

This ACCV 2020 paper, provided here by the Computer Vision Foundation, is the author-created version. The content of this paper is identical to the content of the officially published ACCV 2020 LNCS version of the paper as available on SpringerLink: https://link.springer.com/conference/accv

Single-Image Camera Response Function Using Prediction Consistency and Gradual Refinement

 $\begin{array}{l} \mbox{Aashish Sharma}^{1[0000-0001-9211-3368]}, \mbox{Robby T. } Tan^{1,2[0000-0001-7532-6919]}, \\ \mbox{ and Loong-Fah Cheong}^{1[0000-0001-5259-9113]}, \end{array}$

¹National University of Singapore, ²Yale-NUS College aashish.sharmaQu.nus.edu, {robby.tan, eleclf}Qnus.edu.sg

Abstract. A few methods have been proposed to estimate the CRF from a single image, however most of them tend to fail in handling general real images. For instance, EdgeCRF based on patches extracted from colour edges works effectively only when the presence of noise is insignificant, which is not the case for many real images; and, CRFNet, a recent method based on fully supervised deep learning works only for the CRFs that are in the training data, and hence fail to deal with other possible CRFs beyond the training data. To address these problems, we introduce a non-deep-learning method using prediction consistency and gradual refinement. First, we rely more on the patches of the input image that provide more consistent predictions. If the predictions from a patch are more consistent, it means that the patch is likely to be less affected by noise or any inferior colour combinations, and hence, it can be more reliable for CRF estimation. Second, we employ a gradual refinement scheme in which we start from a simple CRF model to generate a result which is more robust to noise but less accurate, and then we gradually increase the model's complexity to improve the estimation. This is because a simple model, while being less accurate, overfits less to noise than a complex model does. Our experiments confirm that our method outperforms the existing single-image methods for both daytime and nighttime real images.

Keywords: Camera Response Function (CRF); Radiometric Calibration

1 Introduction

In most cameras, the camera irradiance has a non-linear correlation to the image intensities. This is caused by the non-linearity of the camera response function (CRF), which is one of the components of the camera imaging pipeline and is intentionally designed to be non-linear to create more aesthetic effects and perform dynamic range compression [1]. Linearising the image intensities through estimating the CRF is critical for many computer vision algorithms, such as shape from shading [2,3], colour constancy [4,5], photometric stereo [6,7], specular removal [8], shadow removal [9], low-light image enhancement [10], etc.

^{\dagger} This work is supported by MOE2019-T2-1-130.

2



Fig. 1. Results of our single-image CRF estimation method in comparison with those of two existing methods, EdgeCRF [15] and GICRF [18], for both daytime and nighttime image. RMSE is the Root Mean Squares Error between the predicted and true CRF.

Some CRF methods use multiple images of a static scene taken with different exposures [11–13, 7]. To improve the practicality, a few methods utilise a single image, which is considerably more challenging than using multiple images. To accomplish this, the methods require additional constraints, such as: the symmetric nature of the irradiance noise profile [14], or the non-linear distributions of the pixel intensities in the RGB space for pixels around edges [15, 16]. The former constraint is applicable only to images that have a considerable large range of intensity values, which is not the case for many real images. The latter constraint is more applicable to many real images, since it only implies RGB images to have non-uniform coloured patches. Unfortunately, the existing methods that use this constraint are erroneous when noise is present in the input image. Recently, a deep-learning based method [17] using a single image is introduced. This method can work robustly on real images; however, it suffers from the generalization problem. It works effectively only for CRFs represented in the training data.

To address the problem of single-image CRF estimation, we introduce a nondeep-learning method based on the ideas of prediction consistency and gradual refinement. First, unlike the existing methods, our method relies more on patches that provide more consistent CRF predictions. This is based on the idea that if a patch provides more consistent CRF predictions, it is likely to be less affected by noise or any inferior colour combinations, and hence, can be more reliable to estimate the CRF. We put more weight on more reliable patches, which renders more robust and accurate CRF estimation.

Second, unlike the existing methods, our method employs a gradual refinement scheme to improve our CRF result. CRF estimation is generally carried out by optimising a CRF model that can linearise the non-linear pixel distributions with minimum error [15]. There are various models that can be employed to represent CRFs (e.g., [19, 6, 20, 18]). Each model is defined by a set of coefficients, hence the problem of finding the CRF is equivalent to optimizing the coefficients of the CRF model. A model with more coefficients is more accurate, but is also more prone to overfit to noise. Therefore, to gain robustness to noise, our method optimises a simple model first (i.e., one coefficient), and then optimises a more complex model. During the refinement, the CRF result in the current stage is constrained to remain near the CRF result obtained in the previous stage. Fig. 1 shows our estimations for both daytime and nighttime real images.

To summarise, in this paper, our main contributions are as follows:

- Unlike existing methods, our method is designed to deal with noise, making it more applicable to general real images. For achieving this, first, we found that different patches are unequal in terms of noise and how much they carry the CRF information (or degree of non-linearity). For instance, in one input image, patches taken from darker regions are more affected by noise compared to patches taken from well-lit regions; on top of this, in a patch, certain colour combinations carry more CRF information compared to other colour combinations. Second, unlike [15], we found non-edge patches can be used to estimate a CRF. Third, we found noise can significantly alter the degree of non-linearity in a patch, which is critical for a robust method to consider. All these findings are new and important for CRF estimation.
- We propose a new approach to compute the reliability of a patch in estimating CRFs. It is based on the consistency among a patch's CRF predictions. More consistency implies more reliability.
- We introduce a gradual refinement scheme, in which we start from a simple model to obtain results that are robust to noise but less accurate, and then we gradually increase the model's complexity to improve the results. We show that our refinement scheme generates better results than those of oneattempt estimation methods. The idea of gradual refinement is, in various forms, already used in computer vision algorithms, however to our knowledge,

it is the first time it is designed specifically for single-image CRF estimation. To our knowledge, our method is the first method that works robustly in well-lit and dimly-lit scenes. We also show that our method has practical applications such as visibility enhancement under nighttime conditions.

