

Special Section on CEIG 2023

Automated digital color restitution of mural paintings using minimal art historian input

Imanol Munoz-Pandiella^{a,*}, Carlos Andujar^b, Begonya Cayuela^c, Xavier Pueyo^d, Carles Bosch^e

^a Universitat de Barcelona, Spain

^b Universitat Politècnica de Catalunya, Spain

^c Independent Researcher, Spain

^d Universitat de Girona, Spain

^e Universitat de Vic - Universitat Central de Catalunya, Spain

ARTICLE INFO

Article history:

Received 21 May 2023

Accepted 14 June 2023

Available online 22 June 2023

Keywords:

Computer Graphics

Cultural Heritage

Color transformation

ABSTRACT

Digital color restitution aims to digitally restore the original colors of a painting. Existing image editing applications can be used for this purpose, but they require a select-and-edit workflow and thus they do not scale well to large collections of paintings or different regions of the same painting. To address this issue, we propose an automated workflow that requires only a few representative source colors and associated target colors as input from art historians. The system then creates a control grid to model a deformation of the CIELAB color space. Such deformation can be applied to arbitrary images of the same painting. The proposed approach is suitable for restituting the color of images from a large photographic campaign, as well as for the textures of 3D reconstructions of a monument. We demonstrate the benefits of our method on a collection of mural paintings from a medieval monument.

© 2023 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Paintings suffer from physical and chemical changes over the years that degrade their visual color appearance. In the case of mural paintings, damages can also affect the physical support of the paintings, causing in some cases a substantial loss of content. *Digital color restitution* aims at digitally simulating how the colors of a painting looked at a precise time by trying to reverse only the color degradation. It is always approximate (since a display cannot render exactly the appearance of the painting), hypothetical (as the effect of the degradation process is not invertible), and limited to those regions of the painting whose current reflectance properties provide some evidence that a particular color was used; it does not attempt to recover content in severely degraded parts where color information is lost.

In contrast, *virtual restoration* is more ambitious in the sense that it also attempts to reconstruct content in severely damaged parts of the painting [1]. The digital restorer has to rely on contextual information to recover the content. Since virtual restoration must face the problem of missing content, art historians must necessarily use image editing tools for different tasks, such as segmenting the image according to the degree of degradation,

segmenting regions with different presumed original colors, and drawing missing content. These operations require, besides art history research, a considerable amount of image editing work that is pixel-wise specific to a given image and therefore it is not straightforward to transfer these operations to a different region of the same painting.

This paper deals with the problem of digital color restitution of medieval mural paintings. Since color restitution only attempts to restore the colors in reasonably well-preserved areas, without modifying damaged regions, the restitution process should not require the art historians to select specific regions of the image to be edited. Instead, the user input can consist of a collection of color samples, not necessarily from the image to be digitally restituted. Such an approach provides significant advantages in terms of scalability: the same user input can be reused to reconstitute the color of different images of the same painting, and even images of other paintings from the same monument, provided that the hypothesis encoded by the chosen color samples also applies to the other paintings. Notice that this scalability is critical when many different images must be digitally restituted. One example is when photogrammetry software is used to create a 3D reconstruction of a monument. Different executions of the application might produce different mesh parameterizations, thus modifying the resulting textures to be restituted. Alternatively, color restitution can be applied to the photographs used for

* Corresponding author.

E-mail address: imanol.munoz@ub.edu (I. Munoz-Pandiella).

colorizing the model, but this means that multiple overlapping views of the paintings have to be digitally restituted. Finally, some monuments include large collections of mural paintings, which involve capturing thousands of images. As an example, the paintings of the Sta Maria Antiqua church (Rome) cover a surface of about 250 square meters, and some of them decorate non-planar surfaces such as the apses [2]. The imagery for a complete 3D reconstruction of the monument with its mural decorations contains several gigapixels, thus requiring scalable color restitution techniques.

We propose a color restitution workflow that requires minimal input from art historians. For each representative color (e.g. light ochre) identified in the paintings, the user provides the CIELAB values of a few source samples (i.e. samples of the current appearance of the color, showing a moderate level of degradation) and a target sample (the presumed original color, taken for example from a well-conserved sample). In the case of medieval monuments, the number of representative colors (pigment combinations) for a given painting is relatively small (as discussed in Section 5).

The key ingredient of our approach is a deformation of the CIELAB color space, subject to a collection of constraints derived from the user input, in addition to some boundary constraints. In particular, we use a cage-based deformation approach based on generalized barycentric coordinates. The color restitution algorithm proceeds as follows. First, user-provided color samples are converted into a collection of (source, destination) pairs. These pairs will be used as handles to control the deformation. Second, we create an initial cage by slightly modifying a regular tetrahedralization covering the CIELAB space so that all source colors are represented exactly by one vertex of the tetrahedralization. Third, we use the (source, destination) pairs as deformation constraints, and use a shape deformation method to create a deformed cage that satisfies the deformation constraints at the handles, followed by a smooth propagation of these handle deformations to the rest of the cage. The initial and final cages define a deformation of the complete CIELAB space (through barycentric interpolation within each tetrahedron) that can be applied for color restitution of any (color-calibrated) photograph of the painting.

