-), Right: synthetic version based on viewpoint dependent texture mapping on the orange's 3D shape.

The left figure shows an original image of an orange. The 3D shape of it was estimated. The right figure shows the result of mapping viewpoint dependent texture on this shape. The textures for the different orientations of the orange surface were learnt from the left image. The result looks quite realistic. Note that, although we have called this `viewpoint-dependent' textures, the effects of viewpoint and illumination direction are in fact confounded in such experiments. Further work to disentangle these influences is needed.

Under Murale this work will be enhanced by providing additional means to generate textures for different viewpoints that are mutually consistent, i.e. that may originate from a single, physical terrain structure. For now, changing viewpoint dynamically is not possible. The textures for different viewpoints are generated independently from each other. The creation of consistent, viewpoint-dependent textures is a challenging task.

5 Conclusions and future work

We have described the goals of the MURALE project. As it has only just started, the purpose of this paper mainly is to indicate in which directions the consortium plans to carry out its further research. In particular, 3 groups (ETH Zurich, Eyetronics, and the University of Leuven) have presented their research plans in somewhat more detail.

Archaeology poses an interesting mix of challenges to 3D technology. It should be cheap, easy to use and transport, and flexible in the size and type of objects that can be dealt with. Simultaneously the presentation of the results to a wider public will require that the visual quality comes close or even surpasses what people are becoming accustomed to, i.e. at least game quality. The project hopes to offer such solutions on the basis of photo-realistic modeling tools, such as 3D photocams, the shape-from-video technique, and the mapping of viewpoint dependent textures.

But Murale intends to be about more than generating pretty pictures. It is crucial that its technologies become powerful tools in the hands of the archaeologists themselves. To that end they should in the first place be affordable, so that a large number of excavations can benefit.

References

[1] M. Armstrong, A. Zisserman, and P. Beardsley, Euclidean structure from uncalibrated images, Proc. British Machine Vision Conf., 1994.

[2] Besl, P., Active Optical Range Imaging Sensors, Machine Vision and Applications, Vol. 1, No. 2, pp. 127-152, 1988.

[3] K. Boyer and A. Kak, Color-encoded structured light for rapid active ranging, IEEE Trans. Pattern Anal. And Machine Intell., Vol. 9, No. 10, pp. 14-28, 1987.

[4] http://www.eyetronics.com

[5] T. Ezzat and T. Poggio, Visual speech synthesis by morphing visemes, AI memo No. 1658, MIT, May 1999.

[6] A. Gaglowicz and S.D. Ma, Sequential Synthesis of Natural Textures. CVGIP, vol. 30, pp. 289-315, 1985.

[7] A. Heyden and K. Aström, Euclidean reconstruction from image sequences with varying and unknown focal length and principal point, Proc. Conf. On Computer Vision and Pattern Recognition, pp. 438-443, 1997.

[8] Jarvis, A perspective on range finding techniques for computer vision, IEEE Trans. On PAMI, Vol. 5, No. 2, pp. 122-139, 1983.

[9] R. Koch, M. Pollefeys, and L. Van Gool, Multi viewpoint stereo from uncalibrated video sequences, Proc. Eur. Conf. Computer Vision, Vol. I, pp. 55-71, 1998.

[10] M. Maruyama, and S. Abe, Range Sensing by Projecting Multoiple Slits with Random Cuts, IEEE PAMI 15(6), pp. 647-650, 1993.

[11] M. Pollefeys, R. Koch and L. Van Gool, Self-calibration and metric reconstruction in spite of varying and unknown internal camera parameters, Int. Conf. On Computer Vision, pp. 90-95, Bombay, India, Jan. 4-7, 1998.

[12] M. Pollefeys, R. Koch and L. Van Gool, Self-calibration and metric reconstruction in spite of varying and unknown intrinsic camera parameters, Int. Journal of Computer Vision, Vol. 32, No. 1, pp. 7-25, 1999.

[13] T. Tuytelaars, l. Van Gool, L. d'Haene, R. Koch, Matching affinely invariant regions for visual servoing, International Conference on Robotics and Automation, Detroit, pp. 1601-1606, May 10-15, 1999.

[14] P. Vuylsteke and A. Oosterlinck, Range Image Acquisition with a Single Binary-Encoded Light Pattern, IEEE PAMI 12(2), pp. 148-164, 1990.

[15] A. Zalesney and L. Van Gool, A compact model for viewpoint dependent texture synthesis, Proc. Of the 2nd SMILE Workshop, Dublin, July 2000, to be published in the Springer LNCS Series.