Realistic Textures for VR

Alexey Zalesny[1], Dominik Auf der Maur[1], and Luc Van Gool[1, 2]

[1] Swiss Federal Institute of Technology Zurich, Switzerland
{zalesny, aufdermaur, vangool}@vision.ee.ethz.ch
http://www.vision.ee.ethz.ch/~zales
[2] Catholic University of Leuven, Belgium
Luc.VanGool@esat.kuleuven.ac.be

Abstract. In the construction of 3D site models, much more emphasis has been put on the creation of the 3D shapes than on their textures. Nevertheless, the overall visual impression will often depend more on these textures than on the precision of the underlying geometry. The paper proposes a hierarchical texture modeling and synthesis technique to simulate the intricate appearances of building materials and landscapes. A macrotexture or "label map" prescribes the layout of microtextures or "subtextures". The system takes example images, e.g. of a certain vegetation landscape, as input and generates the corresponding composite texture models. From such model, arbitrary amounts of similar, non-repetitive texture can be generated (i.e. without verbatim copying). The creation of the composite texture models follows a kind of bootstrap procedure, where simple texture features help to generate the label map and then more complicated texture descriptions are called on for the subtextures.

Keywords. Texture synthesis, texture analysis, statistical texture modeling, composite textures, texture hierarchy.

1 Textures for Archaeological Sites

With the growing realism that computer graphics has to offer, there also is an increasing interest in the 3D modeling and visualization of archaeological sites. Such techniques hold enormous promise for both the public and the archaeologists. For the former, a much more lively and enticing account of the living conditions in ancient times can be given. For archaeologists, 3D technology allows them to produce 3D records of stratigraphy, to test the validity of hypotheses on city architecture or on road and water supply networks, etc.

So far, much work went into the development of flexible technology for the creation of 3D building and terrain models. This includes methods to capture in 3D the ruins as well as methods to recreate their original state through CAAD. Apart from their shape, the patterns by which the models of buildings and terrain are covered - their textures - are at least as important for the creation of a realistic impression. Just like the developments in 3D shape modeling of the late 90s, one would like to model textures solely from example images. This is the very goal of the work described here.

This paper describes a texture modeling and synthesis method that works from example images of the target texture. Although the approach is generic, the particular goals here are to produce textures for the simulation of building materials and landscape vegetation. Building materials of which ruins have been constructed usually have undergone serious erosion and have lost their original appearance. Hence, it is useful if the original appearance can be simulated as a texture that is mapped onto the CAAD models of the buildings. This is obviously a more veridical visualization than just mapping the ruin's texture onto the model. The same goes for the terrain model. The existing vegetation may be very different from that prevailing in the era for which a site model is produced. Rather than mapping the existing texture onto the terrain model, one would like to cover it with texture that simulates the vegetation of that era.

Texture synthesis is not only useful to recreate appearances related to another period. Vegetation and landscape textures cannot always be photographed with sufficient resolution throughout the site to allow convincing fly-throughs. Mapping realistically looking landscape texture of an appropriate resolution onto the terrain model can solve that problem.

2 Sagalassos as Testing Ground

We currently focus our efforts on the archaeological site of Sagalassos, located in present-day Turkey. The site is lying about 100 km to the north of Antalya. The excavation in Sagalassos is one of the largest ongoing archaeological projects in the Mediterranean. The project is lead by Prof. Marc Waelkens of the University of Leuven.

Sagalassos is lying at the southern flank of the Aglasun mountain ridge (western part of the Taurus-mountains) at a height between 1400 and 1650 meters. In its day, it was one of the three most important cities of Pisidia. The city had thrived for about 1000 years, when it was finally abandoned after an earthquake in the 7th century AD. During this long period, it got under the military, political, and cultural influence of a series of foreign powers. In 333 BC the Sagalassons were defeated by Alexander the Great. Sagalassos was already an important city at that time. In the subsequent period it changed hands at several occasions between the successor kingdoms of Alexander's splintered Macedonian empire. From 189 BC onwards, the Romans directly intervened in the region until in 133 BC it finally became formally part of the Roman state. In 25 BC emperor Augustus created the province of Galatia, which also incorporated Pisidia. It goes without saying that the city has changed substantially during this extended period, with its heydays around the 2nd century AD.