Since we focus on digital color restitution, human color perception plays an essential role. In particular, we need to work in a perceptually-uniform color space to avoid differences between the degree of transformation of any color and its perception. For this reason, we propose to use CIELAB as the working color space, because it was designed so that uniform changes in L^*a^*b components correspond to uniform changes in perception. Moreover, it is the standard color space used in applications where perception is important, such as Cultural Heritage [3]. However, our method should correctly work with any other perceptually uniform color space.

We use as a test dataset a collection of Romanesque-style frescoes from the medieval church of Sant Quirze de Pedret (described in detail in Section 3). The formerly vivid colors of the paintings have been affected not only by the passage of time but also by numerous anthropogenic interventions, including the detachment (and the latter reassembling) of the paintings for preservation. Since the original paintings are now distributed across different museums, the understanding of the monument greatly benefits from a digital reintegration of its paintings. This also illustrates the need for scalable color restitution techniques, since hundreds of photographs of the frescoes were taken to build a 3D reconstruction of the monument.

2. Previous work

Virtual restoration Digital or virtual restoration of paintings might involve varied operations such as completing the drawings of figures and other pictorial elements, filling detached or deteriorated parts, and recovering the original colors and texture of the paintings [1]. Such operations are performed after capturing the paintings and unfolding them onto an appropriate 2D representation. For color restoration, experts might rely on clean areas that have not suffered the same damage or which have been previously restored, in order to replace the current ones [4]. Otherwise, it might be necessary to prepare the same pigment mixtures as in the original painting after analyzing the underlying materials. The most accurate way to analyze and identify materials is by taking micro-samples of the paint layer, which is invasive and limited to specific portions [5]. Other approaches such as Multispectral Imaging (MSI) and Hyperspectral Imaging allow the identification of pigments and materials in a non-destructive way [6]. These techniques have been used extensively for analyzing paintings' pigments and their degradation, as well as for revealing retouching or restoration campaigns. Exact identification is however a hard problem, since pigments are often mixed in layers to make a desired color and effect, and the acquired spectra may refer to more than one layer. In addition, such approaches are limited to specific portions of the paintings and to planar shapes. Colorimetric analysis, such as using spectrophotometry devices, is often applied to validate the captured images (MSI or color-calibrated photographs). It has been shown that such images provide a similar accuracy as long as the lighting is uniform enough [7]. After having identified the appropriate colors, most authors finally rely on photo editing software to manually replace colors or complete the overall painted surface, which is a time-consuming task [1,7]. In our case, we focus on color restitution and propose an automatic approach that only requires the selection of a set of representative colors.

Automatic color restoration Pappas and Pitas first evaluated a set of models to automatically transform colors from dirty to clean datasets, applied to old paintings suffering from varnish oxidation [3]. Such models included sample mean matching, linear, ICP (Iterative Closest Point), white-point, and RBF (Radial Basis Functions), all applied in CIELAB color space. These transformations were uniformly applied onto the whole color space as global operators, and hence had limited flexibility. Other subsequent works suffered from the same limitation [8,9]. Color transformation can also be learned from provided samples, either using Support Vector Machines [10] or Neural Networks [11,12]. These methods, however, often require a considerable amount of samples for the learning and might require resorting to synthetic data [13]. Our method relies on a more flexible transformation that deforms the color space to adapt to the selected representative colors, without any learning involved.

Color transfer and correction In computer graphics and vision, color transfer and color correction are two areas that focus on the general problem of transferring colors between images, either for correcting colors, blending images, or stitching them [14]. Color transfer approaches can be divided into parametric and non-parametric ones, where the former rely on explicit models for estimating the transfer [15]. Most of these methods focus on global transformations [16], but local approaches tend to provide more control [17]. Hwang et al. proposed a non-parametric solution relying on a set of control points in RGB space for computing the transform [18]. The transformation between the two sets is done using Moving Least Squares and a probabilistic measure to improve robustness. The method was later extended to spatially varying color transfer [19]. Our approach uses a similar idea but uses a cage-based deformation with constraints derived from

user samples and boundary constraints specifically tailored to the color restitution of paintings.

Colors in medieval mural paintings Unfortunately, the study of the colors of Romanesque paintings is often a neglected topic among experts [20]. In the Middle Ages, the palette of colors that could be applied with the mural technique of fresco was restricted, due to the use of lime as a support and binder [21,22]. The most common pigments were obtained from the nearby environment and came from easily available natural soils, compatible with lime, that provided white, black, red, and yellow colors. Green was rarer in nature and was therefore used more sparingly, in specific areas such as shadows on the flesh tones, highlights, etc. Blue pigments were even more difficult to obtain and apply, because they were of a mineral nature, incompatible with lime. Moreover, they were rare in nature and usually imported, which made them very expensive pigments. However, in the Romanesque frescoes of the Pyrenees, the use of aerinite, a clay mineral that gives a less intense but cheaper blue, was used more frequently [23]. In our approach, we take into account these studies to decide the representative set of colors for color restitution.

