Methods and Workflows for Colour Change Modelling and Reconstruction

Yoko Arteaga[1], Cristiana Barandoni[2], Lucia Burgio[3], David Buti[2], Irina-Mihaela Ciortan[1](✉), Brenda Doherty[4], Jon Yngve Hardeberg[1], Catlin Langford[3], Donata Magrini[2], Letizia Monico[4], Catarina Pinto[4], Irina Crina Anca Sandu[5], Faïçal Selka[6], Saptarshi Neil Sinha[7], Panagiotis Siozos[8], Sophia Sotiropoulou[8,9], and Giorgio Trumpy[1]

[1] ColourLab, Norwegian University of Science and Technology, Gjøvik, Norway
{yoko.arteaga,irina-mihaela.ciortan,jon.hardeberg,giorgio.trumpy}@ntnu.no
[2] National Research Council, Institute of Heritage Science (ISPC), Florence, Italy
{cristianabarandoni,david.buti,donata.magrini}@cnr.it
[3] Victoria and Albert Museum, London, UK
l.burgio@vam.ac.uk
[4] National Research Council, Institute of Chemical Sciences and Technologies (SCITEC), Perugia, Italy
{brenda.doherty,letizia.monico}@cnr.it, catarina.pinto@scitec.cnr.it
[5] Munch Museum, Oslo, Norway
irina.sandu@munchmuseet.no
[6] IMKI, Strasbourg, France
f.selka@imki.com
[7] Fraunhofer (IGD) Institute for Computer Graphics Research, Darmstadt, Hesse, Germany
saptarshi.neil.sinha@igd.fraunhofer.de
[8] Institute of Electronic Structure and Laser, Foundation for Research and Technology-Hellas (FORTH), Heraklion, Crete, Greece
{psiozos,sophiaso}@iesl.forth.gr
[9] School of Applied Arts and Sustainable Design, Hellenic Open University, Patras, Greece
sophiaso@eap.gr

Abstract. This chapter presents end-to-end methodologies and workflows for modelling, reconstructing, and communicating colour change in cultural heritage across four scenarios: classical sculpture, paintings, historical textiles, and early colour photographs/films. We couple experimental mock-ups, created with chemically appropriate pigments/dyes, with computational imaging and appearance modelling to realistically capture the evolution of the aging process. For sculpture, we propose an atlas-based segmentation and semantic shading pipeline that transfers measured appearance from mock-ups to 3D models, enabling virtual reconstructions of lost polychromy. For paintings, we introduce a dual pathway to predict and visualize light-induced discolouration: a photochemical model linked to radiant exposure and a data-driven spectral unmixing approach; both are demonstrated on Edvard Munch's *The Scream (ca. 1910)*, and further supported by corrections of archival

I. M. Ciortan et al. (eds.), Chromatic Visions, LNCS 16450, pp. 1–56, 2026.
https://doi.org/10.1007/978-3-032-17700-1_1

colour film records. For textiles, we implement appearance-aware 3D restoration and fading simulation, integrating spectral unmixing with multi-temporal photographic documentation to visualize past states and forecast future fading. Finally, for early colour photographs, we present high-magnification multispectral imaging for the recovery of colour in faded autochromes, neural methods for repairing the bleeding of the green dye in some autochromes, and lastly, a refined digital decoding strategy for Kodacolor lenticular film. Together, these methods bridge analytical evidence, appearance capture, and perceptual visualization, providing reproducible, decision-support tools for conservation, display, and public engagement.

Keywords: Photodegradation · Spectral unmixing · Appearance modelling

1 Introduction

Collections of cultural heritage assets that incorporate colour are a priority for museums. Due to their high fragility, the complexity of their study and the significance of their proper communication to future generations, special attention and shared methods are required for curating, preserving, and exhibiting them.

The PERCEIVE project (Perceptive Enhanced Realities of Colored Collections through AI and Virtual Experiences), focuses on coloured collections and aims to empower researchers and museum professionals to achieve their objectives regarding the preservation, authenticity and perception of colour. Colour is one of the most fragile yet meaningful features of cultural objects and heritage collections.

A major objective is to deepen and enrich our knowledge of the original conception of coloured objects and collections. In most cases, depending on the artwork's typology and current preservation state, the colour has changed, faded, or become discoloured compared to the original. Therefore, it is crucial to investigate and identify the original colours, ascertaining not only the chemical composition and provenance of the constituent materials (pigments and binders), but also the original appearance of the coloured surfaces.

Communicating research and interpretation, and the resulting knowledge and understanding, to present and future generations is an equal priority. This knowledge informs debates about authenticity, particularly with regard to reconstructions. The notion of what constitutes the "true" or "faithful" appearance of an object is informed by both material evidence and cultural interpretation. Ultimately, this knowledge enhances perception and interpretation, broadening accessibility to these cultural heritage collections and strengthening public engagement with them, thereby enabling personal and social interaction. However, scientific data is not immediately understandable, interpretable and exploitable for creating pathways that allow us to connect the colour preserved today to its original appearance.

PERCEIVE has explored state-of-the-art technologies to develop digital solutions, tools and services that empower professionals to address their objectives related to the challenges of colour preservation. These challenges include the fragility, change, loss or misrepresentation of colour in artworks and historical objects. The developed tools are presented in Chap. 2 of this book and are based on methods for exploring scientific data to model colour change. The aim is to simulate the effect of deteriorating factors over time in order to either track changes back and reconstruct the original appearance, or predict further change or loss, with the ultimate goal of preventing or slowing down any ongoing degradation process. As well as being used directly, simulations and colour change models allow us to visualise changes and use this visual material to open dialogue and support decision-making among scholars, researchers, and stakeholders.

This chapter is divided into sections, each dedicated to one of the PERCEIVE scenarios, which refer to specific colour change challenges. Colour change can manifest in several ways, such as colour loss, fading and alteration, and the different cases depend on the typology of the object as well as various deteriorating factors. Light has the most significant impact on colouring materials, but this is often exacerbated by the synergistic influence of other weathering parameters or conditions. This dictates a different approach to reconstructing the original appearance in each case. When the colour has barely been preserved on the artwork, identifying the original materials guides the fabrication of mock-ups representing the original painting techniques. The analytical data acquired from the mock-ups or the preserved colours in the case studies were used to develop data-driven models.

2 Polychromy of Classical Sculpture

The persistent myth of a monochromatic aesthetic in classical antiquity has been systematically challenged by modern research.

The study of lost polychromy extends beyond mere aesthetic interest, representing a critical scientific field that focuses on the study of primary and secondary sources, drawing on and combining approaches and methodologies from different disciplines with the common goal of understanding and scientifically re-tracking the effects of material deterioration in order to recover the original visual language of ancient art.

This research combines advanced scientific analysis with rigorous archaeological investigation. Beyond simply identifying pigment residues, this interdisciplinary approach enables a deeper understanding of how colours were applied to statues and architecture. This is essential for comprehending and subsequently communicating meaning, status, and emotion. Ultimately, this research reveals a more complex and dynamic representation of ancient civilizations than was previously thought possible.

The digital reconstruction of lost polychromy on ancient artefacts represents a significant subfield of virtual archaeology and cultural heritage preservation. Early attempts at digitally reconstructing lost polychromy on ancient artefacts were often criticised for lacking authenticity and appearing unrealistic. These models frequently relied on colours that were spread out evenly, resulting in a flat approach that failed to capture the subtle visual appearance properties of real-world materials like gloss and roughness.

Our current approach aims to overcome the limitations of standard reconstruction methods by focusing on achieving more authentic colour representations.

2.1 Material Appearance of Lost Polychromy

Original Colours. As part of the PERCEIVE project, a selection of statues from the National Archaeological Museum of Naples (MANN) were chosen for a case study. The museum's extensive collection provided a significant number of examples, creating a rich iconographic context for comparison with related frescoes and mosaics. This comparative analysis, supported by scientific research and historical documentation, formed the basis for reconstructing the original conception of these artefacts in colour (with polychromy).

Two analytical campaigns were conducted at the MANN on selected statues, following a standardized protocol detailed in previous works [17,18].

2.2 Reconstruction Approach

The workflow for digitally reconstructing lost polychromy on classical sculptures adopted by PERCEIVE combines an experimental and computational approach based on five steps (Fig. 1):

1. Identification of polychromy traces
2. Experimental fabrication of mock-ups
3. Capturing measurements and digital modelling of mock-ups
4. Atlas-based segmentation of a sculpture modelled digitally
5. Application of mock-up models as shaders

Fabrication of Mock-Ups. Although highly fragmented, some residual polychromy remains visible on the statues of the selected iconography, i.e. Venus, with particular interest in the Anadyomene type. The results indicate a wide colour palette used for the hair and drapery. To replicate appearance, a series of mock-ups were created using the identified materials, guided by ancient historical sources like Pliny the Elder (Naturalis Historia), Vitruvius (De Architectura), and Theophrastus (De Lapidibus) who provided essential descriptions of ancient painting techniques and pigment recipes [23,31,34].

The discovered colour palette consisted of a few key pigments, used either individually or in binary and ternary mixtures. Iron-based yellow and red ochres were identified as the primary pigments for the hair. Egyptian blue was used to create blue shades, sometimes mixed with another pigment or applied as a preparatory layer with a white pigment. In the case of the Anadyomene Venus

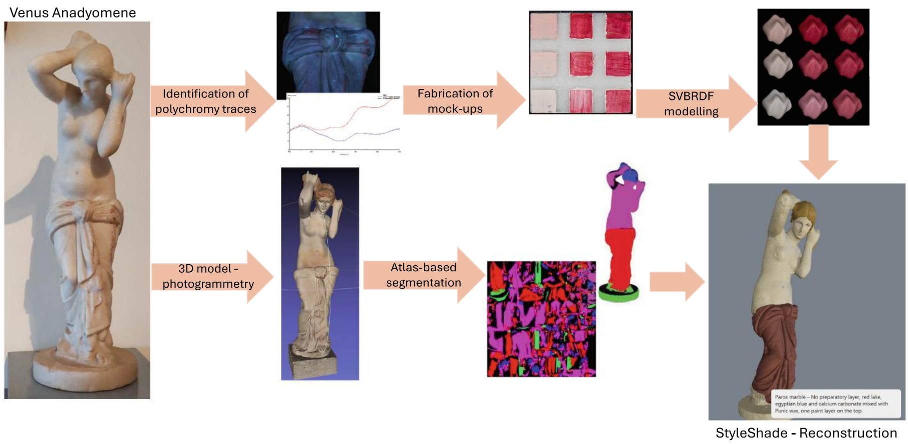

Fig. 1. Reconstruction workflow based on five steps: identification of polychromy traces, fabrication of mock-ups and SVBRDF modelling, 3D model photogrammetry and Atlas-based segmentation. (Color figure online)

from the Iseum of Pompeii (inv. 6298), the research allowed to discover on her the drapery Egyptian blue combined with a madder lake, likely to achieve a pinkish hue. Green earth was also identified on the Lovatelli Venus (inv. 109608).

To recreate the appearance of original polychromy, two replicas of the same set of paint mock-ups were fabricated on Pentelic and Paros marble (Fig. 2). Selecting a substrate similar to that in the original case study aimed to reproduce the texture, translucency, and overall impact of the marble on the final colour. The replicas were designed to simulate the complex interplay of light between the marble surface and the overlying paint layers, as well as the effect of the natural veins in the marble. To this end, a variety of pigment combinations and stratigraphies were created, adhering as closely as possible to the descriptions found in ancient sources.

The paints were applied to 36 marble cubes, each measuring $5 \times 5 \times 3$ cm. The surface of each cube was lightly sanded to create an optimal texture for pigment application. Each surface was then divided into nine distinct squares, each measuring approximately 1.25×1.25 cm. This allowed different combinations of pigments to be applied to separate and limited areas (Fig. 3).

The painting layers were applied using either a ready-to-use Punic wax (from Kremer) or egg white as binding media. The selected pigment was mixed with egg white to create a priming layer.

The pigments chosen for the replicas represent the colours commonly used in the polychromy of classical Roman sculptures. These colours include red ochre, yellow ochre, green earth, Egyptian blue, madder lake, cinnabar, lead white, and calcium carbonate. Lead white was used exclusively for the preparatory layer, while calcium carbonate was also added to various paint mixtures.

Fig. 2. Fabrication process of the mock-ups.

The marble cubes were painted following a structured approach simulating different application methods:

- **Ground Layer Options**: Four different preparatory layers were used to create the replicas: (a) no ground layer; (b) a calcite (calcium carbonate) ground layer; (c) a lead white ground layer and (d) a mixture containing a small amount of highly dispersed Egyptian blue in calcite.
- **Layering:** Pure pigments were dispersed in Punic wax and applied in single, double, or triple superimposed layers on the ground layer options.
- **Lightness variations:** In addition to pure pigments, mixtures of pure pigments with varying ratios of calcium carbonate were also applied.
- **Pigment Combinations:** Various pigment combinations of pigments were also tested, including madder lake mixed with calcium carbonate, green earth mixed with Egyptian blue, and red ochre layered over yellow ochre.

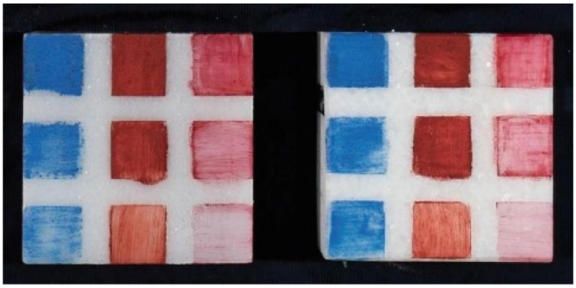

Fig. 3. Example of some mock-ups. The dimensions of the marble cube are 5×5 cm, and each patch has a dimension of approx. 1.25×1.25 cm. (Color figure online)

Material Appearance Modelling. To capture the appearance of the mock-ups' materials, we measure and model the spatially varying bidirectional reflectance distribution function (SVBRDF). The mock-ups were measured using an X-Rite MA-T12 multi-angle spectrophotometer. The MA-T12 instrument measures at six illumination sources and two pick-up points. Surface reflectance is measured in the range of 400–700 nm at 10 nm intervals. The illumination source is a polychromatic white LED with enhancement in the blue region of the spectrum, and its spot size is 9 mm by 12 mm. The MA-T12 also measures six texture images at an angle of incidence of 15° (Fig. 4).

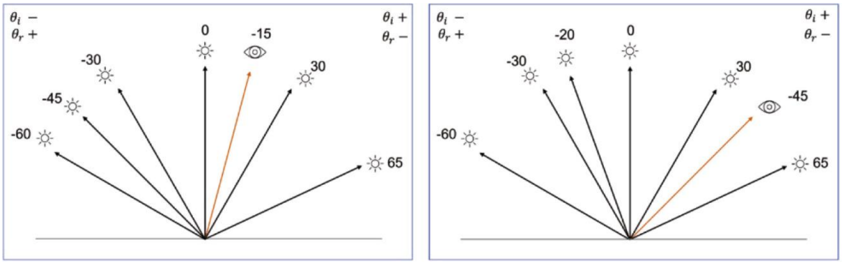

Fig. 4. Schematic of the MA-T12 illumination and acquisition angles. Taken from [1].

The twelve spectra together and the six texture images, are used in the X-Rite Pantora software [36] to obtain the material appearance models. Pantora performs a non-linear fitting to model the SVBRDF to the data and obtains the model parameters in the form of material maps. These are a three-channel diffuse colour map, a normal map, and a specular colour map, and a greyscale roughness map. Pantora also generates a Material Exchange Format (AxF) file [21] which can be used to render the appearance of the material (Fig. 5).

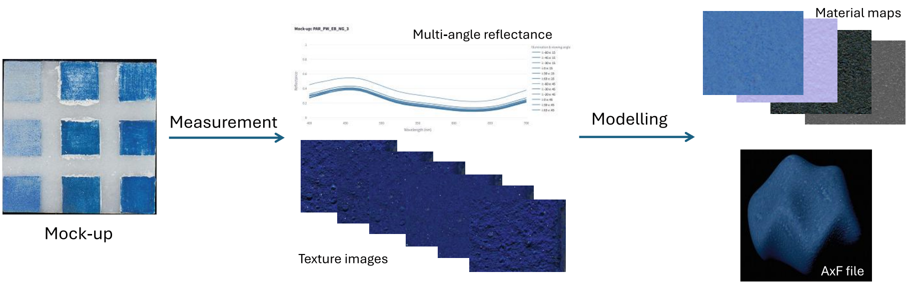

Fig. 5. Pipeline for SVBRDF modelling. The multi-angle reflectance and texture images of each mock-up are measured to obtain the SVBRDF in the form of an AxF file and a set of material maps.

