

An Interactive Local Flattening Operator to Support Digital Investigations on Artwork Surfaces

Nico Pietroni, Massimiliano Corsini, Paolo Cignoni, Roberto Scopigno




Fig. 1. A graphical example of residual marks detection and characterization: the user interactively selects a portion of the surface to be studied; the system produces a parameterized image patch; then, high frequencies encoding residual marks are evaluated and stored in an image; finally, image based processing helps the user to identify and classify the different chisel marks.

Abstract— Analyzing either high-frequency shape detail or any other 2D fields (scalar or vector) embedded over a 3D geometry is a complex task, since detaching the detail from the overall shape can be tricky. An alternative approach is to move to the 2D space, resolving shape reasoning to easier image processing techniques. In this paper we propose a novel framework for the analysis of 2D information distributed over 3D geometry, based on a locally smooth parametrization technique that allows us to treat local 3D data in terms of image content. The proposed approach has been implemented as a sketch-based system that allows to design with a few gestures a set of (possibly overlapping) parameterizations of rectangular portions of the surface. We demonstrate that, due to the locality of the parametrization, the distortion is under an acceptable threshold, while discontinuities can be avoided since the parametrized geometry is always homeomorphic to a disk. We show the effectiveness of the proposed technique to solve specific Cultural Heritage (CH) tasks: the analysis of chisel marks over the surface of a unfinished sculpture and the local comparison of multiple photographs mapped over the surface of an artwork. For this very difficult task, we believe that our framework and the corresponding tool are the first steps toward a computer-based shape reasoning system, able to support CH scholars with a medium they are more used to.

Index Terms— surface characterization, interactive inspection, Cultural Heritage, mesh parameterization, image processing.

1 INTRODUCTION

In the last decade, many archaeologists or CH curators decided to adopt 3D scanning techniques to obtain digital copies of the artworks of interest, moving from 2D to 3D representation. This digital revolution in the CH domain ensures considerable advantages in terms of accessibility to the artwork, documentation of the conservation status and dissemination of knowledge. However, the use of digital 3D models is often limited to visualization, presentation and documentation, while using them for numerical or shape-based analysis is still not a common practice. Cultural Heritage scholars usually perform detailed analysis on the field, by taking a number of characterizing measures over the artwork surface. Measuring can be very easy, (e.g. computation of height and width), or could become very complex when we have to produce measures of complex shape elements. Taking curved paths over a sculpture's surface is a fundamental operation for many analysis tasks: for example to classify features, to supervise the status of a restoration process or, in general, to formulate or validate hypothesis. The classic approach is to perform such measures directly on the original sculpture, using manual measuring in-

struments (e.g. calipers) and photographs as a supporting medium to document the measurements. This approach is time consuming and cumbersome when adopted to characterize many features and requires the researcher to be physically close to the real object. Moreover, the nature of the measurement process can easily introduce errors or severe approximations.

The focus of this paper is to propose and implement a framework for the characterization and analysis of 3D data which supports complex measuring, and shape analysis over high-resolution surfaces through its conversion to a local 2D domain. Taking advantages from mesh parametrization techniques, we designed a new interactive operator, called *Local Flattening*, that constitutes the basis of a novel framework for the analysis of 2D information distributed over 3D geometry. Our *local flattening operator* is designed to be interactive and user friendly. The whole framework is built upon a simple idea: reducing complexity of the characterization process by moving from 3D to 2D space and thus resolving shape reasoning to the easier evaluation of image-based processing computations. This is convenient in our specific context because:

- Comparing different objects (or different regions of a same object) become easier when mapped on a common 2D space;
- Many efficient image processing algorithms for features extraction and analysis are available;
- Despite the research effort, some mathematical tools are specifically designed to work in the image domain, for example the Fast Fourier Transform (FFT) for frequency analysis;
- We can effectively visualize and document the information distributed over an object's surface by just providing an image;

• Nico Pietroni, Massimiliano Corsini, Paolo Cignoni and Roberto Scopigno are with CNR-ISTI, Pisa, Italy E-mail: pietroni@isti.cnr.it

Manuscript received 31 March 2011; accepted 1 August 2011; posted online 23 October 2011; mailed on 14 October 2011.

For information on obtaining reprints of this article, please send email to: tvcg@computer.org.

- Finally, measuring lengths in 2D is more intuitive and generally faster than in 3D. In general, non expert users are more familiar with analyzing images rather than a 3D shape.