3. Cultural heritage requirements regarding color restitution

We have developed our research by taking the church of *Sant Quirze de Pedret* in Cercs, Catalonia, as a case study (Fig. 1). This small Romanesque church was built in the late 9th century as a single-nave church with a square, horseshoe-vaulted apse. Later, in the middle of the 10th century, it was enlarged and transformed into a three-nave building with corresponding apses. It was decorated twice with wall paintings. An early medieval phase shows rustic and enigmatic figures [24]. Around 1100, a more ambitious campaign of Romanesque-style frescoes adorned the eastern third of the church, covering the earlier decoration [25]. The Romanesque phase survived only until the end of the 13th century when a fire destroyed some parts of the building. In addition, two-thirds of the southern nave collapsed for reasons that are not yet clear, and a major renovation of the church was undertaken. The resulting building hid a considerable part of the Romanesque frescoes between the walls and under a new layer of non-figurative decoration. It was only in the last quarter of the 19th century that these frescoes were noticed and appreciated by a wider public.

Unfortunately, the discovery of all these medieval Catalan frescoes attracted private collectors and international art dealers who began to buy the paintings and remove them from their original walls. The Board of Catalan Museums reacted quickly and negotiated with the dealers to buy the paintings for the Barcelona Art Museum, now *Museu Nacional d'Art de Catalunya* (MNAC). Nevertheless, the official operation did not spare the detachments of these Catalan medieval frescoes, as those of Pedret in 1921. During the turmoil of the Spanish Civil War, new paintings were discovered under the walls of the central nave and the white layer that covered the main apse. After a bizarre sequence of events, in 1939 the frescoes that had just been removed eventually ended up in the collections of the *Museu Diocesà i Comarcal de Solsona* (MDCS).

The fact that the paintings were removed (and subsequently reassembled) at different times, by completely different teams, and that they have been preserved in two different locations, with their respective conditions and conservators, has resulted in the two paintings undergoing two sufficiently differentiated aging processes, causing the current appearance of their colors to differ from each other. This characteristic makes this church an ideal case study, not only for the digital color restitution of mural paintings but also in the extreme case where such paintings have undergone different aging and conservation processes.



Fig. 1. Current situation of the interior of St. Quirze de Pedret church (central nave and central and north apse, left), as well as the exhibition of its paintings in MDCS (central nave and apse, center) and MNAC (north and south apse, right). The fragmentation of the artworks across three different locations has resulted in varying appearances of originally equivalent colors, due to different weathering and curation processes.

For a thorough understanding of the Romanesque decoration of Pedret, digital tools can play an invaluable role when a physical restoration is unfeasible. One of the goals is to perform a digital color restitution without trying to fill in all the gaps to reconstruct the presumed original state. This phase involves estimating what the colors of paintings looked like at the time of creation in the 12th century, by attempting to reverse the current state of decay that affect the visual color image. Digital color restitution means that the scope is limited to those areas of the paintings whose reflective properties (at the time the painting was digitized) indicate the use of a particular color. No attempt is made to restore the content in severely degraded areas where color information has been lost nor to fully remove the degradation effects of the colored parts themselves, in order to maintain the painting's history.

4. Color acquisition

The paintings were captured in photographs by an expert photographer from our team (see Fig. 2left). During the process, he followed a methodology that attempted to homogenize the paintings' lighting and maximize the detail in the images, while guaranteeing an important overlap between images. To achieve this, the photographer used a Canon 6D Mark II DSLR camera with a full-frame sensor and various lenses (from normal to telephoto) to maximize the sharpness of the images, taking into account the accessibility to the paintings. In addition, a lighting scheme with two Godox 600 W flashes placed at 45° (one to the left and one above) was used to try to maintain constant lighting of the paintings and enhance their texture. To ensure color accuracy, the images were calibrated using x-rite ColorCheckers, in groups of a small number of images with the same lighting configuration. In this way, the post-processing of the images was minimalist: only white balance and exposure were adjusted, and the corresponding ICC profile was applied.

We used the obtained images to generate textures for a 3D architectural model of the church. To achieve this, we used a photogrammetric software to reconstruct the 3D model of each exhibited set of paintings. For each model, we then obtained a texture with a resolution of less than 0.2 mm per texel. Initially, the architectural model of the church was textured using Adobe Substance, which displayed the hypothetical historical materials. To incorporate the paintings, we aligned the photogrammetric models of the paintings with the architectural model and created a new texture by baking the paintings onto the hypothetical materials. Although thanks to this procedure we obtained a model that shows the paintings again unified (see Fig. 2right), they still show aged colors, the appearance of which varies according to their current location.



Fig. 2. The acquisition process of the paintings done by the photographer (left) and the result of transferring the paintings to the 3D model of the church (right).

5. Color selection

Our method requires the selection of a color table that is enough representative of the colors used in the paintings to be restituted. For each of these colors, we need its current appearance and the presumed original one, in order to transform from one to another.

In order to fix an estimation of the original colors we used the calibrated high resolution images. To select the colors with greater accuracy, we compared them with the schemes available in previous technical reports preserved in the archives of the Conservation and Restoration Department of the MNAC Museum [26,27]. These schemes indicate the points where the micro-samples were collected for the study of pigments used by the medieval artists. Then, to correct the obtained RGB values, we compared them with colorimetric data provided by databases of historical pigments, when the information was available.

According to the technical reports, the basic palette of Pedret's frescoes corresponds to the colors white, black, red, and yellow, which were obtained with cheap pigments. Green was also used quite extensively, but its composition has not yet been established confidently. With these basic colors, medieval painters were able to obtain various shades of other hues, such as gray, pink, skin color, and so on. For this reason, we have also added grays, pinks, and some ochres to the color chart. The resulting selected color table is as follows (see Fig. 3): white, pale ochre, yellow light ochre, middle ochre, yellow ochre, pink, orange, red, deep red, light green, green, light gray, middle gray, deep gray, and black.

