

Revisiting and Optimising a CNN Colour Constancy Method for Multi-Illuminant Estimation

Ghalia Hemrit^{1,2} and Joseph Meehan¹

¹Huawei Technologies, Nice, France; ²University of East Anglia, Norwich, United Kingdom

Abstract

The aim of colour constancy is to discount the effect of the scene illumination from the image colours and restore the colours of the objects as captured under a 'white' illuminant. For the majority of colour constancy methods, the first step is to estimate the scene illuminant colour. Generally, it is assumed that the illumination is uniform in the scene. However, real world scenes have multiple illuminants, like sunlight and spot lights all together in one scene. We present in this paper a simple yet very effective framework using a deep CNN-based method to estimate and use multiple illuminants for colour constancy. Our approach works well in both the multi and single illuminant cases. The output of the CNN method is a region-wise estimate map of the scene which is smoothed and divided out from the image to perform colour constancy. The method that we propose outperforms other recent and state of the art methods and has promising visual results.

Introduction

Performing colour constancy is an important pre-processing step of the digital camera pipeline. It consists of removing the colour bias introduced by the scene illuminant from the colours in the image. This way, a digital camera gives images of the scene where the objects have the same colours independently of the scene illumination conditions. Colour constancy is important for various computer vision tasks like object detection and tracking but can also be important to improve the visual aesthetics of an image.

Most colour constancy methods rely on a first step which is estimating the colour of the scene illuminant before this colour can be discounted from the image. Very often, the single illuminant case is assumed and only the predominant light colour in the scene is estimated (global estimate) expressed as an RGB vector and used for colour constancy. This is the case for state of the art and most commonly used methods like statistics-based methods [6][11][8] and more recent learning-based methods [3][2].

This approach succeeds in solving for the illuminant colour in many cases and it has been largely adopted. There exist a wide variety of methods but the answer to illuminant estimation varies considerably with the method that is used [10] and can give unsatisfactory results, in particular when there are multiple illuminants in the scene.

In fact, most real life scenes have more than one ambient light. We consider in this work the multi-illuminant case and propose instead of having a single (global) estimate for the scene illuminant colour, to generate multiple local estimates in the form of a region-wise map. This estimate map can be then divided out from the image to perform colour constancy. We use the fully convolutional neural network SqueezeNet-FC4 introduced by Hu et al. [2]. Eponymously, we call our method the Multi-Estimate Map CNN (MEM CNN). We evaluate our method on the MIMO public multi-illuminant dataset for colour constancy [1] and we

compare it to other methods using the angular error metric that we adapt to the multi-estimate case. We show that MEM CNN is effective in the multi and single illuminant cases and that it outperforms the other methods when evaluated on this dataset of images.

The rest of the paper is organised as follows. In the next section, we present some related work on colour constancy and deep learning methods. Then, we describe the proposed approach and introduce the dataset that we used for training. We also present the network architecture in another section. This is followed by the experiments and results section, and a conclusion.

Related Work

Some recent learning-based methods address the question of the multiple illuminants in the scene and how to give an estimate of the illuminant colour in this particular case [17]. Often, multiple local estimates are produced then combined into a global estimated illuminant. Additionally, when local estimates are used for colour constancy, the global estimate remain a reference in the colour difference evaluation.

In [2], local estimates are produced by the FC4 trained model as well as confidence values where a confidence expresses the value of the corresponding patch for estimating the illuminant based on its semantic content. The confidence-weighted local estimates are combined into a one global estimate.

In [17], the input image is sampled into patches and local estimates are produced with an edge-based method [7] then a clustering technique gives the global estimate. In [18], a user preference study allows to choose the global estimate. Here, estimates are predicted with a learning-based method from 4×8 sub-images and used to classify the image as single or two-illuminant. Then, two illuminants are calculated with a k-means clustering.