Changes were not only of a political nature, of course. Over time, architectural styles and techniques have changed, and so did the gamut of building materials that were at the builders' disposal and that they preferred to use. Probably even more noticeable are the changes in vegetation found near the site in different periods. Nowadays the mountain slopes are not covered by what one would consider an abundant vegetation (in contrast to the well-irrigated valley itself). They are covered by "thorn-cushion steppe", a result of overgrazing. But at some point, the slope to the north of the city was at least partially covered by cedar woods, for instance. The first signs of human

influence on the vegetation (anthropogenic pollen indicators) appear around 2000 BC. Human activity and its influence have increased since, with an expansion of open terrain at the expense of forest and with large scale replacement of natural vegetation by agriculture and horticulture. This evolution is not monotonous though. For instance, during Hellenistic to middle Imperial times, indicators of anthropogenic activity decrease again. This may have had to do with the political unrest during the beginning of this period and the cutting of the enemy's slow growing olive trees as part of standard warfare. Towards the end of Sagalassos' existence as a living city, pollen finds show a pattern consistent with a return to pine forests and low level grazing, probably as a consequence of a reduced population. In fact, much more detailed knowledge is available about how these factors changed during the centuries, thanks to thorough, multidisciplinary research within the scope of the Sagalassos project [VERMOERE]. Hence, when one creates 3D models of the site, the choice of the right textures for the simulation of building materials and vegetation is an important one and depends on place and time. In summary, in our archaeological applications texture synthesis is used for two purposes:

1. mapping *landscape (esp. vegetation) texture* onto the terrain model of the site, where the texture is dependent on the chosen era (incl. modern times).

2. mapping *building material textures* onto 3D CAD models of buildings, simulating their state in the absence of erosion, or at different stages of erosion.

That even modern landscape texture will be synthesized rather than photographed has to do with the resolution that would be required. One of our goals is to allow a user to experience a virtual walkthrough of the site. The ruins are modeled on the basis of close-range photogrammetry techniques, whereas the overall terrain is modeled based on photographs taken from a long distance. As a consequence, there is a large discrepancy between the resolution of the texture on the landscape and the texture on the ruins. Although visitors can be expected to focus their attention on the ruins/buildings, the visual quality of the landscape texture in between should match that of the buildings. Otherwise, a disturbing, perceptual contrast appears, as shown in Fig. 1.

The resolution of the landscape texture does not match that of the ruin models (this is actually the case for both the geometry and the texture, but the geometrical aspect is not dealt with here and is less conspicuous). A way out would be to take more detailed pictures of the landscape texture throughout the complete area. This is not feasible however, as the site is tens of square kilometers in size. It would simply take too much time and computer memory, spent on non-essential information (not to mention that at some spots taking photographs would be a dangerous undertaking). 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 terrain model is then covered with similar textures of the right type.
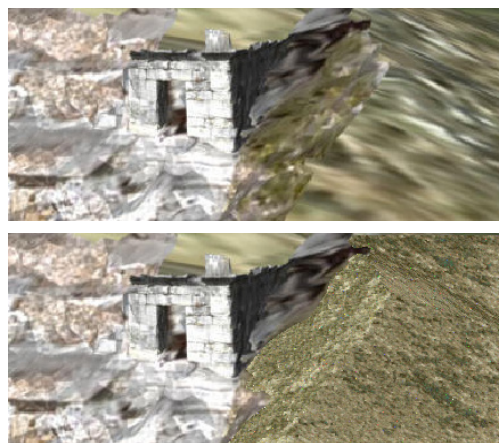

Fig. 1. *View of part of a Roman bathhouse and surrounding landscape at Sagalassos. Top: view with the original landscape texture. As this is a view strongly zoomed in onto part of the landscape, the texture is of insufficient quality when compared to that of the ruin. Bottom: the landscape texture has been replaced by synthetic texture in the bottom right part of the scene.*

For all the applications of texture synthesis that were mentioned, a texture model is learnt from example images. The texture model is very compact and can be used to synthesize arbitrarily large patches of the texture. Section 3 describes the basic texture analysis and synthesis approach used for the relatively simple cases like this one. Then the paper moves on to the description of "composite textures" in section 4, which are used to deal with the synthesis of more complicated material and landscape patterns. Section 5 concludes the paper.