The southern apse was chosen as the reference for the selection of the original colors because most of the samples for pigment analysis carried out in the past were extracted from this area. The remaining frescoes were used to complete the color chart, although they have not received the same attention from the experts. The project's art historians were responsible for color selection to meet the engineers' requirements for digital restitution. The original requirement was simply to select the most representative colors, verify the values in the RGB system of digital colors, and to create a table with the chromatic data and location of the selected points.

The MNAC reports [26,27] include a description of the state of deterioration of the frescoes, an identification of some pigments used by the medieval painters, the technique employed when playing with thin layers of different pigments to obtain either deeper or lighter shades, and a mapping of the altered areas. With this information, we tried to ensure that the candidate colors in the calibrated photos were not overpainted or retouched. Nevertheless, the available colors in the calibrated photos represent the current state of the frescoes. Therefore, we must keep in

Colors palette of Pedret			
white			
photo	x,y	RGB color	
	148,148,148	255,255,255	
pale ochre			
photo	x,y	RGB color	
	245,227,205	245,227,205	
yellow light ochre			
photo	x,y	RGB color	
	232,198,160	232,198,160	
middle ochre			
photo	x,y	RGB color	
	235,186,133	235,186,133	
yellow ochre			
photo	x,y	RGB color	
	235,186,133	235,186,133	
pink			
photo	x,y	RGB color	
	235,186,133	235,186,133	
orange			
photo	x,y	RGB color	
	235,186,133	235,186,133	
red			
photo	x,y	RGB color	
	235,186,133	235,186,133	
deep red			
photo	x,y	RGB color	
	148,148,148	148,148,148	
light green			
photo	x,y	RGB color	
	235,186,133	235,186,133	
green			
photo	x,y	RGB color	
	235,186,133	235,186,133	
light gray			
photo	x,y	RGB color	
	235,186,133	235,186,133	
middle gray			
photo	x,y	RGB color	
	235,186,133	235,186,133	
deep gray			
photo	x,y	RGB color	
	235,186,133	235,186,133	
black			
photo	x,y	RGB color	
	235,186,133	235,186,133	

Fig. 3. Table of reference colors.

mind that they cannot be assumed as the exact original colors. Since no clean or undamaged areas are actually available for these paintings, as happens with most medieval mural paintings, we found this method as the most appropriate one.

With these constraints in mind, the first step was to select the reference color palette using the most representative points in the calibrated photos. The preferential points were near those taken to extract the samples for the physic and chemical analysis, when available. The measurement of these colors was fine-tuned by the art historian to hypothetically invert their weathering further.

After defining the reference palette, the next step was to collect the samples representing the current state for the same colors, hence generating the second palette or color table. For this, additional points were taken with similar colors as those referenced, using all the calibrated photos available for the different spaces of the fresco's decoration. This was done to account for different variations from the original, reference colors, which might be due to different aging or curation processes.

The art historians developed a custom app to select different points per color in the calibrated photographs and immediately validate their suitability before storing the x, and y coordinates of the sample and its RGB values. This was used for the selection of the two color tables. The main procedure was to generate samples of different sizes (e.g. 3×3 , 5×5) per color to obtain the mean chromatic value (Fig. 4). It was important that the selected pixels did not include pixels with different hues.

The chromatic analysis of medieval paintings based on calibrated photographs raises several questions. The main problem is the lack of a free and complete reference color table in RGB for all the historical pigments used by medieval painters. Moreover, the medieval process of applying different layers of pigment to change the color shades, as shown in the stratigraphies analyzed in the technical reports, hinders the task to assess the color values. For example, to make the green look deeper, the painter




Fig. 4. Sampling of reference colors.

sometimes put a black layer underneath [27]. On the other hand, we have no references for color blends, like the green of Pedret, which is a mixture of yellow and blue. These medieval techniques made it difficult to find the appropriate RGB reference values.

6. Color transformation

Our proposal relies on the simple but powerful idea that a color transformation can be considered analogous to a geometrical transformation of a volume. This way we consider both color sets (current and original) as the handles that guide the deformation of the whole color space to revert the colors' aging process. Following this idea, we explain below the chosen geometrical transformation, the subdivision of the color space we use and the settings the user can modify.

6.1. Deformation model

We propose to transform colors by adapting the linear shape deformation technique by Wang et al. [28] available at the [igl](https://github.com/libigl/libigl) library ([libigl.github.io](https://github.com/libigl/libigl)). This approach unifies linear blending skinning and generalized barycentric coordinates. After defining a set of handles \mathbf{H} , which can be either points or regions, the method computes a set of weights \mathbf{W} for each handle that allows to linearly compute the positions of the mesh vertices \mathbf{V} from the positions of the handles \mathbf{H} :

$$\mathbf{V} = \mathbf{WH} \quad (1)$$

where $\mathbf{V} \in \mathbb{R}^{n \times d}$, $\mathbf{W} \in \mathbb{R}^{n \times m}$, $\mathbf{H} \in \mathbb{R}^{m \times d}$, d is the dimension of the space ($d = 3$ in our case), n is the number of vertices of the mesh, and m is the number of handles. The weights \mathbf{W} are computed by optimizing a quadratic smoothness energy problem.

