

CLOE: A CONFIDENCE-BASED LOCAL-TO-GLOBAL ESTIMATION FRAMEWORK FOR MULTISPECTRAL ILLUMINANT RECOVERY

Matteo Kolyszko, Alessio Mognato, Marco Buzzelli, Simone Bianco, Raimondo Schettini

University of Milano-Bicocca, Viale Sarca 336, Milan 20126, Italy

ABSTRACT

This paper introduces CLOE, a confidence-based local-to-global fully convolutional framework for multispectral illuminant estimation, which jointly captures spatial structure and spectral correlations through adaptive confidence-weighted pooling. The architecture integrates a Spatial Feature Extractor and a Spectral Branch through a confidence-based fusion mechanism that produces robust global illuminant predictions. Evaluations on the KAUST and BeyondRGB datasets demonstrate that CLOE consistently outperforms both classical statistical methods and recent learning-based approaches, achieving state-of-the-art performance in terms of angular error between the estimated and ground-truth illuminant spectra (mean- Δ AMS and std- Δ AMS). Qualitative analyses further show that CLOE provides interpretable intermediate representations, enabling visual inspection of spatial-spectral cues. These results highlight the effectiveness and robustness of the proposed method for real-world multispectral imaging.

Index Terms— Multispectral Imaging, Illuminant Estimation, Color Constancy, Deep Learning

1. INTRODUCTION

Illuminant estimation is a fundamental component of color constancy, enabling the recovery of scene reflectance independently of illumination conditions. While traditional approaches rely on RGB images [1, 2, 3, 4, 5], their limited spectral resolution often leads to ambiguities between illumination and surface properties. Multispectral imaging (MSI), by contrast, provides richer spectral information that can significantly improve illuminant recovery. However, effectively exploiting MSI remains challenging: high-dimensional data increases computational complexity, while real-world scenes introduce noise, spectral variability, and heterogeneous lighting conditions. Existing methods for multispectral illuminant estimation [6, 7, 8, 9] typically focus either on spatial features or on spectral structure, but fail to jointly integrate both sources of information in a coherent, interpretable framework. Rather than introducing a radically new estimation paradigm, this work focuses on a principled integration of spatial and spectral cues within a single confidence-driven framework, aiming at improved robustness and interpretability.

In this paper, we introduce CLOE, a multispectral illuminant estimation model that integrates spatial structure with wavelength-dependent spectral features. The framework is built on a dual-branch design composed of a Feature Extractor Branch (FXB) and a Spectral Branch (SB), whose outputs are combined through a confidence-based Fusion Module (FM) to produce accurate and robust illuminant predictions.

We summarize our main contributions as follows:

- We propose a confidence-driven fusion mechanism that integrates spatial and spectral reliability cues, enabling stable and robust illuminant estimation across diverse imaging conditions.
- We design a spectral grouping strategy that models spectral correlations while suppressing unreliable wavelength regions, improving both estimation accuracy and interpretability.
- We conduct extensive experiments on the KAUST and BeyondRGB datasets, showing that CLOE consistently outperforms classical and recent learning-based methods in both mean- Δ AMS and std- Δ AMS, and generalizes effectively across controlled and real-world scenarios.

Overall, CLOE establishes a new state-of-the-art in multispectral illuminant estimation, achieving highly accurate spectral predictions while providing interpretable intermediate representations.

2. RELATED WORKS

2.1. Illuminant Estimation

Color constancy is generally accomplished in two stages: first estimating the illumination and then applying chromatic adaptation. Earlier research primarily targeted illumination estimation using RGB images. These works introduced statistical approaches [1, 2] as well as deep learning-based solutions [10, 11, 12] to improve estimation accuracy. However, these techniques relied exclusively on RGB data to disentangle illumination from surface reflectance.

More recent studies adopted multispectral (MS) images to address color constancy, leveraging their richer spectral

content to obtain more reliable results. Several methods expanded established statistical models from RGB to the multispectral domain [13, 14, 15]. Other approaches separated illumination spectra from surface reflectance by applying low-rank matrix factorization techniques [3, 4, 5], which are effective particularly when only a few objects dominate the scene.

