

VisNIR pigment mapping and re-rendering of an experimental painting

Federico Grillini[†], Jean-Baptiste Thomas and Sony George

The Norwegian Colour and Visual Computing Laboratory, IDI, NTNU, Gjøvik, Norway

[†]Email: federico.grillini@ntnu.no

Pigment mapping allows the classification and estimation of the abundances of pigments in paintings. The information learned becomes extremely important for conservators, who are then able to decide the best strategies in the conservation of the artefacts. When the goal is to restore a painting, it is also important to know what the effects of the newly introduced materials are. To fulfil this purpose, a proper mixing model must be defined. We propose a framework to perform pigment mapping on the hyperspectral image of an experimental painting realised for the occasion, with the goal of rendering a colour image using the concentrations retrieved from the mapping. Contrarily to spectral unmixing tasks, where subtractive models prevailed, hybrid models have the advantage of outputting more accurate colours in this workflow.

Received 15 January 2021; accepted 28 January 2021

Published online: 16 February 2021

Introduction

Imaging and analytical methodologies can help conservators during the non-invasive examination of historical paintings [1]. Such artefacts represent a complex category of objects, as they contain many mixtures of materials which are hard to be discerned visually. However, recently developed non-invasive techniques can be adopted to detect and separate the components of mixtures. Hyperspectral Imaging (HSI) is one of such techniques [2-3]. It allows the acquisition of images which are highly resolved in the spatial and spectral domains. The image cubes produced by hyperspectral cameras have a spectral sampling of a few nanometres, thus allowing the pixel-wise examination of the acquired object's radiance spectra. HSI allows to estimate the reflectance factors of the pixels, by including in the scene appropriate targets that serve to discard the illuminant's effect [4]. One of the advantages of HSI is that the sensitivity is not limited solely to the visible range: depending on the type of sensor deployed, it is possible to retrieve the reflectance spectra in the UV, IR, and short wave IR (SWIR) as well. This is very relevant in conservation science, as varnishes (the usually outer layer of a painting) reveal interesting properties in the UV range [5], and underdrawings and pentimenti can be detected only at longer wavelengths, exploiting the transmittance properties of pigments at those frequencies [6].

The abundances, or concentrations of the endmembers, i.e. the material that compose a mixture, can be retrieved when their spectra are known, by inverting a properly designed optical mixing model, in a process that is known as spectral unmixing [7]. When it is applied to areas rather than points in the context of paintings, spectral unmixing takes on the name of pigment mapping [8]. An optical mixing model combines the reflectance factors of the endmembers, included in the matrix E , with the abundances α assigned to each spectrum, to finally yield a resulting reflectance spectrum Y . The entries of the concentration vector C must comply with two constraints: non-negativity and sum-to-one, as Equation (1) shows.

$$\begin{aligned}
 Y_{[\lambda \times 1]} &= f(E_{[\lambda \times q]}, C_{[q \times 1]}) \\
 \text{with } C &= (\alpha_1, \alpha_2, \dots, \alpha_q) \\
 \text{s. t. } \alpha_i &\geq 0 \text{ and } \sum_{i=1}^q \alpha_i = 1
 \end{aligned}
 \tag{1}$$

Different optical mixing models can be defined and classified accordingly to their nature. When the mixture of pigments is addressed, depending on how the paints are lay down on the canvas, a model can be more appropriate to correctly describe a specific configuration. The additive model explains optical mixing at the camera level: due to limits in spatial resolution, two materials that are physically separated, end up being represented in a single pixel value, as it is common in remote sensing. This model is characterised by linearity. On the other extreme, the subtractive model can better describe the configuration of intimate mixing, by employing a geometrical mean-based formula that combines the endmembers and their concentrations [9]. In between the additive and subtractive mixing, a series of hybrid models can be formulated [10].

In a previous study by Grillini *et al.* [11], the spectral unmixing performances of a series of optical mixing model was investigated on a set of oil painted mockups. Subtractive-based models outperformed the others in terms of spectral accuracy and estimation of the concentrations.

This article revisits a previous conference paper published at the AIC2020 conference [12], and reports a full investigation of an experimental painting, rather than mockups. We use the same pigments (Kremer¹) than in the AIC2020 conference paper to realise the oil painting. The optical mixing models are then evaluated both for unmixing (considering only a subset of the previously evaluated models) and for colour rendering. Colour rendering after a step of unmixing is indeed of major importance for conservators who are interested in the final appearance of an artifact after the introduction or removal of a material with known reflective properties, without of course acting directly on the precious painting.