Parametrization is commonly used in Computer Graphics for producing texture mapped meshes, where quadrilateral meshes or simplified triangle meshes are enriched by detail (color, normals). However, our target application is slightly different. The main requirements for our application domain can be synthesized as follows:

- **Controllability** The user should be able to define and control the parametrization in an interactive manner. Enough degrees of freedom should grant the user the ability of choosing the *orientation* of the parametrization over the surface;
- **Intuitiveness** We should not expect the user to have a solid math background on mesh parametrization (we are focusing on the CH application domain), but he/she probably has an insight on how a given surface region would unfold on a plane. Hence, the graphical user interface should be as intuitive as possible;
- **Compact Representation** The parametrization produced should be easy to process; possibly, we should be able to represent it by a set of rectangular patches;
- **Locality** We can suppose that the analysis is performed region by region, by taking into account each time a subset of the entire surface. Working on sub-regions brings performance gains (usually scanned artworks are composed by several million triangles) and helps to cope with human perception and analysis limitations. Moreover, producing a global parametrization at an interactive rate on highly complex shapes could be even impossible;
- **Accuracy** Mapping a complex 3D geometry onto a 2D domain usually introduces a certain amount of distortion. We are looking for *nearly isometric parameterizations*, i.e. preserving as much as possible areas and angles. Working on local regions rather than on the entire mesh is also crucial to fulfill this goal.

To fulfill the above requirements we introduce a local flattening operator that allows to easily build a parameterization of a curved portion of the 3D surface onto a simple 2D rectangular image, producing as little distortion as possible.

1.1 Application scenarios

We describe here the three major CH scenarios for which the proposed approach has been designed, but it can be used also in other application domains:

- Shape analysis and matching of the *residual traces* left by carving instruments over unfinished sculptures. Residual traces are important both for understanding the sculpting technique adopted by a given artist and for assessing possible attribution hypothesis. We show how the proposed methodology allows to characterize and compare the residual traces on the digital 3D models instead of working on the real artifacts, increasing the speed and accuracy of this characterization process. The specific case we are investigating is the attribution of a disputed artwork, the *Pietà di Palestrina* (Museum Galleria dell'Accademia, Florence, Italy), commonly ascribed to Michelangelo.
- Monitoring and assessing a restoration action, by checking differences among multiple photographs that document the pre- and post-restoration status. The specific case presented is the restoration of Michelangelo's *David* (Museum Galleria dell'Accademia, Florence, Italy).
- An innovative characterization of the surface and of the sculpting technique, obtained by comparing the low resolution *main curvature directions* of the sculpted surface with the main axial direction of the *residual traces*. The preliminary results (see results concerning Michelangelo's *St. Matthew*) seem to be quite interesting, even if they must be still interpreted by art scholars.

1.2 Main contributions

The main contributions of this paper are:

- An intuitive and interactive local parametrization tool that allows to specify almost isometric parameterization of surface portions through simple sketching actions;
- A new shape characterization approach, which moves the analysis phase from 3D to 2D space thanks to the above parametrization framework, providing CH scholars with a new digital procedure that allows to make measurements and comparisons in a radically new way;
- The applicability of our framework to several CH domains, with clear advantages over more classical approaches.

2 STATE OF THE ART

Presenting an exhaustive state of the art of all the technologies or methodologies mentioned in this paper is not easy given the length limitations. We describe only a few major works.

2.1 Mesh parametrization

Mesh parametrization is an extensively studied problem, for a complete discussion please refer to [1]. The main approaches can be characterized into four main classes:

Methods based on Simplification, which compute a globally smooth parametrization by applying a sequence of local simplifications to produce a coarse base mesh representing the parametric domain [2, 3, 4]. These methods may eventually provide adaptivity of patch size over the mesh [5]. However, they are not controllable by the user and the alignment with the shape features is limited.

Global Harmonic or Conformal Parametrizations, able to produce globally smooth mesh parametrizations [6, 7, 8]; these methods do not directly align with geometric features, moreover they preserve angles at the expense of areas.

Field-aligned methods, which conform a parametrization to a smooth cross field capturing geometric features [9, 10, 11, 12]; these methods generally ensure a good degree of isometry.

Manifold learning/dimensionality reduction methods can be also applied for surface flattening, such as the local tangent space alignment (LTSA) algorithm [13] or [14], that is based on LTSA for explicit pose of mesh parametrization, and the locally linear embedding (LLE) algorithm [15]. Even if these methods are designed to work in a very different context (manifold is unknown and the number of dimension is usually very high), they also could be used for our purposes, however they don't provide feature alignment.

Our *local flattening operator* is based on a field-aligned parameterization method. However, with respect to [9, 10, 11], we don't aim to produce a global parametrization. Since we should provide the user with a high degree of controllability, we have designed an interactive tool to define and customize parametric rectangular patches, allowing the intuitive and interactive selection of the desired alignment direction.

