

Pigment Mapping for Tomb Murals using Neural Representation and Physics-based Model

Mayuka Tsuji¹ Yuki Fujimura¹ Takuya Funatomi¹ Yasuhiro Mukaigawa¹
 Tetsuro Morimoto² Takeshi Oishi² Jun Takamatsu³ Katsushi Ikeuchi³
¹Nara Institute of Science and Technology ²The University of Tokyo ³Microsoft

{tsuji.mayuka.td0, fujimura.yuki, funatomi, mukaigawa}@is.naist.jp

tetsuro.morimoto@toppan.co.jp, oishi@cvl.iis.u-tokyo.ac.jp, {jun.takamatsu, katsuike}@microsoft.com

Abstract

This paper presents a pigment mapping in tomb murals for digitization. In order to separate pigments from their substrates, we utilize the Kubelka-Munk (KM) model. However, these murals are drawn on rocks, and the pigments have deteriorated and thinned over time. As such, the challenge is to cancel the impact of the rocks' heterogeneous patterns; previous studies using the KM model either ignored the substrate or assumed it to be constant. We introduce unsupervised learning based on neural representations and physics to facilitate pigment mapping, even on a heterogeneous substrate. The model takes an image of the spectral reflectance data at a specific position of a tomb mural image and the corresponding position as inputs and outputs the pigment thickness, pigment class, and substrate class. For physically-consistent estimation, the input reflectance is reconstructed using the Kubelka-Munk model and the output. This allows unsupervised training via the calculation of the reconstruction loss. While the Kubelka-Munk model operates on a pixel-by-pixel basis, the utilization of neural representation by the input position facilitates highly accurate estimation, all the while maintaining spatial continuity.

1. Introduction

Decorated ancient tombs, the burial mounds of ancient rulers, are recognized by their patterned decorations that adorn the rock grave chambers. These patterns offer valuable archaeological insights into the religious beliefs, art, and societal progression during that period. However, despite their significance, preservation concerns often keep these tombs closed to the public. Consequently, a critical task arises: mapping these patterns and subsequently digitizing them for wider accessibility.

When analyzing the patterns within decorated ancient tombs, distinguishing the pigment from the substrate is es-

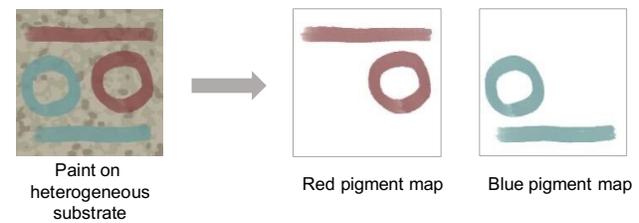


Figure 1: Our goal is to estimate pigment maps even on heterogeneous substrates.

sential. This process involves viewing the mural as a multi-layered surface object, comprising pigment and substrate layers. Two optical models have proven particularly effective in this regard: Lambert-Beer's law and the Kubelka-Munk model [15] (KM model). The latter, introducing more complex assumptions than the former, offers a more precise approximation of this layered structure.

The challenge of understanding the pigment distribution in tomb murals stems from the heterogeneous patterns inherent in the rock substrate. The pigments used in these murals resemble semi-transparent optical models, similar to watercolor paints, indicating that the influence of the substrate surface is non-negligible. The substrate pattern's complexity significantly amplifies the task of pigment separation.

This study proposes a method for pigment mapping in tomb murals, even with heterogeneous rock patterns (Figure 1). Specifically, we analyze which pigments are of what thickness. Our approach employs neural representation and physics-based unsupervised learning. The inputs are the spectral reflectance data at a specific position of a tomb mural image and its position, while the outputs are pigment thickness, pigment class, and substrate class. For loss calculation, we estimate the spectral reflectance from these outputs in line with the KM model and calculate its error from the input. That enables the model to make infer-

ences based on the physics model.

The significance of our study can be summarized in the following key points:

- To the best of our knowledge, this paper presents the first pigment mapping method that adjusts for heterogeneous substrates, thus reducing the impact of substrate patterns.
- The neural representation, utilizing the input position, facilitates a comprehensive calculation considering spatial continuity, extending beyond a mere pixel-by-pixel estimation.
- By integrating the results of reflectance reconstruction with a physics-based loss function model, we can effectuate pigment mapping without depending on supervised data.

This study not only facilitates public exhibitions, such as through AR but also provides a vital reference for restoration work should these historical sites suffer damage from natural disasters or other unforeseen circumstances.

2. Related work

2.1. Pigment mapping using KM model

We can divide the variables of the KM model into two categories: pigment layer parameters and a substrate parameter (see Section 3). Prior studies often made the assumption that the substrate parameter was either negligible or constant. For example, in the case of opaque paintings, such as oil or acrylic paintings, the pigment layer is so thick that it completely obscures the canvas substrate [24, 7, 8, 11, 21, 23]. The substrate parameter is supposed to be a constant value in [1, 22] because the substrate is often homogeneous paper (white) in the watercolors and printings. However, when we assume a non-negligible heterogeneous substrate, the problem becomes more complicated; we need to estimate the substrate parameter on each pixel while analyzing pigments.