2 Related Work

Several methods have been proposed for CRF estimation. Many previous works rely on a sequence of images of the same scene taken with different exposure times [11–13] or varying illumination [6, 1]. In addition to the requirement of multiple images, these methods also assume that the camera position is fixed. Some other methods [21–25] have relaxed the fixed camera condition by performing image alignment. In contrast to these multi-image methods, our method requires a single image, and hence, is more applicable.

Lin et al. [15] are the first to demonstrate that the CRF can be recovered from a single colour image. Their method is based on the property that pixels around color boundaries in the camera irradiance form linear distributions in the RGB space. However, due to the non-linear CRF, the pixel distributions become non-linear for the intensity image in the RGB space. Thus the CRF is recovered 4

by finding a CRF model that can linearise the distributions with minimum error. Later, the method is extended to work on gravscale images [16]. The ideas of these two methods are elegant, yet unfortunately they are prone to suffer from noise that is commonly present in real images. Matsushita and Lin [14] use noise profile for CRF estimation, based on the observation that the profile is symmetric in the irradiance domain but due to the non-linear CRF, it becomes asymmetric in the intensity domain. However, this method does not work well if there is insufficient data available for generating the noise profile. Li et al. [26] propose to use the low-rank structure of skin pixels to recover the CRF from an image that contains a human face. This method, therefore, is limited to work on images that consist of human faces. Recently, Li and Peers [17] propose a network to estimate the CRF from a single image by predicting 11 coefficients of the basis functions model [19], for which they use the 201 CRFs from [19] during training Kao et al. further propose [27], a stacked version of [17], which shares a similar idea of iterative refinement with our method. Both [17] and [27] are supervised learning-based methods and suffer from the generalization problem. This means they cannot work properly for estimating the CRFs that are not represented in the training set. In contrast, our method is not a learning-based method, and hence, does not suffer from the generalization problem of learning-based methods.

Similar to Lin et al. [15], our method uses the non-linear pixel distributions to estimate the CRF from a single image. However, in contrast to the method, our method uses a consistency based metric to take into account the reliability of the patches (containing the distributions) for estimating the CRF. Additionally, unlike the one-attempt estimation approach of their method, our method uses a gradual refinement scheme that provides more robust and accurate CRF results.

3 Proposed Method

The camera imaging pipeline can be represented by [28]:

$$\mathbf{I} = f(T(\mathbf{E})),\tag{1}$$

where **E** is the camera irradiance in the RGB colour channels, T is a linear operator for colour transformations (for e.g., white balance), f is a non-linear function representing the CRF, and **I** is the RGB image outputted by the camera. Given an image **I** (normalised in [0, 1]) as input, our goal is to estimate the inverse CRF $g = f^{-1}$, such that the image $g(\mathbf{I})$ becomes linearly related to the camera irradiance **E**. Figs. 2a and 2d shows an example of **I** and $g = f^{-1}$ respectively.

3.1 Non-Linearity of Pixel Distributions

From the input image \mathbf{I} , we can obtain a set of non-uniform patches. Let Ω represent that set of the patches, and $|\Omega|$ is the number of the patches in the set. A patch drawn from this set is represented by $\mathbf{I}^p \in \Omega$, where p is the index of the drawn patch, and $p \in [1, |\Omega|]$. Moreover, let the size of the patch \mathbf{I}^p be $s \times s$; then, $\mathbf{I}^{p,k}$ is a set of pixels in the patch (which we call a pixel distribution) where k is



Fig. 2. (a) Input image **I**. (b) Patches around edges in the image **I**. Each patch is of size $s \times s$. The coloured line in a patch shows a pixel distribution scanned horizontally or vertically from that patch. (c) Pixel distributions of the patches plotted in the RGB space. (d) True CRF g.

the index of this set of pixels $(k \in [1, s])$. Fig. 2a shows an example of an input image **I**. Fig. 2b shows an example of some patches, \mathbf{I}^p , selected from **I**. Fig. 2b also shows the patches with a line in each patch, where the line represents the locations of pixels that form a pixel distribution, $\mathbf{I}^{p,k}$.

Let $\mathbf{E}^{p,k}$ be the corresponding pixel distribution in the camera irradiance, and the three colour channels of $\mathbf{E}^{p,k}$ be represented by $\mathbf{E}_{\mathbf{r}}^{p,k}$, $\mathbf{E}_{\mathbf{g}}^{p,k}$ and $\mathbf{E}_{\mathbf{b}}^{p,k}$. When we transform $\mathbf{E}^{p,k}$ into the RGB space, it will form a straight line; since $\mathbf{E}_{\mathbf{r}}^{p,k}$, $\mathbf{E}_{\mathbf{g}}^{p,k}$ and $\mathbf{E}_{\mathbf{b}}^{p,k}$ are linear to each other. Hence, we can express the linear correlation between two colour channels of the camera irradiance as:

$$\mathbf{E}_{\mathrm{g}}^{p,k} = m_{\mathrm{gr}}^{p,k} \mathbf{E}_{\mathrm{r}}^{p,k} + b_{\mathrm{gr}}^{p,k}, \qquad (2)$$

$$\mathbf{E}_{\rm b}^{p,k} = m_{\rm br}^{p,k} \mathbf{E}_{\rm r}^{p,k} + b_{\rm br}^{p,k}, \qquad (3)$$

where $(m_{\rm gr}^{p,k}, b_{\rm gr}^{p,k})$ and $(m_{\rm br}^{p,k}, b_{\rm br}^{p,k})$ are the parameters of the line equations.