2.2 Chisel marks characterization: the classic approach

A milestone work on chisel marks characterization has been proposed by Prisca Giovannini [16], focused on the chisel marks left over the S. Matthew marble statue (another example of unfinished work of Michelangelo). She studied the marks of the carving instruments by adopting a classical approach based on visual inspection and measures taken by hand on the marble surface. The work was organized as follows (Figure 2): an initial visual analysis allowed to select twenty-four sections of the surface, considered as the most important for the carving traces characterization; for each section, the visible traces were documented with macro-photography and analyzed; direct measurements were taken of each individual mark, using a caliper and taking direct measures of the length and sections and initial and ending width of each mark; finally, she classified all marks in six classes, corresponding to single-point or multiple-points chisels.




Fig. 2. (a) The characterization (24 sections) of the St Matthew surface; on each section, all chisel marks have been documented and measured; (b) an image of the chisel marks; (c) chisel marks were classified by recognizing six classes corresponding to different tools [16].

It is easy to imagine how time consuming is that manual process. The need to take measures of each single mark requires a lot of time and attention (use of metallic calipers could be harmful for the marble); it requires a close examination of the surface and thus the availability of a scaffolding for a long time, that is very hard to obtain when the artwork is exposed in a museum. Moreover, to be practical (even if very slow) the marks characterization has to be based on just a few parameters: one length and two widths, according to the presented approach.

2.3 Restoration Monitoring

Restoration of artworks is usually paired by a massive photographic documentation that allows to document the pre- and post-restoration conservation status. But taking photographs of a complex artifact from exactly the same set of views in different time frames is not an easy process. Therefore, images parameters are usually different and it is not easy to compare the content of the different images in an objective manner. One possible solution can be to re-project the photos onto the 3D model (by finding intrinsic and extrinsic calibration parameter of each photo), and then render the two models in a coordinated manner (from the same viewpoint) [17]. Cons of this approach are the need to render two very high resolution models at interactive speed and the complexity of doing comparisons using just our visual matching capabilities. We will show in Section 5 how a local flattening can be used to produce a common 2D reference mapping that allows an easy and objective comparison of different photographs.

3 THE LOCAL FLATTENING OPERATOR

The goal of our *local flattening operator* consists of mapping a parameterization of a curved portion of the 3D surface onto a 2D rectangular image, producing as little distortion as possible. We refer to Figure 3 for a more intuitive explanation. The local flattening operation can be subdivided into the following steps:

Sketch. The user sketches two polylines on the object’s surface (Figure 3.a). The first one (which is shown in red and is called *guidance polyline*) defines the patch length and the principal orientation of the cross field that, by definition, the new local parametrization should be aligned with. A second polyline (which is shown in blue and called *interval polyline*) defines the patch width;

Selection. A subset of faces from the original mesh is selected (Figure 3.b), such that their geodesic distance to the *guidance polyline* is less than the *interval polyline*’s length;

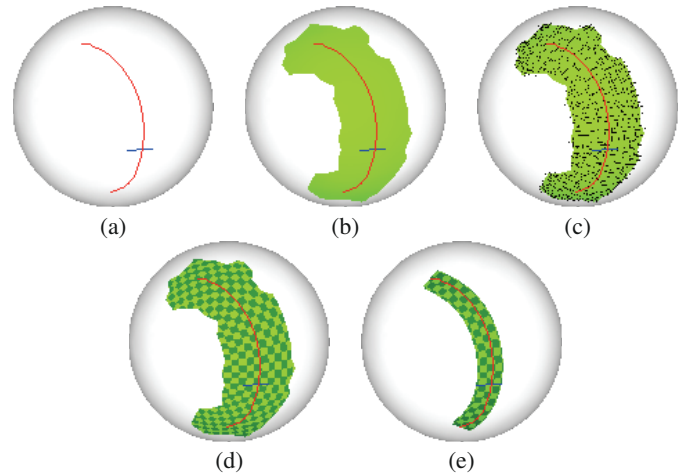


Fig. 3. The steps composing the local flattening operator: (a) user sketches the extents; (b) selection; (c) direction field interpolation; (d) parametrization; (e) cropping to a rectangle.

Direction field interpolation. Upon the selected mesh section, a smooth *cross field* is extrapolated from the *guidance polyline* (see Figure 3.c and Section 3.1);

Parameterization. The selected mesh portion is parameterized according to the field evaluated in the previous step (see Figure 3.d and Section 3.1);

Crop to a rectangle. Finally, the selected mesh is cut to be a rectangle in UV space (Figure 3.e). This is implemented by simply considering the axis-aligned bounding box defined by the UV coordinates of the polylines’s extremes.

The system rejects the user-selected sketch if the selected region is not homeomorphic to a disc. Since we guarantee that the parametrized patch is always homeomorphic to a disc, then we may implicitly avoid discontinuities in UV space (which are usually referred as *cuts* or *jumps* in the literature). This is possible since our method produces a *local* parametrization, while the majority of *global* parametrization methods (e.g. [9, 11]) must necessarily introduce cuts to maps generic shapes (even non homeomorphic to a disk) onto a bi-dimensional space.