Various learning-based illumination estimation methods have also been specifically designed for MS images. In [6], illumination estimation was formulated as a constrained matrix factorization problem, though it required per-image calibration using known references or measurement devices [7]. Their training relied on synthetic datasets with simulated light sources, however, the limited diversity of true reflectance spectra available for data synthesis, despite the wide range of illumination conditions, can hinder generalization and lead to reduced performance in complex scenes.

The work in [8] introduced an unsupervised method that segments regions mainly influenced by illumination, assuming the illumination spectrum is piecewise constant: this simplification however struggles to capture the diversity of real illuminants. Similarly, [9] used an image segmentation strategy to localize illumination-dominant regions, but mismatches were found between segmentation results and the regions crucial for accurate illumination estimation.

Another method [16] split the input MS image into multiple channel triplets and processed them through a pyramid-based convolutional block for color constancy. However, this design did not fully exploit the spectral relationships linking illumination and surface reflectance across different wavelengths.

Finally, [17] proposed a deep framework that jointly models illumination and surface reflectance by integrating a learnable spectral unmixing module with a cross-attention feature-mixing block. While effective in coupling spectral components, the method relies on accurate unmixing and attention mechanisms, which may be sensitive to complex material distributions and scene variability.

3. PROPOSED METHOD

In this section, we first formulate the problem of multispectral illuminant estimation. We then introduce the overall CLOE architecture and finally provide a detailed description of its three main components: the Feature Extractor Branch, the Spectral Branch, and the Fusion Module. An overview of the complete framework is shown in Fig. 1.

3.1. Problem Formulation

Let $\mathbf{I} \in \mathbb{R}^{N \times H \times W}$ denote a multispectral image with N spectral channels and spatial size (H, W) . Each pixel spectrum $\mathbf{I}(x, y) \in \mathbb{R}^N$ can be modelled according to the image-formation process:

$$\mathbf{I}(x, y, \lambda) = E(\lambda) R(x, y, \lambda) S(\lambda), \quad (1)$$

where $E(\lambda)$ is the illuminant spectral power distribution (SPD), here assumed to be global and thus depending only from wavelength λ , $R(x, y, \lambda)$ is the surface reflectance, and $S(\lambda)$ is the camera spectral sensitivity.

The objective of multispectral illuminant estimation is to recover the global illuminant spectrum

$$\mathbf{E} = [E(\lambda_1), \dots, E(\lambda_N)]^\top \in \mathbb{R}^N, \quad (2)$$

from the single multispectral observation $\mathbf{I}$. We formulate the task as learning a parametric function $\mathcal{F}_\theta(\cdot)$ such that

$$\hat{\mathbf{E}} = \mathcal{F}_\theta(\mathbf{I}), \quad (3)$$

where $\hat{\mathbf{E}}$ is the predicted illuminant.

3.2. Proposed Framework

CLOE is a fully convolutional illuminant estimation architecture that jointly exploits spatial structure and spectral correlations. As illustrated in Fig. 1, the network processes the input through two parallel pathways and merges their outputs using a confidence-based fusion mechanism.

The overall forward pass of the proposed CLOE model is defined as:

$$\hat{\mathbf{E}} = \text{FM}\left([\text{FXB}(\mathbf{I}), \text{SB}(\mathbf{I})]\right), \quad (4)$$

where $\text{FXB}(\cdot)$ denotes the Feature Extractor Branch, which extracts spatial features from the multispectral input and produces both local illuminant hypotheses and a spatial confidence map, reflecting the reliability of different image regions for illuminant estimation, as spatial information content varies across the scene; $\text{SB}(\cdot)$ denotes the Spectral Branch, which processes grouped spectral bands to model inter-band correlations and estimate spectral confidence maps, capturing the varying informativeness and stability of different wavelength ranges; $[\cdot, \cdot]$ denotes concatenation along the channel dimension and $\text{FM}(\cdot)$ denotes the Fusion Module, which integrates spatial and spectral confidence features and computes the final global illuminant through confidence-weighted pooling.

3.3. Feature Extractor Branch

The Feature Extractor Branch (FXB) processes the multispectral image $\mathbf{I}$ using a convolutional backbone and produces two complementary outputs: a single-channel spatial confidence map, which identifies the regions of the scene that provide reliable cues for illuminant recovery; and a set of local illuminant estimates, where each spatial location is associated with a predicted illuminant spectrum.