Unlike $\mathbf{E}^{p,k}$ that forms a linear distribution, due to the non-linearity of the CRF, $\mathbf{I}^{p,k}$ forms a non-linear distribution in the RGB space, an observation introduced in [15]. Fig. 2c shows an example of a few pixel distributions, $\{\mathbf{I}^{p,k}\}_{p=1}^{4}$, plotted in the RGB space. We can observe the non-linearity in the distributions, which is due to the non-linear CRF shown in Fig. 2d. The reason of the nonlinearity of $\mathbf{I}^{p,k}$ in the RGB space can be explained as follows. Since the three colour channels can have the same CRF and the linear operator T can be a linear identity mapping, from Eq. (1), we can obtain: $\mathbf{E}_{\mathbf{r}}^{p,k} = g(\mathbf{I}_{\mathbf{r}}^{p,k}), \mathbf{E}_{\mathbf{g}}^{p,k} = g(\mathbf{I}_{\mathbf{g}}^{p,k})$ and $\mathbf{E}_{\mathbf{b}}^{p,k} = g(\mathbf{I}_{\mathbf{b}}^{p,k})$. Substituting these terms in Eqs. (2) and (3), we can obtain the following non-linear equations:

$$\mathbf{I}_{g}^{p,k} = g^{-1}(m_{gr}^{p,k} g(\mathbf{I}_{r}^{p,k}) + b_{gr}^{p,k}), \qquad (4)$$

$$\mathbf{I}_{\rm b}^{p,k} = g^{-1}(m_{\rm br}^{p,k} g(\mathbf{I}_{\rm r}^{p,k}) + b_{\rm br}^{p,k}).$$
(5)

No-Noise Case Given any pixel distribution $\mathbf{I}^{p,k}$ from a non-uniform patch \mathbf{I}^{p} , where all the pixels are free from any noise and thus follow Eqs. (4) and (5) strictly, we can estimate the CRF by optimizing a CRF model \hat{g} that can linearise the distribution with minimum error. The subsequent paragraph discusses this idea in more detail.

 $\mathbf{6}$



Fig. 3. For the no-noise case, the CRF can be estimated from any non-uniform patch. For the patches shown in (a), the linearisation error profiles without and with the normalization operation are shown in (c) and (d), respectively. We can observe that all the profiles show the correct global minima, and with the normalization operation, we can remove the local minima observed in (c) that occurs due to the scale problem.

Fig. 3a shows four non-uniform patches taken from a synthetic image with no noise, which is generated using a simple CRF of a gamma function: $\mathbf{I} = f(\mathbf{E}) = \mathbf{E}^{\gamma}$ with $\gamma = 0.4$. Thus, for this case, the CRF model \hat{g} , is the gamma model, where we need to find a gamma value $\hat{\gamma}$ to estimate the CRF. One of the simplest ways to find $\hat{\gamma}$ is to try all possible values of $\hat{\gamma}$ in a certain range; and for each of them, we generate the linearised distribution represented by $(\mathbf{I}^{p,k})^{1/\hat{\gamma}}$. Then, we can do a line fitting, and compute the error of each point from the line. This is basically the linearisation or line fitting error for $(\mathbf{I}^{p,k})^{1/\hat{\gamma}}$. Fig. 3c shows the linearisation error profiles for the four pixel distributions taken from the four patches shown in Fig. 3a, respectively. The error profile for a pixel distribution is obtained by taking $\hat{\gamma}$ values in the range of [0, 1] with increments of 0.02, and then computing the linearisation error for the distribution for all the $\hat{\gamma}$ values. From the results in Fig. 3c, we can observe that all the error profiles show the correct global minima at $\hat{\gamma} = 0.4$, with zero linearisation error. Note that, in this discussion, for the sake of clarity, we use only one pixel distribution, $\mathbf{I}^{p,k}$, for every patch, \mathbf{I}^{p} . In our actual algorithm, we use all the pixel distributions, $\{\mathbf{I}^{p,k}\}_{k=1}^{s}$, for every patch, \mathbf{I}^p , in the set Ω (see Sec. 3.3).

Scale Problem and the Normalization Operation From Fig. 3c, we can also observe that all the error profiles show local minima at $\hat{\gamma} = 0.1$. The reason why this occurs is because the line fitting error is affected by the scale of $(\mathbf{I}^{p,k})^{1/\hat{\gamma}}$. As $\hat{\gamma}$ becomes small, the scale of $(\mathbf{I}^{p,k})^{1/\hat{\gamma}}$ also reduces; which in turn reduces the line fitting error. While this is not a significant problem for synthetic images with no noise, for real images where noise is inevitable, these local minima can possibly become global minima (see Fig. 4c). This is one of the reasons that some existing methods [15, 16] are erroneous.