This technique produces subspaces with many desirable properties: interpolation, smoothness, shape awareness, locality, and both constant and linear precision. Moreover, speed is another important advantage, since the method requires solving a sparse linear system. Just as the vast majority of deformation techniques, this technique is designed for placing control points and, thus, deforming the surface of 3D models. However, we propose to use it with another aim. We would like to preserve the shape steady (i.e. the boundaries of the color space) and transform the inside of the volume to modify the color space.

6.2. Colors as handles

Handles are one of the main components of the geometrical deformation technique used. They allow the user to control the deformation of the mesh. Just as points and regions are used in that case, we propose using colors as the handles H_c of our color space transformation. More specifically, we propose to use the sets of colors defined by the user and described in Section 5 as the handles that guide the transformation.

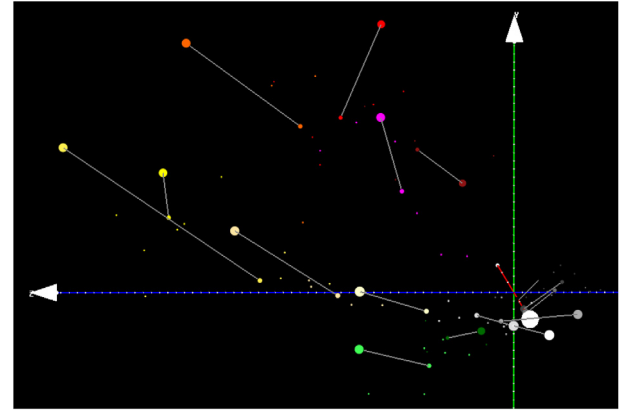


Fig. 5. Color samples S_c (unconnected dots) and the corresponding transformations from the median of the samples C_c (smallest endpoint of each line) to the hypothetical original color O_c (biggest endpoint of each line).

These sets of colors describe the palette of colors used by the artist in the mural paintings and that we aim to reconstitute. For each considered color c , we have, on one hand, a color sample that is used as the hypothesis of its original appearance O_c and, on the other hand, a set of (often degraded) samples that describe its current state of preservation $S_c = \{s_{c,1}, s_{c,2}, \dots, s_{c,N}\}$. It can be easily noticed that we can consider the hypothetical color sample O_c as the target position of each handle H_c after the transformation. However, how to relate the set of samples of its current preservation state S_c to the resting position of each handle is more challenging. Transforming all the samples of the current state of a color S_c to the corresponding hypothetical color O_c would result in a restitution that would lose much of the texture of the paintings, since all those color variations would be transforming to the same O_c color. We thus propose to represent the rest position of each considered color using a unique sample C_c that is computed from the provided samples S_c . In particular, we use the median ($C_c = \text{median}(S_c)$) of the S_c samples, due to its robustness to outliers. Our color transformation handles can be summarized as:

$$H_c : C_c \xrightarrow{\text{transforms to}} O_c \quad (2)$$

Fig. 5 shows in 3D CIELAB space, the samples of the colors S_c and the corresponding transformation handles H_c from C_c to O_c . In our experiments, we have seen that one sample is enough to correctly handle each considered color. It is important to notice that, thanks to the smoothness, locality, and precision properties of the deformation technique, a single sample will also transform similar colors in a smooth and precise way. However, in the case the user needs to have higher control of the transformation of similar colors, additional samples or regions can be used as extra handles.

6.3. Color space division

Our method does not only modify the colors of the samples provided by the user S_c , but the whole color space. To be able to deal with the transformation, we propose to divide the CIELAB color space in a 3D grid of N^3 cells that will serve as the proxy geometry to deform (see Fig. 6). Notice that only a subset of the CIELAB space can be represented as sRGB values (in Fig. 6 we show the closest sRGB color for those cells outside the sRGB




Fig. 6. Division of the CIELAB color space into a uniform 3D grid (here we show the cells slightly downscaled for clarity).

gamut). For each cell, we create a vertex in its center. We use it as a pivot point to create 12 tetrahedra (2 for each face of the cell) that cover all the space. This configuration guarantees the continuity of the edges between cells which is a constraint of the deformation technique. Once this structure is created, we slightly modify it to guarantee that exactly one vertex is lying in the L^*a^*b position of each considered color c . This is performed by, for each color c , moving the nearest grid vertex to the corresponding color position. This strategy of adapting the tetrahedralization to the control points is a common practice in cage-based deformation techniques. In our experiments, we have seen that $N = 30$ is enough to ensure good quality in the transformations. This number can be increased when the specified colors are closer to the resolution of the color space division.