The SVBRDF provides a comprehensive representation of how a material reflects light, accounting for both spatial and angular variations in reflectance. This level of detail is crucial for accurately reproducing the visual properties of real-world materials, particularly those with complex surface features such as gloss, texture, and roughness. By capturing both spectral and spatial information, the MA-T12 enables the creation of digital material models that closely match the appearance of the physical samples under a wide range of lighting conditions.

Once the data is processed in Pantora, the resulting material maps provide a versatile digital representation of the surface. The diffuse colour map encodes the material's base colour, while the normal map captures fine-scale surface geometry that affects how light interacts with the surface.

The specular colour map and roughness map together define the glossiness and specular highlight behaviour, enabling realistic rendering in virtual environments. The AxF file format is a proprietary format and is compatible with a range of rendering engines and design software, facilitating the integration of these material models into digital workflows.

The measurement data, as well as material maps and AxF files, have been uploaded as part of a publicly available Polychromy Appearance database with over 200 mock-up materials. This database has been implemented on Streamlit [29] and it allows the user to search for the mock-up they are interested in, in an interactive way. With drop-down menus, the user can narrow their search by querying different categories such as type of marble, type of binder, pigment, ground layer and number of paint layers. The user can view and download this data to generate their own shaders based on the material maps or directly use the AxF file for their desired application.

2.3 Atlas-Based Segmentation and Semantic Shading

The segmentation (see Fig. 6) and semantic shading of the different segments of the 3D model involves the following steps:

- 3D digitization of the statue using photogrammetry or other AI/non-AI approaches. In this case study, we utilized photogrammetry to digitize the Anadyomene Venus.
- Then, inspired by path planning approaches used in robotic applications, we render the minimum number of multi-view images required to cover the entire 3D surface. The entire implementation was done in Blender.
- We select a few images from the rendered multi-view images, which are then segmented using Segment Anything. The segmentation masks generated by the AI are then refined using user guidance before being tracked to all the segments using an AI-based zero-shot tracker.
- The segmented multi-view images are then projected back to generate a 2D parametrised segmentation map over the 3D surface. We name this process 'segmentation atlas'.

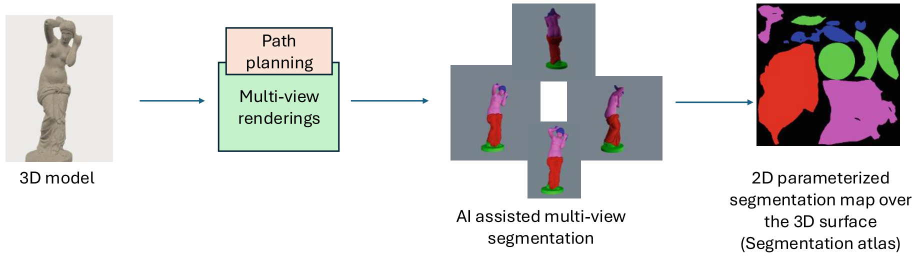

Fig. 6. Pipeline for generation 2D parametrised segmentation map from the 3D model.

This 2D representation of the 3D surface (segmentation atlas) is used to apply the different materials captured by the spectrophotometer as previously explained. We use a modified Cook-Torrance shading model to apply the different material-maps generated by the material capture process. This allows us to apply different materials to the different segments and perform various simulations of the restoration process with each segment. Figure 7 shows a particular example using our approach:

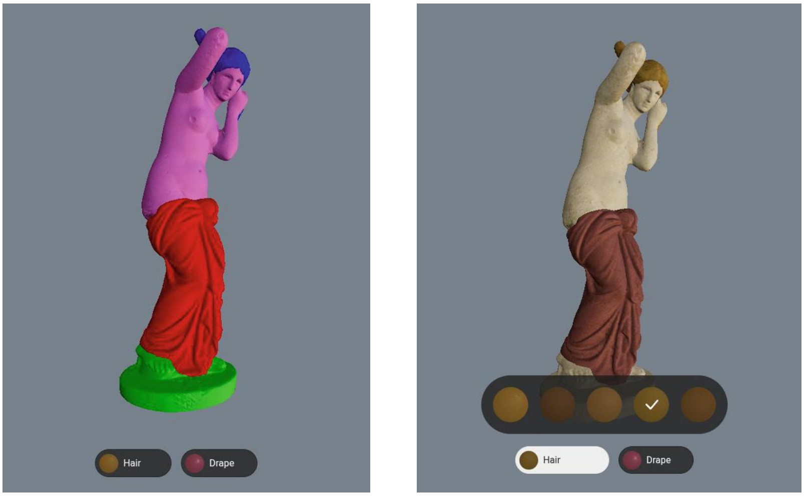

Fig. 7. Segmentation and Reconstruction Hypothesis. On the left the segmentation of the statue using non-representative colours to define distinct segments. On the right one possible reconstruction hypothesis, utilizing material modelled from the paint mock-ups, applied to the segments.

3 Paintings and Works on Paper

This section focuses on methods and workflows developed in the Perceive project based on the case studies provided for Scenario 2: Colour changes in Paintings and Works on paper. Simulation of light-induced colour change had to be applied on pigment mock-ups first, via accelerated aging. Then, one of the approaches used for predicting discoloration in paintings, models changes of the mockups, in the colour domain as linear functions, and then transfer them to a real painting. The other approach models the changes in the spectral domain, using hyperspectral images of the painting as input, where spectral unmixing is used to correlate the mockups with the pigments distribution in the real painting. Beside mock-ups, both methods have been successfully tested on *The Scream (ca. 1910)* by Edward Munch.

Another complementary computational method is presented for studying the temporal evolution of *The Scream (ca. 1910)* and its changing appearance due to the effects of light, moisture, and material ageing. This approach focuses on colour reconstruction using four historical colour photographs from 1971, 1989, 1993, and 2003. To address the challenges of film degradation, the method employs stable reference areas of the painting for advanced colour correction, enabling more faithful digital renderings of its appearance at several points in the 20^{th} century.

3.1 Colour Change Simulation Informed by Artificially Aged Mock-Ups

Photochemical Colour Change Model for Predicting Pigments Discolouration

General Description. The photochemical model presented in this section is one of the two approaches integrated into the Light Damage Estimator (LDE) tool. It offers a mechanistic framework for predicting pigment colour change by linking light absorption to the underlying chemical reactions that drive degradation. In this approach, colour shifts in the CIE L*a*b* space are correlated with the gradual reduction of pigment concentration caused by photochemical processes. The model incorporates both the spectral distribution of the light source and the kinetics of pigment degradation, allowing long-term changes to be expressed as a function of radiant exposure [28]. Unlike the statistical model, which relies on empirical fitting of observed data, the photochemical model offers a process-based interpretation that connects measurable colour alterations directly to the physicochemical mechanisms of light-induced pigment transformation.

Model Workflow. The colour change methodology is organized into two complementary workflows (Fig. 8). The first workflow focuses on the experimental determination of model parameters using artificially aged pigment mock-ups,

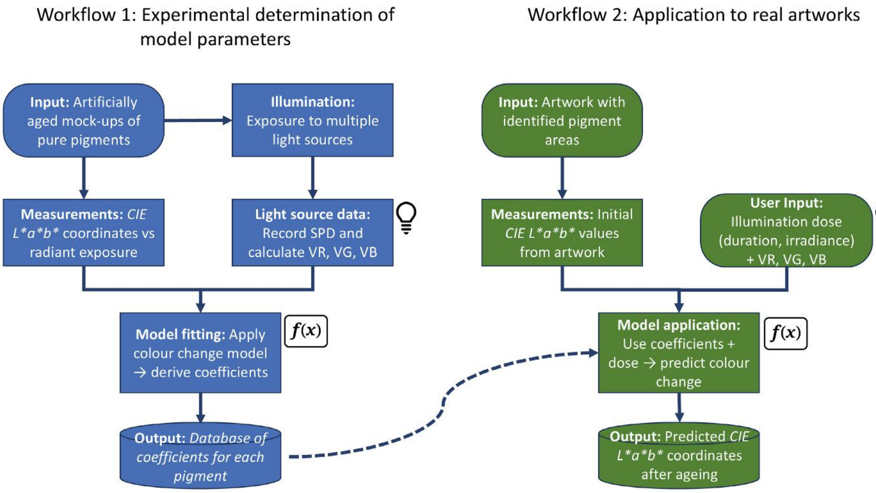

Fig. 8. The two workflows of the photochemical colour change methodology.

while the second applies these parameters to real artworks to estimate colour alterations after prolonged illumination.

Workflow 1: Experimental Determination of Model Parameters. The initial workflow involves the systematic collection of experimental data from artificially aged mock-ups prepared with pure pigments. Each pigment is exposed to controlled artificial ageing using several light sources, each characterized by distinct spectral features. For every light source, the CIE L*a*b* chromatic coordinates are measured at successive ageing intervals and plotted against the corresponding radiant exposure values. In parallel, the spectral power distribution (SPD) of each light source is recorded, and the associated light source classification values (VR, VG, VB) are calculated [28]. These datasets provide the foundation for fitting the colour change model. By applying the model to the CIE L*a*b* coordinates of the artificially aged mock-ups, linear regression yields the coefficients a and b of Equation (2). These coefficients quantitatively describe the colour evolution of each pigment under specific lighting conditions. Once derived, they are stored in a database and subsequently used for predictive simulations on real artworks.

Workflow 2: Application to Real Artworks. The second workflow applies the model directly to cultural heritage objects. The process begins with the identification of specific-coloured regions in the artwork that correspond to pigments previously characterized in the laboratory. From these regions, the initial CIE L*a*b* chromatic coordinates are extracted. To simulate long-term colour changes, the user specifies the illumination dose that the artwork will experience. This dose is calculated based on:

- the total number of years of display,
- the average number of illuminated days per year,
- the daily exposure duration (hours per day), and
- the irradiance of the chosen lighting source.

Additionally, the user provides the light source classification values (VR, VG, VB), ensuring that the illumination is characterized consistently with the experimental conditions. By integrating these parameters with the model coefficients derived in Workflow 1, the methodology predicts the expected chromatic alterations. The outcome is expressed as new CIE L*a*b* coordinates, representing the pigment colour after the simulated ageing period.

Light Source Classification Based on Spectral Power Distribution. We adopt a simplified approach for classifying light sources according to their spectral characteristics. This method is designed to be accessible to users without specialist expertise. The classification scheme categorizes light sources based on the distribution of their emission across spectral ranges corresponding to the three primary colours of light as perceived by the human eye within the visible spectrum.

Table 1. Spectral ranges selected for the calculation of a light source percentage of power distribution.

Classification	Low (nm)	High (nm)
Red spectral range (R)	380	500
Green spectral range (G)	500	600
Blue spectral range (B)	600	780

This approach is consistent with museum lighting practices, where emissions outside the visible range are typically filtered out. As a result, each light source can be described using only three values, which represent the percentage of power emitted in the red, green, and blue spectral regions of the visible range (Table 1). For each light source, the values VR, VG, and VB correspond to the fraction of emission within the respective spectral range relative to the total emission across the 380–780 nm interval:

$$V_{R,G,B} = \frac{\text{Red, Green, Blue spectral range}}{\text{Total emission in visible (380–780 nm)}} \times 100\% \quad (1)$$

The model presumes that colour changes induced by light are determined mainly by the intensity of radiation within the spectral band to which a pigment is most responsive (VR, VG, or VB), rather than by the total irradiance across the visible spectrum. Said otherwise, when a pigment is exposed to light concentrated in its most sensitive wavelength region, the extent of colour alteration will be similar to that observed under full-spectrum illumination, provided the effective portion of irradiance remains equivalent. Based on this, photodegradation arises specifically from the segment of the spectrum that the pigment absorbs.

Photochemical Colour Change Model. The model implemented in the Light Damage Estimator (LDE) tool assumes that colour changes in a material, expressed as shifts in the L*, a*, and b* values (which correspond to variations in the visible diffuse reflectance spectra), are primarily driven by photochemical reactions affecting the pigment. These reactions alter the pigment's chemical composition, leading to a reduction in its original concentration. The model therefore correlates colour change directly with the decrease in pigment concentration. It is important to note, however, that other factors—such as loss of surface gloss or yellowing of the binding medium—may also influence colourimetric measurements. While these effects are not accounted for in the model, they should be considered when interpreting results. When exposed to light, the material absorbs radiation, triggering photochemical reactions that gradually reduce pigment concentration. At low levels of light absorption (zero-order kinetics), the rate of this reduction is proportional to the absorbed energy. As pigment concentration decreases, the fraction of incident radiation interacting with pigment molecules also changes, influencing the rate of degradation.

Assumption 1. Photochemical reactions in pigments are generally slow, producing only minor changes in pigment concentration during typical exposure times. This assumption is valid for gradual degradation processes, where pigment concentration remains nearly constant over the observation period. It also provides a useful metric for assessing the sensitivity of colour or reflectance shifts to small pigment losses. By tracking reflectance or colour values over time, even subtle degradation can be detected and quantified.

Assumption 2. Any colorimetric change C, whether expressed in the CIE Lab* colour space, as ΔE_{1976}, or as diffuse reflectance variations, is assumed to be proportional to the pigment concentration n in the paint. From this assumption, the following expression is obtained:

$$C = -\frac{1}{A}\ln\left(\frac{\varepsilon n_0}{\Phi}\right) + \frac{1}{A}\ln(E_0) = a + b\ln(E_0) \tag{2}$$

In this formulation, the quantum yield Φ represents the efficiency of the photochemical process, defined as the number of molecules degraded per photon absorbed. Although Φ could, in principle, be experimentally determined through chemical or spectroscopic analysis, in the present model it functions as a proportionality factor embedded within the slope of the linear relationship between colour change and the logarithm of radiant exposure (see Eq. 2). Consequently, both Φ and the molar absorptivity ε are not measured directly but are incorporated implicitly through empirical validation, thereby supporting the model's applicability to light-induced colour change data derived from mock-up experiments. As currently formulated, the model is optimized for single pigments or simple binary mixtures in which one pigment exerts dominant control over the observed colour. Extending the model to more complex systems, such as multi-pigment mixtures, would require accounting for the distinct spectral sensitivities, degradation kinetics, and potential interactions among different pigments. This includes modeling how each pigment contributes to overall reflectance and how their individual photodegradation pathways evolve—either concurrently or competitively—over time. The model also assumes a constant oil-to-pigment ratio and attributes colour changes primarily to pigment degradation. In practice, however, variations in binder-to-pigment ratios can significantly affect optical properties such as surface gloss, scattering, absorption, haze, and translucency. Under these conditions, degradation of the binder may have an equally or even more pronounced impact on colour change than pigment alteration alone.

Calculating the Model's Coefficients for a Pigment. Light-induced colour change in the paint is determined not by the total irradiance over the entire emission spectrum, but by the irradiance in a specific spectral region that corresponds to the part of the spectrum that the pigment absorbs, i.e. the spectral region that is most sensitive. In this context, it is not necessary to consider the total irradiance that illuminates the material. Instead, only the fraction corresponding to the spectral component that affects the pigment should be considered. To determine the spectral component affecting the pigment, we plotted the values of ΔL^*, Δa^*, and Δb^* as a function of each of the spectral components R, G, B of the radiant exposure for each light source. The three spectral components of radiant exposure were calculated using the following equations:

$$H_{B,G,R} = \frac{V_{B,G,R}}{100} H_{\text{Total}} \tag{3}$$

where H_{total} is the total radiant exposure, V_B, V_G and V_R the percentage of emission in the specific spectral range derived from Eq. 1, and H_B, H_G, H_R the blue, green and red radiant exposures, respectively. To quantify these observations, the colour-change model described in this chapter was applied to establish predictive relationships for colour evolution. Following Eq. 2, the natural logarithm of both the total radiant exposure and the spectral component-specific radiant exposures was calculated. The resulting ΔL^*, Δa^* and Δb^* values obtained under different light sources were plotted against the logarithm of radiant exposure for Prussian blue–lead white paint mock-ups (Fig. 9). In all cases, the colourimetric shifts exhibited a pronounced linear dependence on radiant exposure when represented on a logarithmic scale, thereby supporting the validity of the proposed model. The fitting-derived coefficients a and b for each colour coordinate are subsequently stored and employed to predict colour changes in artworks.