To address this problem, we propose to normalise $(\mathbf{I}^{p,k})^{1/\hat{\gamma}}$ in each colour channel before computing the line fitting error. Namely, if N represents the normalization operation, we normalise $(\mathbf{I}^{p,k})^{1/\hat{\gamma}}$ such that $\min(N((\mathbf{I}^{p,k}_{c})^{1/\hat{\gamma}}))=0$ and $\max(N((\mathbf{I}^{p,k}_{c})^{1/\hat{\gamma}}))=1$ for $c \in \{r, g, b\}$. The normalization operation makes the line fitting error independent of the scale variations caused by $\hat{\gamma}$ (since it is now the same scale for all values of $\hat{\gamma}$). Fig. 3d shows the profiles after the normalization operation, which now clearly indicate the correct global minima.



Fig. 4. For the noisy case, the linearisation error gets affected by both noise and non-linearity. For the noisy patches in (a), the profiles without the normalization operation shown in (c) show an incorrect global minima. As shown in (d), the normalization operation can alleviate this problem, but depending upon the noise and non-linearity, some distributions (like the purple and blue) can still provide an inaccurate result.

3.2 Effect of Noise on CRF Estimation

For the input image that has no noise, we can recover the CRF from any nonuniform patch, \mathbf{I}^p . However, this is not true when the image has noise. To illustrate this, we generate noisy patches (by adding random noise to the noiseless patches in Fig. 3a). Fig. 4a shows the generated noisy patches, and Fig. 4b shows the pixel distributions from the noisy patches plotted in the RGB space. Following the same procedure in the no-noise case, we compute the linearisation error profiles for the noisy pixel distributions. Fig. 4c and Fig. 4d show the error profiles obtained without and with the normalization operation, respectively.

As shown in Fig. 4c, we can observe that for the noisy patches, most profiles do not show the correct global minima. This is because under the presence of noise, both noise and non-linearity affect the linearisation error. Moreover, we can observe that all the profiles show the global minima at $\hat{\gamma} = 0.1$, which occurs because of the scale problem. Fig. 4d shows the profiles after the normalisation operation that removes the scale problem. As a result, some profiles show the correct global minima, and some do not. This depends on the level of noise and the degree of non-linearity in the pixel distribution in the RGB space.

Eqs. (4) and (5) show that the curviness of a patch (or its distribution $\mathbf{I}^{p,k}$) depends on its colour combination, and thus not all patches carry the CRF information equally. For instance, if the irradiance is achromatic or nearly achromatic, i.e. $\mathbf{E}_{r}^{p,k} \approx \mathbf{E}_{g}^{p,k} \approx \mathbf{E}_{b}^{p,k}$, implying $m_{gr}^{p,k} \approx m_{br}^{p,k} \approx 1$ and $b_{gr}^{p,k} \approx b_{br}^{p,k} \approx 0$, Eqs. (4) and (5) reduce to $\mathbf{I}_{r}^{p,k} \approx \mathbf{I}_{g}^{p,k} \approx \mathbf{I}_{b}^{p,k}$. This means that the line becomes linear for achromatic patches (or achromatic pixel distributions), and they have no information about the CRF. As can be observed in Fig. 3d, the purple coloured profile of the pixel distribution that belongs to a nearly achromatic patch has a considerably shallow minima basin, which means that the distribution is close to being linear. This is what we mean by the degrees of non-linearity of patches (or pixel distributions) are unequal.

Moreover, from Eqs. (4) and (5), we can also observe that due to the presence of both g and g^{-1} in the same equation, there is a "cancellation" effect, which means that in general, the degree of non-linearity in the RGB space is reduced. In other words, the curviness of the pixel distribution in the RGB space is lessened, causing the global minimum to be less obvious. Hence, when the noise comes into the picture, the error profile can be more significantly influenced by the noise, rendering incorrect CRF estimation, which can be observed for the blue and purple coloured distributions in Fig. 4d. Therefore, we need some reliability measure to know the reliability score (which can indicate the degree of genuine non-linearity) of every patch in our set, so that we can rely more on the patches with higher reliability scores to estimate the CRF robustly and accurately. Note that most of the existing methods, particularly [15, 16], assume insignificant noise in the input image and also ignore the patches' varying degrees of non-linearity, and thus treat all the processed patches equally.

3.3 Prediction Consistency for Reliability

To estimate the reliability score of a patch, we propose to use the consistency of the CRF predictions of all the pixel distributions in the patch. In a patch \mathbf{I}^p of resolution $s \times s$, we have s horizontally or vertically scanned pixel distributions, which can provide s estimations of the CRF. If the estimations are more consistent, the reliability score will be higher.

By employing the GGCM model [18] as our CRF model, we can express:

$$\mathbf{I} = f(\mathbf{E}) = (\mathbf{E})^{\gamma_1 + \gamma_2 \mathbf{E} + \dots + \gamma_c(\mathbf{E})^{c-1}},\tag{6}$$

where c is the number of coefficients, and $\{\gamma_1, ..., \gamma_c\}$ are the coefficients of the model. Thus by using the estimated coefficients, $\{\hat{\gamma}_1, ..., \hat{\gamma}_c\}$, we can obtain the estimated CRF represented by: $\hat{g}(\mathbf{x}) = (\mathbf{x})^{\frac{1}{\hat{\gamma}_1 + \hat{\gamma}_2 \mathbf{x} + ... + \hat{\gamma}_c(\mathbf{x})^{c-1}}, \forall \mathbf{x} \in \mathbf{x}$, where x is a set of 100 equidistant values in the range of [0, 1]. Let $\hat{g}^{p,k}(\mathbf{x})$ be the estimated CRF from a pixel distribution $\mathbf{I}^{p,k}$ in a patch. Then, we can compute the consistency between the s estimations of the CRF, $\{\hat{g}^{p,k}(\mathbf{x})\}_{k=1}^{s}$, using:

$$\hat{\sigma}^{p} = \frac{1}{|\mathbf{x}|} \sum_{\mathbf{x}\in\mathbf{x}} \left(\frac{1}{|s|} \sum_{k=1}^{s} \left(\hat{g}^{p,k}(\mathbf{x}) - \frac{1}{|s|} \sum_{k=1}^{s} \left(\hat{g}^{p,k}(\mathbf{x}) \right) \right)^{2} \right).$$
(7)

The reliability score of the patch, $\hat{\alpha}^p$, where: $\hat{\alpha}^p = \exp \frac{-\hat{\sigma}^p}{0.05}$. This score will be used in our gradual refinement scheme to weigh the CRF estimate of the corresponding patch (Sec. 3.4).

To compute the CRF estimate of a single patch, $\hat{g}^{p}(\mathbf{x})$, instead of using the mean of the *s* predictions, we use the mode of the *s* predictions. Since, from our investigation, in the present of noise, the mean is more influenced by the variations in the predictions than the mode. Here are the details. For every pixel distribution, $\mathbf{I}^{p,k}$ in a patch \mathbf{I}^{p} , we can obtain a CRF estimated curve, $\hat{g}^{p,k}(\mathbf{x})$. If we have *s* pixel distributions in the patch, then we have *s* CRF estimated curves, $\{\hat{g}^{p,k}(\mathbf{x})\}_{k=1}^{s}$. We discretise the CRF space (as shown in the examples in Fig. 5) by creating a grid in the space. The grid is a $\Delta \times \Delta$ grid. Meaning, there are Δ columns and Δ rows (where in our implementation $\Delta = 20$, hence our grid has 20×20 cells). Subsequently, we count how many $\{\hat{g}^{p,k}(\mathbf{x})\}_{k=1}^{s}$ that fall into each of the cells. For each column in the grid, we choose the row that has the highest count. If we do this for all columns, we can have our discretised CRF estimated.



Fig. 5. Examples showing the voting mechanism used for obtaining the mode of the CRF predictions. The mode CRF being discretised looks like a staircase function.

Mathematically, we express this discretised CRF estimate as:

 $\hat{g}^{p}(\mathbf{x}) = H\left(\{\hat{g}^{p,1}(\mathbf{x}), \hat{g}^{p,2}(\mathbf{x}), ..., \hat{g}^{p,s}(\mathbf{x})\}, \Delta\right), \forall \, \mathbf{x} \in \mathbf{x},\tag{8}$

where H represents a function that chooses the best CRF for every column or intensity value in the discretised CRF space. In this space, the more correct cells contain higher values (see the examples shown in Fig. 5).

3.4 Gradual Refinement

As we discussed, a CRF model with more coefficients can be more accurate in representing the CRF. However, under the presence of noise, a model with more coefficients is also more prone to overfit to noise. The existing methods ([15, 16, 18]) use a one-attempt optimization approach to directly optimise a complex model (model with a high number of coefficients) which under the presence of noise can cause instability. To address this problem, we propose to initialise the CRF estimation using the one-coefficient model (i.e., a model with only one coefficient). The initial result is more robust to noise but less accurate. We then increase the coefficients of the model to estimate more refined CRF result, which we constrain to remain near the previous CRF result. This is the core idea of our gradual refinement scheme. Also, note that while limiting the values of higher-order coefficients can also act as a regularizer, it can suppress the CRF representation capability, since some CRFs can require high values of the coefficients. Our gradual refinement, however, does not have this problem.

At stage t of the refinement process, where $t \in [1, T]$, for a patch \mathbf{I}^p , we obtain the CRF prediction $\hat{g}_t^{p,k}(\mathbf{x})$ for its pixel distribution $\mathbf{I}^{p,k}$ by optimising the following objective function:

$$\hat{g}_{t}^{p,k}(\mathbf{x}) \equiv \{\hat{\gamma}_{1}^{p,k}, ... \hat{\gamma}_{c}^{p,k}\} = \underset{\{\hat{\gamma}_{1}^{p,k}, ... \hat{\gamma}_{c}^{p,k}\}}{\operatorname{argmin}} \left(\mathcal{L} \left[N \left((\mathbf{I}^{p,k})^{\frac{1}{\hat{\gamma}_{1}^{p,k} + ... + \hat{\gamma}_{c}^{p,k} (\mathbf{I}^{p,k})^{c-1}} \right) \right] + \lambda \left((\mathbf{x})^{\frac{1}{\hat{\gamma}_{1}^{p,k} + ... + \hat{\gamma}_{c}^{p,k} (\mathbf{x})^{c-1}} - \hat{g}_{t-1}(\mathbf{x}) \right)^{2} \right), \forall \mathbf{x} \in \mathbf{x} \right)$$

$$(9)$$

where c is the number of coefficients at stage t, N is the normalization operation, \mathcal{L} is a function that computes the line fitting error, and λ is a parameter that controls the closeness of the CRF prediction $\hat{g}_t^{p,k}(\mathbf{x})$ to the CRF estimated in the previous stage $\hat{g}_{t-1}(\mathbf{x})$. We keep $\lambda=0$ at t=1. Having obtained the predictions

