Optimization Techniques for Photogrammetry Applied to Cultural Heritage and the Action of Transformation Groups

Fabrizio Bazzurri[1,2] and Massimo A. Picardello[2]

[1] Nu.M.I.D.I.A. srl, Via Giacomo Peroni 130, 00131 Roma, Italy
fabrizio.bazzurri@gmail.com
[2] Department of Mathematics, University of Rome "Tor Vergata", Via della Ricerca Scientifica 1, 00133 Roma, Italy
picard@mat.uniroma2.it
ORCID: https://orcid.org/0000-0002-4281-0429

Abstract. Technology could sensitize end-users to preservation of cultural heritage while stimulating more interest based on a greater involvement. Our approach to this goal is based on a novel gamification and story-telling based on scientific foundations. At the same time, this application can offer the pro-user an interactive visualization suitable to study and verify new theories about the artifacts both on their current state and on its 3D reconstruction.ata, but, at the same time, design a workflow capable to integrate different optimization techniques to adjust data depending on different scenarios. This leads to a unified system able to adapt itself to external requirements in a semi-automatic fashion. Finally, new ways of playful interaction, based on simple mathematical operations, were investigated.
The second-named author is partially suppported by MIUR Excellence Department Project funds, awarded MIUR to the Department of Mathematics, University of Rome Tor Vergata, CUP E83C18000100006.

1 Introduction and State of the Art

Studying and preserving cultural heritage have always been primary goals. These goals are highly preeminent in Italy, a country with rich historical background and endowed with a large number of archaeological finds and historical artifacts of outstanding quality. It is necessary not only to study and collect data related to artifacts, but also to make them widely known with in-depth scientific analysis and with attractive exhibitions and presentations. In order to give the witness clear and accurate information, but at the same time to encourage him to undertake further investigations and to make him feel as an active part of the cultural environment, these two sides must be always kept in balance.

In recent years, technological progress and continuous growth of ICT gave rise to completely new types of man/machine and man/man interfaces. From the state of the art of just a few years ago, giant steps have been made, above all in modeling and rendering. Henceforth we shall focus on new technologies

Fig. 1. A frame output by our real-time app

for 3D-modeling and visualization of cultural heritage. For the former aspect, there is a significant progress in making use of photogrammetry for 3D-model acquisitions of real-world artifacts: this progress is made possible by the constant increase of computing power even in end-user equipment and non-professional workstations. The latest CUDA programming techniques on graphic processors have significantly shortened the computing time to obtain dense point clouds starting with photographic images. This has promoted photogrammetric techniques to the status of strong competitor of the classical methods based on laser scans, even more so because of higher and higher quality of texturing.

On the other hand, the visualization of historic artifact has profited of significant advances in real time rendering, that used to be an almost exclusive monopoly of the gaming world, particularly evident in the development of increasingly more refined physically based models for the interaction of different materials with light. This growing interest led to new rendering techniques aimed to more accurate photorealistic output and faster processing time. Visualizations that used to be available only through offline rendering are now achieved in real time with minor compromises on quality. Moreover, we are witnessing a demand for user-experience based on immersive interface and interaction with real world. This push towards virtual reality interfaces led to the marketing of new and advanced Virtual Reality viewers (Oculus and HTC Vive) that can project a virtual world where the viewer is fully immersed and can naturally interact. At the same time, the industry is developing new devices able to generate mixed reality, that is, mixing the true world with virtual elements: these are the augmented reality viewers. These viewers need a much higher computing speed, to support real time endering of more frames (two instead than one for VR viewers), and higher frame rates. Indeed, a frame rate of 90 Hz instead of the usual 60 Hz is the minimum in order to avoid motion sickness due to poor synchronization of real movement of the head and the related movement of the virtual camera: at a lower refresh frequency, the projected images are not fluent enough.