3 Image-Based Texture Synthesis

Several powerful texture descriptions have been proposed in the literature, all with their pros and cons ([DE BONET], [EFROS], [GAGALOWICZ], [GIMEL'FARB], [LEUNG], [PORTILLA], [WEI], [ZHU]). The approach proposed here is in line with the cooccurrence tradition (also see [GIMEL'FARB]), which seems to offer a good compromise between descriptive power and model compactness. Textures are synthesized to mimic the pairwise statistics of their example texture. This means that the joint probabilities of the colors at pixel pairs with a fixed relative position are approximated as closely as possible. Such pairs will be referred to as *cliques* and pairs of the same type (same relative position between the pixels) as *clique types*. This is illustrated in Fig. 2.

Cliques of the same type   Cliques of different types

Fig. 2. *Dots represent pixels. Pixels connected by lines represent cliques. Left: cliques of the same type, right: cliques of different types.*

The texture model consists of statistics for a set of different clique types. Just including all pairwise interactions (all clique types as in [GAGALOWICZ]) in the model is not a viable approach and a good selection needs to be made. We have opted for an approach that makes a selection to keep this set minimal but that on the other hand brings the complete clique statistics of the synthesized textures very close to that of the example textures, i.e. also for these clique types that are not

included in the model [ZALESNY]. Clique type selection follows an iterative approach, where clique types 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 synthesized texture is analyzed to decide which further addition to make. The set of selected clique types (from which textures are synthesized) is called the *neighborhood system*. The complete texture model consists of this neighborhood system and the *statistical parameter set*. The latter contains the joint probabilities for the selected relative, pairwise pixel positions. In fact we do not keep the complete, joint probabilities of colors at the clique pixels, but rather the histogram of intensity differences within and between the color bands. A sketch of the texture model extraction algorithm is as follows:

step 1: Collect the complete 2nd-order statistics for the example texture, i.e. the statistics of all clique types. 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 synthesized texture, to be refined in subsequent steps.

step 3: Collect the statistics for all clique types from the current synthesized image.

step 4: For each clique type, compare the statistics of the example texture and the synthesized texture and calculate their "distance". For the statistics the intensity difference distribution (normalized histograms) were used and the distance was simply Euclidean. In fact, the intensity histograms pure were added also, where "singletons" played the role of an additional interaction.

step 5: Select the clique type with the maximal distance (cf. step 4). If this distance is less than a threshold, stop. Otherwise add the clique 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 neighborhood system and texture parameter set and go to step 3.

After the 8-step analysis algorithm we have the final neighborhood system of the texture and its statistical parameter set. A more detailed description of this texture modeling approach is given elsewhere [ZALESNY]. In that paper it is also explained how the synthesis step works. In this section we demonstrate the use of this basic algorithm for the synthesis of Sagalassos textures, and in the next section we propose an extension towards "composite textures", which we use for the synthesis of more complicated textures.

In the case of Sagalassos, the method has mainly been used for colored textures. For the modeling of colored textures clique types are added that combine intensities of the different color bands.

The shortest 4-neighborhood system within the color bands and the interband interactions between identically placed pixels were always preselected because experiments showed that they are important for the vast majority of the texture classes.

The proposed algorithm produces texture models that are very small compared to the complete 2nd-order statistics extracted in step 1 and also compared to the example image. Typically only 10 to 40 clique types are included and the model amounts from a few hundred to a few thousand bytes. Nevertheless, these models have proven effective for the synthesis of realistically looking textures of a wide variation.