[4] is a learning patch-based method. It also produces multiple local estimates. A detector decides whether the image has a single or multiple illuminants using a statistical analysis of the local estimates (the difference between the estimates in the chromaticity space is compared to a threshold). If the difference is small, it is a single illuminant case and a local-to-global regression is performed with a supervised neural network. In the multi-illuminant case, the local estimates are used with no further processing. The method was not evaluated in terms of colour differences in the multi-illuminant case. In [13], the local illuminant vectors are estimated from the whole image and not from patches. Instead of learning the estimates, the model predicts local kernels. These kernels are then used to calculate the estimated illuminant. A global-to-local aggregation gives a region-wise estimate map using clustering results. The evaluation in terms of colour differences uses the global estimate.

We propose to use the estimate map from SqueezeNet-FC4 [2] to perform colour constancy. The FC4 trained model generates 60 patch-based local illuminants organised in a 6×10 RGB estimate map. We upsample this map and process it to smooth

the transitions between the colours. The map is then divided out from the image to restore the objects colours. Other methods also assume a smooth transition of the illumination in the scene image in the multi-illuminant case like [16][15].

Method Overview

We present in this paper a framework based on a learning-based method for performing colour constancy in the multi and single illuminant cases. Our method, MEM CNN, uses the fully convolutional network SqueezeNet-FC4 introduced by Hu et al. [2]. The original FC4 method gives a global estimate which is a single 3-colour vector.

The novelty of our work is that we do not estimate a single illuminant in the scene. Instead, we use multiple estimates to solve for colour constancy. Our method generates an estimate map which is an image where the illuminant colour is estimated for every region of the image.

The model generates an estimate map which is a thumbnail image of size $6 \times 10 \times 3$ (60 patches). By upsampling this image to the scene image size (by duplicating the patches), it becomes a region-wise estimate map. This map is linearised and more importantly, it is processed in a way that smoothes the edges between the different patches. To do that, we filter the map with a kernel of area 50×50 . It is finally divided out from the image to remove the illuminant colour bias, as shown in Equation 1 for every pixel k from the image.

$$\mathbf{I}_c(k) = \mathbf{I}_o(k) \text{diag}(\hat{\boldsymbol{\rho}}(j))^{-1} \quad (1)$$

where k refers to the k -th pixel from the scene image, in the j -th region. \mathbf{I}_o is the image of the scene under illuminant E_o and \mathbf{I}_c is the restored image of the scene under white light E_c . $\hat{\boldsymbol{\rho}}(j)$ is the estimated illuminant RGB vector for the j -th region of the image, and diag is the diagonal matrix operator.

Training Dataset

To train the network, we use the MIMO dataset introduced by Beigpour et al. [1][9]. This dataset has 78 RGB images of real life scenes and images taken in a controlled lab environment. The scenes have multiple illuminants: two main non-uniformly distributed illuminants in every scene. The dataset is different from other public sets for colour constancy as every image is provided with a pixel-wise ground-truth illuminants image which is also an RGB image that has the illuminant 3-vector for every pixel of the corresponding captured scene. Scenes images and pixel-wise ground-truth images have size $302 \times 452 \times 3$. Samples from the dataset are presented in Figure 1.

The scenes contain simple and complex diffuse and specular objects. The pixel-wise ground-truth is calculated following a methodology described in [1] where the authors present a procedure to collect the data and calculate the ground-truth. This methodology is not in the scope of this paper.

The methodology is given for two illuminants scenes which is the case of the lab-based images. For these scenes, three coloured light sources were used: a reddish, a blueish and a 'white' illuminants. Every scene was lit by a distinct pair of light sources from two angles: left and right. The real life scenes contain two major illuminants: the ambient light which is the overall illumination and a direct light which was added to the scene. For this reason, the given ground-truth calculation is also considered valid for these scenes and pixel-wise ground-truth for these scenes were calculated the same way.

The neural network takes inputs of size $216 \times 325 \times 3$. The dataset images are cropped (at random positions) to fit the input size requirement, as part of a data augmentation step. The




Figure 1. Samples from the MIMO dataset [1]: the captured scenes (gamma-corrected) and the corresponding pixel-wise ground-truth images.

images cannot be resized because the ground-truth image pixel values should not change. Data augmentation also includes random vertical and horizontal flipping of the training sample. The ground-truth image is processed accordingly. As the FC4 network was pre-trained on gamma-corrected data for display, we also apply a gamma to the scenes linear RGB images.