Application of the Colour-Change Model to Artworks. The calculation of colour change in artworks is based on the user-provided input of radiant exposure (light dose) together with the selection of the pigment of interest. These two parameters form the starting point for the prediction process. Using Eq. 2,

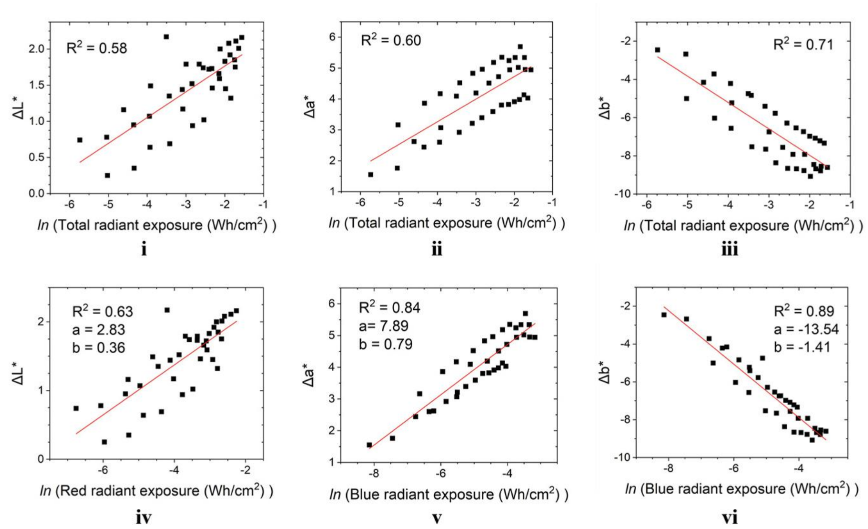

Fig. 9. Plots of ΔL^*, Δa^* and Δb^* values obtained with the colour-change model for Prussian blue–lead white oil paints, shown as a function of total radiant exposure (i–iii) and of the spectral radiant exposure components (iv–vi). The coefficients a and b are derived from the linear fits and subsequently stored for use in colour-change prediction. (Color figure online)

the model determines the expected colour shifts, expressed as ΔL^*, Δa^* and Δb^*, for the chosen pigment under the specified illumination conditions. To translate these shifts into perceptible changes, the calculated ΔL^*, Δa^* and Δb^* values are added to the current CIE Lab* coordinates that describe the pigment's colour within the artwork. This operation results in a new set of colour coordinates that represent the predicted appearance of the pigment after exposure to the given light dose. In practice, this approach allows users to simulate the progressive alteration of colour in a specific pigment region of an artwork, thereby providing a quantitative basis for visualising and anticipating the impact of light-induced degradation over time.

Data-Driven Pigment Photo-Degradation Model in the Hyperspectral Domain. In the context of pigment degradation modelling, previous knowledge on the material palette is very meaningful as it facilitates the creation of mock-ups, where swatches of fresh pigments are created to resemble the chemical composition of the artwork, and then artificially aged in a weathering chamber where environmental parameters such as light and relative humidity can be controlled. In this colour change simulation approach, we combine artificially aged mock-ups with a hyperspectral image of the artwork under investigation. More precisely, as shown in the diagram in Fig. 10, we perform spectral linear unmixing in the hyperspectral image, using least-squares minimization and assuming known spectral signatures of the pure pigments (also known as endmembers in the spectral imaging field). After unmixing, we obtain the abundance maps for each endmember. Then, we can reconstruct the painting by re-mixing the abundance maps with the artificially aged spectra of those endmembers included in the mock-ups, following the same linear mixing model. Nonetheless, the mixing techniques in paintings are commonly of non-linear nature, as there could be superimposed layers or chemical mixture between various materials. To account for the limitations of linear spectral unmixing and mixing models when it comes to paintings, we operate in the absorbance space, $-log(R)$, where R is the reflectance. This is based on the findings of Valero et al. [33] who proved that linear unmixing in the absorbance space, $-log(R)$, is as efficient as non-linear unmixing in the reflectance space for paintings' decomposition.

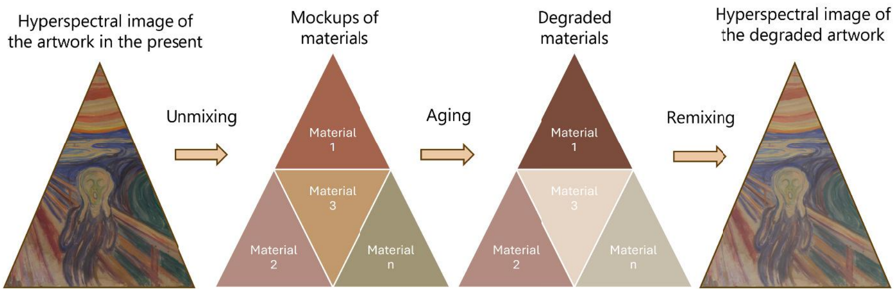

Fig. 10. Our data-driven photodegradation models needs as input a hyperspectral image of the artwork and a set of spectral measurements of the material palette, at various aging levels.

The material palette of *The Scream (ca. 1910)* is known from previous analytical studies and it includes the following pigments: vermilion, cadmium yellow, chrome yellow, viridian, cobalt blue, ultramarine blue, two types of red lake, zinc white, lead white and possibly red lead. To assess light-induced damage and predict future appearance of the artworks, in the context of PERCEIVE several such mock-ups were created by CNR and MOLAB that were later artificially aged. A subset of the palette was chosen for the *The Scream (ca. 1910)* case study according to the transience level of the pigments and so we created mockups of vermilion, chrome yellow, cadmium yellow, zinc white, lead white, and red lead. In addition, the painting was scanned with a hyperspectral imaging system under a MOLAB campaign in 2017.

Thus, to pursue the spectral unmixing task, the spectral signatures of the endmembers in *The Scream (ca. 1910)* have a diverse provenance: some are given by the mock-ups (that only partly cover the painting palette), others are extracted from the spectroscopic measurements on the painting (cardboard substrate, red lakes) and the remaining ones from a privately shared spectral library, with pigments mixed in egg-tempera medium. The reflectance spectra of the endmembers are plotted in Fig. 11. The accelerated aging experiment exposed the mockups to 33.06 $W/cm^2 \cdot hr$ of total light irradiation under a UV-filter Xenon lamp. The spectral measurements were collected at 8 discrete time steps, including t_0. To access unmeasured time steps within the monitored time-span, we interpolate in-between these discrete 8 aging steps. Specifically, an aging polynomial is fitted for each pigment using spline interpolation in the spectral domain, where the dependent variable is the total exposure (in either irradiance or illuminance units) (Fig. 12).

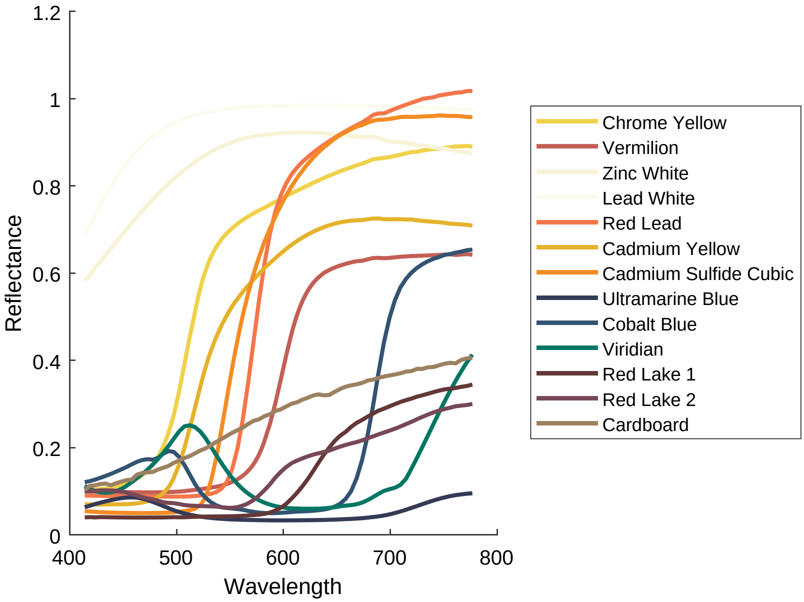

Fig. 11. Endmembers considered for the spectral unmixing of *The Scream*.

The spectral unmixing is carried out in the spectral range from 400 to 850 nm and yields the distribution of pigments over the entire surface of the artwork. The pigment maps (see Fig. 13) show a reasonable distribution and are comparable with the results of analytical techniques. For instance, in the case of the inorganic pigments, such as vermilion the abundance maps are compatible with the XRF maps of the associated chemical elements, mercury (Hg) and chromium (Cr), respectively, and both indicate the presence of the pigments in similar locations, as can be noticed in Fig. 14. Nonetheless, the two abundance maps that represent cadmium yellow in two chemical states (including the cubic form) do not manage to fully capture the cadmium element. This is especially obvious in the depiction of the lake area in the painting, where the yellow ellipse-shaped strokes are erroneously and partially attributed to chrome yellow instead, by the spectral unmixing model. One reason for this could be limited spectral range (visible-near

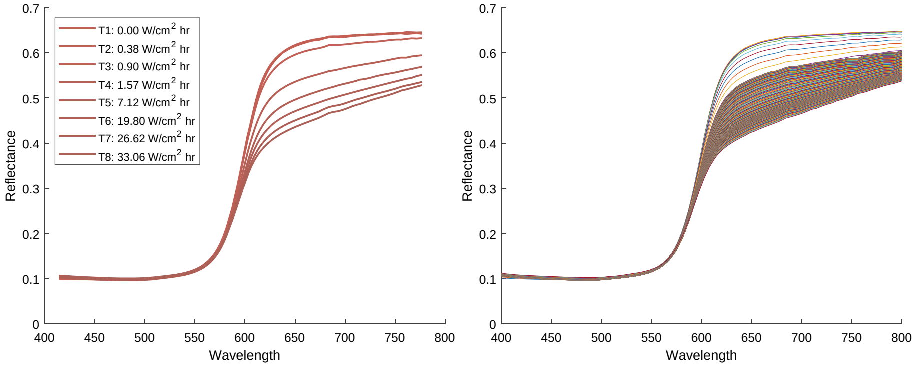

Fig. 12. The aging of vermilion monitored at discrete time steps (left) and the interpolation of intermediate steps, sampled every 0.15 $\mathrm{W/cm^2}(right) \cdot \mathrm{hr}$. As the light exposure increases, the reflectance decreases in the right tail of the spectrum, meaning that the red colour of vermilion becomes darker and less saturated. (Color figure online)

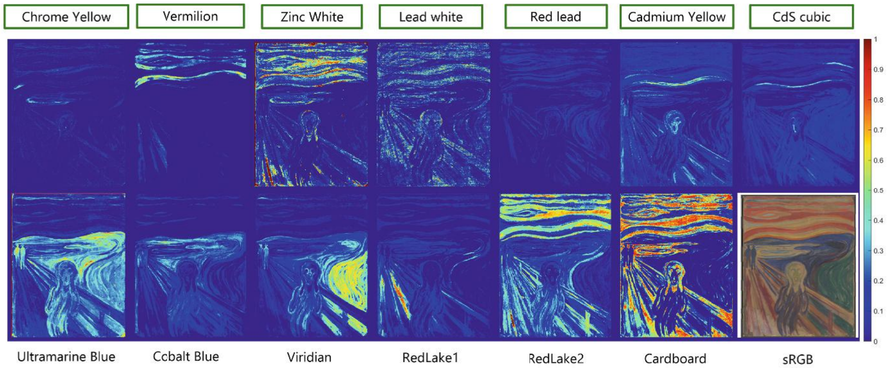

Fig. 13. The pigments maps of *The Scream*, where the red false colour indicates areas with a higher concentration of a material. The photodegradation of the pigments in the top row can be predicted, as aging data was collected for their corresponding mockups. (Color figure online)

infrared) that we operate in, where the yellows have less discriminative features as opposed to the ultraviolet range.

Given the abundance maps and the aging polynomial for some of the pigment, we can now simulate further damage by recombining the pigments maps with the aged spectra. The pigments whose aging behaviour we haven't measured are left unchanged, as in their current state. Figure 15 proposes a degraded version of the painting, equivalent to 34.5 $\mathrm{W/cm^2 \cdot hr}$. In the simulation, the brushstrokes in the sky that correspond to vermilion get overall darker, which is a good result, and consistent with our other simulation methods, presented later in this chapter.

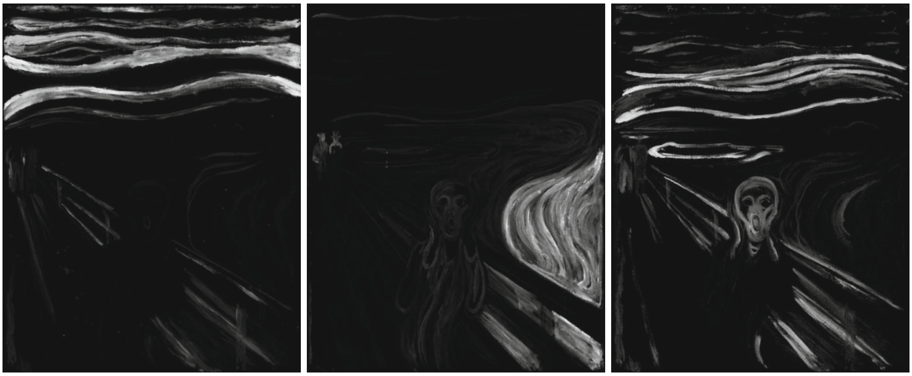

Fig. 14. XRF maps for mercury (Hg), chromium (Cr) and cadmium (Cd) chemical elements, associated with vermilion, viridian and cadmium yellow pigments. (Color figure online)

One limitation of our approach is that not all the pigments in the painting have had aging data collected. Nonetheless, this is a data capture limitation, and not a method limitation, because our method can accommodate additional aging data for all the endmembers.

3.2 Modelling Colour Change Using Film Archival Records as Temporal Reference

Colour Film Modelling. Based on Munch museum's archival records, *The Scream (ca. 1910)* was documented with 4 colour photographs from 1971, 1989, 1993 and 2003, shown in Fig. 16. The challenge is that some of these photographs were themselves subject to degradation, as it is obvious with the one from 1971, where the cyan dye heavily faded. In this approach, we aim to apply a colour correction so that the photographs become more reliable in revealing the past colours of the painting.

In the first step of our approach, we adjust the spectral densities of the cyan, magenta, and yellow dyes so that they match the original concentrations, taking the black level in the film as reference [4]. To this purpose, we need the spectral densities of the photograph, so we captured it with a hyperspectral transmittance imaging setup. The dye adjustment result, displayed in Fig. 17 is a more chromatic balanced film record of the painting, where the magenta colour cast is toned down and the cyan dye is slightly reinjected into the image.

In the second step, we develop a film to painting transformation, where we correct the colours of the film record using the painting as reference. From prior chemical, analytical and accelerated aging investigations, we are aware of which materials and areas in the painting have changed. More precisely, it is known that the vermilion has darkened due to light-exposure, and that cadmium yellow has faded due to humidity. From the XRF analysis, we have the accurate distributions of these pigments in the artwork (see Fig. 14). In addition, the painting

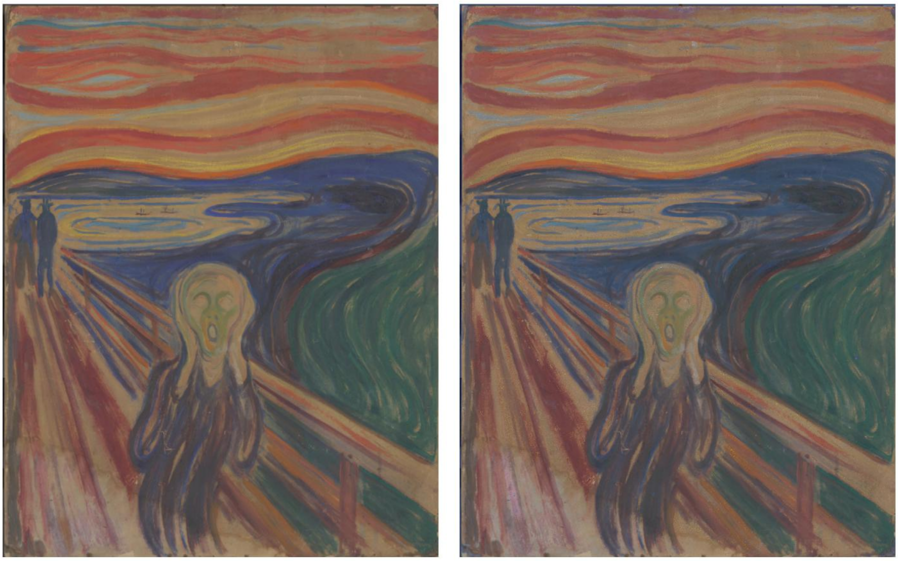

Fig. 15. Current state of The Scream, as scanned in 2017 (left), and the future version, simulated for a total exposure of 33.06 $W/cm^2 \cdot hr$ under a Xenon light source (right). The darkening of vermilion is evident in the brushstrokes in the sky.