Another important advantage of the method is that – in contrast to some of the most interesting alternatives [DE BONET], [EFROS] - it avoids verbatim copying, i.e. no pieces of the example texture are taken as such and then spatially reorganized. Verbatim copying becomes particularly salient when large patches of texture need to be created. Then the repeated appearance of the same structures quickly becomes salient to the human eye. In our case we certainly have to produce extended texture patches. Verbatim copying would e.g. easily lead to the "same" stone or bush reappearing time and again in the landscape. This problem often not transpires very clearly from publications, as the extent of the texture that can be shown there has to be kept small in any case. In the field of cultural heritage avoiding verbatim copying is especially important as one should not convey a false impression of symmetry or regularity. If stones of e.g. the Roman bath in Sagalassos were all of different sizes, the repetition of the same stones creates a wrong impression that reaches beyond the pure aesthetic level, as it suggests an architectural feature.

Fig. 3 shows a few examples of textures synthesized with our method for different building materials used at Sagalassos. The upper images are original textures of these materials (limestone). The images underneath show the results of texture synthesis based on models extracted from the originals. Fig. 4 shows results for two types of vegetation found at modern Sagalassos. From left to right we show an original image, an image with a part cut out, and an image with this part filled in with synthetic texture for the particular vegetation. As can be seen, the synthetic texture blends in well.
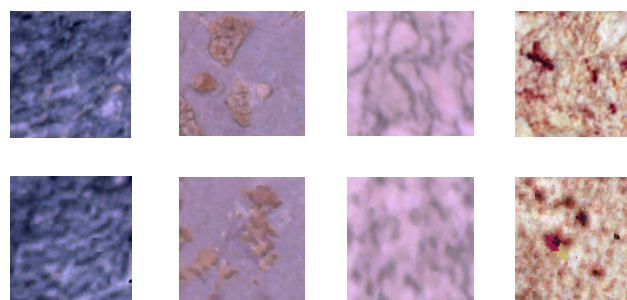

Fig.3. *Example images of building material (top) and synthetic textures (bottom) based on models extracted from these examples.*

But we have glossed over a problem that pops up in the fill-in example of Fig. 4. If no special measures are taken, the boundaries between the surrounding texture and the filled-in texture show up as clearly visible seams. Textures need to be knitted together in a seamless way. A knitting procedure has been developed, and is described in detail elsewhere [ZALESNY]. 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, based on the border zone texture model. In the case of boundaries between the same texture (as in Fig. 4), the seams between

separately generated patches can be removed by resynthesizing a region near the boundary applying that texture's own model.

Fig 4. *The left column images were taken at Sagalassos. The top image shows bush, the bottom image grass. The middle images are the same but with a part cropped out. The right column images have these parts replaced by synthetic texture.*

This knitting procedure can also be used to iron out seams between different textures, as shown in Fig. 5. The image on the left of Fig. 5 shows a combination of rock and grass texture images, both taken at Sagalassos and with a sharp boundary in between. The right image shows the result obtained with the texture knitting algorithm.
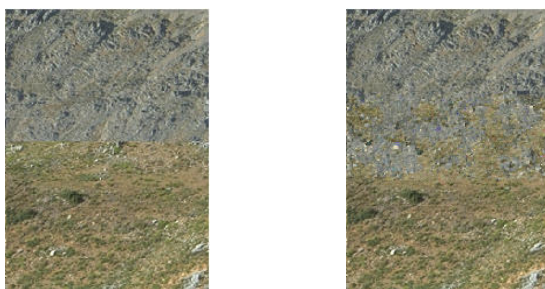
Fig. 5. *Example of texture knitting. Left: image comprising two types of Sagalassos texture, with rocks on top and grass below. Right: knitted rock and grass textures*

Texture knitting does not solve problems of photometric inconsistency. For instance, if the texture to be implanted in a scene has been learned under different illumination conditions than those of the target scene, this is another cause of seams. Fig. 6 (a) and (b) show example images of cedar woods, taken under bright sunlight. Fig. 6 (e) shows the mountain to the north of Sagalassos, photographed under an overcast sky. In order to obtain figure (f), where synthetic wood under similar illumination has been generated at hand picked places, a photometric correction is needed. We have followed a simple gamut transformation technique. It extracts the histograms of similarly colored regions in both type of images and determines the transformation of the three HSV components needed to realize the transition from the example histograms to the target histograms. Figures (c) and (d) show the effect on the example forest texture.