Algorithm 1 Single-Image Camera Response Function using Prediction Consistency and Gradual Refinement

1: Input: Image I. 2: Generate a set of non-uniform patches Ω following the selection process in Sec. 3.5. 3: for stage t where $t \in [1, T]$ do for patch \mathbf{I}^p of resolution $s \times s$ where $\mathbf{I}^p \in \Omega$ do 4: for patch distribution $\mathbf{I}^{p,k}$ where $k \in [1,s]$ do 5:Obtain the CRF estimate for the distribution, $\hat{g}_t^{p,k}(\mathbf{x})$, using Eq. (9). 6: 7: end for 8: Obtain the CRF estimate for the patch, $\hat{g}_t^p(\mathbf{x})$, using Eq. (8). 9: Compute the reliability for the patch, $\hat{\alpha}_t^p$, using Eq. (7). 10: end for Obtain the CRF estimate at stage t, $\hat{g}_t(\mathbf{x})$, using Eq. (10). 11: Update Ω . Remove the patches whose $\hat{\alpha}_t^p < \tau_{\rm re}$. 12:13: end for 14: Obtain the final CRF, $\hat{g}(\mathbf{x})$, from $\hat{g}_T(\mathbf{x})$ using Eq. (11). 15: **Output:** Final CRF $\hat{g}(\mathbf{x})$.

from the patch's pixel distributions, we compute the reliability score of the patch, \hat{a}_t^p , and the CRF estimate from the patch, $\hat{g}_t^p(\mathbf{x})$, (as described in Sec. 3.3).

The CRF estimate at stage t from all the patches, $\hat{g}_t(\mathbf{x})$, is then computed by considering the reliability scores of the patches:

$$\hat{g}_t(\mathbf{x}) = H\left(\{\hat{\alpha}_t^1 \hat{g}_t^1(\mathbf{x}), \hat{\alpha}_t^2 \hat{g}_t^2(\mathbf{x}), ..., \hat{\alpha}_t^{|\Omega|} \hat{g}_t^{|\Omega|}(\mathbf{x})\}, \Delta\right), \forall \mathbf{x} \in \mathbf{x},$$
(10)

where we use the same voting mechanism, except instead of counting the number of predictions for each grid cell, we sum the reliability scores of the predictions for each grid cell. For each column, the row that contains the largest sum is selected.

Practically, to improve the CRF estimation accuracy in the next stage and to keep the runtime of the method small, we remove the patches in Ω whose reliability scores are lower than a certain threshold, $\tau_{\rm re}$. After completion of stage T, i.e. at the end of the refinement process, the estimated CRF $\hat{g}_T(\mathbf{x})$ is a staircase curve. To generate a smooth CRF curve, we fit our *c*-coefficient CRF model on $\hat{g}_T(\mathbf{x})$:

$$\hat{g}(\mathbf{x}) \equiv \{\hat{\gamma}_1, .. \hat{\gamma}_c\} = \underset{\{\hat{\gamma}_1, .. \hat{\gamma}_c\}}{\operatorname{argmin}} \left((\mathbf{x})^{\frac{1}{\hat{\gamma}_1 + \hat{\gamma}_2 \mathbf{x} + ... + \hat{\gamma}_c(\mathbf{x})^{c-1}}} - \hat{g}_T(\mathbf{x}) \right)^2, \forall \mathbf{x} \in \mathbf{x}, \quad (11)$$

where $\hat{g}(\mathbf{x})$ is the final CRF result from our method. See Algorithm 1 for our entire CRF estimation process.

3.5 Selection of Patches

We add a patch into the set Ω if it meets the following criteria:

1. The patch has no under-saturated or over-saturated pixels, i.e. for every pixel in the patch, its magnitude (mean of the three colour channel values) is above $\tau_{\rm us}$ and below $\tau_{\rm os}$, where $\tau_{\rm us}$ and $\tau_{\rm os}$ are the thresholds set for the under-saturated and over-saturated pixels, respectively.

Single-Image CRF Using Prediction Consistency and Gradual Refinement

- 2. The patch is not uniform and contains a mixture of colours, i.e. the variance of all the three colour channel values in the patch is above τ_{un} , where τ_{un} is the threshold set for the uniformity of a patch.
- 3. The patch's pixel distributions are not narrow and are well spread in the RGB space, i.e the variance of a colour channel values in the patch is above τ_{na} , where τ_{na} is the threshold set for narrowness of the distributions.

To select the pixel distributions in a patch, we select the distributions by scanning in either horizontal or vertical direction, depending upon which direction the distributions have more variance.

Note that: (1) unlike [15], which uses solely edge patches, our patch selection method also includes patches with any mixture of colours beyond edges. This gives us a more rich set of patches; (2) The selection criteria above are learned empirically, and can be adjusted depending on the conditions of the target images; (3) Our horizontal-vertical scanning technique is not a hard requirement, as any other techniques can also be used. The only reason we opt for the simple scanning technique is to keep our method efficient.