In the previous studies the subtractive-based models showed the best results in terms of spectral reconstruction and in terms of pigment classification and estimation. In this rendering-oriented framework, that is no longer the case. Indeed, the hybrid model proposed by Yule and Nielsen [13] achieved more accurate rendering results. However, rendering yet represents a challenge in this type of framework, as it is posed in contrast with accurate pigment detection and concentration estimation.

¹Kremer Pigmente GmbH & Co.KG, <https://www.kremer-pigmente.com/en/>. [Online; last accessed 21 August 2020]

Materials and methods

The target used for this article is an oil painting realised on a pre-primed stretched canvas made of linen. An ulterior double layer of gesso is applied for priming, in order to facilitate the adherence of the paint layer. Seven historical pigments except Kremer White, which is adopted as best substitute for Lead White (a toxic pigment, not commercially available), are used to depict the famous *Baby Yoda meme* (Figure 1a). The list of the pigments used can be found in Table 1. Pigments are bound with linseed oil and then applied on canvas, and no varnish is applied, in order to observe directly only the paint layer.

Pigment	Code	Label
Kremer White	46360	W
Ultramarine Blue	45030	B
Naples Yellow	43125	Y
Novoperm Carmine Red	23403	C
Vermilion	42000	V
Viridian Green	44250	G
Gold Ochre DD	40214	O

Table 1: Pigments used in the realisation of the painting of Baby Yoda. The codes refer to the Kremer nomenclature, while the labels are semantically assigned by the authors.

A ground truth reporting the exact composition of each area of the painting was not recorded, due to the complexity of the realisation. However, the painting is later divided into 12 macro-areas in which the approximated concentrations of the present pigments are known (Figure 1b).



Figure 1: Experimental painting of Baby Yoda (a). The painting is realised using the set of pigments reported in Table 1, bound with linseed oil. The 12 macro areas of (b) indicate zones with generally similar and known ratios of concentrations.

To capture the hyperspectral cube, a push-broom HySpex VNIR-1800 hyperspectral camera produced by Norsk Elektro Optikk is deployed. The illumination is provided by a halogen Smart light 3900e by Illumination Technologies, guided via optic fibre onto the scene, forming a standard CIE 45°/0° acquisition geometry. The focus is set at 30cm, with a field of view of 16°, yielding a pixel resolution of approximately 50 μm . The final image has a spatial dimension of 2688 \times 3488 pixels. The total number of pixels would imply very long computation times in our framework; therefore, the image

is spatially down-sampled by a factor 10. The sensor of the HS camera is sensitive in the range 400nm to 1000nm, at steps of about 3 nm, resulting in 186 bands recorded in the final image cube. A Spectralon® target with known reflectance factor is included in the scene, in order to perform flat-field correction and compute the reflectance map of the spectral cube.

The ground truth for comparing the rendering results is the image reported in Fig. 1a. This image is obtained from the reflectance data, assuming a D65 standard illuminant and the 2° standard observer. The XYZ values are then transformed to sRGB [14] using the Matlab function *xyz2rgb*. The same workflow is followed for all the conversions from reflectance data to sRGB that are presented hereafter.

Three optical mixing models, namely linear, subtractive, and Yule-Nielsen [13] are considered for this work. The models selected are simple but significant for the category they represent (Table 2).

Model	Equation	Category
Linear	$Y = \sum_{i=1}^q \rho_i \alpha_i$	A
Subtractive	$Y = \prod_{i=1}^q \rho_i^{\alpha_i}$	S
Yule and Nielsen	$Y = \left(\sum_{i=1}^q \rho_i^{\tau} \alpha_i \right)^{\frac{1}{\tau}}$	H

Table 2: List of investigated models. The category labels are as follows: A-additive, S-subtractive, H-hybrid. The mixing constant τ used in the Yule-Nielsen model is set to 0.5.

Pigment mapping is performed on the down-sampled reflectance image pixel-wise, each time inverting a different mixing model. During this phase, the concentration vector is retrieved optimising a function that tries to maximise the Peak Signal-to-Noise ratio (PSNR) between the reference spectrum, and the reconstructed one. As demonstrated by Grillini *et al.* [11], the PSNR correlates with an accurate pigment detection and estimation. At the end of this process, 3 image cubes containing 7 concentration maps each are produced. For a thorough investigation, each concentration cube is combined with the 3 mixing models, to produce 3 different renderings. In total there will be 9 rendered versions of *Baby Yoda*.