Fig. 6 shows an additional use of texture synthesis. The reader may have noticed that the addition of trees is not the only difference between images (e) and (f). Several objects in the foreground have disappeared, like the blue crane with a car on the left and another car in the bottom-right part of the image. This was achieved by extracting texture models from the surrounding area, and by replacing the area occupied by these

object by synthetic texture generated on the basis of these models. Such ability to remove objects is interesting for archaeological applications, as in order to create convincing visualizations one needs not only to add high quality models of objects that are missing, but also to remove those that represent anachronisms.

The basic texture synthesis approach described in this section can handle quite broad classes of textures. Nevertheless, it has problems with capturing "*composite textures*": complex orderings of patches which themselves show textures (sometimes referred to as microtextures in the literature). This is why an extension towards a hierarchical approach – proposed in the next section – is necessary.
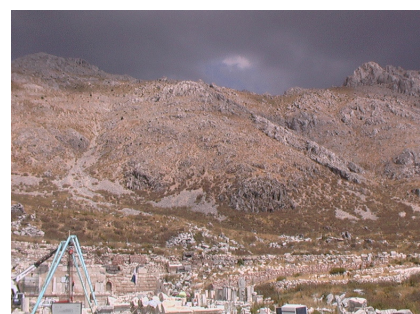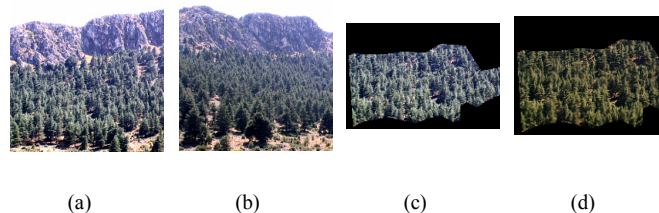

(a)   (b)   (c)   (d)

(e)

(f)

Fig. 6. *(a) and (b) example images of cedar wood texture; (c) forest before the photometric correction, (d) photometrically corrected forest; (e) mountain slope behind Sagalassos in its current state, (f) virtual reforestation and irrelevant object removal.*

4 Composite Textures

Fig. 7 shows part of the modern "thorn-cushion steppe" landscape found around Sagalassos (left). It consists of several ground cover types, like "rock", "green bush", "sand", etc., for which the corresponding segments are drawn in the right figure. If one were to directly model this composite ground cover as a single texture, the basic texture analysis and synthesis algorithm of the last section would not be able to capture its

complexity (Fig. 8). Therefore, a hierarchical version of the texture modeling approach is propounded, where a scene like this is first decomposed into the different composing elements, as shown in Fig. 7 (right). The segments that correspond to different types have been given the same intensity (i.e. the same label). This segmentation has been done manually. Fig. 9 shows the image patterns corresponding to the different segments.
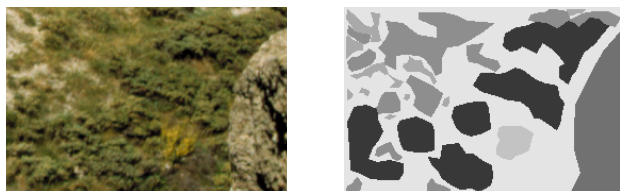

Fig. 7. *Left: an example of modern Sagalassos texture, "thorn-cushion steppe". Right: manual segmentation into basic ground cover types (also see Fig. 9).*
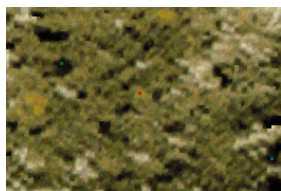

Fig. 8. *Attempt to directly model the scene in Fig. 7 as a single texture.*
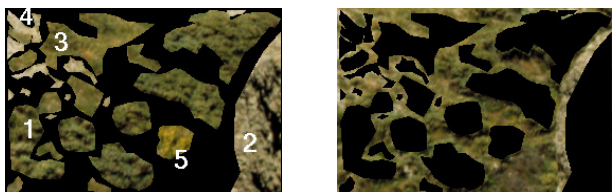

Fig. 9. *Manual segmentation of the Sagalassos terrain texture shown in Fig. 7. Top: segments corresponding to 1-green bush, 2-rock, 3-grass, 4-sand, 5-yellow bush. Bottom: left-over regions are grouped into an additional class corresponding to transition areas.*