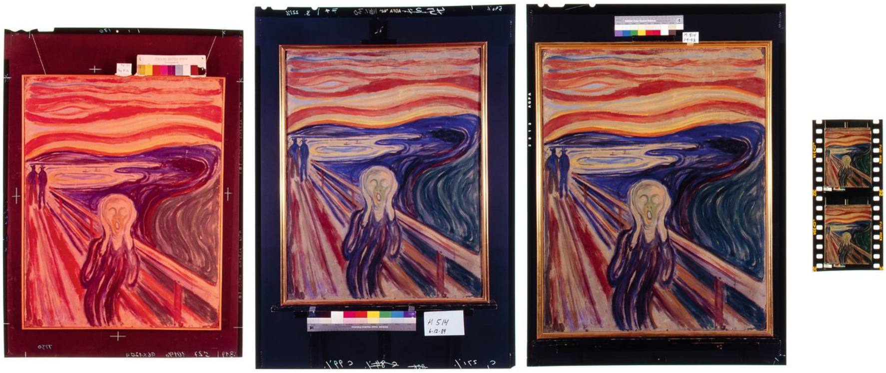

Fig. 16. The analogue photographs of The Scream in 1971 (Kodak Ektachrome), 1989 (Agfa), 1993 (Agfa), and 2003 (Fuji 35 mm).

suffered a water damage in the bottom left corner. We assume that apart from these changes, the rest of the painting is stable and has not changed during the last decades. These stable areas become anchor points to define a transformation from the colours of the slides to the recent colours of the painting last captured in an imaging campaign in 2017, as portrayed in Fig. 18.

Mathematically, we formalize the film-to-painting correction with least-squares regression of polynomial models (first-order, second-order, third-order)

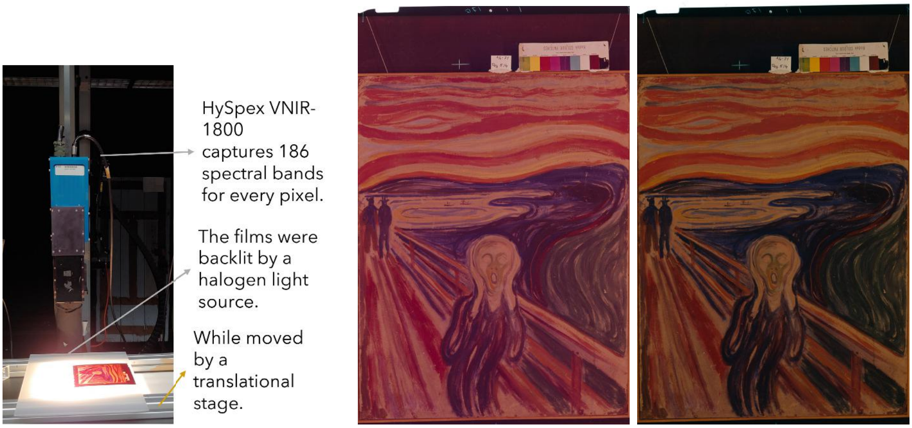

Fig. 17. The hyperspectral transmittance imaging setup (left), the film record from 1971 in its original version (middle) and after pre-processing step with chromogenic dye adjustment (right) that reduces the magenta cast. (Color figure online)

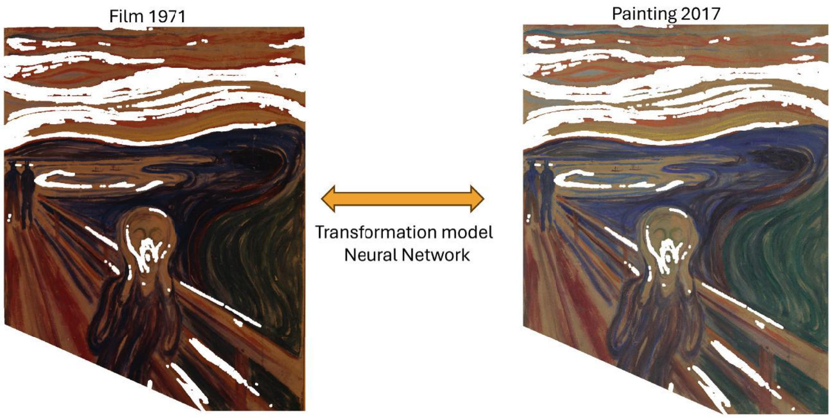

Fig. 18. The stable areas in the painting (not covered by white in the above images) become anchor points for establishing a film-to-paint correspondence.

and neural network fitting with a shallow feed-forward neural network with one hidden layer [7]. We choose the best model according to the mean absolute error given a train-test split of the data, which in our case is the shallow neural network [6]. Finally, we apply the resulting transformation to each film (see Fig. 19), which ultimately enables us to peek into the painting's past with improved colour accuracy.

Generally, one of the challenges of digital restoration approaches resides in the lack of ground truth. A work-around the lack of ground truth is to check the consistency of well-known colour change behaviours of some pigments. We

Fig. 19. The corrected film records (in order: 1971, 1989, 1991, 2003) following the neural network transformation based on the anchor points between the recent colours of the painting (2017) and each film.

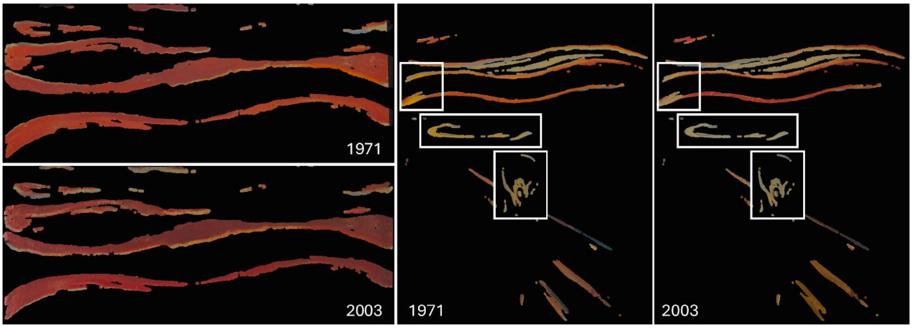

Fig. 20. In the colour corrected versions of the films from 1971 and 2003, we can notice the darkening of the vermilion brushstrokes in the sky, and the fading of the cadmium yellow especially in the lake area, on the face of the character and on the upper left side. (Color figure online)

expect the vermilion to darken and cadmium yellow to fade with the passing of time, and this is perceivable in our results displayed in Fig. 20.

4 Virtual 3D Restoration and Simulation of Fading in light-sensitive Textiles

Heritage textiles are known for containing dyes particularly sensitive to photodegradation. For this reason, their virtual restoration and prediction of further discoloration are of especially high importance towards their documentation and preservation. Digital computation techniques introduce innovative approaches to realistically depict aging processes without intervention on the real objects. In this sub-chapter, we present an approach for modelling and visualizing colour changes in two textiles housed by the Victoria and Albert museum: a 20th-century kimono and 19th-century printed cotton Victorian dress, dyed with natural and synthetic materials, respectively. Our simulations are rendered in the three-dimensional space providing a realistic and interactive environment for scrutinizing the different hypotheses regarding the past and future appearance of the textiles.

Fig. 21. Snapshots of the digitized 3D models of the kimono and Victorian dress (from the V&A collection) Produced by the V&A Photographic Studio.

Our model combines an *atlas-based material segmentation* with advanced colour restoration techniques to transfer the material appearance from textile mock-ups to the 3D models of the garments [26]. Fading effects are simulated through the accelerated aging of these mock-ups that have the same chemical composition as the historical textiles. The aging process is monitored both colorimetrically and spectrally, while the full appearance attributes are captured using a total appearance capture device (TAC7). We then create a colour-based segmentation map over the 3D surface of the textiles to transfer the appearance of the mock-ups on the segmented regions. The fading simulations are further refined by tracing a correspondence in the spectral domain, between the mock-ups and the current appearance of the textiles. Lastly, we also harness available photographic documentation to speculate on the past colour and texture of the fugitive dyes in the textiles.

We embedded our appearance-aware fading simulations in an interactive application, entitled SimTex (**Sim**ulation of colour changes in **Tex**tiles), where the user can toggle between the current state of the two textiles and the three realistic scenarios for depicting the colour changes (see Figs. 22 and 23), as follows:

1. **Mock-up Appearance**: This mode allows for the application of material maps from both aged and unaged mock-ups to the fugitive colors, facilitating appearance-aware construction.

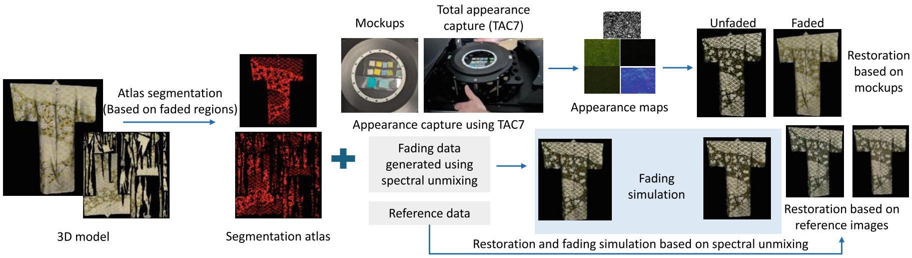

Fig. 22. The segmentation of fugitive colors in the UV-parameterized 3D model is achieved using a 2D parameterized atlas. This segmented atlas enables fading simulations and appearance-aware restoration of fugitive colors, guided by mockups, fading data, and reference images. Image source [26].

2. **Spectral Unmixing**: This mode combines accelerated aging data from the mock-ups with spectral unmixing techniques, revealing the correlation between the fading of the mock-ups and the historical object. It is represented by texture images that illustrate a sequence of fading applied to the segmented fugitive colors.
3. **Photographic documentation**: This mode enables the application of textures from reference documentation images that depict the object in the past, allowing for visualization of potential restoration. The museum's archive contains multiple digital images showcasing details of the kimono. The earliest image is from 2006 (front view), followed by one from 2019 (front view), and a 2024 capture that reveals the inner surface of the inner folds, previously shielded from light exposure. Additionally, there is a 2024 digital image of the inner surface of the Victorian dress.

The interactive apps for the two textiles can be found at the following permanent URLs: https://perceive-tools.igd.fraunhofer.de/kimono/ and https://perceive-tools.igd.fraunhofer.de/dress/.

In the remainder of this sub-chapter, we will elaborate on all the stages in our methodology.

4.1 3D Model Capture and Atlas-Based Segmentation

In this section, we describe our process for digitising the textiles in 3D and conducting automatic segmentation of the fugitive regions.

Photogrammetry Capture for 3D Model Generation: The Victorian dress and the kimono (see Fig. 21) were positioned on a mannequin and a T-bar structure, respectively, within a photographic studio. Continuous halogen lights were employed, remaining fixed in their positions throughout the capture process to ensure consistent lighting. The lighting was kept soft and non-directional

Fig. 23. The figure above illustrates the three visualization modes available in the SimTex interactive visualization tool: Mock-up Appearance (left), where the material appearance from the mock-ups is directly transferred; Spectral Unmixing (middle), which simulates fading based on the current concentrations of the fugitive dyes; and Photographic Documentation (right), which utilizes colour data from various archival photographic records.

to minimize harsh shadows on the objects. A Sony Alpha 7 IV hybrid camera paired with a Sony 50 mm f1.4 lens was utilized for the capture, ensuring over 60% overlap between images in both the x and y directions. The focus was set and maintained consistently throughout the entire capture process. Agisoft Metashape Professional software was used for image processing.

Atlas-Based Segmentation. The 3D modelling of the kimono and Victorian dress using photogrammetry provided us with a texture atlas for both textiles. A texture atlas is a single composite 2D image that packs multiple sub-textures into one image file for efficient rendering. These sub-textures define the UV parameterization of a 3D surface by assigning each mesh vertex (x, y, z) specific UV coordinates (u, v) within the atlas. When the atlas additionally encodes semantic regions, it is referred to as a segmentation atlas.

To create the segmentation atlas, we apply segmentation in the HSV (Hue, Saturation, Value) colour space, with thresholding to identify and isolate the faded regions within the parameterized texture atlas, as exemplified in Fig. 24 for the kimono. For the *i-th* faded region, the thresholds are defined as $H_{min}^{(i)} \leq H_{max}^{(i)}$, $S_{min}^{(i)} \leq S_{max}^{(i)}$, and $V_{min}^{(i)} \leq V_{max}^{(i)}$ for hue, saturation, and value, respectively. The overall condition for segmenting a pixel in the *i-th* faded segment can be expressed as:

$$(H_{\min}^{(i)} \leq H \leq H_{\max}^{(i)}) \wedge (S_{\min}^{(i)} \leq S \leq S_{\max}^{(i)}) \wedge (V_{\min}^{(i)} \leq V \leq V_{\max}^{(i)}) \tag{4}$$

Finally, we developed a custom shader that utilized the UV-parametrized segmentation atlas to achieve the desired appearance for the various segments of the heritage garments.

Improved Segmentation Based of Structural Features. Segmenting fine-grained features is challenging when colours are similar or structures are very small. In our Victorian-dress example, we therefore relied on structural features

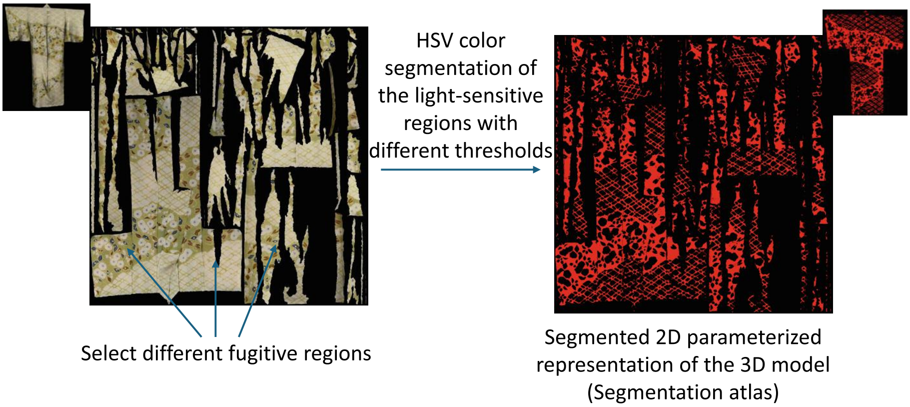

Fig. 24. Segmentation atlas generation through HSV colour segmentation of the faded regions using various thresholds. Image source [26].

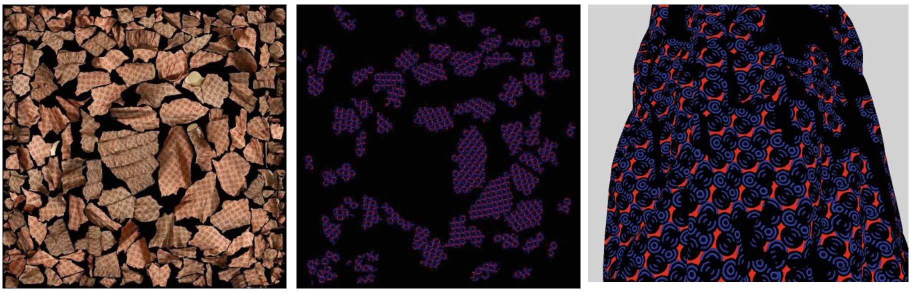

Fig. 25. Segmentation of the Victorian dress using structural features: (left) original atlas, (middle) registered features in the atlas, and (right) resulting partial segmentation when rendered in false-colour for a particular camera view (here blue represents faded purple and red the faded turquoise). (Color figure online)

rather than colour alone. By registering those structural features to a reference atlas, we were able to isolate the faded purple (blue-labeled) and turquoise (red-labeled) regions (see Fig. 25). Although this method yields a partial segmentation of the dress, completing the full segmentation remains future work where we intend to combine colour and structure based segmentation.

4.2 Fabrication of Mockups

The textile mock-ups were prepared as representative models of the historic kimono and dress, designed to address the specific research questions of each case study with regards to their fading and to allow controlled experimentation without risking the original garments. By closely replicating materials, techniques, and visual appearance, the mock-ups provide a platform to study ageing and degradation under defined environmental conditions, particularly light exposure, offering crucial insights for preventive conservation.

The process began with detailed on-site, non-invasive analyses of each textile, including fibre identification, dye composition, and assessment of conservation state. These results guided the selection of historically accurate materials and dyeing techniques. Bibliographic documentation ensured that key variables such as mordanting methods, dye application sequences, and environmental history could be reproduced.

Once completed, the mock-ups were subjected to artificial photoaging under museum-specific lighting conditions. The exposure regimes replicated recorded lux levels, spectra, and durations from actual display periods, making the results directly applicable to the original artefacts. This approach enables prediction of material behaviour, investigation of degradation pathways, and evaluation of preventive conservation strategies in a reproducible, risk-free context.