3.1 Cross field interpolation and parametrization

In this phase we compute a smooth *cross field* defined over the selected portion of the mesh (as shown in Figure 3.c), which should smoothly follow the direction specified by the guidance polyline. A cross field is generally specified by two unit vectors for each face. We search for an orthogonal cross fields, and thus one vector per face defines the whole field, since the other can be simply expressed by cross product with the face normal.

In our settings each polyline (*guidance* and *interval*) is attached to the mesh surface and it is identified by a sequence of pairs (f_i, β_i) , where f_i is a face of the original mesh and β_i are the barycentric coordinates of a point inside such face. Let’s call U and V axis the ones defining the bi-dimensional domain on which our mesh is parameterized. Our idea consists of smoothly parameterizing the mesh, forcing at the same time the parametric positions of (f_i, β_i) to be aligned with axis U . As a consequence, the U direction of the parametric space defines a smooth 2D vector field by itself and it can be simply transformed back to a 3D smooth vector field. In more detail, we rely on discrete harmonic parametrization using the well known cotangent weights (see [18] for details). We align the polyline barycentric positions to the U axis by adding a least squares term:

$$E_{align} = k_{st} \sum_{i \in poly} (\chi_U^i - U_k)^2 \quad (1)$$




Fig. 4. The user may drive the selection with a few gestures.

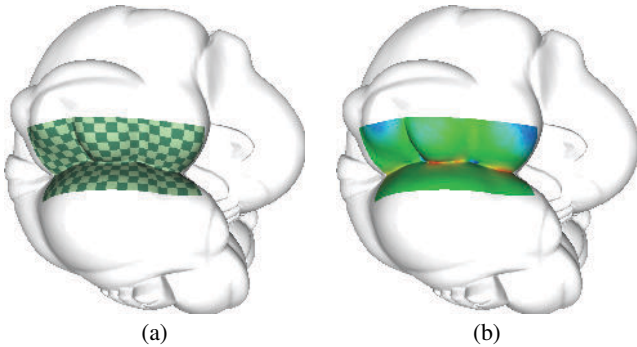


Fig. 5. Parameterizing the extreme case of a portion of surface with high curvature and a crease angle: image (a) shows the parametrization obtained, while image (b) shows the corresponding distortion (blue regions corresponds to expansion, red regions to compression); the edge length distortion is in average 6% and rises to maximum of 27% only in a few small regions.

where $poly$ is the guidance polyline, k_{st} expresses the stiffness of such constraint, U_k is the chosen constant the polyline suppose to be aligned with and, finally, χ_{ij}^U is the U component of the polyline point expressed in barycentric coordinates in UV space.

At this point, it is sufficient to fix the two extremes of the guidance polyline (obviously they should be correctly aligned with U axis) to make the sparse system solvable [19]. Then, once we have computed the cross field, we produce a parametrization aligned with the cross-field as in [11]. Some results are presented in Figures 4 and 5.

4 ANALYSIS OF RESIDUAL TRACES

Methods for visualizing and measuring residual traces over digital models have been developed in the framework of the Digital Michelangelo Project [20] and other more recent works. The classical approach experimented so far in the digital domain has been, first, to produce the outmost accurate digital 3D model of the surface under examination [21] and then to choose either to apply visualization (producing non-photorealistic images that enhance the carvings like [22], to make them more easy to perceive and to compare visually) or to compute the intersection of the surfaces with cut-through planes (section lines give us data on the depth and the 2D shape of sections of the chisel traces, but the selection of those cut-through planes is not easy

and, more important, it is a subjective choice that makes very hard to compare different surface sections).

Our local flattening operator may be conceived as a new digital tool to make measurements and comparisons in a radically new way over the digital representation of the artworks. Our idea is to go beyond the classical approaches by designing a new methodology based on the following three major ideas (see Figure 1):

- Allow the user to select interactively the target surface portions for performing the analysis of the residual traces. This selection stage is based on a very simple sketch-based interaction (just the selection of two polylines, as presented in the previous section); the accompanying video shows the GUI of the interactive tool;
- At the level of the digital 3D representation, split the basic shape description (the overall shape of the sculpture) from the high-frequency detail corresponding to the remains of the chisel marks (the shape detail over the statue corresponding to just the traces of the sculptor's carving tools);
- Transfer high-frequency geometric detail (representing chisel marks) onto an image by exploiting the 3D to UV space mapping specified by the parametrization. Requisite of this mapping is to maintain full control over the metric information encoded in the 2D image;
- Finally, design a set of image processing operations for performing the easy comparisons and analysis of the traces over the 2D representations. These algorithms will work efficiently on the 2D representation and allow the CH expert to compute several different types of measures and matches.

The execution of the second and third step is totally transparent to the user; he just selects the region he wants to work with and he gets as a result the parameterized map onto which he can apply the image processing operations.