The textures within the different segments are simple enough to be handled by the basic algorithm. Hence, in this case 6 texture models are created, one for each of the ground cover types (see caption of Fig. 7). But also the map with segment labels (Fig. 7 right) can be considered to be a texture, describing a typical landscape layout in this case. This "label map" texture is quite simple again, and can be handled by the basic algorithm. Hence, such label maps can be generated automatically as well. The following idea then stands to reason: generate a composite texture, where first a landscape layout texture is generated (i.e. a synthetic label map). Subsequently, the different segments are filled in with the corresponding "*subtexture*", based on these textures' models. As an alternative, a graphical designer or artist can draw the layout, after which the computer fills in the subtextures in the segments that s/he has defined, according to their labels. Similar ideas have independently been proposed [HERTZMANN], but our proposal automates the whole process, including the generation of the label map. Fig. 10 shows one example for both procedures. Note that the right image has been created fully automatically and arbitrary amounts of such texture can be generated, enough

to cover the terrain model with never-repeating, yet detailed texture. As mentioned before, the fact that this approach doesn't use *verbatim* copying of parts in the example images has the advantage that no disturbing repetitions are created.
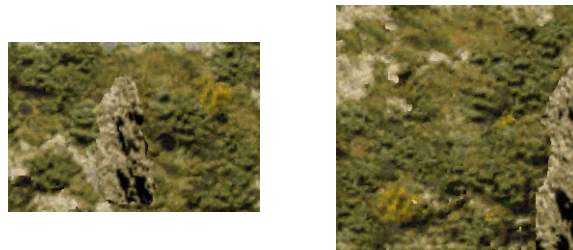

Fig. 10. *Synthetically generated Sagalassos landscape textures. Left: based on a manually drawn label map. Right: based on an automatically generated label map*

A similar approach can generate textures of the more complicated types of marbles and limestones that were used as building materials. In a building like the Nymphaeum, for instance, about 10 differently colored stones were combined to arrive at splendid effects. If one wants to recreate the original appearance of this building and others in the monumental center of the city, such textures need to be shown in their full complexity. Fig. 11 shows an example of a limestone (pink-gray breccia) (a), which has a kind of patchy structure.
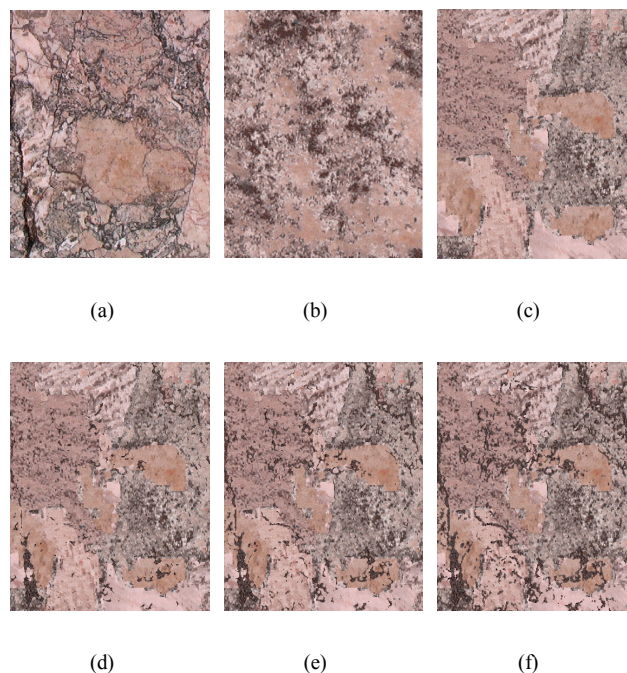

(a)   (b)   (c)

(d)   (e)   (f)

Fig. 11. *Pink-gray breccia. (a) original, (b) synthesized as a single texture, (c) as a composite texture; (d)-(e)-(f) gradually adding erosion effects, where the superimposed, black erosion mask can be modeled as an additional texture.*