All these tools give rise to a new way of interaction with the artifacts, leading to new opportunities and new ethic and formal viewpoints about the individual-vs-community and individual-vs-world approach. These dual approaches give rise to social aggregation and social inclusion of people with physical disabilities. As for the first point, the digitalization process leads to a new way of sharing information and a new way of social aggregation that overcomes geographical and time barriers. However, particular attention must be devoted to the correct use and management of enormous quantity of data produced as time goes by. A rigorous and reproducible workflow is essential for a correct data storage and processing aimed to avoid the inability to adapt to future use of data and to prevent data fragmentation. On the basis of this last point and the available set of technological tools, the adopted procedures should be open-source to favor prototyping and a rich exchange between academic and industrial environments. This leads to different paths: by means of the same data, different applications could be developed for end-users or professionals. A common language (adaptable to various applications) can transfer scientific approach and contents to the end-user. In this way the final application is not based on pure spectacularization neglecting pedagogic aspects (the basis of cultural production). Moreover, the gamification process enables a new interactive communication mode, different from expository or graphic (that is, distinct from narrative or written). This new way of communicating balances the interaction and collaboration respect to the generality of the exposed message. This balance is acceptable because the gamification turns out to be a cognitive tool of significant impact (see [1]. This was also verified on-field by the use of the application, object of this article.

We already observed that the introduction of digital tools in cultural heritage facilitated the social inclusion of people with physical inabilities. Immersive displays and virtual reality allow the access (both in their present state and in historical reconstruction) to site otherwise precluded by physical barriers. Other sites could be easily visited by all the people (with or without physical disabilities) but they are restricted for preservation needs. For this reason, virtual reality is a great tool of social inclusion: the cultural heritage could be shared by all in the same way (no separate paths should be implemented depending on physical abilities) and this identical sharing could destroy (or at least greatly reduce) the psychological and physical barriers and led to social aggregation. Furthermore, on the grounds of preservation of cultural heritage, its digitalization enables two main aspects: studying and archiving. To the former it offers a chance to study and inspect artifacts without damage (and also remotely). The latter aspect is a pivotal starting point for the conservation of the current state of a specific artifact or site.

2 A Case Study and our Workflow

In this Section we present the workflow for interactive rendering of the main hall of the Catacombs of San Senatore at Albano Laziale, developed by the first-named author with the collaboration of Dr. Roberto Libera, Director of Museo Diocesano of Albano, Italy. This archaeological area raises some exciting challenges:

- a correct rendering of light in the scanning process to achieve a neutral texture;
- a suitable split of the processing of surface geometry (normal map) of the fresco and its colored texture.

Moreover, our goal was a virtual reality reconstruction, suitable for full-immersion remote navigation by the visitors, aimed to divulgation, spread of knowledge and user-friendliness. The last issue is a very important challenge with high-impact ethical value, because it makes possible to disabled visitors to enjoy the monument: this is a fundamental reason to design a museum presentation in photorealistic virtual reality.[3]

The flow chart of our procedure, in order to achieve a smooth VR simulation, was subdivided as in Fig. 2.

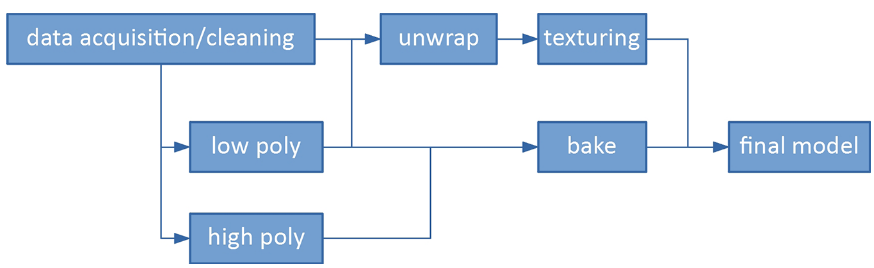

Fig. 2. The flowchart

Clearly, the flowchart in Fig. 2 is aimed to create an individual modeled object to be inserted into a scene (for instance, the frescoed wall or a vase or a lamp laying on a table inside the room to render). The first phase consists of capturing/creating the 3D-model and cleaning the data: it starts with a set of photographic images from several viewpoints and extracts from them a dense

[3] This aspect is particularly important for catacombs. A visitor with walking disabilities will not be able to visit a catacomb not equipped with suitable adjustments of the walking path. but these adjustments are often incompatible with the preservation of the archaeological area, or at least very difficult. Virtual Reality allows these persons to experience a full immersion into the archaeological site, with natural interaction.