Formally, the spatial confidence map is defined as

$$\mathbf{C}_{\text{spatial}} \in \mathbb{R}^{H' \times W'}, \quad (5)$$




Fig. 1. Overview of the proposed CLOE architecture. The multispectral input is processed through two parallel branches: the Feature Extractor Branch (FXB), which generates local illuminant estimates and a spatial confidence map, and the Spectral Branch (SB), which produces spectral confidence maps. These are fused by the Fusion Module (FM) to produce the global illuminant estimate.

while the local illuminant estimates are given by

$$\mathbf{E}_{\text{local}} \in \mathbb{R}^{N \times H' \times W'}. \quad (6)$$

The FXB captures structural features such as edges, textures, and homogeneous regions, providing spatial guidance that helps disambiguate surface reflectance from illumination.

3.4. Spectral Branch

The Spectral Branch (SB) is designed to capture spectral dependencies that cannot be learned purely from spatial processing. To this end, the N spectral channels of the input are partitioned into K contiguous groups, each corresponding to a local spectral neighborhood:

$$\mathbf{I}_{G_k} \in \mathbb{R}^{\frac{N}{K} \times H \times W}, \quad k = 1, \dots, K. \quad (7)$$

In our implementation, we set $K = 4$, resulting in groups of four adjacent spectral bands. This choice provides a good trade-off between modeling local spectral correlations and maintaining computational efficiency, and was empirically validated in our experiments.

Each group is processed by a lightweight spectral module consisting of depthwise-pointwise convolutions followed by Squeeze-and-Excitation recalibration, which adaptively emphasizes the most informative spectral responses within the group. For each group, the module outputs a spectral confidence map:

$$\mathbf{C}_k^{\text{spec}} \in \mathbb{R}^{H' \times W'}. \quad (8)$$

The confidence maps from all groups are then concatenated along the channel dimension to form the full spectral confidence representation:

$$\mathbf{C}_{\text{spec}} = [\mathbf{C}_1^{\text{spec}}, \dots, \mathbf{C}_K^{\text{spec}}]. \quad (9)$$

Through this mechanism, the Spectral Branch identifies the wavelength regions that best reveal the underlying illuminant characteristics and suppresses bands whose responses are unreliable or dominated by reflectance or noise.

3.5. Fusion Module

The Fusion Module (FM) integrates the spatial and spectral confidence cues produced by the two branches and consolidates them into a single confidence map that governs the final illuminant estimation. The spatial confidence map highlights regions that are structurally reliable, while the spectral confidence maps indicate the wavelength groups that best capture the illuminant characteristics. The module combines these signals to produce a fused confidence map and uses it to compute the global illuminant through confidence-weighted pooling.

Formally, given the spatial confidence map $\mathbf{C}_{\text{spatial}}$, the set of spectral confidence maps $\mathbf{C}_{\text{spec}}$, and the local illuminant predictions $\mathbf{E}_{\text{local}}$, the fused prediction is obtained as:

$$\mathbf{C}_{\text{fused}} = \sigma(\text{Conv}_{1 \times 1}([\mathbf{C}_{\text{spatial}}, \mathbf{C}_{\text{spec}}])), \quad (10)$$

and the final illuminant estimate is computed via

$$\hat{\mathbf{E}} = \frac{\sum_{x,y} \mathbf{C}_{\text{fused}}(x,y) \mathbf{E}_{\text{local}}(x,y)}{\sum_{x,y} \mathbf{C}_{\text{fused}}(x,y)}. \quad (11)$$

The Fusion Module therefore suppresses unreliable regions, producing a robust global illuminant estimate that reflects the most informative parts of both the scene structure and the spectral distribution.

4. EXPERIMENTAL SETUP

In this section, we describe the experimental setup used to evaluate our method. We present the datasets, the evaluation metrics, and the implementation details.

4.1. Datasets

We evaluate the proposed CLOE model on two challenging multispectral datasets commonly used for illuminant estimation: the KAUST reflectance dataset [6] and the BeyondRGB real-world dataset [16].