The different parts do not only have different colors, but also a substructure (microtexture) of their own. The overall breccia texture is too difficult to be modeled well by our basic algorithm, as shown in (b). This image is the texture generated from a single model. As for the landscapes, we can follow the composite texture approach. First, the original limestone image is manually segmented, whereupon texture models are generated for the different parts. An automatically generated, synthetic result is shown in (c). We can now drive the composite texture

idea one step further and use it to simulate the evolving effect of erosion. By thresholding the image (a) the cracks and pits that are the consequence of erosion, can be isolated, as they correspond to black regions. These have also been left out from (c) (have not been included into the model used for synthesis), as we wanted to synthesize an intact, polished breccia, as it would have looked like in Roman times. The thresholded image again is a quite simple texture that can be modeled on its own. Image (f) is the result of superimposing the erosion image onto image (c). In fact, a time series of erosion masks can be produced by peeling off layers from this initial erosion mask. Then one can start with the one that has the fewest black pixels as the earliest level of erosion and so on. Parts (d), (e), and (f) of the figure simulate increasing levels of erosion for this breccia. The effect so generated is purely visual, however, and not physics based. Currently, we are working on schemes that stay closer to geological processes. Fig. 12 shows an example of the use of such textures. Several textures have been superimposed on the pillars of the Nymphaeum.

Fig. 12. The Nymphaeum at the upper agora of Sagalassos with differently textured pillars.

Ongoing work aims at an important, further step in the automation of composite texture learning. In the given examples, the textures were segmented by hand to generate an example label map. The potential of the composite texture approach would clearly increase substantially if also this first step could be automated. At first sight this seems like a chicken-and-egg problem, as it seems necessary to have good subtexture models before the segmentation could be done. However, as extensive experiments by Paget [PAGET] have shown, texture features used for segmentation can be simpler than those needed for synthesis. In fact, Paget even demonstrated that the *optimal* complexity of features for segmentation is lower than for synthesis. Hence, a kind of bootstrapping procedure seems possible, where an initial segmentation is based on simpler color and filter bank outputs. These features are segmented on the basis of a clique partitioning algorithm, described in [FERRARI]. This step generates a label map and indicates where different subtextures have to be learned from. Synthesis is then as before, using the much more sophisticated subtexture models. This allows the system to synthesize textures completely automatically from an example image.

Fig. 13 shows some initial results. Both scenes on the left are landscapes photographed in the Sagalassos region. The textures on the right have been generated from these, without any human interaction. The one in the top row was automatically segmented into 2 regions, which in this case amount to

the brownish soil and the green tree canopies. The landscape in the bottom row is more complex. It was automatically segmented into three regions: grass, bush, and rock fragments. Note how the synthetic image for the latter (bottom right image in the figure) manages to keep the overall spatial arrangement correct. This is due to the fact that the basic synthesis method makes a distinction between head and tail pixels in the handling of the cliques. As a consequence, the label map will concentrate the same subtextures at the top and the bottom of the image as is the case in the example image.
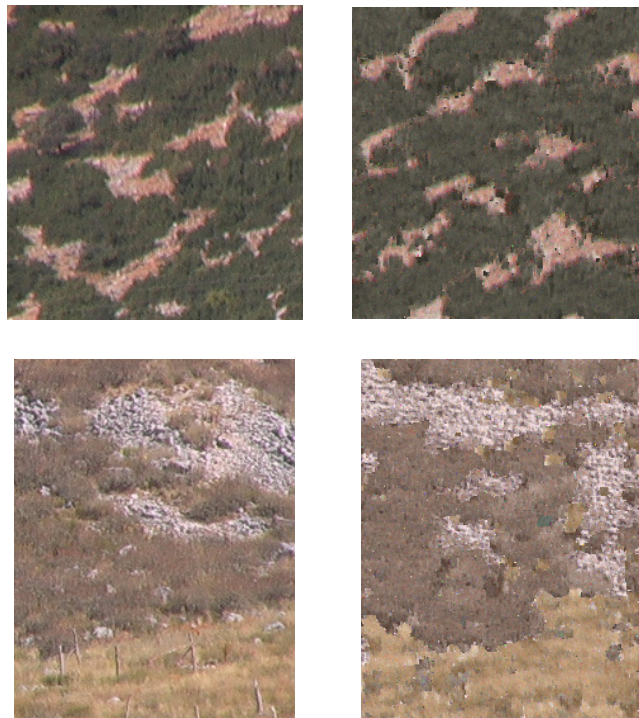

Fig. 13. Top row, left: example landscape texture, top row right: completely automatically generated texture, with only the image on the left as example; bottom row: same, where the texture on the right has been generated based on the example on the left.

