

Object Color Relighting with Progressively Decreasing Information

Luca Cogo¹, Marco Buzzelli¹, Simone Bianco¹ and Raimondo Schettini¹

¹Department of Informatics, Systems and Communication (DISCo), University of Milano-Bicocca, Milan, Italy

E-mail: luca.cogo@unimib.it

Abstract. Object color relighting, the process of predicting an object's colorimetric values under new lighting conditions, is a significant challenge in computational imaging and graphics. This technique has important applications in augmented reality, digital heritage, and e-commerce. In this paper, we address object color relighting under progressively decreasing information settings, ranging from full spectral knowledge to tristimulus-only input. Our framework systematically compares physics-based rendering, spectral reconstruction, and colorimetric mapping techniques across varying data regimes. Experiments span five benchmark reflectance datasets and eleven standard illuminants, with relighting accuracy assessed via ΔE_{00} metric. Results indicate that third-order polynomial regressions give good results when trained with small datasets, while neural spectral reconstruction achieves superior performance with large-scale training. Spectral methods also exhibit higher robustness to illuminant variability, emphasizing the value of intermediate spectral estimation in practical relighting scenarios.

1 Introduction

The appearance of surface colors depends not only on the intrinsic reflectance properties of the materials but also on the spectral characteristics of the illumination. This dependency poses a central challenge in color science and imaging applications: how to predict or estimate the color of a surface when the illuminant changes. Accurate modeling of this relighting process is fundamental for tasks such as color correction, image rendering, and visual perception modeling.

This work addresses the relighting problem of surface colors under progressively decreasing levels of available information. We begin with the most complete scenario, where both the spectral reflectance of the object and the Spectral Power Distributions (SPDs) of the source and target illuminants are known. In such cases, relighting can be formulated via physically-based rendering equations. However, in most practical scenarios, spectral information is either unavailable or too costly to acquire. As information becomes more limited, the relighting task must be reformulated using only partial data, such as colorimetric descriptors (e.g., CIE XYZ values) measured under a given illuminant. We therefore consider how to estimate the colorimetric output under a new illuminant by only knowing the XYZ values under a source illuminant with a known SPD.

This descending perspective, from full spectral knowledge to purely tristimulus-based relighting, reveals the inherent trade-offs between data availability, model complexity, and achievable accuracy. As we descend this information ladder, we explore corresponding computational models, ranging from physics-based rendering to chromatic adaptation transforms and data-driven methods.

We survey and experimentally compare these approaches using several benchmark reflectance datasets (including Munsell Color Chips [1], Macbeth ColorChecker [2], and real-object spectra [3, 4, 5]) under



widely used illuminants [5]. Our goal is to clarify what can be gained (or must be sacrificed) as relighting is attempted with less information. By organizing the problem along this decreasing-information axis, we offer a unified framework for evaluating relighting techniques and contribute to the development of practical pipelines for color correction in systems lacking spectral input.

2 Object Color Relighting Methods

In practical settings, object color relighting refers to the task of predicting an object's color appearance, typically in the CIE XYZ color space, under a target illuminant, given observations of that object under a source illuminant. In the following subsections, we describe and compare various techniques for addressing the task. We begin with the most informed scenario, in which full spectral information for both the object and the source illuminant is available. Next, we consider the case where only the source illuminant's Spectral Power Distribution is known, and the object's spectrum must be estimated from its CIE XYZ values. Finally, we examine the scenario in which only the XYZ values of both the object and the source illuminant are available, and the mapping must be performed in the tristimulus space using chromatic adaptation transforms or learning-based methods.

2.1 Full spectral rendering

When full spectral information of both the object and the source illuminant is available, relighting can be achieved through physics-based rendering. Given the surface reflectance function $R(\lambda)$ and the Spectral Power Distribution of the source illuminant $E(\lambda)$, the reflected radiance L at wavelength λ is given by:

$$L(\lambda) = E(\lambda)R(\lambda). \quad (1)$$

To compute colorimetric values, this radiance is integrated with the standard observer color matching functions $S(c, \lambda)$ (e.g., CIE 1931 XYZ):

$$I(c) = \int_{\lambda_{\min}}^{\lambda_{\max}} E(\lambda)R(\lambda)S(c, \lambda) d\lambda. \quad (2)$$

Under this model, if $L(\lambda)$ and $E(\lambda)$ are known (properly sampled or interpolated), one can recover the surface reflectance $R(\lambda)$, and thereby re-render it under any desired target illuminant. In theory, access to full spectral data enables perfect relighting.