The PSNR-based optimisation deployed in this phase utilises all the points of the reflectance spectra, from 400 to 1000 nm. However, colour information is coded only up to 780nm, utilising the CIE 2° standard observer. It is possible that cost functions limited solely to the visible range may be able to produce more accurate renderings. For this reason, two more rounds of optimisation are computed: the first round deploys the PSNR but with a limitation to the visible range, whereas the second round deploys a cost function based on the colour difference ΔE_{2000} [15]. All three optimisations are conducted adopting the simplex downhill algorithm proposed by Nelder and Mead [16].

The results of these tests are compared to the ground truth by means of PSNR and ΔE_{2000} .

Results and discussion

Pigment mapping on the *Baby Yoda* painting is conducted utilising the spectra of the seven pure endmembers, the reflectance measurements of the painting, and the three selected mixing models in turns.

After the 3 concentration cubes are learned, 9 reconstructed hyperspectral images are composed and rendered to sRGB (2° observer, D65 illuminant). Figure 2 reports such rendered images. At a first glance, none of the 9 rendered version has a close resemblance to the ground truth, especially if the colourfulness² factor is considered. However, there are internal differences that give indications on which model is more appropriate for the re-rendering. To aid the evaluation, Table 3 provides the values of PSNR and ΔE_{2000} for the re-rendered versions.

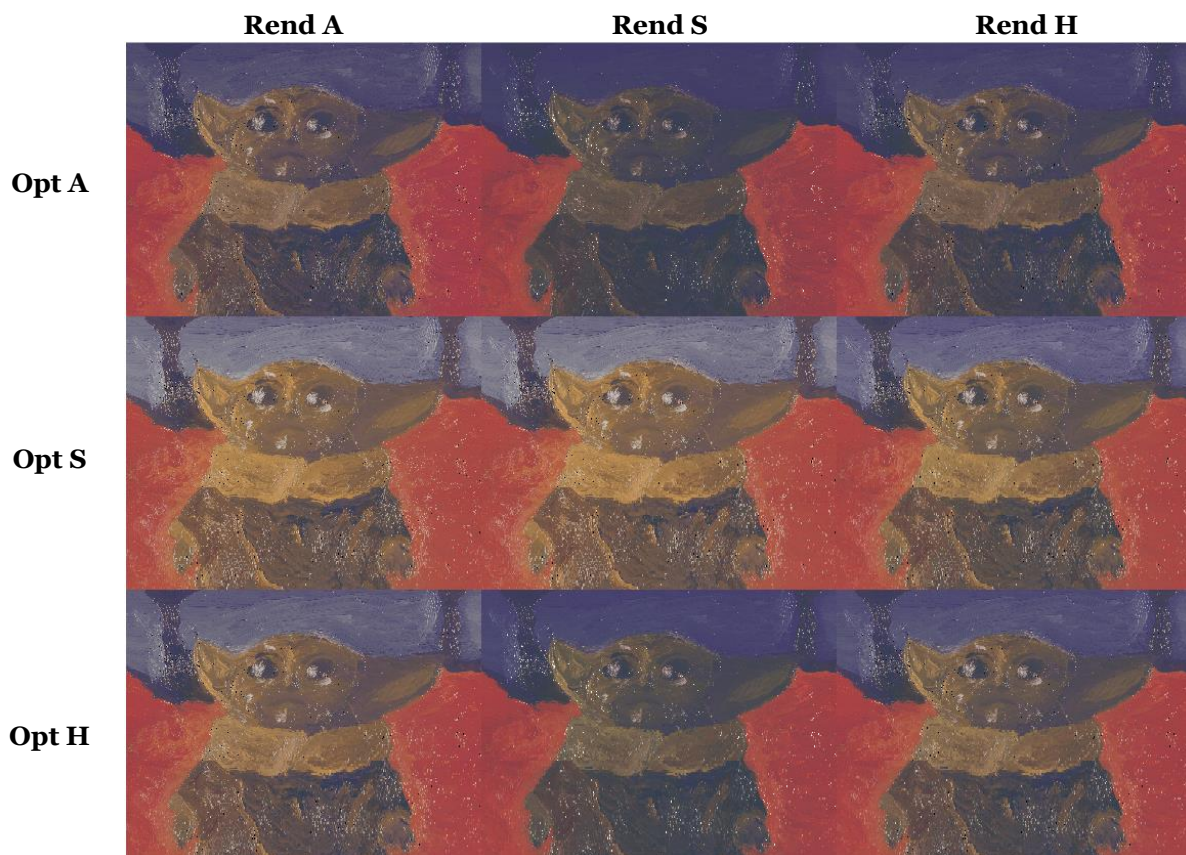


Figure 2: Nine re-rendered versions of Baby Yoda. The rows report what model is inverted in the spectral unmixing step, while the columns report which model is used to combine the matrix of endmembers with the learned concentrations. These images lack of colourfulness and brightness when compared to the ground truth. The rendering conducted with the hybrid (H) Yule-Nielsen model is the one that approaches the ground truth more closely, followed by the linear model (A) and the subtractive (S).