The KAUST dataset provides multispectral reflectance measurements captured under controlled imaging conditions. Each scene is represented as a hyperspectral cube spanning 31 bands between 400 nm and 700 nm at 10 nm intervals. To simulate realistic illuminant variations, the authors relit each reflectance image using a set of 59 standard illuminants. Following prior work, the dataset is split into 80% training and 20% testing. This dataset offers clean spectral measurements and controlled variability, making it suitable for evaluating the spectral fidelity of illuminant estimation.

The BeyondRGB dataset contains real-world multispectral radiance images acquired with a mobile camera prototype under diverse indoor and outdoor lighting conditions, covering the spectral range from 380 nm to 730 nm. The dataset is divided into the Lab subset, which includes controlled laboratory scenes under mixed artificial illuminants, and the Field subset, which contains complex real-world scenes under natural illumination. Unlike KAUST, BeyondRGB exhibits significant spatial and spectral variability, sensor noise, and challenging illumination distributions, providing a rigorous test of robustness.

4.2. Evaluation Metrics

Following prior work on multispectral illuminant estimation [16, 17], performance is quantified using the mean- Δ AMS and std- Δ AMS metrics, where Δ AMS (Angular error between Multispectral Spectra) measures the angular discrepancy between the predicted illuminant spectrum $\hat{\mathbf{E}}$ and the ground-truth illuminant spectrum $\mathbf{E}$, and is defined as:

$$\Delta\text{AMS} = \arccos\left(\frac{\hat{\mathbf{E}} \cdot \mathbf{E}}{\|\hat{\mathbf{E}}\|_2 \|\mathbf{E}\|_2}\right). \quad (12)$$

4.3. Implementation details

The network was implemented in PyTorch and trained on a single NVIDIA TITAN X GPU. Each experimental run consisted of 500 training epochs and required approximately one and a half hours.

Model parameters were optimized using the Adam optimizer, with a weight decay set to 1×10^{-5} . A OneCycleLR scheduler [18] was employed to dynamically modulate the learning rate throughout training, ensuring stable convergence.

The backbone selected for the Feature Extractor Branch (FXB) is the EfficientNetB0-E [19], which represents the best tradeoff between performance, interpretability, and efficiency in terms of model size and computational cost.

To guarantee consistency and reproducibility across all experiments, the same random seed and hardware configuration were maintained.

The loss function $\mathcal{L}$ is defined as the average angular error between the estimated and ground-truth illuminant spectra

over a batch of size B :

$$\mathcal{L}(\hat{\mathbf{E}}, \mathbf{E}) = \frac{1}{B} \sum_{b=1}^B \theta_b, \quad (13)$$

where the angular error θ_b is computed as:

$$\theta_b = \arccos\left(\frac{\hat{\mathbf{E}}_b \cdot \mathbf{E}_b}{\|\hat{\mathbf{E}}_b\|_2 \|\mathbf{E}_b\|_2}\right) \times \frac{180}{\pi}. \quad (14)$$

This angular error corresponds to the evaluation metric used throughout the paper, referred to as Δ AMS.

5. RESULTS

We evaluate the performance of the proposed frameworks against a range of existing methods, including both statistic-based and learning-based approaches. For the KAUST dataset, comparisons are conducted with GrayEdge [14], LRMF [5], ISNL [20], PWIR [21], DUN [6], BeyondRGB [16] and ISS [17]. For the BeyondRGB dataset, the evaluation includes PWIR [21], BeyondRGB [16] and ISS [17].

Table 1 reports a comparative evaluation of illuminant estimation methods on the KAUST and BeyondRGB datasets using mean- Δ AMS and std- Δ AMS.

Across both datasets, a clear trend emerges: traditional statistical methods produce substantially higher errors, while learning-based and multispectral-aware approaches achieve better accuracy.

On the KAUST dataset, classical models such as GrayEdge, LRMF, and ISNL show the highest errors (11° – 24°). Recent learning-based methods, including DUN and BeyondRGB, reduce the error to 9.17° and 7.30° respectively, with ISS further improving to 5.55° . Our method achieves the best overall performance, with a mean- Δ AMS of 4.60° and a std- Δ AMS of 2.96° , confirming the effectiveness of the proposed fusion mechanism.