However, acquiring spectral measurements is often impractical due to equipment cost, limited portability, or other application-specific constraints. Consequently, relighting must often be performed using only tristimulus measurements. In this case, relighting is formulated as a mapping between CIE XYZ values under different illuminants:

$$\begin{bmatrix} X^{(t)} \\ Y^{(t)} \\ Z^{(t)} \end{bmatrix} = \phi \left(\begin{bmatrix} X^{(s)} \\ Y^{(s)} \\ Z^{(s)} \end{bmatrix}, \begin{bmatrix} X_w^{(s)} \\ Y_w^{(s)} \\ Z_w^{(s)} \end{bmatrix}, \begin{bmatrix} X_w^{(t)} \\ Y_w^{(t)} \\ Z_w^{(t)} \end{bmatrix} \right) \quad (3)$$

where $[X^{(s)}Y^{(s)}Z^{(s)}]^T$ and $[X^{(t)}Y^{(t)}Z^{(t)}]^T$ are the tristimulus values under source and target illuminants, and $[X_w^{(*)}Y_w^{(*)}Z_w^{(*)}]^T$ denotes the white point of the respective illuminants. In the following subsections, we describe various techniques for estimating this mapping.

2.2 Relighting Through Spectral Reconstruction

In cases where the Spectral Power Distribution of the source illuminant is known but only the CIE XYZ values of the object are available, it is possible to leverage a reflectance dataset to learn a mapping that reconstructs the surface's spectral reflectance $R(\lambda)$ from its tristimulus values. This inverse problem is inherently ill-posed, as multiple reflectance spectra can produce the same XYZ values. However, with appropriate priors or learned models, approximate reconstruction becomes feasible.

A common method assumes that reflectance spectra lie in a low-dimensional subspace spanned by basis functions obtained via Principal Component Analysis (PCA) of a reflectance dataset [3, 6]. The spectral radiance can then be approximated as:

$$L(\lambda) = \begin{bmatrix} X^{(s)} \\ Y^{(s)} \\ Z^{(s)} \end{bmatrix} MB + b \quad (4)$$

where B is a matrix of basis functions, b is a bias term (e.g., the mean spectrum of the dataset), and M is a transformation matrix learned via least-squares optimization to best reproduce training spectra.

A nonlinear variant employs a shallow neural network [7] to directly regress the radiance spectrum from XYZ inputs:

$$L(\lambda) = f^{(L)} \left(W^{(L)} f^{(L-1)} \left(\dots f^{(1)} \left(W^{(1)} \begin{bmatrix} X^{(s)} \\ Y^{(s)} \\ Z^{(s)} \end{bmatrix} \right) + b^{(1)} \right) \dots + b^{(L)} \right) \quad (5)$$

Once $L(\lambda)$ is estimated, the surface reflectance can be computed using the known source illuminant:

$$R(\lambda) = \frac{L(\lambda)}{E^{(s)}(\lambda) + \epsilon} \quad (6)$$

Subsequently, the relit tristimulus values under a target illuminant can be computed by:

$$\begin{bmatrix} X^{(t)} \\ Y^{(t)} \\ Z^{(t)} \end{bmatrix} = \phi_{\text{spec} \rightarrow \text{XYZ}} \left(R(\lambda) E^{(t)}(\lambda) \right) \quad (7)$$

where $\phi_{\text{spec} \rightarrow \text{XYZ}}(\cdot)$ denotes the rendering function defined in Equation 2.

Alternatively, some models aim to directly estimate $R(\lambda)$ from XYZ values without first recovering $L(\lambda)$ [8, 9]. In such cases, the model implicitly accounts for the source illuminant, enabling spectral relighting under arbitrary target illumination.

Several pixel-based spectral recovery methods originally developed to reconstruct reflectance spectra from RGB values could, in principle, be adapted to the scenario addressed in this work, where only a single XYZ triplet under a known illuminant is available. Techniques such as sparse coding (e.g., A+ [10] and A++ [11]) and learning-based regression models reviewed in recent surveys [12, 13] are inherently pixel-based and could be reparameterized to use XYZ inputs instead of RGB, leveraging the known spectral power distribution of the source illuminant. However, the intrinsic ambiguity in recovering reflectance from tristimulus values remains a key limitation, and the added complexity may not yield significantly improved relighting accuracy in the single-color case. For these reasons, we do not investigate such adaptations further in this paper, though they represent an interesting avenue for future exploration.