6.4. Image transformation

Once the tetrahedral structure that divides the color space has been created and its deformation computed, considering the sets of colors introduced by the user, we can transform any image straightforwardly. Algorithm 1 depicts this process:

Algorithm 1 Transform an image:

```

for each  $p_{rgb} \in \mathcal{I}$  do
   $p_{lab} \leftarrow \text{rgb2lab}(p_{rgb})$ 
   $i_t \leftarrow \text{look4Tetra}(p_{lab})$ 
   $bc_{p,t} \leftarrow \text{barycentricCoords}(T[i_t], p_{lab})$ 
   $p'_{lab} \leftarrow \text{interpolate}(T'[i_t], bc_{p,t})$ 
   $p'_{rgb} \leftarrow \text{lab2rgb}(p'_{lab})$ 
end for

```

For each pixel p_{rgb} of the image \mathcal{I} , we first need to transform it to CIELAB space p_{lab} (observer 2° , illuminant D65). Then, we look for the tetrahedron i_t that contains the point p_{lab} . This search is fast as the grid division serves as an acceleration structure. Then, we compute the barycentric coordinates $bc_{p,t}$ of the point p_{lab} with respect to the tetrahedron $T[i_t]$. After that, obtaining the new L^*a^*b values of the pixel p_{lab} is as easy as interpolating these barycentric coordinates $bc_{p,t}$ for the tetrahedron in the deformed grid $T'[i_t]$. Finally, we only need to convert the color back to RGB space p'_{rgb} .

6.5. Additional constraints

The chosen deformation technique can generate color transformations that do not preserve the boundaries of the grid. This can cause more colors to go beyond the color space gamut. We have observed that, due to the size of the color space, these problems can be often solved by avoiding overpassing the gamut of the lightness component. For this reason, we have defined several constraint points on these limit planes ($L = 0$ and $L = 100$) that are kept steady by default during the transformation. This simple approach forces the transformation to keep most colors inside the gamut, avoiding undesired clamps on the boundaries. Furthermore, besides removing these default constraints, we allow the user to define new constraints, which can be points or regions, and hence have more control over the transformation of the color space. These user-defined constraints show to be useful to avoid undesired transformations of colors not present in the user-provided color set.

6.6. Fast preview mode

Since we are dealing with the transformation of the whole color space, the chosen deformation method is considerably fast. However, it cannot provide interactive results, which can hinder the quick testing of different hypotheses, particularly in the early stages. To address this limitation, we have developed a fast preview method that allows the user to obtain interactive feedback while roughly defining the hypothesis (original colors). In a second stage, the user can then refine the selection using the precise method explained before.

This fast mode is generated by removing complexity to the deformation. We convert the 3D transformation problem to 2D by removing the lightness component from the model. This reduction is justified by the fact that, in the initial step, the choice of the representative colors of artwork is dominated by chromaticity rather than by lightness. In this way, the deformation problem is simplified to only deform two components, which implies a conversion of the division into tetrahedra of the color space, to a division of the chroma plane into triangles. This results in an important speed-up in the computations and a reduction of the memory consumption of the needed data structures. Fig. 7 shows the resulting simplification of a color transformation when this mode is active.

7. Results and discussion

We have developed an app to test our results. The app has been developed with an intuitive GUI to facilitate its use by art historians. The main window shows an image in the central layer that can be zoomed and panned. Aside, it shows the defined set of colors (and their samples) that will guide the transformation. The app allows defining, saving, and loading this set by selecting regions of pixels in images or directly defining the components through a menu (see Fig. 8). Once defined, the app lets the user compute the transformation in the chosen mode (3D or 2D) and save it for future tests. Then, the user can apply the transformation to the image in the central layer. Our app allows the user to define or load a mask in case the user is interested in only modifying some parts of the image (see Fig. 9). At this point, the user can choose to visualize in the central layer the original image or its transformed version and save the transformed image. Finally, when the user is satisfied with the results, (s)he can launch a process to apply (and save) the transformation to a collection of images.

We have tested our technique in an ASUS ROG laptop (i7-10750H CPU, 32 GB RAM, NVIDIA GeForce RTX 2070 with Max-Q




Fig. 7. Color transformation is simplified after being converted to 2D.

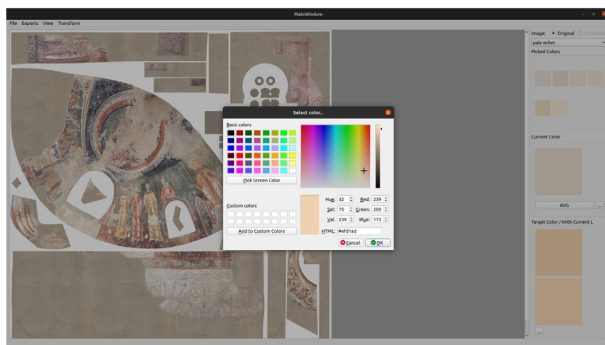


Fig. 8. Snapshot of the developed app with its intuitive GUI.

Design 8 GB GPU). Using a color space grid division of 30^3 cells and 16 handles, computing the transformation takes about 222 s (2.1 s for the 2D mode). Each time the user modifies the value of a hypothetical original color O_c , the transformed grid is updated in 1 ms regardless of the mode. Transforming a pixel p_{rgb} to its corrected version p'_{rgb} takes 1.44 ms (0.7 ms for the 2D mode). Thus, in an unoptimized and non-parallelized code, processing an image with a resolution of 2048×1024 takes 3.0 s (1.5 s for the 2D mode). These timings allow the experts to fine tune color hypotheses (original palette) interactively for the 2D mode and in a reasonably fast way for the 3D mode. This fine tuning is sometimes necessary, as the original palette is not exact and might require testing different hypotheses, as described in Section 5. As transformations can be saved, processing groups of images is fast and easy, even in batches.