cloud of 3D-points. The building of those point clouds is based on photogrammetry (without laser or structured light scanners). We limited attention to this technique to verify, on-field, the quality of the algorithms and the software currently available. We shot 570 photos (Nikon D7000 equipped with 35mm lens) of the main hall following some best practices [3]. By making use of Reality Capture as reconstruction software, we obtained in a few hours a dense point cloud (15 millions triangles) on consumer hardware (CPU Intel i7-5820K, 32Gb RAM, GPU NVIDIA Titan). The number of images as well as the polygon density of the dense point cloud should be higher for final use. For this target, we plan to improve the final visual performance by trying to optimize both the snapshot step and the point cloud management. Our research provided a good enough performance when the data are used only for gamification, but a greater resolution (consequently, a greater number of photos) should be used in order to build up a unified workflow (end-user and pro-user) .

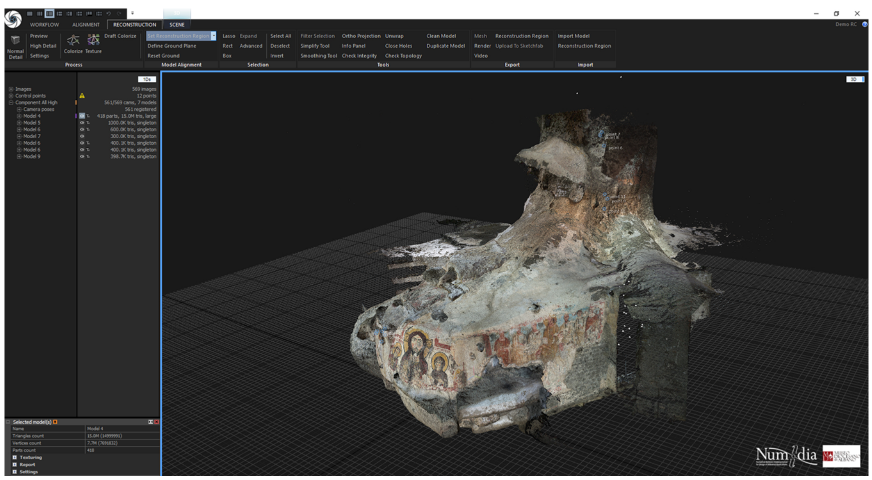

Fig. 3. View of Reality Capture screenshot, with the dense point cloud of the main hall of the Catacombs of San Senatore

Note that, during the snapshot step, controlled temperature lamps were used in order to maintain an adequate colorimetric level. Furthermore, we move them as to have a good illumination response. This is really important to get a correct capture of the paintings and to manage the texture during real-time rendering.

A neutral and uniform illumination without shadows allows to insert artificial lights in the rendering phase with photorealistic material response. Furthermore, uniform and “clean” data speed up image post-processing considerably.

This dense cloud is the starting step to generate a *high-poly* (high-resolution) and *low poly* (low-resolution) meshes. These two variants of the meshes allow us to use the same captured 3D-model at the high precision needed for a scientific

Fig. 4. Lamps used for snapshot step

presentation and at the lower precision but faster processing speed suitable for a virtual visit and interactive remote navigation. The projection of the meshes onto the texture plane perfomed by the reconstruction software is highly dependent of the required output quality, and can be scarcely adapted to provide an adequate rendering performance. For this reason we performed an appropriate semi-automatic unwrap capable of an optimal compromise between the conflicting goals of preserving the quality and speeding up the rendering. Furthermore we transferred (baked) the high-resolution topological information to the texture of the low-poly in a way that allows to add surface details without impacting on performances (see more details below in the example of the wall).