Network Architecture

FC4 [2] uses a SqueezeNet set of layers [12]. The model generates two types of outputs: an estimate map and a corresponding confidence map which has the confidence weights of the patches in the illuminant estimation. The whole estimate map is used for training (to calculate the loss) and for testing. We do not use the confidence map in this work.

The network is implemented with Tensorflow and it is trained end-to-end by back-propagation. Adam optimiser is used with $\text{learning rate} = 3 \times 10^{-4}$ and $\text{batch size} = 16$. All the convolutional layers are fine-tuned from the pre-trained network (for image classification). The ReLU function is used for non-linear activation except for the final layer 7.

In MEM CNN we use the FC4 neural network. We present the architecture in Figure 2. The network is trained with the scenes images and ground-truth illuminants images. The input size is $216 \times 325 \times 3$.

Angular Error Map

We define a new cost function to take into account the multiple ground-truth illuminants in one image and the multi-estimate map. To calculate the loss for every training sample, the ground-truth illuminants image used for training is downsampled to the size of the estimate map ($6 \times 10 \times 3$): the image is divided into 60 distinct regions then the median over every region is calculated. For the 6×10 pairs of colour vectors, the traditional angular error is calculated. The loss L is an average of the 60 angular errors, as shown in Equation 2.

$$L(\hat{\boldsymbol{\rho}}_i) = \text{mean}(\text{err}(\hat{\boldsymbol{\rho}}_i(1), \boldsymbol{\rho}_i(1)), \dots, \text{err}(\hat{\boldsymbol{\rho}}_i(60), \boldsymbol{\rho}_i(60))) \quad (2)$$

where i refers to the i -th training sample, $\hat{\boldsymbol{\rho}}_i(j)$ is the j -th estimated illuminant which is the RGB vector of the j -th patch from the estimate map and $\boldsymbol{\rho}_i(j)$ is the j -th region-wise ground-truth

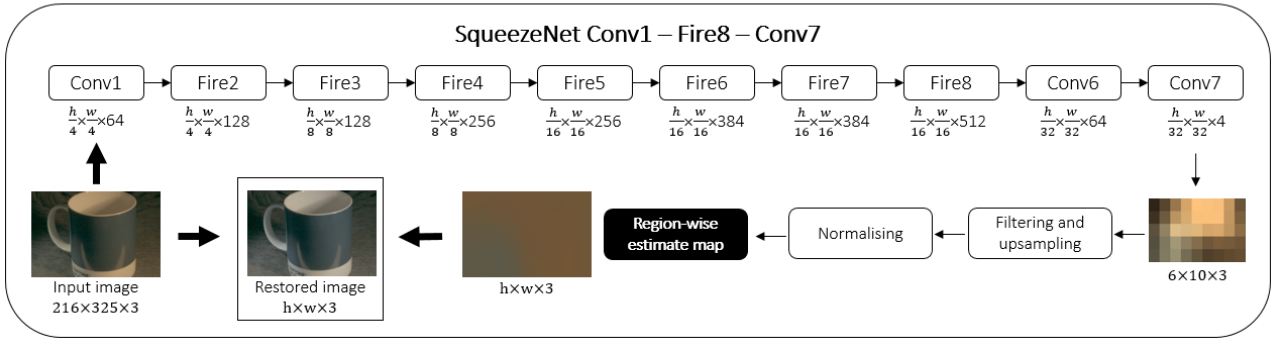


Figure 2. Architecture of MEM CNN.

illuminant. It is defined as the median of the j -th region of the ground-truth image (see Equation 3).

$$\rho(j) = \text{median}(\mathbf{R}_j) \quad (3)$$

where \mathbf{R}_j is the j -th region of the ground-truth image which is divided into 60 equal regions of 36×32 pixels except the right side regions of the image which have more pixels.