Case Study: Myriad Green Leaves Kimono

The silk kimono, created in 1992 by the Japanese artist Furusawa Machiko, is part of the Victoria and Albert Museum's textile collection (accession number FE.422-1992). It was crafted using shibori, a traditional Japanese resist-dyeing technique involving precise stitching, gathering, and binding to protect areas of fabric before successive dye applications. The dominant green hue was achieved through a two-step process: initial dyeing with yellow Kariyasu (Miscanthus tinctorius), followed by over-dyeing with indigo (Indigofera tinctoria), as recorded in the artist's notes.

To replicate this textile, silk samples were prepared using artisanal dyeing methods with natural extracts of Kariyasu (yellow), indigo (blue), and a sequential application of both dyes to produce green. Both mordanted and non-mordanted fabrics were prepared to reflect possible historical variations in dyeing practice, as shown in Fig. 26. An undyed white silk control sample was also included for baseline measurements.

For artificial aging, one set of each colour (white, yellow, blue, and green) was stored in darkness as a reference, while the other set underwent 259 h of controlled light exposure, simulating the museum's halogen and LED display conditions at 50 lx. Colour changes were measured periodically using colorimetry and reflectance spectroscopy, enabling the monitoring of fading and the detection of underlying chemical transformations.

Case Study: Victorian Printed Cotton Dress

The second case study focuses on a cotton dress printed by Edmund Potter & Co. in 1887, a leading calico printer of the 19th century (accession number T.2-1926). The design features a turquoise-blue background with motifs inspired by Japanese kamon and Indian chintz. Over time, sections of the fabric exposed to ambient light have faded markedly, while areas protected by folds retain most of their original brightness, as can be noticed in Fig. 27.

Material analysis suggests the use of early synthetic dyes from the triarylmethane family, common in late 19th-century textile production. To reproduce the dress's colour palette, cotton samples were mordanted and hand-dyed (see

Fig. 26. Left: The *Myriad Green Leaves* silk kimono (Furusawa Machiko, 1992), Victoria and Albert Museum (FE.422-1992). The layered dyeing process using Kariyasu and indigo produces the characteristic green hue. © Victoria and Albert Museum. Right: Silk mock-ups replicating the kimono's dye palette and methods, prepared for artificial aging tests. (Color figure online)

Fig. 27) with diamond green B (Basic Green 4, C.I. 42000), diamond green G (Basic Green 1, C.I. 42040), and yellowish light green SF (Acid Green 5, C.I. 42095). An undyed white cotton sample served as a control.

The dyed cotton mock-ups were subjected to 203 h of artificial photoaging under outdoor light exposure. Regular spectroscopic measurements tracked the rate and nature of colour loss, permitting visual fading with chemical changes to be investigated in the synthetic dye molecules.

By combining historically faithful fabrication with controlled artificial aging, these mock-ups provide a reproducible way to evaluate how natural and synthetic dyes respond to photoageing under both indoor museum lighting and outdoor conditions, covering the full relevant spectral range from UV to visible light. The integration of analytical techniques both non-invasive and laboratory-based ensures a detailed understanding of degradation processes.

This methodology not only safeguards the original textiles but also produces data directly applicable to conservation decision-making, including the optimisation of lighting protocols and storage conditions. Ultimately, the study of these mock-ups contributes to the development of evidence-based preventive conservation strategies for both natural-dye and synthetic-dye textiles.

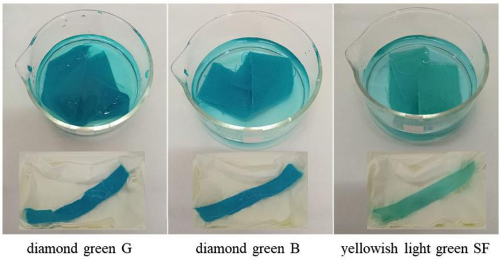

Fig. 27. Left: Victorian printed cotton dress (Edmund Potter & Co., 1887), Victoria and Albert Museum (T.2-1926), showing fading in the background. Right: Dyed cotton mock-ups prepared to match the original dress's colours, for use in artificial aging experiments.

4.3 SVBRDF Capture and BRDF Modelling

To accurately capture the SVBRDF of our mockup samples, we utilize the Total Appearance Capture (TAC7) [14,20,35], developed by X-Rite. This commercially available scanner is specifically designed for capturing material appearances and operates with Pantora software [36], which facilitates postprocessing and optional editing of the resulting SVBRDFs.

The TAC7 features a hemispherical array comprising 29 white LEDs and four monochrome cameras arranged in an arc above a rotating turntable. Five of the LEDs are equipped with colour filter wheels that include 10 filters covering the visible spectrum. To ensure detailed sampling of reflectance lobes, a linear light source can be rotated in the space above the sample holder. For each material analyzed, we gather 100 colour images, 348 monochrome point-lit images, and 280 linear light source images [19]. Additionally, a structured light projector is employed to capture low-resolution surface geometry, and transparent materials are illuminated from below using an LED that passes through a diffuser plate in the sample holder. The TAC7 generally requires about 30 min to process low-gloss anisotropic fabrics [19]. While it can detect fluorescence and alerts users to its presence, it does not incorporate fluorescence into the measurements; however, a mask can be used to identify sections of fluorescent fabric.

The captured appearance of textile mockups (kimono silk and Victorian dress) using the TAC7 device was used to extract the material maps with Pantora [36]. This software fits the appearance data to a spatially varying

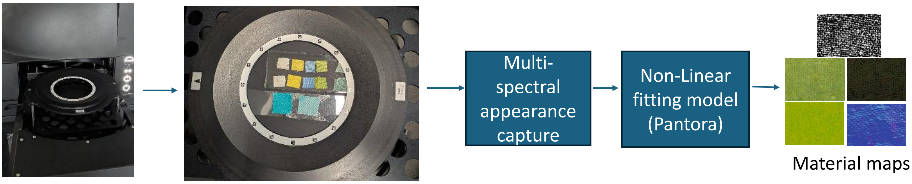

Fig. 28. Appearance capture was conducted using TAC7 [14], and fitting to the appearance model was performed using Pantora [36] on the silk mockup of the kimono and the cotton mockup of the Victorian dress, both created by a conservation scientist.

bidirectional reflectance distribution function (SVBRDF) and generates the corresponding material maps, as illustrated in Fig. 28. To the best of our knowledge, the fitting model provided by Pantora is regarded as state-of-the-art for achieving high-quality material fits. The anisotropic SVBRDF obtained from a TAC7 measurement consists of 13 channels: 3 channels for diffuse colour, 3 channels for normal maps, 3 channels for specular colour, 3 channels for anisotropic roughness, and 1 channel for Fresnel effects (see Fig. 28).

4.4 Fading Simulation

Our approach to simulating colour fading relies on two main assumptions: either we have access to mockups that accurately replicate both the visual appearance and chemical composition of the dyes present in the studied textiles, and that are subjected to controlled accelerated aging, or we possess a series of visual records (photographics documentation) that document the objects' colours over time. The spectral reflectance measurements collected from the mockups at each stage of the accelerated aging process can be integrated with the appearance modeling framework outlined in the previous section, enabling a continuous simulation of colour changes in the textiles.

By measuring how light reflects off the surface of these mockups at different stages of aging, we can build a model that shows how the colours in the original textiles progresses throughout the aging timeline. This allows us to "rewind" the visual appearance of an object to its original state and then "play forward" its gradual fading, almost like watching a time-lapse of its life. This method presumes that the mockups fully and faithfully represent the appearance of the actual objects, which is rarely the case. As we can see in Figs. 29 and 30, the colours resulting from the direct appearance transfer from the mockups, although plausible, are rather saturated and vivid, in stark contrast with the real heritage object that is characterized by a natural patina and softer colours.

To improve accuracy and convey a more natural looking effect, we compare the mock-ups' appearance with the current condition of the textiles using hyperspectral imaging. Thus, for each textile, we capture two details with a hyperspectral camera. The details were selected in such a way so that they contain the fugitive dyes in a good condition. With the spectral unmixing technique,

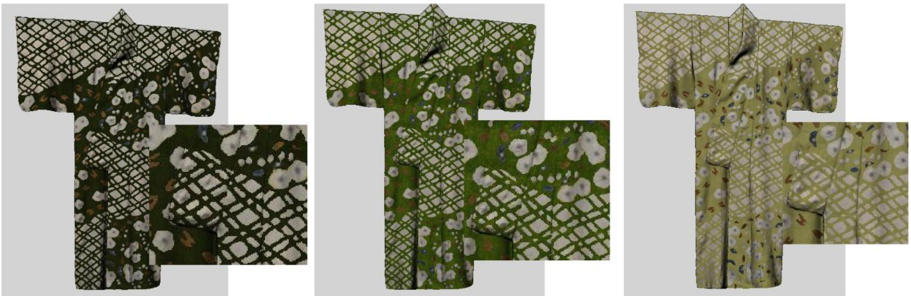

Fig. 29. Past and future simulations of the kimono in Fig. 21, generated with direct appearance transfer from the mock-ups: unaged with reduced exposure (left), unaged as recorded by TAC7 (b), and aged as recorded by TAC7 (c). All renderings utilize a spotlight source with a colour temperature of 6600K.

Fig. 30. Past and future simulations of the turquoise dye in the Victorian dress in Fig. 21, generated with direct appearance transfer from representative mock-ups: unaged as recorded by TAC7 (a), and aged as recorded by TAC7 (b). Additionally, the faded violet regions have been rendered with a representative colour derived from analytical studies. All renderings utilize a spotlight source with a colour temperature of 6600K. (Color figure online)

we then break down the hyperspectral data into the constituent pure materials in the object, known as endmembers. For materials recreated as mock-ups, the spectral signatures of the endmembers are derived directly from the mock-ups; for those not represented by mock-ups, the signatures are extracted from the hyperspectral images themselves. Following spectral unmixing, we generate abundance maps that indicate the spatial distribution of each dye across the textile surface (Fig. 31).

These maps can then be recombined with either aged or unaged spectral data from the mock-ups to simulate the fading process within the hyperspectral image, offering a dynamic and data-driven visualization of colour degradation over time. For the Victorian dress, we first define the turquoise endmember with the spectral signature of the corresponding unaged mock-up (top row in

Fig. 31. Dye abundance maps obtained with spectral unmixing from the hyperspectral images of the details in the Victorian dress.

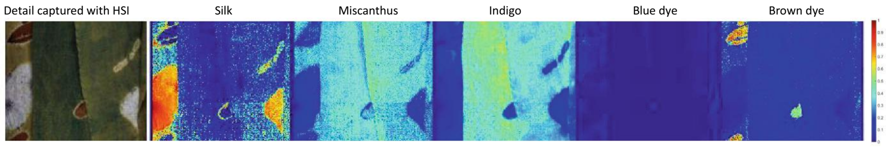

Fig. 32. Dye abundance maps obtained with spectral unmixing from the hyperspectral images of the details in the kimono.

Fig. 31). This is used for fading simulation from present to future. Then, in the bottom row, we extract the signature of the turquoise endmember from the hyperspectral image itself, as this dye is applied purely and unmixed. This is used for attempting to virtually restore the original colour in the dress.

The green in the kimono is a mixture of miscanthus and indigo, so we perform the unmixing with the two spectra of the dyes taken from the unaged mock-ups. Because the detail scanned with HSI contains both green in good condition (in the top-central part) and faded green where indigo has severely depleted (left margin), we attempt the restoration of the past colors with former signal and the present to future discoloration with the latter signal. This simulation correlates any observed colour changes with actual chemical alterations, enabling a meaningful assessment of the progressive loss of colour that is directly linked to the behavior of the original materials. As a matter of fact, the simulations look more natural compared to the mock-up appearance mode, as can be noticed in Figs. 33 and 34.

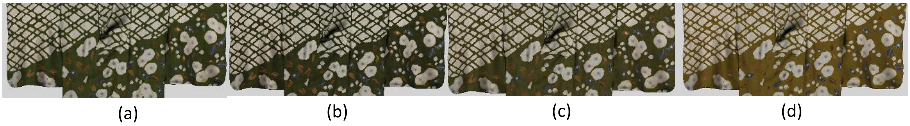

Fig. 33. Fading simulation for the kimono at various time intervals – hypothesized original colour (a) and fading expected in the future (b, c and d) – considering correlation of the mock-ups with the actual heritage clothings through spectral unmixing.

The reconstruction of the kimono using photographic documentation (Fig. 35) reveals significant changes in color and detail from the front view in

Fig. 34. Fading simulation for the turquoise dye in the Victorian dress at various time intervals – hypothesized original colour (left) and fading expected in the future (middle and right) – considering correlation of the mock-ups with the actual heritage clothings through spectral unmixing.

2006 (a) to the front view in 2019 (b), highlighting the effects of time on its appearance. Likewise, the inner surface of the kimono (c), photographed in 2024, exhibits less fading compared to its current appearance captured in the 3D model (see Fig. 21), providing valuable insights into how the colors may have originally appeared. However, the reconstruction based on the reverse side of the kimono appears to be intermediate between the 2006 and 2019 reconstructions. This is expected, as the inner surface, although shielded from light exposure, likely experienced discoloration due to other untracked aging factors such as humidity and temperature affecting the inner folds. Similarly, when comparing the inner surface reconstruction of the Victorian dress (Fig. 35(d)) with its present-day appearance captured in the 3D model (Fig. 21), it is evident that the former is less faded than the latter.

Reconstructing color changes from chronologically archived photographs allows for retrospective modeling of degradation, offering important qualitative insights into the temporal evolution of dye fading and the long-term behavior of historic textile materials, though it does not provide evidence of the environmental factors influencing color alterations. A significant limitation of using documentation pictures taken over various years is the ambiguity introduced by

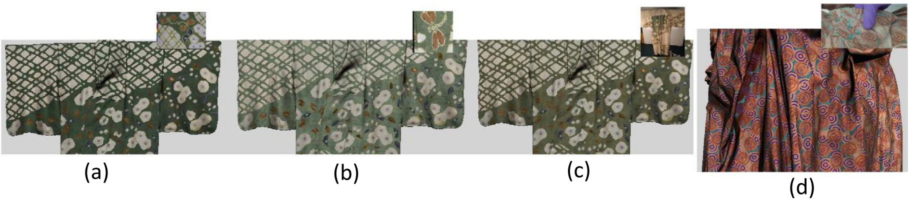

Fig. 35. Reconstruction of the kimono's historical colors using reference documentation images shown in the top corners: (a) from 2006, (b) from 2019 and (c) depicting the inner surface of the kimono that was shielded from light exposure. Also included is the Victorian dress (d) based on a photograph from the inner surface. All images are illuminated with a spotlight set to a colour temperature of 6600K.

subjective choices in photography, such as differing view and illumination setups, as well as potential variations in post-processing and color calibration pipelines.

5 Modelling Colour Change in Historical Photo/Film Collections

The investigation of the deterioration of autochromes and Kodacolor films and the reconstruction of their original appearance are essential to appreciate the full scope and impact of these pioneering colour processes. While these early technologies represent significant milestones in the history of photography and cinematography, their unique aesthetic qualities are often lost due to their inherent fragility, chemical instability and physical degradation.

By digitally reconstructing their original colours and appearance, we make it possible for the general public to experience these works as they were intended—vivid, immersive, and historically rich—thereby fostering greater appreciation and enjoyment of early colour imaging, its cultural significance and its societal impact.

An Introduction to Autochromes

The Autochrome Lumière is considered the first commercially successful colour photography process. Publicly released in 1907, the invention by the Lumière Brothers was publicly praised for its production of luminous colour photographs. Autochromes comprise a glass plate, on top of which sits thousands of starch granules dyed red, green and blue, which act as a colour filter. On top of the colour filter is a layer of panchromatic silver emulsion. In essence, the colour filter produces the impression of colour in the photograph, while the silver gelatin emulsion provides the image, or outline.

Autochromes were used primarily in the period 1907 to 1930. During this time, multiple problems with the process became apparent, particularly in regards to the presentation and preservation of colour. Articles in photographic journals from the period discuss the presence of issues including 'greening', which manifest on the autochrome plate as small spots or large areas in shades of green. Other issues included the presence of orange areas or spots on the plate, silvering and, more commonly, the fading of colours. Autochromes were found to be extremely light sensitive, and popular means of exhibiting glass plates in the period like projection caused the colour dyes to quickly fade. Articles dating from the early twentieth century discuss limiting projection and the light exposure of plates to mere seconds to avoid fading, and the potential melting of photographic emulsion.

Today, as well as in the early twentieth century, it was understood that when the colour in autochromes changed or was damaged, there was little action that could be taken to restore the plates to their original appearance. PERCEIVE has worked on new initiatives that offer means to generate an image which provides a sense of what the autochrome may have originally appeared like.