For the BeyondRGB dataset, results are reported separately for the Lab and Field subsets to reflect domain variability. In the Lab setting, PWIR performs poorly (27.10°), while BeyondRGB and ISS achieve lower errors (5.92° and 3.65° respectively). Our method again obtains the best performance, with a mean- Δ AMS of 3.29° and the lowest std- Δ AMS of 2.05° . In the Field subset, all methods face increased difficulty due to uncontrolled lighting. PWIR and BeyondRGB reach 16.31° and 7.22° , whereas ISS achieves 5.97° . Our method remains the most accurate, with a mean- Δ AMS of 5.17° and a std- Δ AMS of 3.43° , demonstrating strong robustness in real-world conditions.

Figure 2 reports qualitative examples of the proposed CLOE model on the BeyondRGB dataset. These qualitative results indicate that CLOE not only achieves low spectral angular errors but also provides interpretable intermediate representations, enabling visual inspection of both spatial and spectral contributions to the illuminant prediction.

Table 1. Comparison of illuminant estimation methods on the KAUST and BeyondRGB datasets. Metrics include mean- Δ AMS and std- Δ AMS.

Dataset	Location	Method	mean- Δ AMS ↓	std- Δ AMS ↓
KAUST	Mixed	GrayEdge	11.46	8.02
		LRMF	24.06	23.49
		ISNL	20.05	11.46
		PWIR	28.07	21.20
		DUN	9.17	6.30
		BeyondRGB	7.30	4.45
		ISS	5.55	3.29
		CLOE	4.60	2.96
BeyondRGB	Lab	PWIR	27.10	10.86
		BeyondRGB	5.92	2.92
		ISS	3.65	2.72
		CLOE	3.29	2.05
	Field	PWIR	16.31	6.07
		BeyondRGB	7.22	5.54
		ISS	5.97	3.90
		CLOE	5.17	3.43

Overall, the proposed framework consistently achieves the lowest errors across all benchmarks and acquisition conditions, confirming its effectiveness and robustness for multispectral illuminant estimation, while also providing interpretable intermediate representations that offer insights into both spatial and spectral contributions to the final prediction.

6. CONCLUSION

In this work, we introduced CLOE, a fully convolutional architecture for multispectral illuminant estimation that jointly leverages spatial structure and spectral correlations.

Extensive experiments on the KAUST and BeyondRGB datasets demonstrate that CLOE consistently outperforms classical statistical approaches and recent learning-based methods, achieving state-of-the-art performance in terms of both mean- Δ AMS and std- Δ AMS. Furthermore, the qualitative analyses highlight the interpretability of the model, revealing meaningful intermediate confidence maps and stable spectral predictions across diverse scenes.

The proposed framework is deliberately modular and can be extended to more complex illumination models without altering its core confidence-based formulation. While conceptually simple, the proposed design choices are shown to be effective across different datasets and acquisition conditions.

Future work will explore extending CLOE to jointly model illumination and reflectance, as well as adapting the framework to varying spectral resolutions and sensor characteristics. We also plan to investigate self-supervised training

strategies to reduce reliance on calibrated datasets and improve generalization in real-world scenarios.

7. REFERENCES

- [1] Edwin H Land, “The retinex theory of color vision,” *Scientific american*, vol. 237, no. 6, pp. 108–129, 1977.
- [2] Joost Van De Weijer, Theo Gevers, and Arjan Gijsenij, “Edge-based color constancy,” *IEEE Transactions on image processing*, vol. 16, no. 9, pp. 2207–2214, 2007.
- [3] Dongsheng An, Jinli Suo, Haoqian Wang, and Qionghai Dai, “Illumination estimation from specular highlight in a multi-spectral image,” *Optics Express*, vol. 23, no. 13, pp. 17008–17023, 2015.
- [4] Xiaochuan Chen, Mark S Drew, and Ze-Nian Li, “Illumination and reflectance spectra separation of hyper-spectral image data under multiple illumination conditions,” *Electronic Imaging*, vol. 29, pp. 194–199, 2017.
- [5] Yinqiang Zheng, Imari Sato, and Yoichi Sato, “Illumination and reflectance spectra separation of a hyperspectral image meets low-rank matrix factorization,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 1779–1787.
- [6] Yuqi Li, Qiang Fu, and Wolfgang Heidrich, “Multispectral illumination estimation using deep unrolling network,” in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 2672–2681.
- [7] Yuqi Li, Chong Wang, Jieyu Zhao, and Qingshu Yuan, “Efficient spectral reconstruction using a trichromatic camera via sample optimization,” *The Visual Computer*, vol. 34, no. 12, pp. 1773–1783, 2018.
- [8] Lin Gu, Cong Phuoc Huynh, and Antonio Robles-Kelly, “Segmentation and estimation of spatially varying illumination,” *IEEE transactions on image processing*, vol. 23, no. 8, pp. 3478–3489, 2014.
- [9] Donik Vršnak, Ilija Domislović, Marko Subašić, and Sven Lončarić, “Framework for illumination estimation and segmentation in multi-illuminant scenes,” *IEEE access*, vol. 11, pp. 2128–2137, 2023.
- [10] Wu Shi, Chen Change Loy, and Xiaoou Tang, “Deep specialized network for illuminant estimation,” in *European conference on computer vision*. Springer, 2016, pp. 371–387.
- [11] Huanglin Yu, Ke Chen, Kaiqi Wang, Yanlin Qian, Zhaoxiang Zhang, and Kui Jia, “Cascading convolutional color constancy,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020, vol. 34, pp. 12725–12732.