²Attribute of a visual perception according to which the perceived colour of an area appears to be more or less chromatic. For a colour stimulus of a given chromaticity and, in the case of related colours, of a given luminance factor, this attribute usually increases as the luminance is raised except when the brightness is very high. From the CIE vocabulary.

	Rend A		Rend S		Rend H	
	PSNR	ΔE_{2000}	PSNR	ΔE_{2000}	PSNR	ΔE_{2000}
Opt A	19.66	12.24	20.59	12.52	20.37	11.97
Opt S	15.88	16.62	15.88	16.62	17.94	13.11
Opt H	17.89	13.64	21.03	11.00	19.69	11.44

Table 3: Metrics for the evaluation of the re-rendering. Each row refers to a different optimisation, with the letters A, S, and H recalling the three models: additive, subtractive, and hybrid, respectively. Similarly, the results are presented by columns according to the model used to render the colour image. The column of the rendering according to the hybrid Yule-Nielsen model is highlighted in green as it outputs the best pair of rendering metrics. It is noticeable that this model is able to achieve better results even if it is not the one inverted during the unmixing step.

The PSNR and mean ΔE_{2000} values should be evaluated in pairs, as they simultaneously provide an indication of the accuracy of the sRGB values and a perceptual measure. Generally, when the concentration learned are combined with the hybrid Yule-Nielsen model, the best-looking images are obtained. When the unmixing is performed inverting the subtractive model, the rendering results are quite poor, although this model is the best in retrieving more accurately the concentrations. The renderings obtained after inverting the linear model (A) fall right behind the hybrid approach, showing results that are significantly better than the subtractive counterpart.

At this point, in order to improve the rendering from the learned concentrations, two more optimisation cost functions are implemented, focusing solely on optimising the spectra in the visible range. Only the hybrid Yule-Nielsen model will be considered for this last part, as it produced the best renderings in the previous step.



Figure 3: Renderings in sRGB from the pigment mapping inverting the hybrid Yule-Nielsen model. The different cost functions vary from left to right: ground truth, optimisation utilising the full range of wavelengths (400-1000nm), optimisation for PSNR only in the visible range, optimisation for ΔE_{2000} .




Figure 4: Colour difference maps of the renderings illustrated in Figure 3. As the information considered is limited to the visible range, it is possible to obtain more accurate colours.

The visual results of the renderings according to the different cost functions, namely PSNR-full, PSNR-vis and ΔE_{2000} , are reported in Figure 3. For comparison, the ground truth is showed as well. The ΔE_{2000} maps of each image are depicted in Figure 4. As expected, if the cost function focuses only in the visible range, the final sRGB images resemble more closely the ground truth. The metrics reported in Table 4 confirm the visual observation of the colour difference maps.

	Opt PSNR-full	Opt PSNR-vis	Opt ΔE_{2000}
PSNR	19.69	20.67	20.35
ΔE_{2000}	11.44	10.50	6.92

Table 4: Metrics referred to the renderings performed after inverting the Yule-Nielsen model with the different cost functions. Optimising for the colour difference yields better results than optimising for the PSNR in the visible range.

It must be pointed out that even if the colour difference values are smaller when optimising for ΔE_{2000} , the scale suggest that we are still in the domain of large colour differences.

Indeed, this test show that the mixing models adopted are still far from producing good rendering results in this framework.

Moreover, if the portion of the sky is observed, it can be noted that in the rendering performed after the optimisation for ΔE_{2000} it tends towards a greenish hue, showing however a smaller ΔE_{2000} value than in the other two approaches, which still show a bluish tonality in the sky. This hue shift also suggests that during the unmixing step, the green pigment Viridian was selected. The optimisations run for PSNR correctly avoided Viridian in the sky area, concluding the classification more accurately.

This observation might be a hint that an accurate colour rendering and an accurate pigment classification in the unmixing task are found in a trade-off, at the current state of the examined models.

Conclusions

In this article, three different mixing models are inverted in a pigment mapping application, and then used forwardly to reconstruct the hyperspectral cubes. The cubes are then converted to sRGB and compared to the ground truth, in order to assess which model is the more appropriate for the task.