The unwrap step is very important: it is based on a semi-automatic tweaking procedure. The way a texture is unwrapped onto a surface has a direct impact on performance of rendering. To obtain high quality rendering and to take full advantage of the information obtained from the photos, a high pixel density should be used. This could be made possible by "subdividing" a surface mapped on the square $[0,1] \times [0,1]$ of texture coordinates, but this procedure associates many different materials to an individual object, or, in order to achieve better performance, to many objects (stemming from a subdivision based on materials). When the number of objects increases, the number of draw calls also does, with a negative impact on rendering performance. We could counterbalance the negative impact by making use of a texture with lower resolution than to the one used for an individual object. We have three degrees of freedom:

- how many sub-objects stem from an individual one;

– texture resolution for an individual sub-object;
– number of operations to perform to transfer all the data to game engine.

We could make an automatic choice of some of these parameters, but the third one is the most time-consuming as it is mostly manual. There is no general rule to optimize these parameters. We can develop appropriate scripts to speed up some factors over others, but some tweaking is always needed depending on the scene. In the scene shown in the figures, we focused our attention to the main fresco and developedf an unwrap style aimed to maximize the use of this surface. Furthermore, the main surface should be correctly managed in view of a post-processing phase of restoration. Indeed, a fragmentation of the fresco (unwrapping the main surface to different sectors) is a bad choice here, because it makes more difficult to restore the painting during post-processing.

Figure 5 illustrates a ceremony in the catacombs, built in Virtual Reality for the purposes of illustration addressed to visitors of the museum. The scene in-

Fig. 5. A moment in history reconstructed in Virtual Reality

cludes several objects whose geometric shape and texture has been reconstructed via photogrammetric methods, at least in part: for instance, a capital, a lamp, an amphora and a frescoed wall. The modeling was achieved through photogrammetry in the parts where the original surface of the amphora in Fig. 6 was entirely modeled by means of photogrammetric elaboration. The lamp in Fig. 7 was rebuilt for approximately half of its surface, and only about one fourth of the capital in Fig. 8 was obtained starting with photographs.

Fig. 6. An amphora entirely reconstructed via photogrammetry

Fig. 7. A lamp reconstructed via photogrammetry (on the left) and partial rebuilding (on the right): the texture on the left shows the real color captured by the camera, but the color on the right is a superimposed terra-cotta texture

The frescoed wall shows a good example of unwrapping. Figure 9 shows four adjacent slices. From left to right, the first slice has been modeled through the high resolution mesh high-poly. The second is based on low-poly. In the third, the mesh is low-poly, but the normal vectors of each polygon of the mesh have been changed: what we see is the rendering based on Phong illumination

Fig. 8. A reconstructed capital: the red part was obtained through photogrammetry, the rest was remodeled and rebuilt

after replacing at each point the normal vector of the low-poly mesh with the normal vector of the high-poly variant. The data about the normal vectors are embedded in the material as a texture, in order to enhance the speed of the rendering phase. In the fourth slice, we did the same, but added a diffusive component of the lighting, based on Lambert illumination rule: the color at each point used in Lambert model is given by the texture of the fresco extracted from the photographic images.

Beside the simple visual representation, we investigated the chance to use the 3D modeling tool to support the virtual restoration of paintings and artifacts. All the rendering in the figure has been done in Blender. All the geometric, texturing and unwrapping manipulations have been done with proprietary code in Python by the first-named author. We choose Blender because of the following fundamental aspects: its open-source nature, the easiness of its modeling tools, the expandability provided by its API. The mesh obtained from photogrammetry