4 Experimental Results

We evaluate our method on general daytime images taken in indoor and outdoor settings. For this, we use the Color-Constancy dataset [29]. We create a test set of 120 images picked randomly from the dataset. The images are taken from three cameras: NikonD40, Canon600D and SonyA57. In the patch selection process, we use the parameters $\{s, \tau_{us}, \tau_{os}, \tau_{un}, \tau_{na}\} = \{21, 0.15, 0.9, 0.01, 0.065\}$ The parameter for patch resolution, s, is set so that we can have sufficient number of pixels in the pixel distribution to compute the line fitting error. The parameters $\tau_{\rm un}$ and $\tau_{\rm na}$ are tuned to avoid selecting the uniform patches, and inspecting that the distributions are widely spread in the RGB space. We implement our method in MATLAB and use the function fmincon to optimise the coefficients for CRF prediction (Eq. (9)). We set the parameter $\lambda = 0.01$ and $\tau_{\rm re} = 0.3$, that provide the best performance on a separate validation set of 30 images. At stage t, we estimate c = t number of coefficients. We set T=2, as it offers a good tradeoff between our method's accuracy and runtime (as shown in our ablation study). Also, as shown in [18], the GGCM model with two coefficients is more accurate than the other models [19, 6, 20] using the same number of coefficients.

For the baseline methods, we use EdgeCRF [15], CRFNet [17] and GICRF [18]. Similar to ours, EdgeCRF and GICRF are not learning based methods, while CRFNet is a learning based method. Since the codes of EdgeCRF and CRFNet are not available, these methods are based on our implementation. To evaluate the accuracy of the CRF result from each method, we compute Root Mean Squares Error (RMSE) between the method's CRF result $\hat{g}(x)$ and the ground-truth CRF g(x) by: RMSE = $\sqrt{\sum_{x \in x} (\hat{g}(x) - g(x))^2}$ (where the ground-truth CRFs are obtained by using images with Macbeth ColorChecker [15, 16, 18, 19] that are provided by the datasets). For each method, we compute the mean, median, standard deviation, minimum and maximum RMSE values obtained on the entire

12 A. Sharma, R. T. Tan, L.-F. Cheong

 Table 1. Comparisons with the baseline methods on the daytime test dataset. The numbers represent RMSE. Bold font indicates lowest error

Method	Mean	Median	Std	Min	Max
EdgeCRF $[15]$	0.1561	0.1596	0.0402	0.0424	0.2522
CRFNet $[17]$	0.1609	0.1686	0.0538	0.0459	0.2788
GICRF [18]	0.0943	0.0899	0.0340	0.0379	0.2415
Our Method	0.0406	0.0301	0.0308	0.0142	0.2077

 Table 2. Comparisons with the baseline methods on the nighttime test dataset. The numbers represent RMSE. Bold font indicates lowest error

Method	Mean	Median	Std	Min	Max
EdgeCRF $[15]$	0.1738	0.1657	0.0504	0.0593	0.2629
CRFNet $[17]$	0.2407	0.2549	0.1059	0.0681	0.4428
GICRF $[18]$	0.1149	0.1078	0.0533	0.0669	0.2819
Our Method	0.0521	0.0377	0.0350	0.0173	0.1996

test set. The quantitative results are shown in Table 1 and the qualitative CRF results are shown in Fig. 6. From the results, we can observe that our CRF results are more accurate and stable than the baseline methods.

We evaluate our method on nighttime images taken under varying illumination conditions. For this, we use 20 images taken from our own NikonD80 camera, and 30 images from the SID dataset [30] taken from Sony α 7s camera. For patch selection, the same parameters used in the daytime experiment are used, except for $\tau_{\rm us}$ which is relaxed to 0.02 to allow more patches if the images are low-light. We compare our method with the baseline methods, results corresponding to which are shown in Table 2 and Fig. 6 respectively. We can again observe the better performance of our method compared to the baseline methods.

The methods of EdgeCRF and GICRF rely on edges and non-locally planar pixels for their respective algorithms, which can be erroneous under general noisy and nighttime conditions. Especially when the images are near low-light, most of the patches are noisy and nearly achromatic, and they cannot be used for CRF estimation (Sec. 3.2). In contrast, our method estimates the CRF by taking into account the reliability of the patches in generating the CRF estimate. As observed in Sec. 3.2, even if there is a single reliable patch in an image (which can come from the relatively well-illuminated regions in a low-light image), the CRF can be reliably estimated. In addition, contrast to the one-attempt optimisation used by these methods, we use a gradual refinement scheme to gradually improve the CRF results. These factors contribute to our method's better performance for both general daytime and well-to-dim-lit nighttime images.

The results also show the challenges in using a supervised learning based method such as CRFNet for CRF estimation. CRFNet is trained on the 201 CRFs from the DoRF dataset [19]. If the testing CRFs are the same as the training CRFs, then CRFNet performs better than our method (in terms of



Single-Image CRF Using Prediction Consistency and Gradual Refinement 13

Fig. 6. Qualitative comparisons with the baseline methods. As we can observe, compared to the baseline methods, our results are closest to the true CRF, showing our method's robustness in both general daytime and nighttime conditions.



Input Image (11.44) LIME+BM3D (17.23) Ours+SID (24.61) Ground Truth

Fig. 7. We show an application of our CRF estimation method for nighttime visibility enhancement. As we can observe, our enhanced image is closest to the ground truth image (highest PSNR), showing that our enhancement is more physically correct.

mean RMSE, using the same CRF model, CRFNet's score is 0.0201 while ours is 0.0315). However, if the testing CRFs are different (such as the ones used in our experiments), then our method performs better. Since ours is not a learning based method, it can generalise well to different CRFs and imaging conditions.