An Introduction to Kodacolor

Since the emergence of cinema in the 1890s, inventors and researchers sought ways to bring colour photography into motion pictures. However, it took decades before a workable method of recording colour in film was achieved. The lenticular process offered a solution by capturing three colour separations simultaneously within a single film frame. The principles behind this technique originated from earlier work by pioneers such as Liesegang and Lippmann. Liesegang had imagined a pixelated approach using a perforated screen, while Lippmann described in detail the use of lenticular film to achieve more realistic colour reproduction [22]. In the early 1900s, Berthon built upon these ideas and applied them to moving pictures [3]. It took a long time before Kodak managed to refine and commercialise the technology quickly, launching Kodacolor in 1928 as a 16 mm reversal film designed for amateur use. For any new colour process to thrive, it needed widespread adoption, compatibility with existing recording and projection equipment, and the ability to produce distributable copies. Kodacolor struggled in the latter area. Moreover, it suffered from the common drawbacks of additive colour methods, including low image brightness, reduced sharpness, and a limited colour range. Although Kodacolor enjoyed some popularity among amateur filmmakers, its appeal quickly faded when Kodak introduced Kodachrome, and lenticular film disappeared by the late 1930s.

Today, many lenticular films remain hidden or misidentified, as they appear to be ordinary black-and-white reels at first glance. Raising awareness of these films is crucial to ensure their proper identification, preservation, and eventual rediscovery for future generations.

5.1 Modelling Colour Change in Autochromes

We present two distinct methodologies for the colour restoration of autochromes, both of which follow a shared processing pipeline illustrated in Fig. 36, previously published in [5]. This pipeline begins with the separation of the colour mosaic from the silver emulsion and proceeds by replacing the silver-decoupled dyes with unfaded reference dyes to achieve colour restoration. The key difference between the two approaches lies in the strategy employed for silver/dye separation. In this section, we will only focus on one method: spectral unmixing. For more details about the imaging set-up and silver/dye separation, please refer to the original publication.

This method capitalizes on the spectral resolution capabilities of our imaging system by formulating the silver/dye separation as a spectral unmixing problem. Following the unmixing process, we replace the spectral signatures of the degraded dyes in the dye mosaic with those of unfaded reference dyes. The restored colour mosaic is then recombined with the silver emulsion. To the best of our knowledge, this represents the first application of spectral unmixing techniques to the restoration of autochromes.

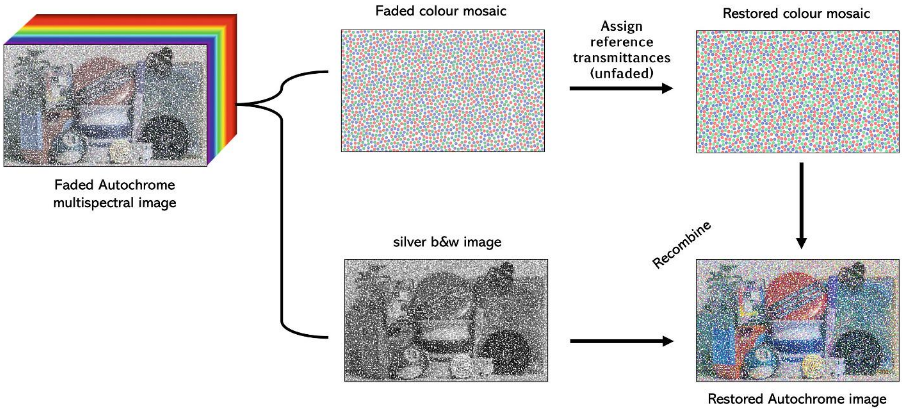

Fig. 36. Our restoration pipeline is defined by the processing of the colour mosaic after its uncoupling from the silver emulsion layer.

Test Autochromes. The significant light sensitivity of autochromes means that many museums choose to limit public access and restrict exhibiting original plates. To avoid degradation on autochromes from public collections, we purchased four autochromes from antiques sites and private collectors for the purpose of the project, each exhibiting varying degrees of fading and typical deterioration such as greening (Fig. 37).

High Resolution Multispectral Image Capture. To capture the fine structural and spectral details of autochromes, we employed a high-magnification multispectral imaging system. Autochromes are composed of microscopic dyed potato starch granules, each approximately 10 µm in diameter. To resolve individual granules, our system achieves a spatial resolution of 1.84 µm on the object plane, enabling precise analysis of the colour mosaic.

The imaging setup is a modified version of the active multispectral system described in [32], consisting of a monochrome 16-bit CCD camera paired with a 5× magnification lens (Fig. 38a). Illumination is provided by an integrating sphere equipped with ten LEDs, each emitting at distinct spectral peaks in the visible spectrum.

To ensure uniform illumination, the light from the integrating sphere is collimated using a condenser lens and diffused through a translucent slab, creating a bright-field setup. The autochrome plate is placed on a custom holder with a black cover (Fig. 38b), exposing only the area being imaged to minimize stray light and reduce unnecessary exposure. The holder is mounted on an XYZ translation stage, allowing for precise positioning and tile-based image acquisition, which can later be stitched into a complete image.

Exposure times are optimized individually for each spectral band to maximize the dynamic range of the 16-bit sensor without saturation. Due to the inherently low transmittance of autochrome glass plates [10], we use a neutral density (ND)

(a) Couple by the lake. Autochrome used by permission of the collector, Arthur Clay.

(b) Gruyere, 1914

(c) Harrogate shop window

(d) Poppy field

(e) Umbrella. Autochrome used by permission of the collector, Arthur Clay.

Fig. 37. Test autochromes acquired for this research. (Images not to scale)

filter with 0.4 optical density as the white reference, rather than an empty gate. Each captured image undergoes flat-field and dark-current correction.

Given the extreme light sensitivity of autochromes, we also assessed the potential impact of our imaging process on material degradation. This monitoring enables us to compare current and future captures, providing a framework to evaluate any fading induced by the digitization process.

Spectral Unmixing. In our spectral unmixing approach to autochrome restoration, we treat the three dyes—red-orange (R), green (G), and blue-violet (B)—as endmembers. Silver and carbon, which primarily modulate image contrast rather than contribute distinct spectral signatures, are captured in the residuals of the unmixing model. This rationale allows us to focus on separating

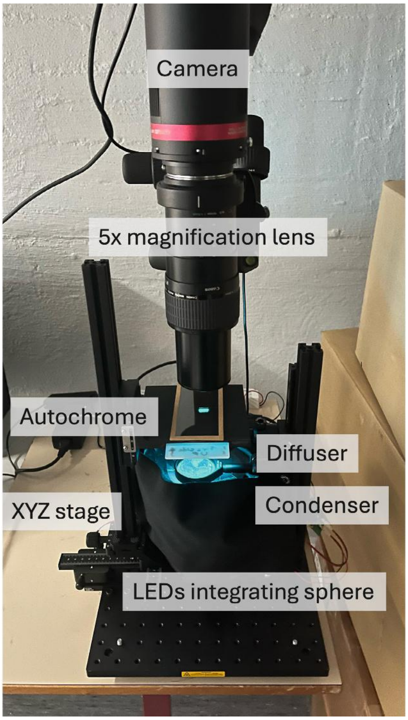

(a) Multispectral Imaging Setup.

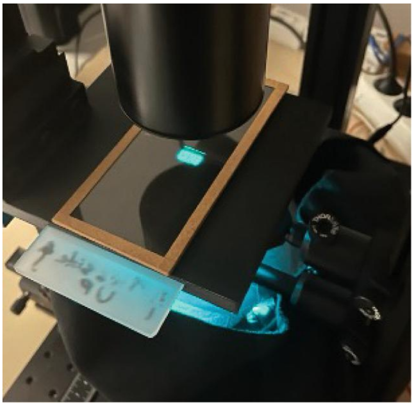

(b) Setup closeup - Autochrome plate on holder, with black cover to minimize stray light and exposure.

Fig. 38. Our autochrome digitization setup is non-invasive and offers flexibility in adjusting the distance between the photograph and the camera.

the dye components from the monochromatic signals introduced by silver and carbon.

For each autochrome, we construct a multispectral image cube using only the bands within the visible spectrum. To identify the spectral signatures of the dyes, we select a region of the image that is free from silver interference. Within this no-silver area, we sample the spectra of individual dye granules and average them to obtain representative spectral profiles for each dye's optical density.

Operating in the additive absorption domain, we assume a linear spectral mixing model. That is, each pixel ij in the multispectral cube is considered a linear combination of the three dye endmembers:

$$OD_{cube} = ab_R^{ij} * OD_R + ab_G^{ij} * OD_G + ab_B^{ij} * OD_B + residual^{ij} \tag{5}$$

with $i = 1...N$ and $j = 1...M$, where N is the number of rows and M is the number of columns in the multispectral image. This assumption simplifies the unmixing process and enables us to estimate the spatial abundance of each dye across the image, forming the basis for accurate colour restoration.

The objective of the spectral unmixing process is to estimate the abundance ab of each dye at every pixel (i, j) in the multispectral image cube. These abundance values represent the spatial distribution of each dye across the image, forming what is known as an abundance map. While abundance is proportional to dye concentration, it does not correspond to absolute quantities and lacks a standardized unit of measurement.

To compute the abundance maps, we solve the linear mixing model using a least-squares minimization approach [13]. The optimization minimizes the squared difference between the observed multispectral cube and its reconstruction based on the linear combination of the dye endmembers:

$$|OD_{cube} - \widehat{OD}_{cube}|^2$$

This minimization is subject to the constraint that all abundance values remain non-negative, ensuring physical plausibility. The resulting abundance maps provide a robust estimate of the spatial presence of each dye, enabling accurate colour restoration while implicitly accounting for the influence of silver and carbon through the residuals of the model.

Once the abundance maps for each dye are computed, we reconstruct the multispectral image cube by multiplying these maps with the corresponding spectral signatures of the dyes:

$$\widehat{OD}_{cube} = ab_R^{ij} * OD_R + ab_G^{ij} * OD_G + ab_B^{ij} * OD_B. \tag{6}$$

This yields a modelled version of the original multispectral cube, representing the contribution of the dye mosaic alone.

The residual error is then calculated as the difference between the original multispectral cube and the reconstructed cube:

$$residual^{ij} = OD_{cube} - \widehat{OD}_{cube} \tag{7}$$

This residual captures the spectral information not explained by the dye components—primarily the influence of the silver emulsion and carbon elements. In this way, the abundance maps effectively isolate the spatial distribution of the dyes, while the residual serves as a segmentation of the silver layer and other non-dye components.

Colour Restoration. To restore the original colours of the autochrome dyes, we rely on the spectral transmittance measurements of unfaded dyes previously obtained by Lavédrine and Gandolfo [16] from a well-preserved historical autochrome. These spectral profiles represent the optical characteristics of the red-orange, green, and blue-violet dyes used to colour the potato starch granules.

Using the spectral unmixing approach, we decompose the multispectral cube of each autochrome into abundance maps for the three dyes, as described by Eq. 5. Faded dyes typically exhibit lower and flatter optical density curves compared to their unfaded counterparts. To reconstruct the unfaded appearance, we multiply the abundance maps by the reference optical densities of the unfaded dyes, resulting in a restored colour mosaic.

To reintroduce spatial detail, we add the residual component which captures the silver and carbon contributions back to the restored mosaic. This yields the final restored multispectral cube, expressed as:

$$OD_{restored\ cube} = ab_R^{ij} * OD_{Ref\ R} + ab_G^{ij} * OD_{Ref\ G} + ab_B^{ij} * OD_{Ref\ B} + residual^{ij} \tag{8}$$

The restored cube is then rendered using the CIE 1931 2° standard observer and the D50 standard illuminant. For final image encoding, we apply the Rec. 2020 colour profile, which follows the ITU-R BT.1886 [25] recommendation and includes gamma encoding with $\gamma = 1/2.4$.

Colour Restoration of Partially Faded Autochromes. The test autochromes were imaged using the high-resolution multispectral setup previously described. Figure 39 presents a selected area from each autochrome where the silver emulsion is assumed to be transparent. The spatial resolution of the imaging system is sufficient to resolve individual dyed starch granules, allowing for a detailed assessment of their current state of conservation. The optical densities of the dyes were computed by averaging pixel values within the selected area shown in Fig. 39. Compared to the reference optical densities of autochromes, the faded dyes exhibit lower maximum densities, indicating reduced light absorption, and higher minimum densities indicating reduced light transmission. This dual effect results in diminished contrast and desaturated colours, characteristic of aged autochromes.

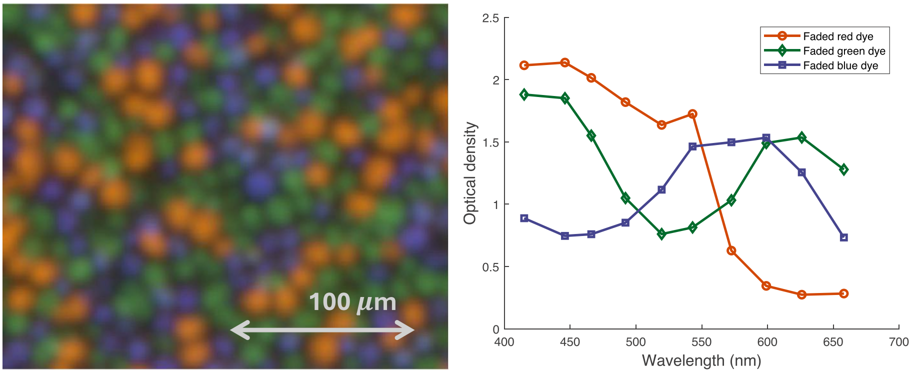

Fig. 39. Left: Dyes uncovered by the silver emulsion, cropped from the colour image of the autochrome Couple by the lake. Right: Spectral signatures of the pure dyes in the "Couple by the lake" autochrome.

For each autochrome in our study, we selected a region devoid of silver emulsion which served as the basis for extracting the average spectral signatures of the red, green, and blue dyes within the visible range Fig. 39. These spectral profiles were used as endmembers in the spectral unmixing method, allowing us to compute the spatial distribution of each dye across the entire image in the form of abundance maps.

Beyond dye separation, we also decomposed the silver and carbon signals. Because carbon is optically mixed with the adjacent dyes, its contribution is embedded within the abundance maps. The silver layer, however, is isolated by calculating the residual between the original multispectral image and the image reconstructed using the unmixing model.

To reconstruct the original multispectral cube, we multiply the abundance maps by their corresponding endmember spectra, accounting only for the dyes and carbon elements. The resulting colour rendering of the faded dye mosaic is shown in Fig. 40b. A faint ghost-like artefact of the woman's silhouette is only slightly visible, likely due to scattering effects not captured by our current model.

For the virtual restoration, we replace the faded dye spectra with the reference spectra of unfaded dyes, combining them with the abundance maps to generate a restored dye mosaic (Fig. 40d). Finally, we add the silver layer, represented by the residual image (Fig. 40e) to the restored dye mosaic, producing the fully restored autochrome shown in Fig. 40c.

Colour Restoration of Autochromes with Completely Faded Blue Dye. The autochromes previously discussed exhibit varying degrees of fading across their dyes. In this subsection, we focus on a particularly degraded autochrome in which the blue-violet dye has completely faded, namely the "Umbrella" autochrome shown in Fig. 37e. Figure 41 shows a magnified view of its dyed potato starch granules, revealing that the blue dye granules are nearly transparent, while only the red-orange and green dyes remain visibly present. Further analysis of the spectral curves confirms this observation: not only is the blue dye absent, but the remaining red and green dyes also show significant fading. Their optical density curves are both flattened and reduced in magnitude, indicating diminished absorption and transmission properties. This results in a severely desaturated colour mosaic with low contrast, characteristic of advanced dye degradation in autochromes.

We applied the spectral unmixing approach, following the same procedure previously outlined. We extracted the spectral signatures of the faded dyes directly from silver-free regions of the image. This allowed us to estimate the spatial distribution of the remaining dye components through abundance maps, even under conditions of extreme fading. Although the spectral contrast between the dyes was minimal, the unmixing method proved sufficiently robust to isolate the residual dye information, as shown in Fig. 42. The resulting abundance maps enabled us to reconstruct the faded dye mosaic and proceed with virtual restoration, despite the near-complete loss of the blue-violet dye and the significant fading of the red-orange and green dyes.