2.3 Relighting in Tristimulus Space

In scenarios where only tristimulus values of the object and the source illuminant are available, spectral reconstruction is not feasible. Instead, relighting must be performed directly in the tristimulus space. This section presents two classes of approaches suited to this setting: analytical methods based on Chromatic Adaptation Transforms (CATs), and data-driven methods that learn a mapping from source to target conditions.

2.3.1 Chromatic Adaptation Transforms Chromatic Adaptation Transforms (CATs) are widely used to approximate human color constancy by adjusting for changes in illumination [14]. These methods typically transform XYZ values into a cone response domain (e.g., LMS), apply a diagonal scaling to adapt to the new illuminant, and transform back to the XYZ space. The general form of a CAT is:

$$\begin{bmatrix} X^{(t)} \\ Y^{(t)} \\ Z^{(t)} \end{bmatrix} = M_A^{-1} \begin{bmatrix} R_w^{(t)}/R_w^{(s)} & 0 & 0 \\ 0 & G_w^{(t)}/G_w^{(s)} & 0 \\ 0 & 0 & B_w^{(t)}/B_w^{(s)} \end{bmatrix} M_A \begin{bmatrix} X^{(s)} \\ Y^{(s)} \\ Z^{(s)} \end{bmatrix} \quad (8)$$

Here, M_A is a 3×3 matrix defining the transformation to the cone response space, while $[R_w^{(t)} \ G_w^{(t)} \ B_w^{(t)}]^T$ and $[R_w^{(s)} \ G_w^{(s)} \ B_w^{(s)}]^T$ are computed from the target and source illuminants by multiplying their XYZ tristimulus values $[X_w^{(t)} \ Y_w^{(t)} \ Z_w^{(t)}]^T$ and $[X_w^{(s)} \ Y_w^{(s)} \ Z_w^{(s)}]^T$ by M_A .

Different CATs use different matrices: for example, the von Kries [15] method uses the LMS space, while CAT02 [16], Bradford [17], and Bianco & Schettini [18] employ empirically optimized matrices. A simple variant is XYZ Scaling, where M_A is the identity matrix, applying direct scaling in XYZ space [14].

2.3.2 Regression from Tristimulus Values When a dataset of surface colors observed under multiple illuminants is available, supervised learning techniques can be employed to empirically learn the relighting transformation. These models aim to approximate the mapping between tristimulus values under the source and target illuminants [19, 20].

A linear approach uses a polynomial expansion $\phi_d(\cdot)$ of the input data followed by a learned linear transformation:

$$\begin{bmatrix} X^{(t)} \\ Y^{(t)} \\ Z^{(t)} \end{bmatrix} = M \phi_d(i)$$

with:

$$i = \left[X^{(s)} Y^{(s)} Z^{(s)} X_w^{(s)} Y_w^{(s)} Z_w^{(s)} X_w^{(t)} Y_w^{(t)} Z_w^{(t)} \right]^T \quad (9)$$

Here, M is a regression matrix learned via least-squares fitting, and d is the degree of polynomial expansion. For $d = 1$, this corresponds to applying only first order monomials.

Alternatively, a shallow neural network [7] can be trained to learn a nonlinear mapping:

$$\begin{bmatrix} X^{(t)} \\ Y^{(t)} \\ Z^{(t)} \end{bmatrix} = f^{(L)} \left(W^{(L)} f^{(L-1)} \left(\dots f^{(1)} \left(W^{(1)} i \right) + b^{(1)} \right) + \dots + b^{(L)} \right) \quad (10)$$

where i is the concatenated input vector, $W^{(l)}$ and $b^{(l)}$ are the weights and biases of layer l , and $f^{(l)}(\cdot)$ is the corresponding non-linear activation function. Training is typically performed via backpropagation using mean squared error loss.

In scenarios like the one considered in this work, where the source and target illuminants are fixed, the white point information can be implicitly encoded during training, allowing simplification of the input vector to just $i = [X^{(s)} Y^{(s)} Z^{(s)}]^T$.

3 Datasets

To evaluate the performance of surface color relighting methods under varying illumination conditions, we employed a diverse set of spectral reflectance datasets. These datasets are categorized into two groups based on their origin and structure: (i) *standardized reflectance datasets*, designed for perceptual uniformity, reproducibility, and benchmarking; and (ii) *measured real-world datasets*, comprising reflectance spectra captured from physical materials using spectrophotometry or spectral imaging. Additionally, we employed a set of illuminant Spectral Power Distributions (SPDs) [5] to render the reflectance spectra under different lighting conditions.