2.2. Other Pigment mapping

Spider model [18] is an effective physics model for separating pigments from the substrates. This model is primarily based on Lambert-Beer’s law, a one-flux model focusing on light attenuation in a singular direction shown in Figure 2(a). However, crucial to note that Lambert-Beer’s law does not consider the non-negligible scattering caused by pigment particles as a parameter. While the Spider model is reliant on RGB images, hyperspectral data for pigment mapping has been widely reported [6, 5, 4, 12, 20]. Hyperspectral data provides more detailed spectral information, which could potentially enhance the accuracy of pigment mapping. In recent years, with the rise of deep learning,

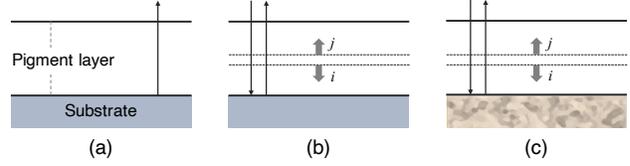


Figure 2: (a) Lambert-Beer’s low based model [18]. (b) General KM model, assuming simple substrate. (c) KM model for our study, assuming heterogeneous substrate.

research has emerged that combines hyperspectral data and deep learning to perform pigment mapping of cultural heritage [9, 13, 3].

3. Kubelka-Munk model

The Kubelka-Munk model [15] is predicated on light attenuation due to scattering and absorption within the pigment layer, considering two fluxes: incident and reflected directions (Figure 2(b)(c)). The two-flux attenuation model parameters include a scattering coefficient S and an absorption coefficient K , each of which indicates the degree to which light gets attenuated by either absorption or scattering. We use i and j to represent the diffuse light flux in the top-to-bottom (incident) and bottom-to-top (reflected) directions within the pigment layer, respectively. We express the change in i and j for an infinitesimal thickness dx as follows.

$$di = (S + K)idx - Sjd, \quad (1)$$

$$dj = -(S + K)jdx + Sid. \quad (2)$$

Solving the differential equations of this two-flux model allows us to determine the reflectance R_{km} of layered-surface objects:

$$R_{km} = \frac{\frac{1}{R_\infty}(R_b - R_\infty) - R_\infty(R_b - \frac{1}{R_\infty})e^{SX(\frac{1}{R_\infty} - R_\infty)}}{(R_b - R_\infty) - (R_b - \frac{1}{R_\infty})e^{SX(\frac{1}{R_\infty} - R_\infty)}} = KM(X, R_\infty, S, R_b), \quad (3)$$

where X signifies the thickness of the pigment layer, R_b denotes the substrate reflectance, and R_∞ represents the reflectance when thickness $X = \infty$. As stated by Kubelka [14], we can obtain the scattering coefficient S and absorption coefficient K for each pigment from measurements. Moreover, we can compute R_∞ using these coefficients.

$$R_\infty = 1 + \frac{K}{S} - \sqrt{\left(\frac{K}{S}\right)^2 + 2\frac{K}{S}} \quad (4)$$

Note that the thickness X of the pigment layer represents the relative optical path length, and as such, it doesn't have any units.

4. Neural pigment representation

4.1. Problem setting

Pigment layer: Existing studies suggest that specific types of pigments have been identified, and the pigments used in tombs are not mixed [10]. This introduces the need to undertake a two-pronged estimation process. Firstly, we must identify which pigment is used at each point on the mural. Secondly, we need to estimate the thickness of the pigment application. Within the context of the KM model variables, this study aims to classify the values of S and R_∞ , and estimate the value of X .

Substrate layer: The substrate layer, comprised of a variety of minerals, allows us to theoretically classify the substrate's spectrum based on the number of minerals present in the rock. In this study, we use K-means clustering on the spectral data from parts of the substrate where no pigment was applied. This method will assist in determining the precise color of the rock at each location on the mural.

4.2. MLP network

Figure 3 provides an overview of our model. To accomplish pigment mapping on a heterogeneous substrate, we utilize a network architecture similar to NeRF [17], wherein the network input comprises coordinates from the input image to account for spatial continuity.

We use the 2D position of the hyperspectral image $\mathbf{p} = (x, y)$, which are normalized to $[-1, 1]$, and the spectral reflectance at the position $R_{km}(\mathbf{p}) \in \mathbb{R}^\Lambda$ as the network inputs, where Λ represents the number of bands. Our network generates the following outputs: the estimated thickness of the pigment layer $\hat{X} \in \mathbb{R}$, the probability of each pigment class $C^n \in [0, 1]$, ($n = 1, 2, \dots, N$), and the probability of each substrate class $B^m \in [0, 1]$, ($m = 1, 2, \dots, M$).

Please note that $\sum_{n=1}^N C^n = 1$ and $\sum_{m=1}^M B^m = 1$. For ensuring non-negative output, we use a ReLU as the activation function for \hat{X} , and for achieving maximally sparse outputs, we apply Sparsemax [16] to C^n and B^m .

Based on the above, the MLP model F_Θ can be defined as follows:

$$F_\Theta : (\mathbf{p}, R_{km}(\mathbf{p})) \mapsto (\hat{X}, C^m, B^m), \quad (5)$$

where Θ represents the parameters of the network.

By using this defined deep learning model F_Θ , we can estimate the parameters of the pigment simultaneously accounting for the substrate.

4.3. Positional encoding

According to NeRF [17], mapping inputs to a higher-dimensional space with a high-frequency function before inputting them into the network improves the fit for high-frequency data. We adhere to this principle by embedding \mathbf{p} within a high-frequency function as

$$PE(\mathbf{p}) = (\sin(2^0 \pi \mathbf{p}), \cos(2^0 \pi \mathbf{p}), \dots, \sin(2^{L-1} \pi \mathbf{p}), \cos(2^{L-1} \pi \mathbf{p})). \quad (6)$$

This formula enables embedding a 2-dimensional vector \mathbf{p} into a $2 \times 2L$ space. For this study, we expand to 24 dimensions with $L = 6$.