4.1 Sampling high frequency detail

The interactive selection of the parameterized patches is performed over a medium resolution model of the artifact; for example, the Pietá is represented at this stage by a 250K faces model⁷. Using a low resolution representation is justified at this stage because we need to obtain a smooth parameterized surface computed at an interactive rate.

Once a rectangular parameterized patch is defined, we can uniformly sample points and normals at the user-selected resolution (specified in pixel per millimeter). The result is therefore a geometry image encoding both shape and normals. But a further goal of this step is to split the high-frequency detail from the basic shape curvature, to allow to represent only the chisel marks in the final image. To produce that, we further smooth geometry and normals belonging to the geometry image by using a gaussian filter. Then sampled points (representing the basic shape in parametric space) are projected along their normal to an high-resolution 3D representation of the artifact (possibly the raster model is produced at the best accuracy allowed by the available 3D scanning technology). Distances between sampled and projected points are finally encoded in a single image. While the recovery of height maps from a 3D model can be done in various sophisticated ways as surveyed in [23], for this class of models having by-definition a smooth base mesh the simple ray shooting approach is quite robust. This process returns a 2D map of scalar values which represent the high frequencies details of the model's surface. Figures 6 and 7 show extraction of different types of residual traces on Michelangelo's Pietá di Palestrina and St Matthew.

To validate the accuracy of the residual traces encoding, sampling a high resolution 3D mesh with a correct inter-sampling distance is not enough. The latter should be paired with a controlled distortion introduced by the parametrization, that should satisfy a given threshold. We quantify the distortion as the difference of the ratio between 3D lengths on the surface mesh and the correspondent length in UV space. We selected by edge lengths preservation since our application domain requires that distances measured in image space should be as




Fig. 6. Example of residual marks extraction over the Michelangelo's Pietà of Palestrina. The residual marks are colored according to their respective orientation, as an example of simple classification. Many other ways of color-coding them (e.g. depth, width) are possible. The sampling resolution of the 2D images (pixel size on the 3D surface) is 0.25 mm.

close as possible to distances measured on the surface embedded in 3D space; therefore, overall edge preservation is a good estimator for isometry.

An analysis of the distortion introduced in our test cases is shown in Table 1. As we can observe, even if a few isolated triangles may introduce up to a 16% of length distortion, the average distortion is usually below 3%. Moreover, we verified for the same experiments that area distortion is always below 31%, while angle distortion is always below 35%. In the experiment shown in Figure 5 we applied our operator to parameterize a region of surface with high curvature and a crease edge, that is a clear worst case. Even in this case the distortion is acceptable (average is 6% while maximum is 27%). Hence, we may conclude that our flattening operator is an accurate instrument to study residual traces, since the potential error is probably much smaller than the one occurring with measures taken manually over the real artifact. Moreover, since we know the distortion, it is also possible to create a map of distortion to compensate per pixel length measured in UV space.

4.2 Image processing of the sampled details

As previously stated in the introduction, there are several practical and theoretical motivations to directly resort to image processing algorithms on the parameterized domain instead of working directly on the 3D mesh domain. We present here the image-based algorithms implemented in a first prototypical version of our system, focusing on efficient scale-invariant chisel traces extraction and automatic chisel marks classification.

Chisel Traces Extraction. The extraction of the residual marks can be performed by using one of the existing methods for the extraction of valleys and ridges from images. From the many existing methods we adopted the one proposed in [24]. It is particularly suitable for our purposes, since it is able to deal with multi-scale extraction: a scale-space representation of the input image is build; then, ridges and valleys are automatically identified according to a differential geometric definition; finally, the identified ones are filtered to keep only the salient ones. This scale-space extraction $L(\cdot)$ is obtained by convolving the input image with gaussian kernels of increasing size:

$$L(x, y, \sigma) = I(x, y) * G(x, y, \sigma) \quad (2)$$

where $I(x, y)$ is the input image and $G(x, y, \sigma)$ is a gaussian kernel with spatial extension of σ . This scale-space representation is particularly efficient to construct, since the convolution can be evaluated efficiently on the GPU [25].

The extraction of the ridges/valleys at a fixed scale (σ) is done according to the following constrains:

$$L_p = 0, L_{pp} < 0, |L_{pp}| \geq |L_{qq}| \text{ for ridges} \quad (3)$$

$$L_q = 0, L_{qq} > 0, |L_{qq}| \geq |L_{pp}| \text{ for valleys} \quad (4)$$

where L_p, L_q, L_{pp}, L_{qq} are the first-order and second-order directional derivatives aligned with the directions of the principal curvature in UV-space. Scaling factor of the features can be achieved by considering the scale-space representation, and thus an additional constrain on $L_t(\cdot)$ is added. This involves the computation of higher-order partial derivatives. Since the partial derivatives of $L(\cdot)$ can be expressed as a convolution with the partial derivatives of the Gaussian kernels:

$$\partial_{x^\alpha} \partial_{y^\beta} L(x, y, \sigma) = I(x, y) * \partial_{x^\alpha} \partial_{y^\beta} G(x, y, \sigma) \quad (5)$$

the numerical evaluation of even higher-order partial derivatives can be done with high accuracy. Therefore, we obtain high accuracy in the extracted ridges/valleys. Figure 8 shows an example of traces computation. Once the user has indicated a trace of interest, other chisel traces are computed having similar *ridges strength* values (see the original paper for further details [24]).