3.1 Standardized Reflectance Datasets

These datasets consist of reflectance spectra from color samples engineered for consistent and reproducible behavior across imaging conditions (see Figure 1). They are widely adopted in color science for algorithm evaluation and calibration.

- **Macbeth ColorChecker** [2]
Contains 24 standardized color patches commonly used in camera calibration and illuminant estimation. The reflectance spectra are well-characterized and span a representative gamut of natural and artificial colors.
- **Munsell Color Chips** [1]
Comprises 1,269 matte samples from the Munsell system, systematically varying in hue, value, and chroma. The spectral reflectances are smooth and uniformly sampled at 10 nm intervals over the visible range (400–700 nm), making them ideal for evaluating colorimetric algorithms.

3.2 Real-World Reflectance Datasets

These datasets contain reflectance spectra measured from real-world objects and materials. Their natural variability makes them suitable for testing the generalization ability of relighting methods.




Figure 1: xy-chromaticity diagrams of the considered reflectance datasets: Macbeth ColorChecker (a), Munsell Color Chips (b), Westland Dataset (c), Vrhel Natural Objects (d), and Zhang et al. Dataset. Each spectrum is rendered to xy values under uniform illuminant. For the Zhang et al. Dataset, we plot only 1 million randomly sampled values.

- **Vrhel Dataset** [3]
Includes 170 measured reflectance spectra from a variety of natural and artificial surfaces (e.g., vegetation, textiles, paint), sampled at 10 nm intervals. These spectra exhibit high variability and spectral structure.
- **Westland Dataset** [4]
Comprises 404 reflectance spectra acquired using a tele-spectroradiometer. Each spectrum was computed by dividing the measured radiance of an object by that of a white reference tile. The dataset primarily consists of natural samples such as leaves, petals, and bark.
- **Zhang's Large-Scale Reflectance Dataset** [5]
Contains approximately 25 million unique reflectance spectra, assembled from multispectral and hyperspectral images, as well as spectrophotometric measurements of real-world surfaces. The spectra are sampled over the visible spectrum (400–700 nm, 10 nm intervals) and span a wide variety of materials including foliage, skin, food, textiles, buildings, paper, artwork, and industrial products. The dataset was deduplicated from an initial pool of roughly 40 million spectra using quantization and filtering techniques to retain distinct samples. Owing to its density and diversity, this dataset is especially valuable for analyzing relighting performance under metameric ambiguity.

3.3 Illuminant Dataset

To simulate relighting under realistic viewing conditions, we used a set of eleven illuminants originally curated by Zhang et al. [5], see Figure 2. This set includes:

- **Standard illuminants:** A, D50, D65, D100, D150, and D200, representing a range of correlated color temperatures from tungsten to various daylight conditions.




Figure 2: Spectral power distribution and xy-chromaticity diagrams of the considered illuminants Spectral Power Distributions (SPDs)

- **Fluorescent sources:** F4, F8, and F11, characterized by spiky spectral power distributions typical of indoor environments.
- **LED illuminants:** Two SPDs representative of white LEDs commonly found in mobile phone displays.

This diverse collection reflects practical lighting scenarios relevant to both human vision and computational imaging systems. All illuminants were normalized such that the CIE Y tristimulus value equals 1 for a perfect reflecting diffuser, ensuring consistent luminance across simulations.

4 Experimental Setting

We evaluate relighting performance under diverse illumination and data conditions by combining multiple illuminant groups, relighting methods, and dataset pairings for training and test. Illuminants are categorized into: (i) natural (A, D50, D65, D100, D150, D200), (ii) artificial (F4, F8, F11, LED-1, LED-2), and (iii) a combined set (including both natural and artificial).

Relighting methods fall into two categories: tristimulus-based and spectral recovery-based approaches. The former includes:

- Diagonal XYZ scaling
- Chromatic Adaptation Transforms: Von Kries, Bianco & Schettini, CAT02
- Polynomial Regression with varying orders (1^{st} , 2^{nd} , 3^{rd} , root- 2^{nd} , root- 3^{rd})
- Shallow Neural Network Regression (XYZ-to-XYZ)

Spectral recovery-based methods estimate surface reflectance or radiance using:

- Principal Component Analysis (PCA) based reconstruction (3 or 6 basis functions)
- Shallow Neural Network for spectral estimation (XYZ-to-Spectrum)

For learning-based methods, experiments are conducted across three dataset pairings: (1) training on Munsell Color Chips and testing on Macbeth ColorChecker, (2) training on Westland Dataset and testing on Vrhel Dataset, and (3) training on a 90% split of the Zhang et al. Dataset with testing on the held-out 10%. Due to computational constraints, non-neural methods use a reduced training set of 500,000 samples from Zhang et al. Dataset. The considered neural networks have two hidden layers of size 64, resulting approximately in 5K parameters, and ReLU nonlinearities [21]. Training is performed for up to 100 epochs using the Adam optimizer [22] with a learning rate of 0.01 and early stopping with patience equal to 5 epochs. This setup enables consistent comparison across controlled and real-world conditions.