4.4. Hidden layer

We combine the two kinds of inputs into a single vector and pass this combined vector through seven fully connected layers with ReLU activation (comprising 300 perceptrons). Additionally, we implement a skip connection linking the input \mathbf{p} , embedded in a $2 \times 2L$ space, to the activation of the fourth layer, as suggested by NeRF [17] and DeepSDF [19].

4.5. Loss using the Kubelka-Munk model

We calculate our loss as the mean squared error between the estimated \hat{R}_{km} and the actual R_{km} used as input, following the method by Shitomi *et al.* [22]. Using the model outputs, we determine the estimated \hat{R}_{km} . During the learning step, the model computes S , R_∞ , and R_b as a weighted sum with the probability of each pigment C^n to maintain differentiability:

$$S(C) = \sum_{n=1}^N C^n S^n \quad (7)$$

$$R_\infty(C) = \sum_{n=1}^N C^n R_\infty^n \quad (8)$$

$$R_b(B) = \sum_{m=1}^M B^m R_b^m, \quad (9)$$

where $S^n \in \mathbb{R}^\Lambda$ and $R_\infty^n \in \mathbb{R}^\Lambda$ are the scattering coefficient and the reflectance of n -th pigment class, and $R_b^m \in \mathbb{R}^\Lambda$ is the spectral reflectance of m -th substrate class. Substituting Eqs. (7)–(9) to Eq. (3), we obtain \hat{R}_{km} as

$$\hat{R}_{km} = KM(\hat{X}, S(\hat{C}), R_\infty(\hat{C}), R_b(\hat{B})). \quad (10)$$

The mean squared error function is

$$\mathcal{L} = \frac{1}{\Lambda} \sum_{\mathbf{p} \in \mathcal{R}} \|\hat{R}_{km}(F_\Theta(\mathbf{p}, R_{km}(\mathbf{p}))) - R_{km}(\mathbf{p})\|_2^2, \quad (11)$$

where \mathcal{R} represents the set of coordinates in each batch. By back-propagating the error through the KM model, we can calculate the gradients of the weights.

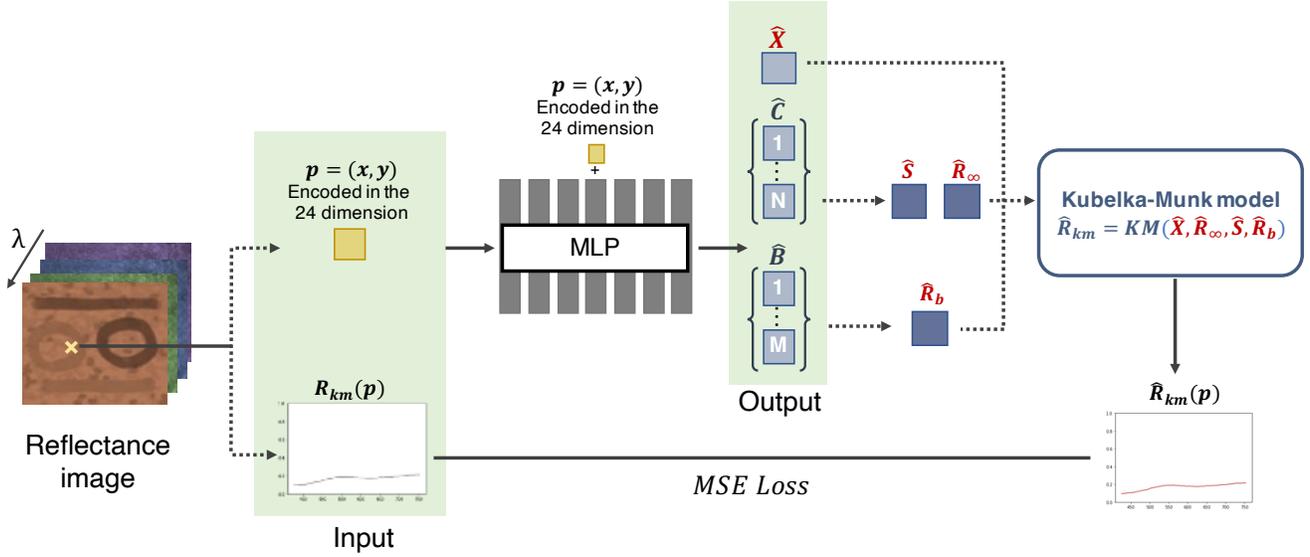


Figure 3: An overview of our deep learning model: We input the position of a reflectance image and its corresponding reflectance (R_{km} in the KM model) into the MLP. The MLP’s output is used to determine the variables of the KM model. As the KM model is differentiable, we can conduct learning by minimizing the error between the estimated reflectance (\hat{R}_{km}) and the input reflectance (R_{km}) used in the process.

4.6. Pigment mapping

In the testing phase, we define pigment mapping as the product of the pigment class and its thickness, calculated for each region. We select the pigment class as $\text{argmax } C^n$. Even though the restriction of $\sum_{n=1}^N C^n = 1$ leads to estimating the pigment class in the area without pigment, this doesn’t cause any issues if the thickness estimation correctly results in zero.

5. Experiment results of simulation data

We conducted a simulation experiment to validate our pigment mapping method on heterogeneous substrates. Due to the challenge of obtaining a sample with a known thickness, we initially synthesized a spectral image for verification purposes.

5.1. Dataset and settings

Figure 4(a) displays an RGB visualization of the synthesized spectral data, encompassing 65 bands within the 430–760nm range. With an image resolution of 200×200 , we utilized 40,000 pixels for training. The experiment was configured with $N = 2$ pigment classes and $M = 3$ substrate classes.