Please note how our final extraction may considerably simplify the measurement task: the measure of length is immediate while, due to the extraction of ridges, the depth of a trace can be simply evaluated as the difference in height with respect to the closest ridge.

Similarity Analysis of Chisel Marks. The possibility to evaluate chisel marks *similarity* leads to useful applications. For example it might allow the scholar to characterize a particular sculpting technique, or it may be useful in the classification of carving tools. As starting point, we measured similarity in terms of ridges orientation and scale of the extracted traces. A result is shown in Figure 9. More sophisticated chisel marks characterizations are possible, for example




Fig. 7. Example of residual marks extraction over the Michelangelo's St. Matthew. Refer to Figure 6 for a description of the color coding. Similarly to the Pietà the sampling resolution is 0.25 mm.




Fig. 8. Computation of chisel traces: user selection of a sample trace (left); valleys (in red) and ridges (in blue) extracted automatically from the user example (right).

taking into account depth, length, gradients changes in the surrounding of the residuals and so on. Extensions in this direction are leaved as a matter of interesting future research.

A completely different approach to analyze the surface's characteristics, that we find interesting here to mention, consists in identifying similar surface regions by applying some of the recent approximate nearest neighbors algorithms for image patches (e.g. the *patch-match* [26] algorithm). In this way, the surfaces characterization is turned into a problem of texture segmentation. Obviously, each parameterized patch should be pre-processed in order to compensate the

user selection and to put it in a common framework. This is another example of image-based processing that it is still not easy to extend to 3D meshes.

5 MEASURABLE EVALUATION OF THE RESTORATION RESULTS

A second example of the capabilities of the proposed framework is the visual comparison of the appearance of an artwork before and after some critical events like a restoration. Currently, such comparison is difficult to be performed because the status is documented by photographs that, even when taken with great accuracy, are never perfectly aligned.

We show an application of our new approach to data from the Michelangelo's David restoration (2003-2004). Two photographic campaigns were carried out, before and after the restoration, but even if these acquisitions were carefully planned the resulting photos were taken from slightly different positions. Those photographs were calibrated and re-projected onto the David 3D model using the approach described in [17]. By using the flattening operator approach we can define in a simple way a common domain of comparison between the photos of the two campaigns (Figure 10). The common parametric domain is used to flatten the corresponding section of two photos into a common rectangular 2D space where they result unwrapped and perfectly aligned. This makes quite easy both visual and numerical analysis of the changes occurred after the restoration. Figure 10 shows two examples focused over the left arm and the right shoulder: the bottom images present corresponding subset of some photographs pairs (presented in the upper row) unwrapped on a common space through the local flattening operator; the middle row presents the locations of the flattening region onto the statue. As an example, the flattened images of the left arm present very clearly the removal of the white plaster fillings of the cracks resulting from the arm fracture happened in the 16th century. Similarly, the extensive cleaning on the back of the right shoulder is very evident on the other image pair.




Fig. 9. On the left, an example of residual traces extracted at multiple scales (two scales are used, thin lines are associated to fine scale traces, while thick lines corresponds to the larger ones); on the right, color-coded orientation of the extracted traces.

Test Num	Min Ratio (optimum=1)	Max Ratio (optimum=1)	Average Diff (% of real length)
1	0.9275	1.0891	3.06
2	0.9396	1.0937	2.79
3	0.9156	1.1696	3.22
4	0.8276	1.1563	3.99
5	0.8716	1.1328	4.79
6	0.9034	1.0742	1.36
7	0.9533	1.0716	1.93
8	0.9470	1.0929	2.50
9	0.9305	1.0889	3.07

Table 1. Numeric evaluation of the length distortion introduced by the flattening operator: the test cases are those presented in Figures 6 and 7; we show respectively minimum and maximum ratio between 3D length and UV length; the last column is the average ratio, in percentage.

This application can be seen as a generalization to 3D shapes of the ortho-rectification tools commonly used in topography for the production of architectural drawings.

6 CORRELATE MARKS WITH CURVATURE

The proposed framework allows also to compare in a common parametric space different shape characteristics, possibly also at different level of scale. In this sense, another application is to compare the main surface curvature direction, evaluated on a low resolution and smoothed representation, with the orientation of the residual traces. This idea comes from a simple intuition: might the sculpting technique (direction of chisel traces) be somehow connected with the shape (i.e. the surface curvature) the artist wants to emerge from the stone mass?