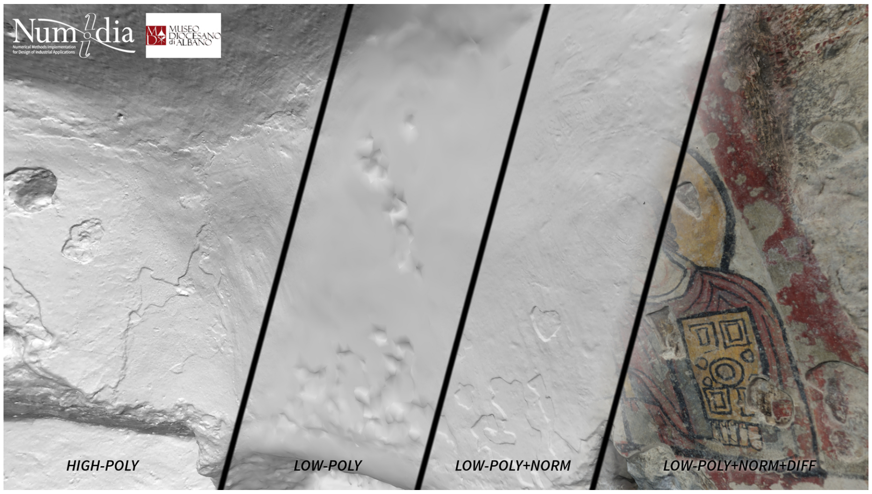

Fig. 9. Mixing meshes of various resolutions to obtain different levels of precision on the reconstruction of an ancient wall

as the base of polygonal modeling allowed the study of new restoration solutions and theories.

The real time rendering was done on Unreal Engine while the VR system was HTC Vive. This choice is due to the quality offered by this game engine over its direct counterpart (Unity). No particular interactions were implemented in the main scene. The main hall dimensions, in fact, are similar to the ones required by the typical installation of room-based virtual reality (4x4 meters). Using this installation, the user could easily reach every point of interest in the virtual scene. To focus the attention on the main fresco we used a “fog shield” sited along directions with no interest. Furthermore, the chance to not use the teleportation system (typical for room scale VR with scene much bigger than 4x4 meters) alleviate the need by the user to learn a new interaction technique and facilitate the approach to the experience.

The user starts his visit in the current state of the main hall (as reconstructed from photogrammetry) and after a while he is carried to the reconstruction of the same hall (around VII century AD). Finally, the user is projected in a totally neutral environment where he could interact with some artifact (the amphora, the capital or the lamps) digitally reconstructed from the photogrammetry step. The interaction (thanks to Vive controller) is deliberately simple to stimulate the exploration and the view of artifacts as they could be seen inside the museum cases, with the added opportunity to virtually handle them and to view them all around.

3 Toward a Mathematical Approach to the Style of Ancient Pottery

We have seen that some artifacts in the catacombs were partially rebuilt for the purposes of museum exhibition, visualization and illustration. This was necessary as a consequence of incomplete preservation of the original surface. It is important, when we rebuild part of an archaeological find, to preserve the original style as much as possible (see for instance the capital in Fig. 8, where the rebuilding was extensive. But how do we define the concept of *s*tyle so that a computer program can understand it? This last section is a mathematical divertissement about this question, with no claims of expertise in art history, but with the hope of providing stimulus to museum exhibitions addressed to children. Except for this purpose, of course, the definition of style that we are about to present is far too simplistic.

Once an artistic artifact is found in good shape, almost intact, like the amphora in Fig. 6, we may clone it into many different variants by acting on the mesh of the modeled object (let us say, an ancient vase) with linear or non-linear operators, maybe even choosing a different operator on different points of the mesh. Of course, the latter choice (varying operators), although very simply implemented in Python, requires interactive intervention of a person, because of many constraints (like, for instance, the fact, to be prevented, that after the manipulations a part of the surface of the vase might penetrate another part of its surface). On the other hand, non-linear operators of this type are essential in producing instances of the vase, for instance in adding handles or a beak or modifying their shape and relative positions. But because of the evident complexity, the approach via varying operators will not be analyzed in this simple sketch, but we shall give some examples later.

We should point out the following fact. There are two ways to modify an object in $\mathbb{R}^3$: by modifying the coordinates of world space, hence in particular the object but as a whole, or by applying local changes to its mesh. The former is expressed by the action of a linear operator, the latter is not, and it offers much more adaptability and variety.

However, the manipulation by linear operators is much simpler, and indeed, it is very congenial to photogrammetry, since the camera models in photogrammetry and computer vision are based on similar tools. There are three types of operators to consider, of increasing generality.

The case of least generality is euclidean operators, that is, similarities. These operators consist of compositions of translations, rotations and dilations. Their action changes the object exactly has we would do by physically handling the object with rigid motions, that is, moving it aside, rotating it and bringing it closer or farther away of our eye. Of course, the modified object is still considered, in our mind, the same as before the modification, not true variant: only position and scale have been changed, but not the relative proportions and shapes of its parts.