The spectral data was synthesized based on actual measurements. For pigments, we used red (Bengara) and green (Terre Verte), both frequently found in ancient tomb paintings. We derived the S and R_∞ values (Figure 4(c),(d)) or these pigments using the method detailed in [14], which in-

involved measuring the reflectance of each pigment painted on both a white and a black substrate. Moreover, we determined the X value (Figure 4(b)) by solving the inverse problem of the KM model from paint applied to a single-color substrate with constant substrate parameters. This process was executed using the technique described in [22],

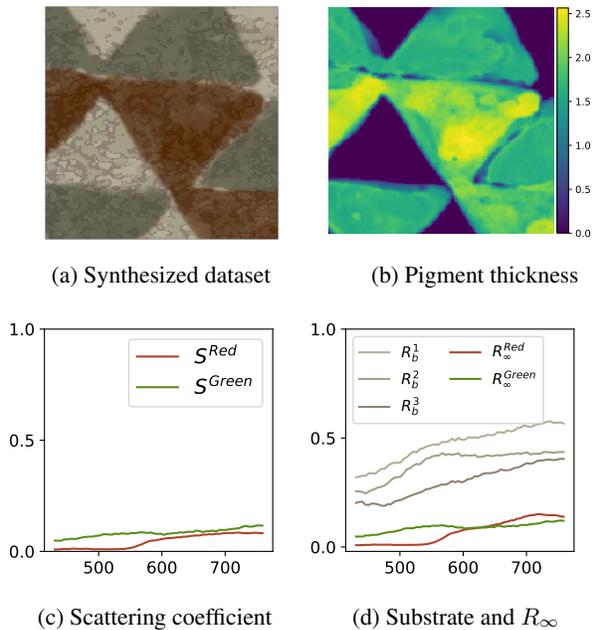


Figure 4: Overviews of synthesized data.

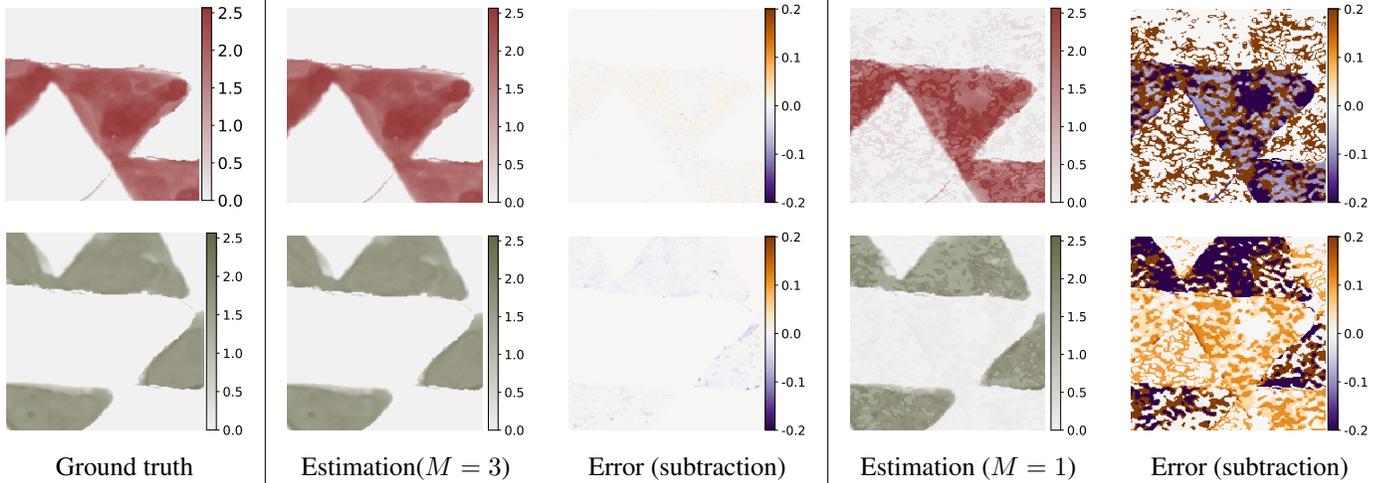


Figure 5: Estimation results of simulation data, presented under both a heterogeneous ($M = 3$) and homogeneous ($M = 1$) assumption.

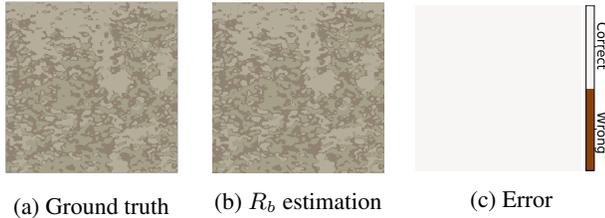


Figure 6: Estimation results of simulation data.

in which we analyzed the thickness of the pattern painted on white paper. For the substrate, we measured the spectral reflectance of a bare rock surface. Using k-means clustering, we identified patterns within the substrate. The spectral component of each class, representing the centroid of its respective cluster in k-means, was used as the reflectance of each class (Figure 4(d)). With the obtained S , R_{∞} , R_b , and X values for each region, we synthesized a spectral image using the KM model.

For our training regimen, we used the Adam optimizer with CosineAnnealingLR serving as a scheduler to adjust the learning rate. The learning rate started at 5×10^{-4} and gradually decreased to 5×10^{-7} during the optimization process. The training was conducted over 1,200 iterations, with CosineAnnealingLR completing one cycle every 25 iterations.

5.2. Results

Figure 5 displays the estimates produced by our pigment mapping. The pigment classification performs well, slightly underestimating the thickness but nonetheless accurately capturing the original thickness distribution. We computed the SSIM and PSNR between the ground truth

and estimated values for each pigment mapping. Our model achieved SSIM= 0.998, PSNR= 48.6 for the red pigment, and SSIM= 0.997, PSNR= 46.9 for the green pigment.