We call the metric associated to the new cost function the angular error map as it uses the error distribution in a downsampled representation of the scene. The output of the angular error map is a single angular error measure. We use the angular error map when testing MEM CNN to evaluate its performance in terms of angular error.

Experiments and Results

We evaluate the performance of our method MEM CNN and compare it to recent and state of the art methods on the MIMO dataset [1]. We use the angular error as a metric. For MEM CNN, we first calculate the angular error maps. For the other single estimate methods, the estimate is an RGB vector, the angular error is the angle between the estimate and the ground-truth. For MIMO scenes images, we use a ground-truth RGB vector that we generate from the ground-truth pixel-wise image as being the average colour of the two RGB vectors with the largest weights in the image. This way, we define an illuminant colour vector per scene to use with the single estimate methods. Of course, the way this vector was defined from the image is questionable and there can be different calculation methodologies that can be adopted. Saying that, we tested other methods and we found very similar results.

Results

Figure 3 shows for three different scenes from the dataset, the estimate maps and the MEM CNN outputs (restored using the estimate maps), the input scenes images and the corresponding ground-truth images, and the images restored using the ground-truth.

Some of the ground-truth images show clipped pixels. These are missed values in the calculation of the provided pixel-wise ground-truth illuminants. We do not have this clipping in the estimated results. The method output gives a good estimate of the illumination distribution in the scene. The changes in the illuminant are visible in the estimate map in the three scenes in the Figure, and these changes are accurately located in the image, even though a high frequency change in the illumination is more difficult to estimate and our solution can be improved in

	mean	median	std	Q1	Q3	max
Cheng et al. (n = 8%) [8]	9.1	7.04	7.01	4.68	11.35	41.96
Grey-Edge 1 st order (p = 3, σ = 9) [7]	7.45	5.86	4.72	5.02	8.81	35.12
Grey-Edge 2 nd order (p = 3, σ = 9) [7]	7.41	5.99	5.12	4.77	8.28	37.04
White Patch [11]	6.89	4.01	7.84	1.62	10.7	40.0
Shades of Gray (p = 3) [5]	6.06	5.13	3.64	3.34	7.75	18.74
CNN [3]	5.95	5.73	3.58	4.86	7.57	12.92
Gray World [6]	5.36	4.53	3.0	3.4	6.9	16.68
FC4 [2]	2.44	1.94	1.77	1.3	3.05	7.31
MEM CNN	1.93	1.34	1.35	0.63	3.2	4.42

Table 1: Summary statistics of angular errors for different colour constancy methods evaluated on the MIMO dataset.

this direction. MEM CNN gives good visual results, compared to the restored images when the ground-truth is used.

Table 1 shows quantitative results of different methods evaluated on the MIMO dataset. They are summary statistics of the angular errors (mean, median, standard deviation, first quantile Q1, third quantile Q3 and maximum) for the whole dataset. In this multi-illuminant case, the results show that MEM CNN outperforms the other learning-based methods like CNN [3] and the original single estimate FC4 method [2], as well as state of the art statistics-based methods: Gray World [6], White Patch [11], Shades of Gray [5], Grey-Edges [7] and Cheng et al.'s method [8].

In Figure 4 we show outputs of three scenes using four illuminant estimation methods: MEM CNN, FC4 [2], CNN [3] and Gray World [6], and the restored images using the ground-truth illuminants (the ground-truth image). MEM CNN gives good results. With the multi-illuminant estimation, in the given examples, our method is capable of recovering the colours of the scene under a white light without a large shift in hue, which is not the case of e.g. Gray World (first scene), FC4 (second scene) and CNN (third scene).



Figure 3. Input data and results of MEM CNN for three scenes from the MIMO dataset. From top to bottom, left to right: image restored with ground-truth, image restored with estimate map, input image, ground-truth image, estimate map, filtered estimate map. The scenes images have the gamma correction applied for visualisation purposes.



Figure 4. Visual results of different colour constancy methods for three scenes from the MIMO dataset. From top to bottom, left to right: images restored respectively with ground-truth, MEM CNN, FC4 [2], CNN [3], Gray World [6]. The images have the gamma correction applied for visualisation purposes.