The case of maximum generality must be expressed in the framework of projective geometry. The vertices of the mesh are vectors $\boldsymbol{v} = (v_1, v_2, v_3) \in \mathbb{R}^3$.

Embed them into $\mathbb{R}^4$, but keep the number of parameters equal to three by identifying any two vectors that are non-zero multiples of each other. Then each non-zero vector in $\mathbb{R}^4$ belongs to a different equivalence class, and each equivalence class has a special representative $\tilde{\boldsymbol{v}} = (v_1, v_2, v_3, 1)$, except the classes with representative $\tilde{\boldsymbol{v}} = (v_1, v_2, v_3, 0)$, that (taken all together) form the *plane at infinity*. The set of equivalence classes is called the *projective space* $\mathbb{P}^3$. A projective transformation is a special type of mapping on $\mathbb{P}^3$, or equivalently, a mapping of $\mathbb{R}^4$ that preserves the equivalence classes, that is, that commuted with dilations. This special mappings are those that are invertible and map collinear triples of points of $\mathbb{R}^4$ into collinear points. It is easy to see that these mappings are expressed by linear operators on $\mathbb{R}^4$ (hence their action is given by a very fast computation in all computer languages), and form a group, called the projective group. Unfortunately, the action of the projective group is transitive on the projective space, and in particular many projective transformations map points at infinity into points at finite (and conversely): that is, the plane at infinity is not preserved. The reason why we wrote "unfortunately" is that transformations that map to infinity a point inside an object give rise to frightening geometric and perspective distortions and shrinkings, that produce unrecognizable variants of the object (say, a vase): nobody would consider these variants as examples of another vase of a similar style. However, we shall soon see that the projective transformations that do not map to infinity points of the vase (or inside its convex hull) provide a large number of interesting and realistic variants of similar type. Once the modeled object is given, the projective transformations that map its convex hull to the convex hull of the transformed object are called *quasi-affine transformations* of the object. The quasi-affine transformations of an object form a group, that we call its quasi-affine group.

The intermediate case between similarities and projective maps consists of those linear operators of $\mathbb{R}^4$ that are invertible and preserve the plane at infinity. These linear operators are called *affine*. When we restrict attention to a three-dimensional subspace $V \approx \mathbb{R}^3$of $\mathbb{R}^4$ that does not contain the plane at infinity, affine operators consist of linear operators on $\mathbb{R}^3$ followed by translations in $\mathbb{R}^4$. The action of an affine transformation on an object is not a similarity: orthogonality of vectors is not preserved, but parallelism is. So an affine deformation of an object is a variant with a distorted but recognizable geometry. However, affine transformations are still rather rigid. For instance, the amphora of Fig. 6 has a handle and a beak on the opposite side: no affine transformation of the amphora can bring handle and beak to the same side, because the three points given by the barycenter of the handle and the beak and the center of the top opening of the vase are collinear, and collinearity is preserved by affine transformation. Actually, the same is true for any linear operator, so also the projective transformations are rigid in this sense: they not yield twists or twirls. In order to make variants of the vase where handles and beaks get closer, we need more handles. We build a vase with four antipodal handles to illustrate the ideas of this Section: see Fig. 10. The variants of a round vase would be much more realistic once deformed, but we choose to design a square vase in order to make the geometric actions clearer.

Fig. 10. A square vase with four handles, designed for this presentation

Linear transformations preserve the number of handles. This statement is a consequence of the following result:

Theorem 1. *Projective transformations preserve the topological genus of a surface.*

So, we can now produce many realistic variants of our vase. Figure 11 shows some examples of quasi-affine transformations of the square vase with four handles. Note that these transformations are not affine: indeed, the four vertical corner lines do not remain parallel.

The viewpoint of a mathematician is that there is something in common between all these variants, and we should call this mysterious entity the *style* of the artpiece. In this sense, the style of an object is its orbit under the group of quasi-affine transformation, or in other words the equivalence classes of the relation that states that two objects are equivalent if each of them can be transformed into the other by a affine (or quasi-affine) transformation.