To capture the full extent of the umbrella depicted in the autochrome and to produce a visually compelling result, we created a mosaic composed of 45

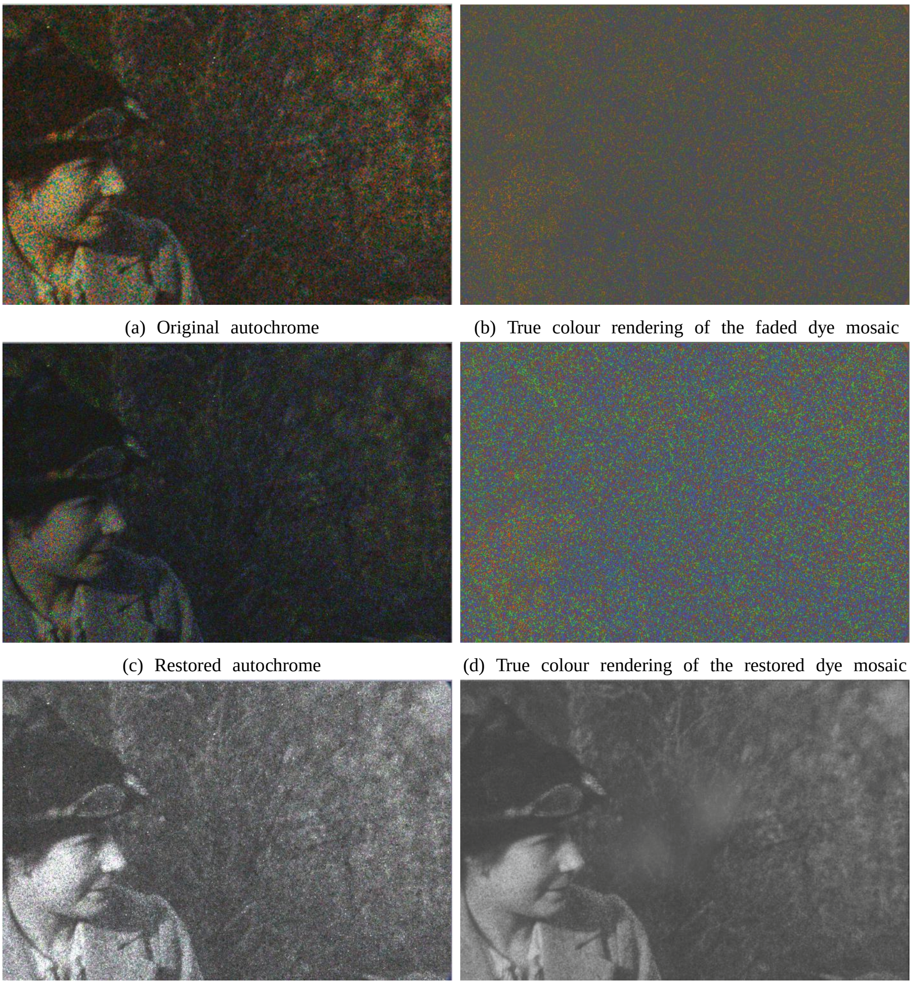

(a) Original autochrome

(b) True colour rendering of the faded dye mosaic

(c) Restored autochrome

(d) True colour rendering of the restored dye mosaic

(e) Residual of the spectral unmixing

(f) Infrared band (940 nm)

Fig. 40. The restoration with the spectral unmixing method (c) pulls up the faded blue colours in the original image, obvious especially in the coat area, and increases the overall contrast. Moreover, the dyes are well decoupled from the silver layer. The restoration is sharper, as the spectral unmixing denoises the image through dimensionality reduction. This is particularly obvious if we compare the silver layer (f) as given by infrared image (that has several artifacts and is affected by blurriness) with the residual of the spectral unmixing (e). (Color figure online)

individual acquisitions. This approach allowed us to preserve fine spatial detail across the entire image while also enhancing its aesthetic impact. The resulting stitched image not only showcases the intricate structure of the dyed starch granules when zoomed in, but also serves as a striking visual demonstration of

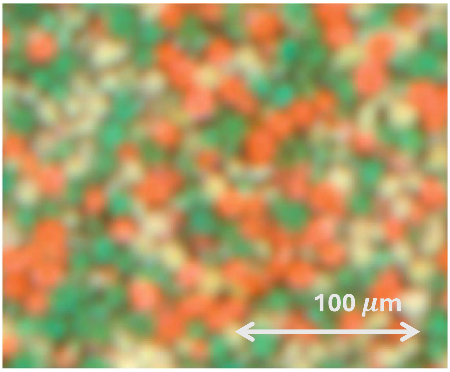

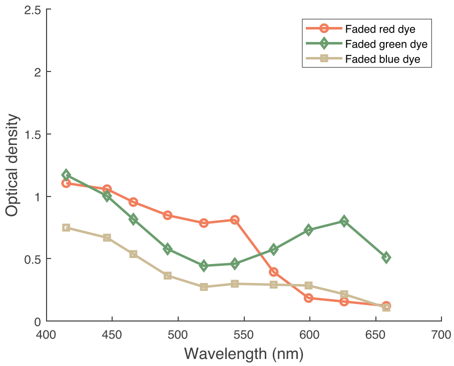

Fig. 41. Left: Dyes uncovered by the silver emulsion, cropped from the colour image of the autochrome "Umbrella". Right: Spectral signatures of the pure dyes in the "Umbrella" autochrome.

the restoration process. By presenting both the original and restored versions, the image effectively illustrates the transformative effect of our method and highlights the contrast between the faded and unfaded colour mosaics (Fig. 43).

5.2 Neural Restoration of Greening Defects in Autochrome Photograph Utilizing Purely Synthetic Data

"Greening" and green spotting (see Fig. 44) are characterized by green areas or spots on the autochrome plate. The historical dyes used to color the starch granules in autochromes are sensitive to water. The green dye, in particular, is vulnerable and can bleed into other areas of the image if moisture penetrates the plate, leading to the greening defect. In severe cases, the entire image can become unreadable.

Since this damage cannot be reversed at all in physical autochrome plates, Sinha et al. [27] proposed using AI models to digitally restore images of digitized autochromes that exhibit greening. To train these AI models for image restoration, we first need to simulate the greening defect, which is referred to as a synthetic defect. The next section describes how we synthetic greening defects [27] were generated.

Synthetic Greening Defect Generation. To address the key issue of insufficient datasets for damaged autochromes, Sinha et al. [27] proposed a method for generating a synthetic dataset suitable for training and testing autochrome restoration techniques. To ensure the quality of this synthetic dataset, we conducted a thorough analysis of actual defects found in autochromes. This analysis revealed that defects can be classified as spot-shaped, widespread, a combination of both, or resulting in complete image loss. We selected seven autochromes for a more detailed evaluation, which demonstrated that spot defects typically

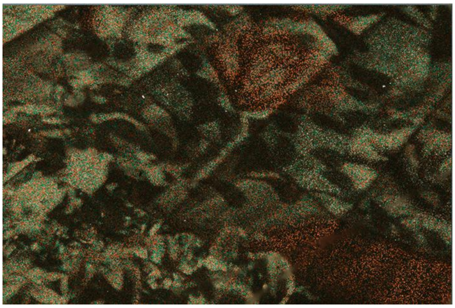

(a) Current state autochrome (detail of "Umbrella")

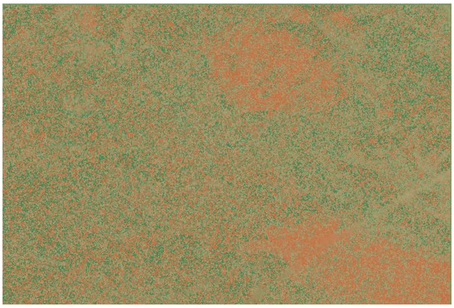

(b) True colour rendering of the faded dye mosaic

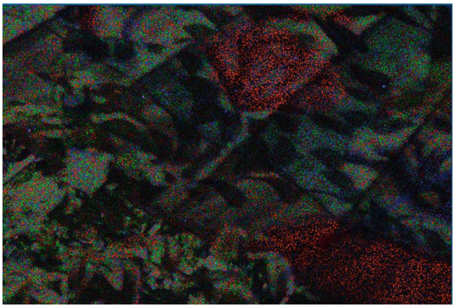

(c) Restored autochrome

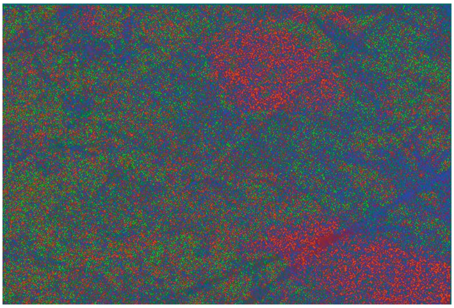

(d) True colour rendering of the restored dye mosaic

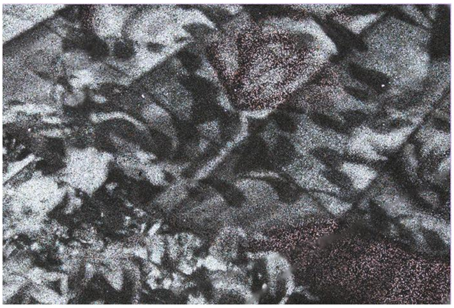

(e) Spatial information contained by the silver layer, computationally obtained as the residual of the spectral unmixing model

Fig. 42. The "Umbrella" autochrome has no blue colours, because the blue dye has faded completely to the point of becoming white. For this reason, its restoration is more dramatic than in the case of the other autochromes. Apart from the blue dye, the red and green are also rather desaturated, more so than in the other test autochromes. (Color figure online)

range from 1% to 5% of the overall image width, while larger defects can occur independently, covering up to one-third of the image. However, the merging of larger area defects with other spot defects is relatively rare.

In an initially generated dataset, defects were simulated using color filters commonly employed in photography, with a particular focus on green filters to

Fig. 43. The full umbrella (as result of stitching 45 captured images) detail in the "Umbrella" autochrome, split between its current, faded state (left side of the diagonal), and its restoration with the spectral unmixing method (right side of the diagonal). We can notice how the umbrella regains not only its blue colour, bleached in the current state, but also the saturation and vividness of its red and green granules. The two magnifier peepholes illustrate the full resolution of the colour granules in the autochrome's mosaic. At this resolution, it is even more obvious how the granules that were once blue are white in the current state as consequence of fading. (Color figure online)

Fig. 44. Examples of real greening defects: small area or spotting defects (left) and large area defects (right). Image source [30].

Fig. 45. Synthetic greening defects (small area (left) and large area (right) defect) generated using the approach proposed by Sinha et al. [27]. Image source [30].

represent the foreground layer during the defect generation for the autochrome in the background layer. However, this approach ultimately failed to produce realistic results, as the defects exhibited low color variance, rendering them unrealistic. Hence, a second dataset was created, referred to as ChannelGreeningDefects (see Fig. 45), which directly captures observable changes in the color channels within the defective areas. The process of generating this synthetic defects [27] can be described below:

- Images may contain point defects (60%), large defects (30%), or both (10%), with defect masks generated according to the characteristics of these known defects. These masks are created with 1 to 7 spot-shaped defects or 1 to 2 larger defects, using sizes and centers defined based on known size ratios from real defects.
- The generation of area and spot-shaped defects involves creating ellipses distorted by random radius adjustments at their boundaries. Spot defects have

their centers located within the image, resulting in point-shaped origins, while large area defects are represented by significantly larger ellipses with origins outside the image, simulating the leakage of liquids that penetrate the autochrome and cause damage. The parametric equations [27] for the ellipse are given by:

$$x(t) = x_c + a \cdot \cos(t) \cdot \text{irregularity},$$
$$y(t) = y_c + b \cdot \sin(t) \cdot \text{irregularity},$$

where (x_c, y_c) are the coordinates of the center of the ellipse, (a, b) are the semi-major and semi-minor axis lengths, t is an angular parameter in $[0, 2\pi]$, and irregularity is a factor affecting the shape of the ellipse. The center of defect masks may be point-shaped or linear, and the intensity of the defect is interpolated based on the normalized distance d to the origin, with a decrease in intensity computed by $I = -d^2 + 1$.

– For each ring, a change in the percentage of the individual color channels in the defect areas is defined, randomly adjusted by a factor of 0.2 (chosen empirically to avoid drastic changes and maintain the realism of the defects) in both directions for each image. The combined defect pattern is smoothed using a Gaussian filter, and the intensity change ΔI_c for color channel c is calculated as $\Delta I_c = (p_c * I_c) - I_c$, where p_c is the corruption percentage. This change is then applied to the final image. Table 2 shows the definition of the individual modifications per channel and ring in a dictionary.

Table 2. Dictionary for ring of corruption. Source [27]

Label	Blue	Green	Red	Effect in defect
9	0.6	0.85	1.05	Outer orange ring
1	0.5	1.2	0.8	Light green ring
2	0.4	0.8	0.6	Middle
3	0.4	0.8	0.6	Middle
4	0.2	0.6	0.1	Dark green second
99	0.2	0.2	0.1	Dark mid
20	0.4	0.95	0.6	Surface

Training of Restoration Network. Image2Image models, specifically Pix2Pix [12] and CycleGAN [37], were trained on the synthetic defect dataset (see Fig. 45) using the original training configurations. However, these approaches did not yield satisfactory results in restoring the original colors of the defective areas, as they modified the images without effectively addressing the

Fig. 46. Large area greening defects removal utilizing the approach proposed by Sinha et al. [27]. Image source [30].

restoration task. Therefore, the ChaIR (Channel Interactions for Image Restoration) model [8] was selected due to its demonstrated state-of-the-art performance in deraining and dehazing tasks which is similar to our problem. This model utilizes spatial and frequency loss functions to evaluate performance and supports dual-channel attention mechanisms. While it primarily affected the greening regions, it showed promise in its application despite the challenges in overall color restoration. The main challenge for autochrome restoration remains the accurate representation of the original colors in the defective areas. To enhance the restoration process, a modified loss function was implemented [27], prioritizing the impact on defective areas over non-defective regions.

A further step involved the digital in-painting of autochromes which have had been de-greened. When digitally restored, the de-greened area can appear pale or off-colour. Using hundreds of undamaged autochromes, we trained the network [15] utilizing low rank adaptation (LoRA) [11] to understand the unique colour palette and texture of the autochrome. Based on this training, the network is able to inpaint those areas digitally "de-greened" so that they more closely resemble the colour and texture of an autochrome (see Fig. 48).

Results. Figure 46 shows that large area defects are removed utilizing this de-greening model. The greening is significantly reduced and this improves the state-of-the-art image restoration methods as evaluated and shown by the work presented by Sinha et al. [27]. Moreover, even greening defects with significant could be removed partially as shown in Fig. 47. In terms of small area defects (green spotting) the de-greened output followed by inpainting using a finetuned model [15] could produce perceptually restored results as shown in Fig. 48. This approach was most effective in stationary regions—areas within an image where statistical properties such as texture, color, and intensity remain relatively constant. This concept is closely linked to stationary texture, which refers to a uniform and consistent pattern throughout a given area.

Fig. 47. Large area greening defects (with significant damage) removal utilizing the approach proposed by Sinha et al. [27]. Image source [30].

Fig. 48. Removal of small area greening defects (left) using the approach proposed by Sinha et al. [27] (middle), followed by inpainting with a fine-tuned model [15] that employs Low Rank Adaptation [11] (right). Image source [30].

Final Comments. The proposed method effectively identifies and removes some greening defects in old autochrome photographs, regardless of their size or location. By adjusting the colors in the affected areas, the method makes these defects less noticeable. However, it sometimes struggles to accurately reproduce the original colors of the autochrome images, which can lead to bluish tones and the omission of very small defects. To enhance the restoration process, AI-based inpainting methods has been employed to selectively remove small defects using the mask and the de-greened output predicted by the model.

The results of this method can improve editing tools, allowing designers to efficiently achieve similar outcomes while tackling defects that are otherwise difficult to manage, especially in larger areas. Future efforts should aim to expand the synthetic dataset to include digitized autochromes from various collections and to identify effective no-reference metrics—tools that assess image quality without needing a reference image for comparison. These advancements would enable a more comprehensive evaluation of restoration quality across a diverse array of samples. Moreover, the techniques developed could be adapted to address other autochrome defects, such as oranging, and could also be applied to restore various types of artworks and images, broadening their impact in the field of image restoration.

5.3 Reconstructing Colour in Kodacolor Films

Technical Specifications. Lenticular Kodacolor used a form of spatial encoding to capture colour information. The 16 mm film base was embossed with hundreds of fine vertical cylindrical lenses. On the reverse side, a black-and-white panchromatic silver emulsion recorded the image. Figure 49 schematically illustrates the photographic process. During film shooting (left), a special red-green-blue filter (tripartite filter) was placed in front of the camera lens. The lenticular surface focused each colour component onto the emulsion as narrow vertical stripes, enabling colour scenes to be captured on black-and-white film stock. Reversal processing produced a positive silver-based image.