Conclusion

We presented in this paper an effective framework for colour constancy, MEM CNN. It is based on a deep convolutional neural network and it solves for illuminant estimation when there are multiple lights in the scene which is very often the case. MEM CNN generates a region-wise estimate map which is divided out from the image to perform colour constancy. It gives competitive results compared to state of the art and recent learning-based methods when evaluated on the MIMO multi-illuminant dataset.

We are limited by the size of the datasets for training specifically when ground-truth illuminants images are needed. A data generator tool like in [14] can be used to create larger datasets. As a future work we would like to evaluate MEM CNN on images with more than two illuminants, where we expect it to perform well (in this work the two illuminants are not uniformly distributed in the scenes) and we would like to train the CNN on a larger dataset and improve the solution to be able to see the small and high frequency changes in the scene illumination.

References

- [1] Beigpour, S. et al. 'Multi-illuminant estimation with conditional random fields', *IEEE Transactions on Image Processing*, 23(1), pp. 83–96 (2014).
- [2] Hu, Y., Wang, B., and Lin, S. 'FC4: Fully Convolutional Color Constancy with Confidence-weighted Pooling', In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4085–4094 (2017).
- [3] Bianco, S., Cusano, C., and Schettini, R. 'Color constancy using CNNs', In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 81–89 (2015).
- [4] Bianco, S., Cusano, C. and Schettini, R. 'Single and Multiple Illuminant Estimation Using Convolutional Neural Networks', *IEEE Transactions on Image Processing*, 26(9), pp. 1–14 (2017).
- [5] Finlayson, G. D. and Trezzi, E. 'Shades of gray and colour constancy', In *Proceedings of the Color and Imaging Conference*, pp. 37–41 (2004).
- [6] Buchsbaum, G. 'A Spatial processor model for object colour perception', *Journal of the Franklin Institute*, 310(1), pp. 1–26 (1980).
- [7] Van De Weijer, J., Gevers, T. and Gijsenij, A. 'Edge-based color constancy', *IEEE Transactions on Image Processing*, 16(9), pp. 2207–2214 (2007).
- [8] Cheng, D., Prasad, D. K. and Brown, M. S. 'Illuminant estimation for color constancy: why spatial-domain methods work and the role of the color distribution', *Journal of the Optical Society of America A*, 31(5), pp. 1049–1058 (2014).
- [9] Beigpour, S. et al., 'Two-Illuminant Dataset with Computed Ground Truth: www5.cs.fau.de/research/data', last accessed in October 2020.
- [10] Results per Dataset, 'www.colorconstancy.com', last accessed in May 2021.
- [11] Land, E. H. and McCann, J. J. 'Lightness and retinex theory', *Journal of the Optical Society of America*, 61(1), pp. 1–11 (1971).
- [12] Iandola, F. N. et al. 'SqueezeNet: AlexNet-level accuracy with 50x fewer parameters and < 0.5 MB model size', *arXiv:1602.07360* (2017).
- [13] Liu, Y. and Shen, S. 'Self-adaptive single and multi-illuminant estimation framework based on deep learning', *arXiv:1902.04705* (2019).
- [14] Banic, N. et al. 'CroP: Color constancy benchmark dataset generator', In *Proceedings of the International Conference on Vision, Image and Signal Processing*, pp. 1–9 (2020).
- [15] Ebner, M. 'Color constancy based on local space average color', *Machine Vision and Applications*, 20(5), pp. 283–301 (2009).
- [16] Land, E. H. 'The Retinex Theory of Color Vision', *Scientific American*, 237(6), pp. 108–128 (1977).
- [17] Gijsenij, A., Rui Lu and Gevers, T. 'Color Constancy for Multiple Light Sources', *IEEE Transactions on Image Processing*, 21(2), pp. 697–707 (2012).
- [18] Cheng, D. et al. 'Two illuminant estimation and user correction preference', In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pp. 469–477 (2016).