Fig. 11. Quasi-affine variants of square vase

A projective transformation is represented by a four-dimensional matrix,

$$\begin{pmatrix} a_{11} & a_{12} & a_{13} & a_{14} \\ a_{21} & a_{22} & a_{23} & a_{24} \\ a_{31} & a_{32} & a_{33} & a_{34} \\ a_{41} & a_{42} & a_{43} & a_{44} \end{pmatrix}.$$

This matrix gives rise to an affine transformation if its last row is $(0, 0, 0, 1)$: this preserves the plane at infinity in the projective space. The square vase in Fig. 10 is a modification of the unit cube centered at the origin in $\mathbb{R}^3$ with main axes given by the three canonical basis vectors $(1, 0, 0, (0, 1, 0), (0, 0, 1)$.

Its four variants of Fig. 11 are produced by the action of four projective transformation that are quasi-affine with respect to the unit cube. The upper three by three block in all four is the same as the identity matrix. In the fourth column we set $a_{34} = 0$ in all four, and choose a_{14}, a_{24} small and possibly different from matrix to matrix (these entries represent skew, that is, horizontal translations of the top of the vase). The last rows of the four matrices are, respectively, $(a_{41}, a_{42}, a_{43}, a_{44}) = (0, 0, 0.3, 0.8)$, $(0.1, 0.1, 0.3, 0.8)$, $(0, 0.1, 0.5, 0.8)$, $(0, 0, 0.2, 0.25)$. It is easy to see that none of these matrices maps any point of the unit cube to infinity, hence they are quasi-affine on it. By reducing the value of the last entry, we can make the projective transformation become non-quasi-affine, but then a part of the vase is mapped to infinity and Blender is unable to offer an understandable rendering (and so is anybody: the shape becomes unbelievably complex). See [2] for a reference on projective, quasi-affine and affine transformations in photogrammetry.

Since matrix operations are fast, the variants can be created (and rendered) quickly, and will serve a good purpose in interactive museum games meant to lead children to reflect on art and styles. Of course, the game should have bounds in its sliders so as to prevent making the transformations non-quasi-affine, as explained above. But a more fullfilling and satisfactory approach can be obtained by adding local deformations, as explained below.

4 Producing Variants of an Artpiece via Non-linear Operators

It is important to realize that quasi-affine transformations, being linear operators, that is, global changes of coordinates in three-dimensional space, on one side have the great advantage of being extremely fast, but on the other side are a bit rigid: the variants of an artpiece unde these transformations do not have any twists. So let us show examples of other transformations, nonlinear, that apply to the mesh of vertices of the model (or just to a part of it) and produce twirls or stretches or other deformations. These transformations, that we call *local transformations*, are easily implemented in Blender. Twirls and stretches are great tools to modify a modeled vase into some new variants that appear as instances of the same style, but have very different looks, as shown in Fig. 12. Therefore, our mathematical idea of style extends to include orbits under non-linear transformation groups. The application Blender implements twirls and stretches on the whole mesh of a modeled object, but also on local parts of the mesh: it allows, for instance, to twist a handle or stretch another. New handles can be added via specific creation (or *prototyping*) operators that modify the mesh by replacing parts of it with appropriate predefines shapes: these transformations, however, must be built case by case by operating upon a given mesh, ad so they must be constructed interactively in real time, a demanding and time-consuming process.

A very nice fact (evident in all our Figures) is that also the texture of the object is transformed to a new texture, both by projective transformations and