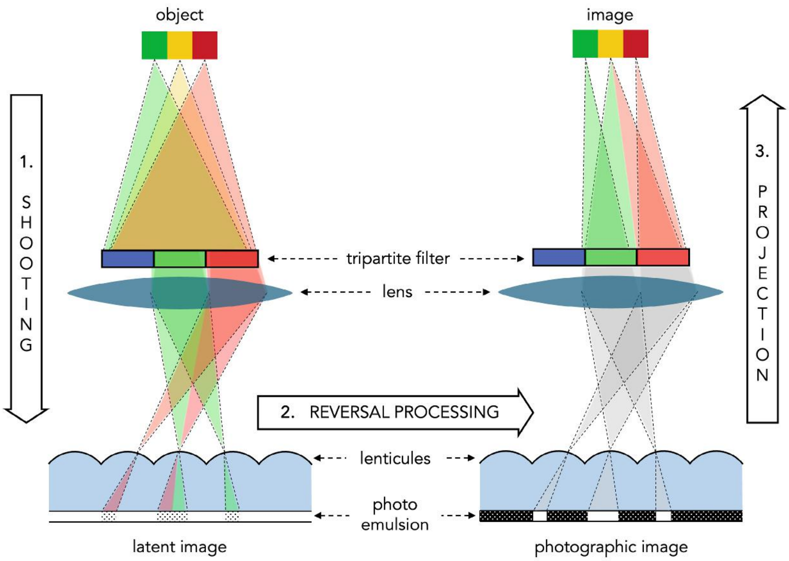

Fig. 49. Schematic illustration of the optical principle underlying image capture and visualization in lenticular colour film, considering a green-yellow-red flag as photographed object. (Color figure online)

The projection (right) required a complementary setup: a projector fitted with a matching red-green-blue filter. When projected, the magnified images recombined to form a full-colour picture on the screen. However, this method of display has since become obsolete. Today, preserving and recovering Kodacolor films requires addressing this issue of technological obsolescence.

Digitisation and Colour Reconstruction. Digitising Kodacolor films provides a practical way to make these works accessible to the public while also safeguarding them against the inevitable decay of the original film material.

Several approaches can be used to extract the embedded colour information. Four main methods can be outlined:

- **Telecine -** This is the most straightforward solution. The original projection system, equipped with the red-green-blue filter, is used, and a digital video camera records the projected colour image on the screen. A synchronization mechanism is required to align the frame rate of the film projector with that of the digital capture device.
- **Colour Digital Capture -** Instead of projecting the image onto a screen (like in "telecine"), a dedicated imaging system can be employed. A lens equipped with a red-green-blue filter–equivalent to the original one–focuses the colour image directly onto a digital colour sensor at low magnification. This process involves frame-by-frame sequential capture using a film transport system.
- **Slit-Scan Capture -** Based on the same principle as *colour digital capture*, the slit-scan eliminates the use of a red-green-blue filter and employs a monochrome digital sensor. A moving slit in front of the lens isolates a specific angular portion of the light, enabling the capture of a single colour component at a time. Three successive exposures per frame, with the slit positioned differently each time, provide the complete colour information. The film transport system is intermittent.
- **Digital Decoding -** In this approach, the film is scanned at high resolution (at 15 or more pixels per lenticule horizontally) with the emulsion facing the imaging system so that the encoded silver-based colour components remain in focus. Specialised software then reconstructs the colour image by:
  1. Locating the lenticular screen, and
  2. Converting the side-by-side silver densities into RGB values.

Among the four methods, *telecine* depends on original Kodacolor projection equipment, which is now extremely rare. Both *colour digital capture* and *slit-scan capture* demand highly specialised, custom-built devices, making them impractical for most institutions. By contrast, *digital decoding* can be considered the most efficient and widely applicable solution. It requires only a conventional high-resolution film scanner for image acquisition, while the reconstruction of the colour information is carried out entirely in software. In 2013, two independent research groups validated the digital decoding approach: Reuteler, Fornaro, and Gschwind within the SNF project doLCE at the University of Basel [24], and Aschenbach in a separate study [2]. The PERCEIVE project draws on the work of Reuteler and Gschwind and on the later advances by D'Aronco and colleagues [9]. The central challenge in digital decoding lies in accurately locating the lenticular screen. The original doLCE method relied on signal processing: it averaged the image rows, then identified the local minima of the pixel value profile to determine the screen's structure. The more recent deep-doLCE approach replaced this strategy with a deep-learning framework. Rather than fine-tuning the existing algorithm with additional parameters, it sought higher robustness and accuracy by letting a neural network learn the structure of the lenticular pattern directly from data.

DoLCE 2.0. Despite the improvement in some cases, deep-doLCE still failed to properly reconstruct some films, therefore, under the auspices of PERCEIVE, the research into this topic has continued. The work focused on a 1930 travelogue to the Norwegian fjords filmed by the amateur filmmaker Arnold Hennings. This 16mm movie is a combination of tinted film (title cards), black-and-white (most of the content), and lenticular Kodacolor (last part). The colour reconstruction using deep-doLCE resulted in image frames that introduced their own set of problems. Many frames displayed artifacts and, in several cases, completely incorrect colours. This issue arises from the lenticule detection model's inability to accurately predict lenticular borders for the provided image set. To improve the colour reconstruction of the travelogue, we refined the original deep-doLCE model using more representative training data aligned with the lenticular boundaries of the scans. This finetuning step enabled the model to more accurately detect boundaries and thereby achieve more reliable colour reproduction. To build this dataset, we employed a semi-automatic labelling process supported by visual analysis, ensuring that reconstructed colours were both technically consistent and historically plausible (e.g., skies appeared bluish, grass greenish, and clothing matched known period tones). We explored two complementary strategies for generating representative data:

- Manual refinement of lenticular boundaries. First, we applied the original deep-doLCE model to colourize the full video sequence. We then selected frames at intervals where colours appeared reasonably plausible. Using these as a starting point, we manually corrected the detected lenticular boundaries to improve accuracy. The corrected frames formed the training data for finetuning the model. If results remained unsatisfactory, additional frames were labelled and incorporated, with the process repeated iteratively until the reconstructed sequence achieved a visually consistent and credible colour balance.
- Automatic selection of plausible frames. To reduce manual intervention, we also adopted a more automated strategy. Instead of adjusting boundaries directly, we identified frames or cropped regions where the original model's output already showed convincing colours. These selections, validated through visual reasoning, were then used to finetune the model. As with the manual approach, additional data were added iteratively until the sequence produced stable and realistic colour results.

Together, these methods demonstrate that user-guided finetuning–whether through manual adjustment or selective frame validation–can substantially improve the colour reconstruction of lenticular film scans, even when source material is of suboptimal quality.

6 Conclusion and Discussion

Colour in heritage objects is both semantically rich and materially fragile. This chapter demonstrates that robust colour reconstruction and change prediction emerge when material science, computational imaging, and visual computing are treated as a single, interoperable workflow. At the foundation are well-designed mock-ups that embody the chemistry and stratigraphy of historical practices—whether Punic-wax polychromy on Pentelic/Parian marble, pigments with linseed oil in paintings, or natural/synthetic dyes on silk and cotton. Colour and spectral imaging, together with appearance capture (MA-T12, TAC7) and SVBRDF/AxF modelling preserve colour, spectral and geometric attributes, supporting photorealistic rendering of instances of the studied artworks in past and future.

classical sculpture, an atlas-based segmentation and semantic shading framework enables the transfer of measured material models from controlled mock-ups onto segmented 3D reconstructions, overcoming earlier "flat colour" restorations by recovering gloss, texture, and micro-geometry cues essential to perceived authenticity of polychromy.

For paintings, we showed that complementary strategies - (i) a mechanistic photochemical model that ties ΔL^*, Δa^*, and Δb^* trajectories to spectral radiant exposure and (ii) spectral unmixing operating in absorbance space - yield predictive maps and renderings of future damage, while historical photographic records can be corrected (via dye density balancing and learned film$\leftarrow$painting transforms) to approximate past appearances. This pluralistic approach reduced the typical "lack-of-ground-truth" problem by cross-validating expected behaviours (e.g. vermilion darkening, cadmium yellow fading).

For textiles, integrating spectral unmixing with appearance-aware rendering and interactive visual tools (SimTex) allows conservators to toggle between hypothesized original states and forecasted fading scenarios, grounded in accelerated aging data and archival photographs. The method balances chemical plausibility with perceptual realism, acknowledging uncertainties from photographic variability.

For early colour photographs and films, we advanced multispectral 'silver/dye decoupling' to restore autochromes (including severely faded blue dyes) and explored neural restoration of "greening" defects trained on synthetic data; for Kodacolor, we improved digital decoding through targeted finetuning of deep learning models for lenticule boundary detection, recovering historically plausible colours from lenticular scans.

Author Contributions. Section 2 was written by Yoko Arteaga, Cristiana Barandoni, Donata Magrini, Saptarshi Neil Sinha and Sophia Sotiropoulou. Section 3 was written by Irina-Mihaela Ciortan, Letizia Monico, Irina Sandu, Panagiotis Siozos, and Sophia Sotiropoulou with contributions from David Buti. Section 4 was written by Irina-Mihaela Ciortan, Lucia Burgio, Brenda Doherty, Catarina Monteiro Pinto, and Saptarshi Neil Sinha with contributions from David Buti and Faïçal Selka. Lastly, Sect. 5 was written by Yoko Arteaga, Lucia Burgio, Irina-Mihaela Ciortan, Catlin Langford, Giorgio Trumpy, and Saptarshi Neil Sinha, with contributions from Jon Yngve Hardeberg.

Acknowledgments. This study was funded by the European Union (grant number 101061157, Project PERCEIVE).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Arteaga, Y., Sole, A.S., Hardeberg, J.Y., Boust, C.: Characterising appearance of gold foils and gilding in conservation and restoration. In: Proceedings of the 11th Colour and Visual Computing Symposium (CVCS), vol. 3271. CEUR Workshop Proceedings (2022). https://ceur-ws.org/Vol-3271/Paper11_CVCS2022.pdf
2. Aschenbach, T.: Digital Color Decoding of Kodacolor Film, November 2013
3. Berthon, R.: A suggested new method of colour photography direct at one exposure. Br. J. Photogr. **57**, 421–423 (1910)
4. Chatterjee, S., Trumpy, G., Ruedel, U.: Digital unfading of chromogenic film informed by its spectral densities. Heritage **6**(44), 3418–3428 (2023). https://doi.org/10.3390/heritage6040181
5. Ciortan, I., Arteaga, Y., Trumpy, G.: Silver/dyes decoupling and colour restoration of autochromes using multispectral imaging. Eur. Phys. J. Plus **140**, 384 (2025). https://doi.org/10.1140/epjp/s13360-025-06284-2
6. Ciortan, I.M., Trumpy, G.: FILM2PAINT: transforming photographic documentation on reversal film into paintings' accurate colors. In: Archiving Conference, vol. 21, pp. 17–22. Society for Imaging Science and Technology IS&T, 7003 Kilworth Lane (2024)
7. Ciortan, I.M., Trumpy, G., Sandu, I.C.A.: *The Scream* (ca. 1910) through the Years: from Photographic Documentation to Digital Rejuvenation. In: Colour Photography and Film: Sharing Knowledge of Analysis, Preservation, Conservation of Analogue and Digital Materials, p. 152 (2024)
8. Cui, Y., Knoll, A.: Exploring the potential of channel interactions for image restoration. Knowl. Based Syst. **282** (2023)
9. D'Aronco, S., Trumpy, G., Pfluger, D., Wegner, J.D.: A deep learning approach for digital color reconstruction of lenticular films (2022). https://arxiv.org/abs/2202.05270, eprint: 2202.05270
10. Hannouch, H.: Three-colour photography: technologies, expeditions, empires. Eur. Soc. Hist. Photogr. **37**, 5 (2022)
11. Hu, E.J., et al.: LoRA: low-rank adaptation of large language models (2021). https://arxiv.org/abs/2106.09685
12. Isola, P., Zhu, J., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: CVPR, pp. 5967–5976 (2017)
13. Keshava, N., Mustard, J.: Spectral unmixing. IEEE Sig. Process. Mag. **19**(1), 44–57 (2002). https://doi.org/10.1109/79.974727
14. Kohlbrenner, A., Rump, M., Frick, B., Schwartz, C.: Method and apparatus for digitizing the appearance of a real material (2018)
15. Labs, B.F., et al.: FLUX.1 Kontext: flow matching for in-context image generation and editing in latent space (2025). https://arxiv.org/abs/2506.15742
16. Lavédrine, B., Gandolfo, J.P.: The Lumière Autochrome: History, Technology, and Preservation. Getty Publications, Los Angeles, California (2013)

17. Liverani, P., Bracci, S., Iannaccone, R., Lenzi, S., Magrini, D.: New evidence about the polychromy of early imperial cycle from the *Augusteum of Rusellae* (Tuscany). Heritage **6**, 3385–3401 (2023). https://doi.org/10.3390/heritage6040179
18. Magrini, D., Bracci, S., Bartolozzi, G., Iannaccone, R., Lenzi, S., Liverani, P.: Revealing Mithras' color with the ICVBC mobile lab in the museum. Heritage **2**, 2160–2170 (2019). https://doi.org/10.3390/heritage2030130
19. Merzbach, S., Hermann, M., Rump, M., Klein, R.: Learned fitting of spatially varying BRDFs. Comput. Graph. Forum (2019). https://cg.cs.uni-bonn.de/svbrdfs/
20. Merzbach, S., Weinmann, M., Klein, R.: High-quality multi-spectral reflectance acquisition with X-rite TAC7. In: Eurographics Workshop on Material Appearance Modeling, MAM@EGSR 2017. Eurographics Association (2017)
21. Müller, G., Tautges, J., Gress, A., Rump, M., Hermann, M., Lamy, F.: AxF - appearance exchange format. Technical report, X-Rite, Inc., 4300 44th St. SE, Grand Rapids, MI 49505, February 2020. https://www.xrite.com/-/media/xrite/files/whitepaper_pdfs/axf/axf_whitepaper_v18_en.pdf
22. Pfluger, D., Garmsen, L., Trumpy, G.: Modern methods for the visualization of lenticular film colors. Color Cult. Sci. J. **14**(1), 64–72 (2022)
23. Pliny the Elder: Natural History, Volumes I-X. Loeb, Cambridge (MA) (1938–1962)
24. Reuteler, J., Fornaro, P., Gschwind, R.: Digital Color Reconstruction of Lenticular Film: Project 'doLCE', June 2013
25. Series, B.: Parameter values for ultra-high definition television systems for production and international programme exchange. In: Proceedings of the ITU-T, Bt. 2020, pp. 1–7 (2012). https://glenwing.github.io/docs/ITU-R-BT.2020-1.pdf
26. Sinha, S.N., et al.: An Atlas-based approach for appearance-aware virtual 3D restoration and simulation of fading in fugitive textiles. In: Campana, S., et al. (eds.) Digital Heritage. The Eurographics Association (2025). https://doi.org/10.2312/dh.20253147
27. Sinha, S.N., Kuehn, P.J., Koppe, J., Kuijper, A., Weinmann, M.: Neural restoration of greening defects in historical autochrome photographs based on purely synthetic data (2025). https://arxiv.org/abs/2505.22291
28. Siozos, P., et al.: Light source classification and colour change modelling for understanding and predicting pigments discolouration. J. Phys. Photonics **7**, 035020 (2025). https://doi.org/10.1088/2515-7647/ade8c2
29. Streamlit Inc.: Streamlit (2019). https://streamlit.io. Accessed 01 Oct 2025
30. Taylor, H.A.: Harold Taylor slide collection. California Revealed (1896). https://californiarevealed.org/search?search_api_fulltext=&f%5B0%5D=search_page_series_title:Harold%20Taylor%20Slide%20Collection. Accessed 26 Nov 2024
31. Theophrastus: De lapidibus. Teubner, Stuttgart (1952)
32. Trumpy, G., Hardeberg, J.Y., George, S., Flueckiger, B.: A multispectral design for a new generation of film scanners. In: Optics for Arts, Architecture, and Archaeology VIII, Munich, Germany, July 2021, vol. 11784, pp. 138–146. SPIE (2021). https://doi.org/10.1117/12.2592655
33. Valero, E.M., Martínez-Domingo, M.A., López-Baldomero, A.B., López-Montes, A., Abad-Muñoz, D., Vílchez-Quero, J.L.: Unmixing and pigment identification using visible and short-wavelength infrared: reflectance vs logarithm reflectance hyperspaces. J. Cult. Herit. **64**, 290–300 (2023). https://doi.org/10.1016/j.culher.2023.10.016
34. Vitruvius, M.P.: De architectura libri decem. Teubner, Stuttgart (1991)
35. X-RITE: TAC7-Scanner (2018). http://web.archive.org/web/20180615015942/, https://www.xrite.com/categories/ap

36. X-RITE: Pantora material hub (2019). http://web.archive.org/web/20190424232441/, https://www.xrite.com/categories/appearance/pantora-software
37. Zhu, J., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: ICCV, pp. 2242–2251 (2017)