Table 1: ΔE_{00} results under natural illuminants (A, D50, D65, D100, D150, D200). Results are shown for the considered dataset pairings. The solid horizontal line separates non-learning-based from learning-based methods, while the dashed line distinguishes tristimulus-based methods from spectral recovery-based ones. Best results are shown in **bold**, second best are underlined, and third best are *italicized*.

Method	ΔE_{00} on Macbeth Dataset (Trained on Munsell)			ΔE_{00} on Vrhel Dataset (Trained on Westland)			ΔE_{00} on Zhang-test (Trained on Zhang-train)		
	Mean	Median	95-P	Mean	Median	95-P	Mean	Median	95-P
XYZ Scaling	4.33	2.94	12.98	4.30	3.39	10.62	3.74	2.84	9.84
Von Kries	2.62	1.70	8.93	2.06	1.37	6.47	2.27	1.50	7.11
Bianco & Schettini	1.95	1.43	5.90	2.14	1.63	5.83	1.76	1.26	4.90
CAT02	1.79	1.19	5.78	1.87	1.34	5.29	1.62	1.10	4.74
3×3 Correction Matrix	1.36	0.82	4.67	1.17	0.78	3.57	1.14	0.74	3.53
2° Ord. Polynomial	1.11	<i>0.49</i>	4.37	<i>1.06</i>	<i>0.66</i>	<i>3.36</i>	1.03	0.64	3.28
3° Ord. Polynomial	<u>0.98</u>	0.55	<u>3.64</u>	0.95	<u>0.57</u>	<u>3.05</u>	<i>0.94</i>	<i>0.56</i>	3.02
2° Ord. Root-Polynomial	<i>1.08</i>	<i>0.48</i>	<i>3.87</i>	1.21	0.76	3.76	0.99	0.62	3.11
3° Ord. Root-Polynomial	0.85	0.39	3.21	<u>0.97</u>	0.54	3.03	<u>0.87</u>	<u>0.52</u>	<u>2.82</u>
ShallowNN (5K)	9.48	5.82	26.01	12.04	10.77	26.77	1.06	0.78	<i>2.93</i>
PCA-Radiance (3 bases)	1.79	0.93	7.02	2.70	1.80	8.52	5.07	3.04	16.45
PCA-Radiance (6 bases)	1.33	0.83	4.55	1.18	0.80	3.56	1.15	0.76	3.56
PCA-Reflectance (3 bases)	1.45	0.96	4.57	2.50	1.95	6.78	1.46	1.11	3.85
PCA-Reflectance (6 bases)	1.33	0.82	4.53	1.16	0.80	3.44	1.14	0.74	3.53
ShallowNN-Radiance (5K)	2.67	1.50	7.88	3.75	2.42	11.64	0.82	0.49	2.64

5 Results and Analysis

The relighting performance of the evaluated methods is reported in Table 1, Table 2, and Table 3, which show the mean, median and 95-th percentile of ΔE_{00} errors under natural, artificial, and combined illuminant sets, respectively. Several observations emerge from these results.

When training is performed on low-cardinality datasets, such as Munsell Color Chips and Westland Dataset, the best-performing methods are those based on polynomial regression, particularly the third-order and root-third-order variants. These models are able to learn the transformation between source and target illuminants while maintaining robustness against overfitting. Spectral recovery methods, on the other hand, tend to underperform in this scenario, as the limited number of training samples does not provide enough spectral diversity to reconstruct reflectances with sufficient accuracy. Among learning-based approaches, neural networks generally perform poorly in these low-cardinality data regimes, especially when evaluated on a different dataset from the one used for training. However, the neural network-based spectral recovery model consistently outperforms its tristimulus-to-tristimulus counterpart. This is likely due to the implicit physical priors embedded in the spectral recovery process, which impose constraints that help guide learning and improve generalization.

In contrast, when training is conducted on a high-cardinality dataset such as the one from Zhang et al. [5], the neural network-based spectral recovery method emerges as the best overall method. The abundance and diversity of training samples allow the model to learn highly accurate mappings from tristimulus values to spectral data, and subsequently produce faithful relighting under novel illuminants. This performance gap highlights the data dependency of learning-based methods and underscores their potential when sufficient coverage of the reflectance space is available.