On the other hand, when the substrate was assumed to be homogeneous ($M = 1$), the effects of the substrate’s pattern were reflected, causing significant errors in the estimation. In this case, SSIM= 0.278, PSNR= 21.1 for the red pigment, and SSIM= 0.359, PSNR= 23.1 for the green pigment.

Figure 6 presents the results of the substrate estimation. In Figure 6(c), we use white to denote pixels where the class has been correctly estimated, and red to denote pixels where the class has been incorrectly estimated. This model succeeds in classifying the class correctly for all pixels.

5.3. Ablation study

Our model has $R_{km}(\mathbf{p})$, $PE(\mathbf{p})$ as inputs to the first layer, and $PE(\mathbf{p})$ as an input to the hidden layer. We have conducted an ablation study to validate our design choices. Table 1 shows each ablation design.

The accuracy was compared using SSIM and PSNR. Table 2 shows the result of the ablation study. Our method

Table 1: Design of ablation study.

Input Layer	$R_{km}(\mathbf{p})$ First	$PE(\mathbf{p})$ First	$PE(\mathbf{p})$ Hidden
<i>Abs-1</i>	—	✓	✓
<i>Abs-2</i>	✓	✓	—
<i>Abs-3</i>	✓	—	✓
<i>Abs-4</i>	✓	—	—
<i>Ours</i>	✓	✓	✓

Table 2: Result of ablation study.

	Red pigment map		Green pigment map	
	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow	PSNR \uparrow
<i>Abs-1</i>	0.897	26.1	0.767	20.9
<i>Abs-2</i>	0.981	42.5	<u>0.986</u>	<u>43.8</u>
<i>Abs-3</i>	<u>0.988</u>	<u>43.1</u>	0.923	32.8
<i>Abs-4</i>	0.785	19.8	0.778	22.0
<i>Ours</i>	0.998	48.6	0.997	46.9

performed the best score for both SSIM and PSNR. This result shows it is essential to input both reflectance $R_{km}(p)$ and its position $PE(p)$. Furthermore, by comparing *Ours* with *Abs-2* and *Abs-3*, the effectiveness of including coordinates twice has been confirmed as presented in previous studies [17, 19].

Figure 7 compares the original and *Abs-4* errors. *Abs-4* is a pixel-by-pixel method with no position input, and the substrate pattern tends to show up more as an error. The position input allows spatial continuity to be taken into account, and the influence of the substrate pattern can be eliminated more. These results indicate that a model based on a neural representation is valid.



(a) Error of *Ours* (b) Error of *Abs-4* (c) Substrate

Figure 7: Comparison of pigment mapping error between *Ours* and *Abs-4*. *Abs-4* is affected by the substrate pattern.

6. Experimental results of real tomb data

6.1. Dataset and settings

We used a spectral image of the Mezurashiduka tomb, located in western Japan. The tomb’s visual representation combines two kinds of pigments ($N = 2$), red (Ben-gara) and gray (Hekikaimatsu), applied over a granite-rock surface (Figure 9(a)). This tombs’ dataset is a collection of spectral images with 81 bands, ranging from 400nm to 720nm. With an image resolution of 420×700 , we used 294,000 pixels for training. Since granite is composed mainly of six minerals, the substrate class was set as $M = 6$ in this study.

In real data, it is difficult to prepare pigments identical to those used in ancient murals. Therefore, S and R_∞ values for each pigment were determined by annotating the pixels

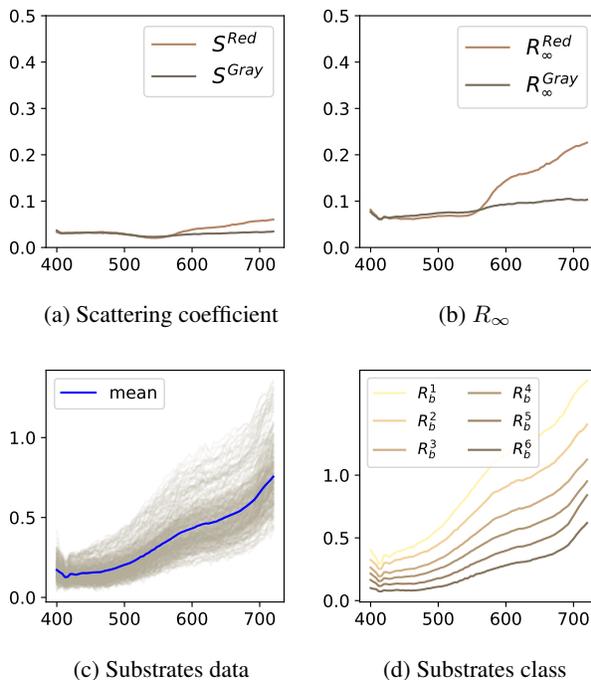


Figure 8: Optical parameters of Mezurashizuka tomb.

where the pigment is applied, and subsequently optimizing using the KM model. We used the reflectance of annotated pixels for the initial value for R_∞ . Also, we used the Adam optimizer to update S and R_∞ to minimize the discrepancy between the measured and modeled reflectance. The optimized S and R_∞ are shown in Figure 8(a) and (b), respectively.

For substrate, R_b was obtained by annotating regions without pigment, clustering approximately 3,000 reflectance data points (seen in Figure 8(c)) using K-means. We adopt the centroid as the reflectance of each substrate class (Figure 8(d)).