5   Conclusions and Future Work

We have focused on a texture synthesis method and its extension towards composite textures. This texture synthesis work is aimed at simulating intricate textures such as those of building materials and vegetation types. An advantage of the methods is that they work from example images as input. In the case of composite textures, a segmentation has to be produced manually once, or this could be generated automatically. Further enhancing the automatic segmentation is another topic of current investigations. So is the enhanced modeling of interactions between subtextures and their positions within label maps, work that should also lead to more geology based erosion, for instance. An advantage of the proposed methods is that they don't yield any *verbatim* copying of parts of the example textures. Such copies quickly become very salient as several of them appear in larger patches of textures. Moreover, the texture models are very compact and storage of the example images is not required.

References

DE BONET, J.S., 1997. Multiresolution Sampling Procedure for Analysis and Synthesis of Texture Images. SIGGRAPH 97, 1997, pp. 361-368.

EFROS, A. and LEUNG, T., 1999. Texture synthesis by non-parametric sampling. *Proc. Int. Conf. Computer Vision (ICCV'99)*, Vol. 2, 1999, pp. 1033-1038.

FERRARI, V., TUYTELAARS, T., and VAN GOOL L., 2001. Real-Time Affine Region Tracking and Coplanar Grouping. *IEEE Computer Soc. Conf. on Computer Vision and Pattern Recognition (CVPR'01)*, December 2001, Vol. 2, pp. 226-233.

GAGALOWICZ, A. and MA, S.D., 1985. Sequential Synthesis of Natural Textures. *Computer Vision, Graphics, and Image Processing*, Vol. 30, 1985, pp. 289-315.

GIMEL'FARB, G., 1999. *Image Textures and Gibbs Random Fields*. Kluwer Academic Publishers: Dordrecht, 1999, 250 p.

HERTZMANN, A., JACOBS, C., OLIVER, N., CURLESS, B., and SALESIN, D., 2001. Image analogies. *ACM SIGGRAPH 2001*, pp.327-340, 2001.

LEUNG, T. and MALIK, J., 1999. Recognizing surfaces using three-dimensional textons. *Proc. Int. Conf. Computer Vision (ICCV'99)*, 1999, pp. 1010-1017.

PAGET, R., 1999. *Nonparametric Markov Random Field Models for Natural Texture Images*. PhD Thesis, University of Queensland, February 1999.

PORTILLA, J. and SIMONCELLI, E. P., 2000. Texture Modeling and Synthesis using Joint Statistics of Complex Wavelet Coefficients. *Int. J. of Computer Vision*, Vol.40, No.1, pp.49-72, Oct. 2000.

VERMOERE, M., VANHECKE, L., WAELKENS, M., and SMETS, E., 2002. Woodlands in Ancient and Modern Times in the Territory of Sagalassos, Southwest Turkey, to appear in: Sagalassos VI. Report on the survey and excavation campaigns of 1998, 1999, and 2000, ed. M. Waelkens, J. Poblome, Turnhout 2002.

WEI, L.-Y. and LEVOY, M., 2000. Fast texture synthesis using tree-structured vector quantization. *SIGGRAPH 00*, pp.479-488.

ZALESNY, A. and VAN GOOL, L., 2000. A Compact Model for Viewpoint Dependent Texture Synthesis. *SMILE 2000, Workshop on 3D Structure from Images, Lecture Notes in Computer Science 2018*, M. Pollefeys et al. (Eds.), 2001, pp. 124-143.

ZHU, S.C., WU, Y.N., and MUMFORD, D., 1998. Filters, Random Fields And Maximum Entropy (FRAME). *Int. J. Computer Vision*, Vol. 27, No. 2, March/April 1998, pp. 1-20.