Fig. 10 shows the result of applying our technique to a set of photos that show mural paintings of *Pedret* and a ColorChecker. The colors of the mural paintings are correctly transformed into

a hypothetical unweathered version of them. It is noticeable how colors recover saturation and brightness after the transformation. Moreover, their hue is homogenized although they correspond to paintings of different spaces, preserved in different museums (left and middle-left, MDCS; middle-right and right, MNAC) and preserved using different techniques. Notice how some patches of the ColorChecker suffer important changes in hue. These changes are not at issue, as they correspond to colors that are not part of, or even close to, the defined color set. In case they would, the transformation will automatically avoid these problems by including them in the transformation, as the other patches show. Fig. 13 shows a visualization of the original and deformed color grids, for different levels of lightness.

Our technique is also suitable to recover colors in Cultural Heritage photogrammetric textures. Fig. 11 shows the result of applying our technique to the mural paintings once texture mapped onto the architectural 3D model of *Pedret's* church. We have used the color transformation defined by the art historian experts of our team. It is very noticeable how the contrast and vibrancy of the colors are increased. This is compatible with unweathering the paintings, as they tend to lose saturation over time. Fig. 12 compares how this transformation affects the paintings of the central nave (left) and south apse (right). We can see how our technique works well in both situations (note improvements in reds, ochers, and dark grays) and, moreover, helps to recover back colors to closer values. This homogenization process of colors helps experts and visitors in a better understanding of the artwork as a whole and the different history of its fragments.

Limitations: an implicit assumption of our technique is that degraded colors of the painting still preserve some of its original chromaticity. If different pigments degrade into similar unsaturated colors, the required deformation handles are likely to move in opposite directions for the same color region, and thus our technique will not reconstitute them appropriately. Note that these unsaturated colors would also pose a challenge for alternative (e.g. selection-based) color restitution techniques.

8. Conclusions and future work

The weathering of mural paintings is a recurrent problem in Cultural Heritage. Although it is irreversible, experts can formulate hypotheses to try to figure out what was their original appearance. Existing techniques for color correction do not fulfill Cultural Heritage requirements. In this work, we have presented a technique for color restitution of mural paintings designed considering these requirements. Our proposal, based on an existing geometry deformation technique, automatically transforms the whole CIELAB color space taking into account a set of weathered and unweathered samples for each color of the artwork palette provided by an art historian using an intuitive GUI. The results, which are suitable for images and photogrammetric textures of 3D models, recover the vibrance and saturation of the colors, features that are normally attenuated during the weathering process of the paintings, even when the paintings have experienced different exhibition and curation conditions. We believe that our technique will also help Cultural Heritage and Computer Graphics communities to improve the understanding of weathering processes in mural paintings.

As future work, we plan to create an adaptive division of the CIELAB color space which will reduce memory consumption and improve the detail of the transformation in the regions of interest of the color space. We would like to include automatic tests to detect important crossings between the colors which



Fig. 9. The result of applying our transformation process to the texture atlas of the south apse of *Pedret*: the original image (left), the defined mask (center), and the transformed image (right).



Fig. 10. Before and after of applying our technique to photos of the mural paintings of *Pedret*: left, central apse (at MDCS); middle-left, central nave (at MDCS); middle-right, north apse (MNAC); and right, south apse (MNAC).



Fig. 11. Restitution of the mural paintings of *Sant Quirze de Pedret* (left: original, right: restituted) by applying our technique to the textures generated using photogrammetry.

can compromise the quality of the color transformation. Another interesting avenue for future work is to compare the output image by replacing Wang's method with alternative handle-based space deformation schemes, including those amenable to GPU acceleration. Finally, we would like to include MSI information to improve the study of the artwork palette of colors and its samples and to be able to create guided masks automatically.



Fig. 12. Details of the restitution of the mural paintings of *Sant Quirze de Pedret* (top: original, bottom: restituted) in the central nave (left) and south apse (right). Our technique is suitable for restituting colors in paintings that have suffered different weathering and curation processes.




Fig. 13. Illustration of the color transformation generated with the constraints shown in Fig. 5. Images in odd rows show the original CIELAB grid and even rows show the deformed grid. For the sake of clarity, we show only the (downscaled) cells instead of the original tetrahedra. Each cell is painted with the RGB color closest to its CIELAB values. The images show different slices for increasing values of perceptual lightness L.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

No data was used for the research described in the article.

Acknowledgments

This work has been partially supported by projects PID2021-122136OB-C21 and PID2021-122136OB-C22 funded by MCIN/AEI/ 10.13039/501100011033 and by ERDF A way of making Europe, by the EU Horizon 2020, JPIC Conservation, Protection and Use initiative (JPIC-0127) and the Spanish Agencia Estatal de Investigación (grant PCI2020-111979).