To try to find some evidence for this hypothesis, we first could characterize the basic shape by applying an estimator of the main curvature directions over a low resolution smoothed model (we rely on the approach proposed in [27]). Then, we overlap the curvature directions with an image showing the chisel marks, that has been produced with the approach presented in Section 4. The final result of this processing is shown in Figure 11. It is interesting to notice how the low-scale surface curvature is somehow related with the direction of the chisel marks. The further analysis of this derived information is obviously a task for Cultural Heritage scholars. We hope that this characterization




Fig. 10. Comparing, in a common parameterized domain, the pre- and post-restoration status of the Michelangelo's David surface. Given some photos taken in different times and positions, registered over the 3D David model, our approach allows to define a common comparison domain that allow immediate quantitative and qualitative comparisons.

may reveal interesting clue about Michelangelo's sculpting style.

7 CONCLUSIONS AND FUTURE WORKS

We presented a novel framework to support CH scholars in the shape-driven characterization and analysis of digital 3D artworks models. Our framework is based on a local surface parametrization operator, called *local flattening operator*. This operator allows to produce a rectangular image corresponding to a portion of the model, with controlled distortion, that allows us to infer shape-related knowledge by means of more simple image processing algorithms. The local flattening operator is interactively controllable by the user, who can specify, by using an intuitive interface, the field that the produced parametrization is aligned with. We successfully adapted our framework to various real application scenarios, concerning the analysis and monitoring of a sculpture's surface. We shown that this approach allows to implement measuring, classification and comparison kernels in an efficient and precise way by simply deploying 2D filters over parameterized 2D encodings of the original 3D surface. The proposed system will be at the base of an extensive study over some selected Michelangelo's artworks, conducted by a multidisciplinary consortium that is using this new digital technology in the framework of consolidated CH methodologies. The goal is to assess the proposed methodology and the tool developed in the framework of a concrete CH research and, hopefully, to extend the investigation and analysis capabilities of CH experts by means of this new digital tool.

Finally, we underline that the proposed local flattening operator can be applied for detail extraction also in other applications scenarios and it is not limited to the CH domain. For example, it could be used




Fig. 11. Correlating low scale curvature (shown with crosses) with residual traces on Michelangelo's S. Matthew (the accompanying video shows effectively the correlation between those two signals, by using a progressive fading of one image into another).

for the analysis of biological surface patterns to drive texture synthesis algorithms or for the characterization of the micro-structure of certain kind of surfaces like skin or brushed/wrinkled metallic surfaces.

Acknowledgments The research leading to these results has received funding from the European Community 7th Framework Programme (FP7/2007-2013) under grant agreement no. 231809 (IST IP "3DCOFORM"). We would like to thank Franca Falletti, director of the Accademia Museum, and the other CH colleagues involved in this activity: prof. Irving Lavin, prof. Alessandro Nova, dr. Ute Dercks, dr. Michael Pfanner).