In training sessions, we employed the Optuna [2] hyperparameter optimization framework. The parameters we focused on were the batch size and the learning rate. The batch size was selected from 512, 1024, 2048, and 4096. On the other hand, the learning rate was searched within a range from 1×10^{-4} to 5×10^{-4} . The number of training sessions was set to 500, with CosineAnnealingLR completing one cycle every 25 iterations.

6.2. Results

We applied our proposed model to actual spectral data obtained from tomb murals, assessing its ability to estimate. For validation purposes, we referred to the restoration sketches produced by Hakkou Kusaka, a noted Nihonga artist (Figure 9(e)). Furthermore, Figure 9(d) and (h) rep-

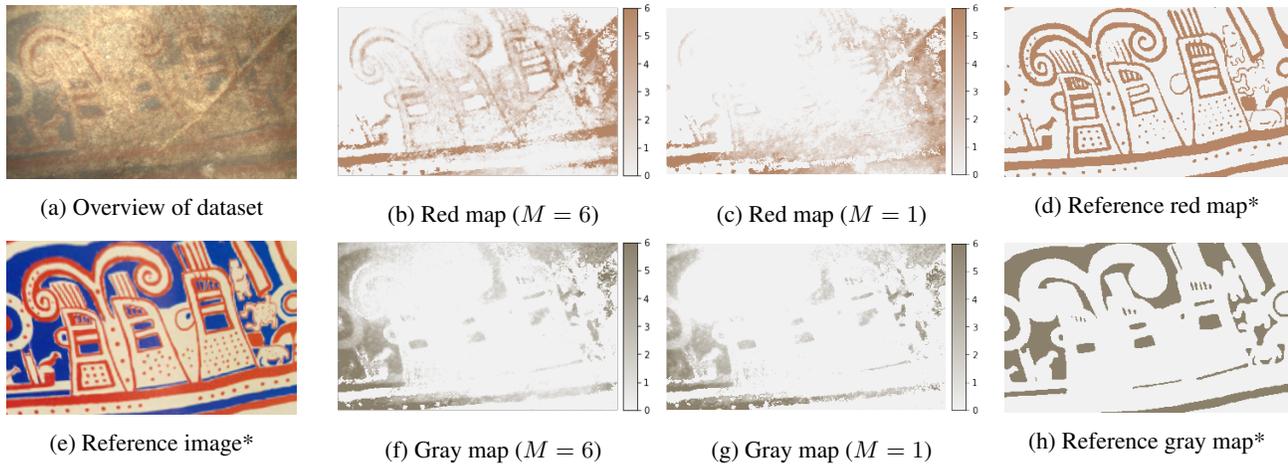


Figure 9: Estimation results for the Mezurashizuka tomb, presented under both a heterogeneous ($M = 6$) and homogeneous ($M = 1$) assumption. The entries marked with an (*) are derived from the restoration reproductions carried out by Hakkou Kusaka in the 1950s. Please note that these reproductions were a subjective task and should not be strictly considered ground truth.

resent the segmentation of Hakkou’s reference image using the K-means method, displaying the regions of each pigment. Commissioned by the Agency for Cultural Affairs between 1953 and 1955, Kusaka created these sketches. Although his artwork is based on subjectivity, it is a valid benchmark against which we can evaluate the accuracy of our model estimates.

The estimation results are included in the Figure 9 (b) and (f). It can be seen that our method captures the tendency of pigment distribution well.

Our model estimates that pigment is applied even in areas not painted on the reference image, and there are three possible interpretations for this. The first factor is that the model might have detected the presence of pigment that has deteriorated and washed away over time. The second one could be a drawback associated with the fact that our research is based on spectral information. Depending on the combination of the substrate and the pigment, it is possible that a spectrum similar to another substrate class could have been reproduced. The third factor pertains to the influ-

ence of the data used for estimation. This includes the error during the optimization of S and R_∞ , the decreased expressiveness due to the classification of the substrate layer, the impact of illumination variations in the spectral image, and the noise when capturing a spectral image. It is reasonable to view our estimation results as a combination of the three phenomena mentioned.

Figure 10 demonstrates the outcome of applying the Difference of Gaussians (DoG) to our estimation results. Given that red pigments are often utilized for delineation, their corresponding values were set to accentuate the edges. Conversely, the grey pigments, frequently used for area-filling expressions, had their values adjusted to encourage smoothing. These processed images illustrate the potential for estimating mural patterns even when past reference images are not available. Also, referring to Hakkou’s inferred figure, the results of this study provide a clue as to how the tumulus wall paintings deteriorated.

6.3. Comparison with homogeneous substrate

The substrate of the rock on which the murals are painted is inherently heterogeneous. However, for the sake of comparison, we also examined the case where we assume the substrate to be homogeneous. When treating the substrate as homogeneous, we adopted the average value of the substrate reflectance (the blue-line in Figure 8(c)) as the substrate spectrum.

The assumption of the homogeneous substrate simplified our modeling process, removing the necessity to classify R_b , which then becomes a constant. The results are shown in Figure 9(c) and (g). Our analysis reveals that when we assume the substrate to be constant, the estimation results exhibit more errors as opposed to when we consider the

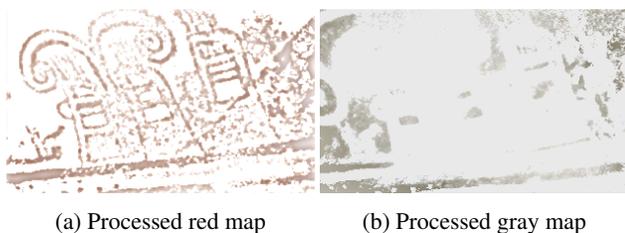


Figure 10: Results of applying the Difference of Gaussian method to our pigment mapping estimation, which was derived from Figure 9(b) and (f).