Another important trend concerns the relative robustness of spectral methods under different illuminant groups. When the illuminant set includes a mix of natural and artificial sources—introducing a wide range of spectral distributions and correlated color temperatures—spectral recovery methods experience a smaller degradation in accuracy compared to tristimulus-based techniques. This suggests that reconstructing the intermediate spectral representation provides a buffer against variability in illumination, as it allows the relighting to be grounded in physical reflectance properties rather than solely in empirical color transformations.

Finally, Chromatic Adaptation Transforms (CATs), while showing some robustness in conditions close to those they were designed for, consistently yield ΔE_{00} values above the perceptual threshold of 1.0. This indicates that although CAT-based methods offer a simple and interpretable approach, they lack the precision necessary for high-fidelity relighting, especially under illuminant conditions not explicitly

Table 2: ΔE_{00} results under artificial illuminants (F4, F8, F11, LED-1, LED-2). Results are shown for the considered dataset pairings. The solid horizontal line separates non-learning-based from learning-based methods, while the dashed line distinguishes tristimulus-based methods from spectral recovery-based ones. Best results are shown in **bold**, second best are underlined, and third best are *italicized*.

Method	ΔE_{00} on Macbeth Dataset (Trained on Munsell)			ΔE_{00} on Vrhel Dataset (Trained on Westland)			ΔE_{00} on Zhang-test (Trained on Zhang-train)		
	Mean	Median	95-P	Mean	Median	95-P	Mean	Median	95-P
XYZ Scaling	3.56	3.22	8.85	2.85	2.60	5.82	2.73	2.37	6.01
Von Kries	3.12	2.48	8.45	3.04	2.74	6.82	2.95	2.54	6.96
Bianco & Schettini	2.44	1.94	6.29	2.40	1.99	5.39	2.41	1.99	5.55
CAT02	2.49	1.99	6.56	2.46	2.11	5.37	2.48	2.10	5.69
3×3 Correction Matrix	1.56	1.13	4.21	1.72	1.14	5.04	1.21	0.89	3.16
2° Ord. Polynomial	<i>1.32</i>	0.87	<i>3.88</i>	1.62	1.14	4.44	1.13	0.82	2.99
3° Ord. Polynomial	1.10	<u>0.78</u>	3.13	1.20	<u>0.82</u>	3.55	1.09	0.79	2.92
2° Ord. Root-Polynomial	1.31	<i>0.79</i>	3.57	<i>1.32</i>	<i>0.99</i>	3.56	1.11	0.81	2.94
3° Ord. Root-Polynomial	<u>1.35</u>	0.71	<u>4.09</u>	1.24	0.77	<i>3.67</i>	<u>1.05</u>	<u>0.75</u>	<i>2.82</i>
ShallowNN (5K)	7.74	4.43	25.32	8.10	6.23	21.93	<i>1.06</i>	<i>0.78</i>	<i>2.76</i>
PCA-Radiance (3 bases)	1.71	1.32	4.34	1.87	1.41	4.84	1.56	1.21	3.93
PCA-Radiance (6 bases)	1.51	1.13	4.15	1.71	1.14	4.97	1.21	0.88	3.16
PCA-Reflectance (3 bases)	1.56	1.19	4.19	2.42	1.68	6.96	1.37	1.07	3.38
PCA-Reflectance (6 bases)	1.52	1.13	4.17	1.71	1.14	4.95	1.22	0.89	3.17
ShallowNN-Radiance (5K)	2.39	1.65	6.92	4.16	2.76	12.60	1.01	0.72	2.68

accounted for during model design.

6 Conclusions

This work explored object color relighting under progressively decreasing levels of information—from full spectral data to minimal tristimulus-only input. By organizing the task along this information axis, we offered a unified framework to compare physically-based, colorimetric, and learning-based approaches. Our experiments demonstrate that method performance depends strongly on data availability. Polynomial regression, particularly of third order or root-transformed variants, performs best in low-cardinality data scenarios. In contrast, neural spectral recovery methods outperform others when trained on large-scale datasets, benefiting from learned priors and higher data diversity. Spectral reconstruction-based methods also show greater robustness to illuminant variability compared to tristimulus-based models. Finally, while chromatic adaptation transforms offer simplicity, their accuracy remains above perceptual thresholds, limiting their use in high-fidelity applications. Among the advanced approaches not yet fully explored in this study, sparse coding over learned reflectance dictionaries represents a promising alternative to PCA- or NN-based spectral recovery. These methods express reflectances as sparse linear combinations of atoms learned from real-world spectra and could yield physically plausible estimates from input data [10, 23]. Nevertheless, a fundamental limitation persists when the Spectral Power Distribution of the illuminant is not known. In such cases, only the CIE XYZ white point may be available, which is generally insufficient to resolve the underlying spectral content. Due to the phenomenon of metamerism, distinct SPDs can produce identical white points while yielding significantly different reflected spectra when interacting with object surfaces [14, 24]. This ambiguity cannot be resolved from tristimulus data alone, imposing a theoretical ceiling on achievable relighting accuracy.