The hybrid Yule-Nielsen model was selected and inverted again, this time with the goal of further improving the rendering results, by changing the cost function in the optimisation. The more accurate colours were obtained when the cost function was based on the ΔE_{2000} colour difference.

Generally, it was observable how the final rendering is highly impacted by the mixing model that is inverted, by the mixing model that is then used forwardly, and by the cost function optimised in the unmixing step.

The best produced renderings minimised ΔE_{2000} , but incidentally showed hue shifts that suggest wrong pigment classifications, which leads to the conclusion that accurate colours and accurate pigment classifications are currently in a trade-off. The limitations of the studied mixing models are showed in this article, as the scale of the colour difference is well above the human just noticeable difference threshold. Those results demonstrate that we are still far to model all the aspects of light-matter interaction present in art-painting with the purpose of analysis and rendering. Further work might focus on finding more adequate models for rendering and mapping, which would be the most useful tool for conservators.

References

1. Bianco S, Colombo A, Gasparini F, Schettini R and Zuffi S (2011), Applications of Spectral Imaging and Reproduction to Cultural Heritage, in *Digital Imaging for Cultural Heritage Preservation: Analysis, Restoration, and Reconstruction of Ancient Artworks*, Stanco F, Battiato S and Gallo G (eds.), 183-209, CRC Press.
2. Liang H (2012), Advances in multispectral and hyperspectral imaging for archaeology and art conservation, *Applied Physics A*, **106**, 309-323.
3. Jones C, Duffy C, Gibson A and Terras M (2020), Understanding multispectral imaging of cultural heritage: Determining best practice in MSI analysis of historical artefacts, *Journal of Cultural Heritage*, **45**, 339-350.
4. Krishna Mohan B and Porwal A (2015), Hyperspectral image processing and analysis, *Current Science*, **108** (5), 833-841.
5. Pelagotti A, Pezzati L, Bevilacqua N, Vascotto V, Reillon V and Daffara C (2005), A study of UV fluorescence emission of painting materials, *Proceedings of the Art '05 8th International Conference on Non-Destructive Investigations and Microanalysis for the Diagnostics and Conservation of the Cultural and Environmental Heritage*, A97, Lecce (Italy).
6. Strojnik, M, Gonzalo P and Antonio O (2011), Near IR diodes as illumination sources to remotely detect under-drawings on century-old paintings, *Proceedings of the 22nd Congress of the International Commission for Optics: Light for the Development of the World International Society for Optics and Photonics*, **8011**, 801177.
7. Bioucas-Dias JM, Plaza A, Dobigeon N, Parente M, Du Q, Gader P and Chanussot J (2012), Hyperspectral unmixing overview: Geometrical, statistical, and sparse regression-based approaches, *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, **5** (2), 354-379.
8. Hilda D, George S and Hardeberg JY (2014), Pigment mapping of the Scream (1893) based on hyperspectral imaging, *Proceedings of the International Conference on image and Signal processing*, 247-256, Cherbourg (France).
9. Burns SA (2017), Subtractive color mixture computation, [arXiv:1710.06364](https://arxiv.org/abs/1710.06364) [last access 15 January 2021].
10. Simonot L and Hébert M (2014), Between additive and subtractive color mixings: intermediate mixing models, *Journal of the Optical Society of America A*, **31** (1), 58-66.
11. Grillini F, Thomas J-B and George S (2020), Linear, subtractive, and logarithmic optical mixing models in oil painting, *Proceedings of the Colour and Visual Computing Symposium (CVCS) 2020*, Paper 7, 1-16, Gjøvik (Norway).
12. Grillini F, Thomas J-B and George S (2020), Mixing models in close-range spectral imaging for pigment mapping in cultural heritage, *Proceedings of the International Colour Association Interim Meeting (AIC2020)*, Avignon (France). [in press]
13. Yule JAC and Nielsen WJ (1951), The penetration of light into paper and its effect on halftone reproduction, *Proceedings of the Technical Association of the Graphic Arts (TAGA)*, **3**, 65-76.
14. Anderson M, Motta R, Chandrasekar S and Stokes M (1996), Proposal for a standard default color space for the internet—sRGB, *Proceedings of the 4th Color Imaging Conference*, 238-245, Scottsdale (USA).
15. Luo MR, Cui G and Rigg B (2001), The development of the CIE 2000 colour-difference formula: CIEDE2000, *Color Research and Application*, **26** (5), 340-350.
16. Nelder JA and Mead R (1965), A simplex method for function minimization, *The Computer Journal*, **7** (4), 308-313.