Fig. 2. CLOE model outputs: (a) RGB visualization of the input image (from Samsung sensor); (b) final fused confidence map; (c) local illuminant estimates for each spectral band, shown as a 6×6 grid corresponding to 36 spectral bands sampled every 10 nm from 380 nm to 730 nm. Each map is independently min-max normalized to emphasize the spatial distribution and consistency of local estimates at each wavelength, rather than their absolute spectral magnitude; (d) predicted versus ground-truth illuminant spectra with angular error in ΔAMS ; (e) fused confidence map overlaid on the RGB image for visual interpretation.

- [12] Yi-Chen Lo, Chia-Che Chang, Hsuan-Chao Chiu, Yu-Hao Huang, Chia-Ping Chen, Yu-Lin Chang, and Kevin Jou, “Clcc: Contrastive learning for color constancy,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 8053–8063.
- [13] Arjan Gijzenij, Theo Gevers, and Joost Van De Weijer, “Improving color constancy by photometric edge weighting,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 34, no. 5, pp. 918–929, 2011.
- [14] Haris Ahmad Khan, Jean-Baptiste Thomas, Jon Yngve Hardeberg, and Olivier Laligant, “Illuminant estimation in multispectral imaging,” *Journal of the Optical Society of America A*, vol. 34, no. 7, pp. 1085–1098, 2017.
- [15] Jean-Baptiste Thomas, “Illuminant estimation from uncalibrated multispectral images,” in *2015 Colour and Visual Computing Symposium (CVCS)*. IEEE, 2015, pp. 1–6.
- [16] Ortal Glatt, Yotam Ater, Woo-Shik Kim, Shira Werman, Oded Berby, Yael Zini, Shay Zelinger, Sangyoon Lee, Heejin Choi, and Evgeny Soloveichik, “Beyond rgb: a real world dataset for multispectral imaging in mobile devices,” in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2024, pp. 4344–4354.
- [17] Hyejin Oh, Woo-Shik Kim, Sangyoon Lee, YungKyung Park, and Je-Won Kang, “Illumination spectrum estimation for multispectral images via surface reflectance modeling and spatial-spectral feature generation,” in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 2215–2225.
- [18] Leslie N Smith and Nicholay Topin, “Super-convergence: Very fast training of residual networks using large learning rates,” *arXiv preprint arXiv:1708.07120*, vol. 5, 2017.
- [19] Mingxing Tan and Quoc Le, “Efficientnet: Rethinking model scaling for convolutional neural networks,” in *International conference on machine learning*. PMLR, 2019, pp. 6105–6114.
- [20] Tong Su, Yu Zhou, Yao Yu, Xun Cao, and Sidan Du, “Illumination separation of non-lambertian scenes from a single hyperspectral image,” *Optics express*, vol. 26, no. 20, pp. 26167–26178, 2018.
- [21] Antonio Robles-Kelly and Ran Wei, “A convolutional neural network for pixelwise illuminant recovery in colour and spectral images,” in *2018 24th International Conference on Pattern Recognition (ICPR)*. IEEE, 2018, pp. 109–114.