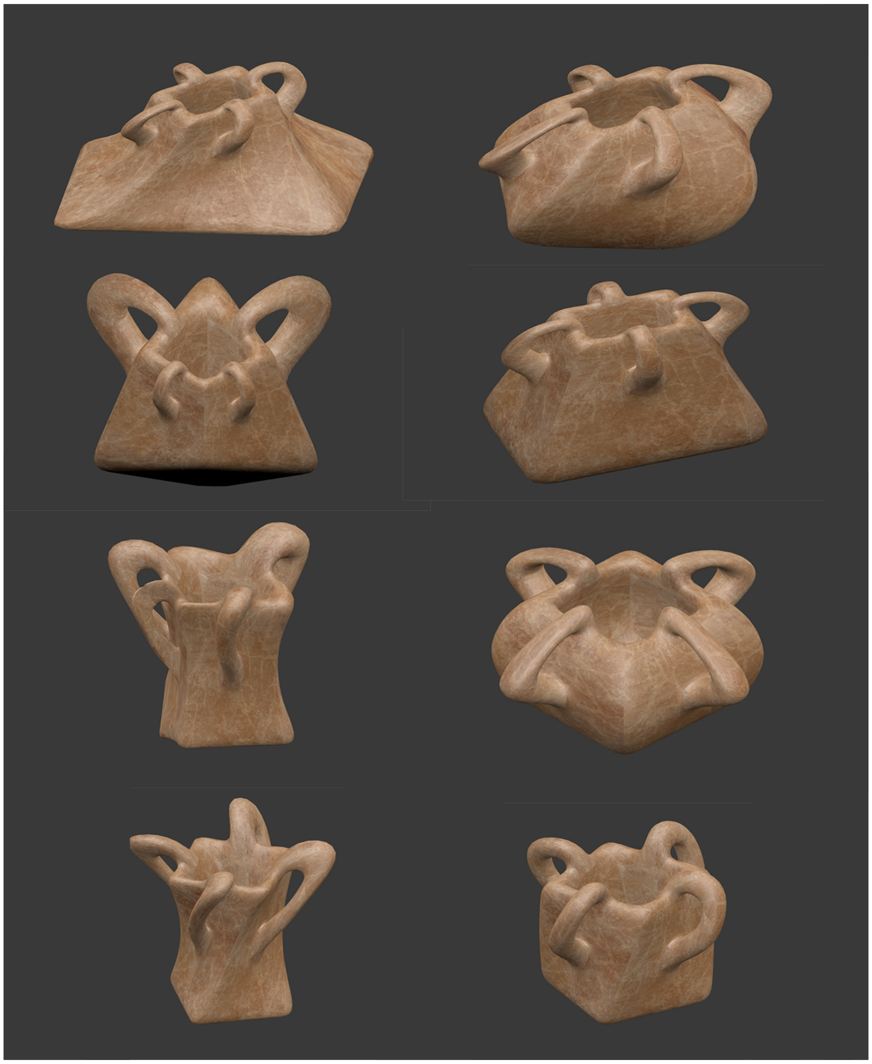

Fig. 12. Twisted or stretched variants of square vase

by local transformations. But here the usual meaning of style requires a much wider group of transformations acting on textures: we must accept as realistic variants of textures those obtained by transformations that are continuous but do not necessarily map collinear triples to collinear triples, hence not linear. For instance, a person painted on a vase may have straight arms and legs, but we would consider variants within the same style those where the painted persons have folded arms and legs. Again here, the texture transformation must preserve the convex hull of (planar) objects, and in particular, for every painted person with a given outline, the transformation not only maps outline to outline, but must also map its interior to the interior of the image. However, there are many more constraints, often different from a style to another. For instance, the decoration of several types of ancient vases often splits the surface into horizontal levels (like floors, that is, horizontal curves, perpendicular to the central vertical axis of the vase), and aligned on the same horizontal level several groups of persons and animals. Therefore the acceptable texture transformations must preserve these horizontal curves. On the other hand, they do not have to preserve the number of persons at each level. Therefore the "style" of the texture or painting can be considered an equivalence class under a group of transformations that, again, must include local creation and cancellation operations. This becomes a very complicated concept or axiomatic definition, and we stop our presentation here.

References

1. de Freitas, S.: What-can-Gamification-add-to-cultural-heritage.pdf, http://www.digitalmeetsculture.net/wp-content/uploads/2013/07/Sara_de_Freitas_What-can-Gamification-add-to-cultural-heritage.pdf
2. Hartley, R., Zisserman, A.: Multiple View Geometry in Computer Vision, 2^{nd} Ed., Cambridge University Press, 2004.
3. Mallison, H., Wings, O.: Photogrammetry in Paleontology - A Practical Guide, Journal of Paleontological Techniques 12 (2014), 1–31.