Future directions should therefore include strategies to mitigate this uncertainty—such as incorporating plausible priors on illuminant spectra, or developing models capable of jointly estimating both reflectance and illumination characteristics when the latter is unknown [25, 26]. Moreover, while this study focused on single-point color inputs, additional gains might be obtained by exploiting contextual or spatial constraints across collections of measurements or temporally coherent data.

References

- [1] University of Eastern Finland, “Spectral database - munsell colors matt (spectrophotometer measured). [online (07/07/2025)],” <https://sites.uef.fi/spectral/databases-software/munsell-colors-matt-spectrophotometer-measured/>.

Table 3: ΔE_{00} results under combined illuminants (A, D50, D65, D100, D150, D200, F4, F8, F11, LED-1, LED-2). Results are shown for the considered dataset pairings. The solid horizontal line separates non-learning-based from learning-based methods, while the dashed line distinguishes tristimulus-based methods from spectral recovery-based ones. Best results are shown in **bold**, second best are underlined, and third best are *italicized*.

Method	ΔE_{00} on Macbeth Dataset (Trained on Munsell)			ΔE_{00} on Vrhel Dataset (Trained on Westland)			ΔE_{00} on Zhang-test (Trained on Zhang-train)		
	Mean	Median	95-P	Mean	Median	95-P	Mean	Median	95-P
XYZ Scaling	4.51	3.46	12.78	3.91	3.14	9.88	3.64	2.88	9.36
Von Kries	3.48	2.57	10.44	3.18	2.53	8.36	3.31	2.56	8.91
Bianco & Schettini	2.39	1.93	6.46	2.30	1.84	5.70	2.23	1.77	5.64
CAT02	2.34	1.75	6.62	2.25	1.77	5.63	2.27	1.79	5.85
3×3 Correction Matrix	1.68	1.11	5.10	1.75	1.14	5.27	1.39	0.99	4.00
2° Ord. Polynomial	1.39	0.80	4.46	1.61	1.08	4.82	1.29	0.89	3.73
3° Ord. Polynomial	1.14	<i>0.76</i>	3.58	1.26	<u>0.84</u>	3.74	1.21	0.81	3.54
2° Ord. Root-Polynomial	<i>1.34</i>	<u>0.75</u>	<u>4.10</u>	<i>1.51</i>	<i>1.06</i>	<i>4.40</i>	1.25	0.87	3.62
3° Ord. Root-Polynomial	1.28	0.64	<i>4.27</i>	1.31	0.78	3.89	<i>1.14</i>	<u>0.77</u>	<i>3.40</i>
ShallowNN (5K)	<u>7.44</u>	4.42	23.88	10.26	8.30	25.07	<u>1.06</u>	<u>0.77</u>	<u>2.91</u>
PCA-Radiance (3 bases)	1.98	1.36	6.01	2.51	1.79	7.18	<u>3.56</u>	<u>1.85</u>	13.12
PCA-Radiance (6 bases)	1.63	1.11	4.93	1.74	1.15	5.14	1.40	0.99	4.01
PCA-Reflectance (3 bases)	1.70	1.19	4.77	2.63	1.95	7.22	1.59	1.19	4.20
PCA-Reflectance (6 bases)	1.63	1.12	4.93	1.73	1.15	5.12	1.39	0.99	4.00
ShallowNN-Radiance (5K)	2.78	1.76	8.45	4.38	2.89	13.56	0.56	0.45	1.34