REFERENCES

- [1] K. Hormann, K. Polthier, and A. Sheffer, "Mesh parameterization: Theory and practice," in *SIGGRAPH Asia 2008 Course Notes*, no. 11. Singapore: ACM Press, Dec. 2008.
- [2] A. W. F. Lee, W. Sweldens, P. Schröder, L. Cowsar, and D. Dobkin, "MAPS: Multiresolution adaptive parameterization of surfaces," in *Proceedings of SIGGRAPH 98*, Jul. 1998, pp. 95–104.
- [3] A. Khodakovsky, N. Litke, and P. Schröder, "Globally smooth parameterizations with low distortion," *ACM Trans. on Graphics (Proc. SIGGRAPH)*, pp. 350–357, 2003.
- [4] N. Pietroni, M. Tarini, and P. Cignoni, "Almost isometric mesh parameterization through abstract domains," *IEEE Trans. on Visualization and Computer Graphics*, vol. 16, pp. 621–635, 2010.
- [5] D. Panozzo, E. Puppo, M. Tarini, N. Pietroni, and P. Cignoni, "Automatic construction of adaptive quad-based subdivision surfaces using fitmaps," *IEEE Trans. on Visualization and Computer Graphics*, vol. 17, 2011.
- [6] M. Jin, Y. Wang, S.-T. Yau, and X. Gu, "Optimal global conformal surface parameterization," in *Proc. Visualization '04*. IEEE Computer Society, 2004, pp. 267–274.
- [7] S. Dong, P.-T. Bremer, M. Garland, V. Pascucci, and J. C. Hart, "Spectral surface quadrangulation," *ACM Trans. Graph.*, vol. 25, pp. 1057–1066, 2006.
- [8] Y. Tong, P. Alliez, D. Cohen-Steiner, and M. Desbrun, "Designing quadrangulations with discrete harmonic forms," in *EG Symp. on Geometry Processing*. Eurographics Association, 2006, pp. 201–210.
- [9] N. Ray, W. C. Li, B. Lévy, A. Sheffer, and P. Alliez, "Periodic global parameterization," *ACM Trans. Graph.*, vol. 25, pp. 1460–1485, 2006.
- [10] F. Kälberer, M. Nieser, and K. Polthier, "Quadcover - surface parameterization using branched coverings," *Computer Graphics Forum*, vol. 26, no. 3, pp. 375–384, 2007.
- [11] D. Bommers, H. Zimmer, and L. Kobbelt, "Mixed-integer quadrangulation," *ACM Trans. Graph.*, vol. 28, no. 3, 2009.
- [12] A. Myles, N. Pietroni, D. Kovacs, and D. Zorin, "Feature-aligned t-meshes," *ACM Trans. Graph.*, vol. 29, pp. 117:1–117:11, 2010.
- [13] Z. Zhang and H. Zha, "Principal manifolds and nonlinear dimensionality reduction via tangent space alignment," *Journal of Shanghai University (English Edition)*, vol. 8, pp. 406–424, 2004.
- [14] Z. Chen, L. Liu, Z. Zhang, and G. Wang, "Surface parameterization via aligning optimal local flattening," in *Proceedings of the 2007 ACM symposium on Solid and Physical Modeling*. ACM, 2007, pp. 291–296.
- [15] S. T. Roweis and L. K. Saul, "Nonlinear dimensionality reduction by locally linear embedding," *Science*, vol. 290, pp. 2323–2326, 2000.
- [16] P. Giovannini, "Il San Matteo di Michelangelo: analisi delle tracce di lavorazione, studio degli strumenti e osservazioni sulla tecnica di scultura," *OPD Restauro*, vol. 10, pp. 205–228, 1998.
- [17] M. Dellepiane, M. Callieri, F. Ponchio, and R. Scopigno, "Mapping highly detailed color information on extremely dense 3D models: the case of David restoration," *Computer Graphics Forum*, vol. 27, no. 8, pp. 2178–2187, 2008.
- [18] M. Eck, T. DeRose, T. Duchamp, H. Hoppe, M. Lounsbery, and W. Stuetzle, "Multiresolution analysis of arbitrary meshes," *Proceedings of SIGGRAPH'95*, pp. 173–182, 1995.
- [19] G. Guennebaud and B. Jacob, "Eigen: a c++ template library for linear algebra." [Online]. Available: <http://eigen.tuxfamily.org>
- [20] M. Levoy, K. Pulli, B. Curless, S. Rusinkiewicz, D. Koller, L. Pereira, M. Ginzton, S. Anderson, J. Davis, J. Ginsberg, J. Shade, and D. Fulk, "The Digital Michelangelo Project: 3D scanning of large statues," in *SIGGRAPH 2000, Computer Graphics Proceedings*, ser. Annual Conference Series. Addison Wesley, July 24-28 2000, pp. 131–144.
- [21] G. Godin, J.-A. Beraldin, J. Taylor, L. Cournoyer, M. Rioux, S. El-Hakim, R. Baribeau, F. Blais, P. Boulanger, J. Domey, and M. Picard, "Active optical 3d imaging for heritage applications," *IEEE Computer Graphics and Applications*, vol. 22, no. 5, pp. 24–36, 2002.
- [22] S. Rusinkiewicz, M. Burns, and D. DeCarlo, "Exaggerated shading for depicting shape and detail," *ACM Transactions on Graphics (Proc. SIGGRAPH)*, vol. 25, no. 3, Jul. 2006.
- [23] L. Szirmay-Kalos and T. Umenhoffer, "Displacement mapping on the gpu state of the art," in *Computer Graphics Forum*, vol. 27, no. 6, 2008, pp. 1567–1592.
- [24] T. Lindeberg, "Edge detection and ridge detection with automatic scale selection," *Int. J. Comput. Vision*, vol. 30, pp. 117–156, November 1998.
- [25] O. Fialka and M. Cadik, "Fft and convolution performance in image filtering on gpu," in *Conf. on Information Visualization*. Washington, DC, USA: IEEE Computer Society, 2006, pp. 609–614.
- [26] C. Barnes, E. Shechtman, A. Finkelstein, and D. B. Goldman, "Patch-Match: A randomized correspondence algorithm for structural image editing," *ACM Trans. on Graphics (Proc. SIGGRAPH)*, vol. 28, no. 3, Aug. 2009.
- [27] Y.-L. Yang, Y.-K. Lai, S.-M. Hu, and H. Pottmann, "Robust principal curvatures on multiple scales," in *EG Symp. on Geometry Processing*, A. Sheffer and K. Polthier, Eds. Cagliari, Italy: Eurographics Association, 2006, pp. 223–226.