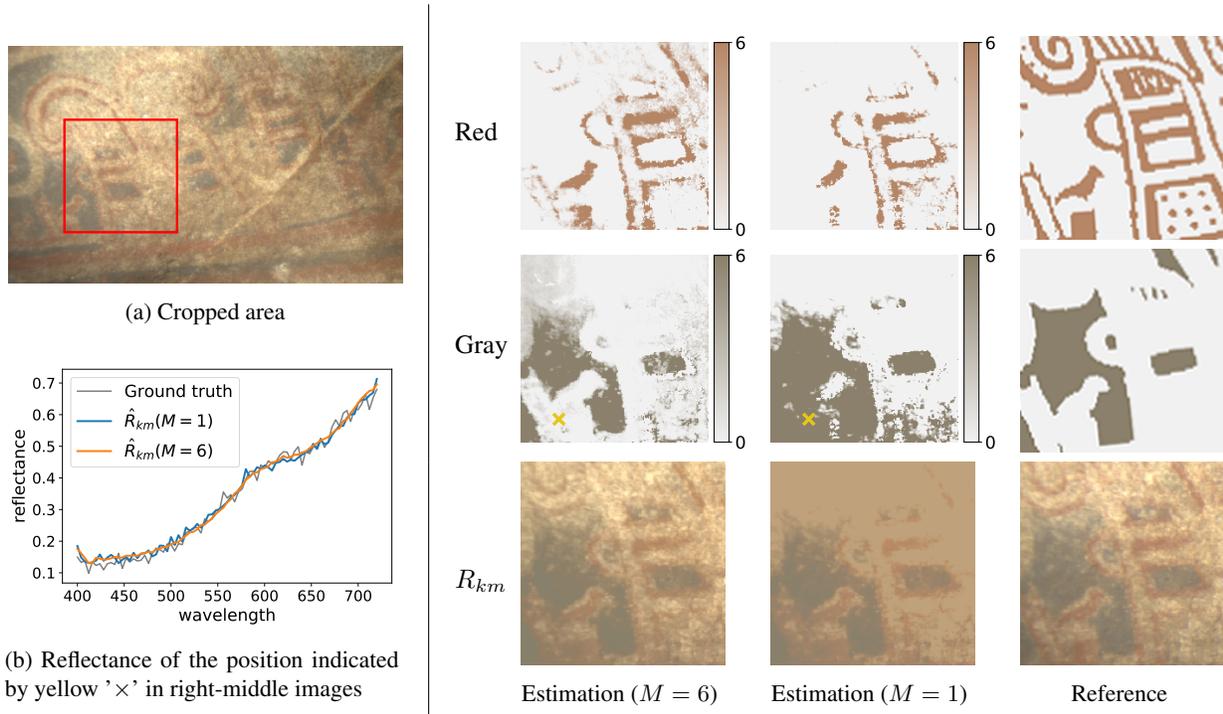


Figure 11: Estimation results in a 200×200 region. We are comparing the case where the substrate is assumed to be heterogeneous (substrate class $M = 6$) and the case where it is assumed to be homogeneous ($M = 1$).

substrate to be heterogeneous. Especially, estimation of red pigment is getting worse.

Furthermore, we have discovered that the results can deteriorate depending on how the images in the dataset are cropped, under the assumption of a homogeneous substrate. The frequency of pattern changes in the dataset is associated with the size of the image area. Specifically, the smaller the image area, the higher the frequency of these pattern changes, especially when cropping occurs in regions with substantial pattern variations. Figure 11 shows estimation results when cropping a 200×200 area from the original dataset. The assumption of a homogeneous substrate clearly reduces the accuracy.

This is due to the fact that assuming a constant substrate reduces the expression of the model and due to physics-based estimation. Figure 11(b) is a plot of the reflectance data at the coordinates indicated by the 'x' marker in the right-middle of Figure 11. Despite the clear inaccuracies in pigment estimation assuming a homogeneous substrate, the reconstructed reflectance \hat{R}_{km} is found to be similar to the original reflectance R_{km} . Such errors are unavoidable as long as reflectance is the basis for estimation, but the results indicate that assuming a heterogeneous substrate and increasing the representational capacity of the model can partially mitigate these issues. These observations underscore the importance of accurately capturing the substrate's heterogeneity in our analyses.

7. Conclusion and Future work

This study realizes the pigment mapping for an ancient tomb mural. Our approach uses unsupervised learning that accounts for the unknown substrate pattern by simultaneously estimating the pigment class, thickness, and substrate class. The experimental results indicate the efficacy of concurrently estimating both the substrate and pigments without disregarding the fact that the substrate is heterogeneous. Moreover, as the ablation study shows, by inputting coordinates and estimation by neural representation, the model takes spatial continuity into account and estimates with less influence on the substrate pattern.

However, this study presents several limitations and future works. While many murals are depicted on relatively smooth rocks, some are located on rocks with significant irregularities. Our current estimation does not account for the three-dimensional structure of murals, potentially resulting in errors when mapping pigments on uneven surfaces, as our future work. Furthermore, we have based our real data on directly obtained pigment parameters from the mural. That is under the assumption that the parameters of the pigments are enough painted at some pixels. Addressing the spectral changes in the mural sections where the pigments have significantly washed out or faded, or have oxidized (undergone chemical changes), is also part of our future work.

Acknowledgement This work was partly supported by JSPS KAKENHI 23H00499 and JST PRESTO JPMJPR2025.