- [2] C. S. McCamy, H. Marcus, J. G. Davidson *et al.*, “A color-rendition chart,” *J. App. Photog. Eng.*, vol. 2, no. 3, pp. 95–99, 1976.
- [3] M. J. Vrhel, R. Gershon, and L. S. Iwan, “Measurement and analysis of object reflectance spectra,” *Color Research & Application*, vol. 19, no. 1, pp. 4–9, 1994.
- [4] S. Westland, J. Shaw, and H. Owens, “Colour statistics of natural and man-made surfaces,” *Sensor Review*, vol. 20, no. 1, pp. 50–55, 2000.
- [5] X. Zhang, B. Funt, and H. Mirzaei, “Metamer mismatching in practice versus theory,” *Journal of the Optical Society of America A*, vol. 33, no. 3, pp. A238–A247, 2016.
- [6] J. P. Parkkinen, J. Hallikainen, and T. Jaaskelainen, “Characteristic spectra of munsell colors,” *Journal of the Optical society of America A*, vol. 6, no. 2, pp. 318–322, 1989.
- [7] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, “Learning representations by back-propagating errors,” *nature*, vol. 323, no. 6088, pp. 533–536, 1986.
- [8] R. Schettini, “Deriving spectral, reflectance functions of computer-simulated object colours,” in *Computer graphics forum*, vol. 13, no. 4. Wiley Online Library, 1994, pp. 211–217.
- [9] R. Schettini and S. Zuffi, “A computational strategy exploiting genetic algorithms to recover color surface reflectance functions,” *Neural Computing and Applications*, vol. 16, no. 1, pp. 69–79, 2007.
- [10] J. Aeschbacher, J. Wu, and R. Timofte, “In defense of shallow learned spectral reconstruction from rgb images,” in *Proceedings of the IEEE International Conference on Computer Vision Workshops*, 2017, pp. 471–479.
- [11] Y.-T. Lin and G. D. Finlayson, “A rehabilitation of pixel-based spectral reconstruction from rgb images,” *Sensors*, vol. 23, no. 8, p. 4155, 2023.
- [12] J. Zhang, R. Su, Q. Fu, W. Ren, F. Heide, and Y. Nie, “A survey on computational spectral reconstruction methods from rgb to hyperspectral imaging,” *Scientific reports*, vol. 12, no. 1, p. 11905, 2022.

- [13] A. N. Fsiian, J.-B. Thomas, J. Y. Hardeberg, and P. Gouton, "Spectral reconstruction from rgb imagery: A potential option for infinite spectral data?" *Sensors*, vol. 24, no. 11, p. 3666, 2024.
- [14] M. D. Fairchild, *Color appearance models*. John Wiley & Sons, 2013.
- [15] J. von Kries, "Chromatic adaptation," in *Festschrift der Albrecht-Ludwig-Universität, Fribourg*, 1902, [Translation: D.L. MacAdam, *Sources of Color Science*, MIT Press, Cambridge, 1970].
- [16] N. Moroney, M. Fairchild, R. Hunt, and C. Li, "The ciecam02 color appearance model," 2002.
- [17] K. M. Lam, "Metamerism and color constancy," *Ph. D. Thesis, University of Bradford*, 1985.
- [18] S. Bianco and R. Schettini, "Two new von kries based chromatic adaptation transforms found by numerical optimization," *Color Research & Application*, vol. 35, no. 3, pp. 184–192, 2010.
- [19] G. D. Finlayson, M. Mackiewicz, and A. Hurlbert, "Color correction using root-polynomial regression," *IEEE Transactions on Image Processing*, vol. 24, no. 5, pp. 1460–1470, 2015.
- [20] G. Hong, M. R. Luo, and P. A. Rhodes, "A study of digital camera colorimetric characterization based on polynomial modeling," *Color Research & Application*, vol. 26, no. 1, pp. 76–84, 2001.
- [21] A. F. Agarap, "Deep learning using rectified linear units (relu)," *arXiv preprint arXiv:1803.08375*, 2018.
- [22] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [23] B. Arad and O. Ben-Shahar, "Sparse recovery of hyperspectral signal from natural rgb images," in *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, the Netherlands, October 11–14, 2016, Proceedings, Part VII 14*. Springer, 2016, pp. 19–34.
- [24] P. M. Hubel, J. Holm, and G. Finlayson, "Illuminant estimation and colour correction," in *The Colour Imaging in Multimedia Conference*, 1998.
- [25] R. M. Nguyen, D. K. Prasad, and M. S. Brown, "Training-based spectral reconstruction from a single rgb image," in *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part VII 13*. Springer, 2014, pp. 186–201.
- [26] Y. Fu, Y. Zheng, L. Zhang, and H. Huang, "Spectral reflectance recovery from a single rgb image," *IEEE Transactions on Computational Imaging*, vol. 4, no. 3, pp. 382–394, 2018.