References

- [1] Limoncelli Massimo. Virtual restoration. paintings and mosaics. <<L'ERMA>> di BRETSCHNEIDER; 2017, p. 1–195.
- [2] Munoz-Pandiella Imanol, Bosch Carles, Guardia Milagros, Cayuela Begoña, Pogliani Paola, Bordini Giulia, Paschali Maria, Andujar Carlos, Charalambous Panayiotis. Digital twins for medieval monuments: requirements from art historians and technical challenges for analysis and restoration. In: 2022 International Conference on Interactive Media, Smart Systems and Emerging Technologies (IMET). 2022, p. 1–8. <http://dx.doi.org/10.1109/IMET54801.2022.9929510>.
- [3] Pappas Michail, Pitas Ioannis. Digital color restoration of old paintings. IEEE Transactions on image processing 2000;9(2):291–4.
- [4] ColourLex.
- [5] BELTRAN Marta PLAZA. Submerged Villages. Recovering Wall Paintings from the Church of Atance (Guadalajara, Spain). Technical Study, Exhibition, and 3D Display. International Journal of Conservation Science 2016;7(4).
- [6] Cosentino Antonino. Identification of pigments by multispectral imaging; a flowchart method. Heritage Science 2014;2(1):8.
- [7] Dardenay A, Mulliez Maud, Mora P. Reconstruction of ancient wall paintings: digital painting and three-dimensional restoration of the House of Neptune and Amphitrite in Herculaneum. In: Virtual Retrospect 2017, Restituer les couleurs/Reconstruction of Polychromy. 8, Ausonius; 2017, p. 67–77.
- [8] Barni Mauro, Bartolini Franco, Cappellini Vito. Image processing for virtual restoration of artworks. IEEE multimedia 2000;7(2):34–7.
- [9] Arora Nidhi, Kumar Ankush, Kalra Prem. Digital restoration of old paintings. In: International conference in central Europe on computer graphics, visualization and computer vision (WSCG). Citeseer; 2012.
- [10] Nemes Paul, Gordan Mihaela, Vlaicu Aurel. Color restoration in cultural heritage images using support vector machines. In: Scientific Computing and Cultural Heritage: Contributions in Computational Humanities. Springer; 2013, p. 51–9.
- [11] Palomero Cherry May T, Soriano Maricor N. Digital cleaning and “dirt” layer visualization of an oil painting. Optics Express 2011;19(21):21011–7.
- [12] Zeng Yuan, van der Lubbe Jan CA, Loog Marco. Multi-scale convolutional neural network for pixel-wise reconstruction of Van Gogh's drawings. Machine Vision and Applications 2019;30(7–8):1229–41.
- [13] Maali Amiri Morteza, Messinger David W. Virtual cleaning of works of art using deep convolutional neural networks. Heritage Science 2021;9(1):1–19.
- [14] Xu Wei, Mulligan Jane. Performance evaluation of color correction approaches for automatic multi-view image and video stitching. In: 2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. IEEE; 2010, p. 263–70.
- [15] Reinhard Erik, Adhikmin Michael, Gooch Bruce, Shirley Peter. Color transfer between images. IEEE Computer graphics and applications 2001;21(5):34–41.
- [16] Finlayson Graham, Gong Han, Fisher Robert B. Color homography: theory and applications. IEEE transactions on pattern analysis and machine intelligence 2017;41(1):20–33.
- [17] Tai Yu-Wing, Jia Jiaya, Tang Chi-Keung. Local color transfer via probabilistic segmentation by expectation-maximization. In: 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05). 1, IEEE; 2005, p. 747–54.
- [18] Hwang Youngbae, Lee Joon-Young, So Kweon In, Joo Kim Seon. Color transfer using probabilistic moving least squares. In: Proceedings of the IEEE conference on computer vision and pattern recognition. 2014, p. 3342–9.
- [19] Hwang Youngbae, Lee Joon-Young, Kweon In So, Kim Seon Joo. Probabilistic moving least squares with spatial constraints for nonlinear color

- transfer between images. *Computer Vision and Image Understanding* 2019;180:1–12.
- [20] Pulliam Heather. Color. *Studies in Iconography* 2012;33(Special Issue Medieval Art History Today—Critical Terms):3–14.
- [21] Giráldez Pilar, Vendrell Màrius, editors. *Romànic de muntanya: materials, tècniques i colors*. Barcelona: Clavell Cultura; 2010.
- [22] Biggam Carole P, Wolf Kirsten. *A cultural history of color in the medieval age*. Bloomsbury Publishing; 2022.
- [23] Palet Antoni, de Andrés Jaime. The identification of aerinite as a blue pigment in the romanesque frescoes of the pyrenean region. *Studies in Conservation* 1992;37:132–6.
- [24] Mancho Carles. *La peinture murale du haut moyen âge en catalogne (ixe-xe siècle*. Turnhout, Brepols.; 2012.
- [25] Cayuela Begoña. *Sant quirze de pedret, un paisatge pictòric romànic*. Barcelona: Edicions UB; 2022.
- [26] Morer Antonio, Font-Altaba Manuel. *Materials pictòrics medievals. investigació de les pintures murals romàniques a catalunya*. Butlletí del MNAC 1993;1:71–116.
- [27] Rapport general n. 887A. Barcelone. MNAC. *Etude des peintures murales déposées des différents eclises du 12e et 14e siècle*. Dossier n. 6, Pedret. Laboratoire de recherche des monuments historiques-Paris. 1992, Informe intern conservat a l'Arxiu del departament Conservació-Restauració del MNAC, Unitat 114, caixa 037511-CJT.
- [28] Wang Yu, Jacobson Alec, Barbič Jernej, Kavan Ladislav. Linear subspace design for real-time shape deformation. *ACM Transactions on Graphics (TOG)* 2015;34(4):1–11.