References

- [1] Elad Aharoni-Mack, Yakov Shambik, and Dani Lischinski. Pigment-based recoloring of watercolor paintings. In *Proceedings of the Symposium on Non-Photorealistic Animation and Rendering*, pages 1–11, 2017.
- [2] Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 2623–2631, 2019.
- [3] Ailin Chen, Rui Jesus, and Márcia Vilarigues. Identification and visualization of pure and mixed paint pigments in heritage artwork using machine learning algorithms. *SN Computer Science*, 4(2):115, 2022.
- [4] Antonino Cosentino. Identification of pigments by multi-spectral imaging; a flowchart method. *Heritage Science*, 2(1):8, 2014.
- [5] Hilda Deborah, Sony George, and Jon Yngve Hardeberg. Pigment mapping of the scream (1893) based on hyperspectral imaging. In *Image and Signal Processing: 6th International Conference, ICISP 2014, Cherbourg, France, June 30–July 2, 2014. Proceedings 6*, pages 247–256. Springer, 2014.
- [6] John K Delaney, Jason G Zeibel, Mathieu Thoury, ROY Littleton, Michael Palmer, Kathryn M Morales, E René de La Rie, and ANN Hoenigswald. Visible and infrared imaging spectroscopy of picasso’s harlequin musician: mapping and identification of artist materials in situ. *Applied spectroscopy*, 64(6):584–594, 2010.
- [7] Daniel W Dichter. Kubelka-munk model of full-gamut oil colour mixing. *Journal of the International Colour Association*, 32:70–78, 2023.
- [8] Kensuke Fukumoto, Norimichi Tsumura, and Roy Berns. Estimating pigment concentrations from spectral images using an encoder-decoder neural network. *Journal of Imaging Science and Technology*, 64(3):30502–1, 2020.
- [9] Bartosz Grabowski, Wojciech Masarczyk, Przemysław Głomb, and Agata Mendys. Automatic pigment identification from hyperspectral data. *Journal of Cultural Heritage*, 31:1–12, 2018.
- [10] Katsushi Ikeuchi, Tetsuro Morimoto, Mawo Kamakura, Nobuaki Kuchitsu, Kazutaka Kawano, and Tomoo Ikeda. Kyushu decorative tumuli project: from e-heritage to cyber-archaeology. *International Journal of Computer Vision*, 130(7):1609–1626, 2022.
- [11] Eric Kirchner, Ivo van Der Lans, Frank Ligterink, Muriel Geldof, Art Ness Proano Gaibor, Ella Hendriks, Koen Janssens, and John Delaney. Digitally reconstructing van gogh’s field with irises near arles. part 2: pigment concentration maps. *Color Research & Application*, 43(2):158–176, 2018.
- [12] Tania Kleynhans, David W Messinger, and John K Delaney. Towards automatic classification of diffuse reflectance image cubes from paintings collected with hyperspectral cameras. *Microchemical Journal*, 157:104934, 2020.
- [13] Tania Kleynhans, Catherine M Schmidt Patterson, Kathryn A Dooley, David W Messinger, and John K Delaney. An alternative approach to mapping pigments in paintings with hyperspectral reflectance image cubes using artificial intelligence. *Heritage Science*, 8(1):1–16, 2020.
- [14] Paul Kubelka. New contributions to the optics of intensely light-scattering materials. part i. *Josa*, 38(5):448–457, 1948.
- [15] Paul Kubelka and Franz Munk. An article on optics of paint layers. *Z. Tech. Phys*, 12(593-601):259–274, 1931.
- [16] Andre Martins and Ramon Astudillo. From softmax to sparsemax: A sparse model of attention and multi-label classification. In *International conference on machine learning*, pages 1614–1623. PMLR, 2016.
- [17] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021.
- [18] Tetsuro Morimoto, Robby T Tan, Rei Kawakami, and Katsushi Ikeuchi. Estimating optical properties of layered surfaces using the spider model. In *2010 IEEE computer society conference on computer vision and pattern recognition*, pages 207–214. IEEE, 2010.
- [19] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 165–174, 2019.
- [20] Roxanne Radpour, Glenn A Gates, Ioanna Kakoulli, and John K Delaney. Identification and mapping of ancient pigments in a roman egyptian funerary portrait by application of reflectance and luminescence imaging spectroscopy. *Heritage Science*, 10(1):1–16, 2022.
- [21] Neda Rohani, Emeline Pouyet, Marc Walton, Oliver Coscairt, and Aggelos K Katsaggelos. Nonlinear unmixing of hyperspectral datasets for the study of painted works of art. *Angewandte Chemie*, 130(34):11076–11080, 2018.
- [22] Ryuta Shitomi, Mayuka Tsuji, Yuki Fujimura, Takuya Funatomi, Yasuhiro Mukaigawa, Tetsuro Morimoto, Takeshi Oishi, Jun Takamatsu, and Katsushi Ikeuchi. Unsupervised learning with a physics-based autoencoder for estimating the thickness and mixing ratio of pigments. *Journal of the Optical Society of America. A, Optics, Image Science, and Vision*, 40(1):116–128, 2023.
- [23] Abu Md Niamul Taufique and David W Messinger. Hyperspectral pigment analysis of cultural heritage artifacts using the opaque form of kubelka-munk theory. In *Algorithms, Technologies, and Applications for Multispectral and Hyperspectral Imagery XXV*, volume 10986, pages 297–307. SPIE, 2019.
- [24] Yonghui Zhao, Roy S Berns, Lawrence A Taplin, and James Coddington. An investigation of multispectral imaging for the mapping of pigments in paintings. In *Computer image analysis in the study of art*, volume 6810, pages 65–73. SPIE, 2008.