Application: Nighttime Enhancement We show that our CRF estimation method has practical applications, such as visibility enhancement under nighttime conditions. Most existing enhancement methods either assume the CRF to be linear (LIME [31]) or require the RAW image (irradiance image) as the input (SID [30]). The former generates enhancement in a physically incorrect manner, and the latter though being physically correct, is limited in application since RAW images are not available in most practical situations. If we can estimate the



Fig. 8. (a) Effectiveness of our gradual refinement scheme as the results improve gradually with the refinement stage. (b) Using the normalization operation and the prediction consistency are important for the better performance of our method.

CRF, we can linearise the intensity image such that it becomes linearly related to the irradiance image. Therefore, we can combine our CRF estimation method with SID to create an enhancement method that accepts an RGB image as input, and enhances the image in a physically correct manner. The results are shown in Fig. 7. We can observe that the enhanced image from our combined method is closest to the ground truth (highest PSNR), thus showing that the enhancement achieved by our method is more physically correct.

5 Ablation Study

Gradual Refinement Fig. 8a shows that our performance improves gradually with the refinement stage t. Since our method's runtime also increases with t, we use T=2 by default as it offers a good tradeoff between our method's runtime and performance. Also, without gradual refinement (i.e. using one-attempt optimisation), our method's performance drops and the mean RMSE increases to 0.0831 from 0.0406 on the daytime test dataset.

Normalization and Prediction Consistency We compare our method with two variants: (1) ours without using the normalization operation; and (2) ours without using the prediction consistency. The results are shown in Fig. 8b. We can observe that both the normalization operation and prediction consistency are important factors for the better performance of our method.

6 Conclusion

In this paper, we have presented a new method for CRF estimation from a single image using prediction consistency and gradual refinement. We showed that under the presence of noise, not every patch is reliable, and it is important to take into account its reliability in estimating the CRF. To handle this problem, we proposed to use consistency between a patch's CRF predictions as a measure of its reliability. Our method puts more weight on the more reliable patches that provides more accurate results. In addition, we employed a gradual refinement scheme in the CRF estimation that gradually improves the CRF results. Compared to the existing learning and non-learning based methods, our experiments confirmed that our method has good generalization capability, and it provides more accurate results for both general daytime and nighttime real images.

References

- 1. Kim, S.J., Frahm, J.M., Pollefeys, M.: Radiometric calibration with illumination change for outdoor scene analysis. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2008)
- Zhang, R., Tsai, P.S., Cryer, J.E., Shah, M.: Shape-from-shading: a survey. IEEE transactions on pattern analysis and machine intelligence 21 (1999) 690–706
- Nayar, S.K., Ikeuchi, K., Kanade, T.: Shape from interreflections. International Journal of Computer Vision 6 (1991) 173–195
- Finlayson, G.D., Hordley, S.D., Hubel, P.M.: Color by correlation: A simple, unifying framework for color constancy. IEEE Transactions on Pattern Analysis and Machine Intelligence 23 (2001) 1209–1221
- Tan, R.T., Nishino, K., Ikeuchi, K.: Color constancy through inverse-intensity chromaticity space. JOSA A 21 (2004) 321–334
- Shi, B., Matsushita, Y., Wei, Y., Xu, C., Tan, P.: Self-calibrating photometric stereo. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2010)
- Shi, B., Inose, K., Matsushita, Y., Tan, P., Yeung, S.K., Ikeuchi, K.: Photometric stereo using internet images. In: International Conference on 3D Vision (3DV). (2014)
- Tan, R.T., Ikeuchi, K.: Separating reflection components of textured surfaces using a single image. IEEE Trans. Pattern Anal. Mach. Intell. 27 (2005) 178193
- Finlayson, G.D., Drew, M.S., Lu, C.: Entropy minimization for shadow removal. International Journal of Computer Vision 85 (2009) 35–57
- Ying, Z., Li, G., Ren, Y., Wang, R., Wang, W.: A new low-light image enhancement algorithm using camera response model. In: IEEE International Conference on Computer Vision Workshops (ICCVW). (2017)
- Debevec, P.E., Malik, J.: Recovering high dynamic range radiance maps from photographs. In: Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH). SIGGRAPH '97 (1997)
- Mitsunaga, T., Nayar, S.K.: Radiometric self calibration. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (1999)
- 13. Mann, S.: Comparametric equations with practical applications in quantigraphic image processing. IEEE Transactions on Image Processing (TIP) **9** (2000) 1389–1406
- 14. Matsushita, Y., Lin, S.: Radiometric calibration from noise distributions. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2007)
- Lin, S., Gu, J., Yamazaki, S., Shum, H.Y.: Radiometric calibration from a single image. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2004)
- Lin, S., Zhang, L.: Determining the radiometric response function from a single grayscale image. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2005)
- 17. Li, H., Peers, P.: Crf-net: Single image radiometric calibration using cnns. In: European Conference on Visual Media Production (CVMP). (2017)
- Ng, T.T., Chang, S.F., Tsui, M.P.: Using geometry invariants for camera response function estimation. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2007)
- Grossberg, M.D., Nayar, S.K.: Modeling the space of camera response functions. IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI) 26 (2004) 1272–1282

- 16 A. Sharma, R. T. Tan, L.-F. Cheong
- Lee, J.Y., Matsushita, Y., Shi, B., Kweon, I.S., Ikeuchi, K.: Radiometric calibration by rank minimization. IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI) 35 (2012) 144–156
- Kim, S.J., Pollefeys, M.: Robust radiometric calibration and vignetting correction. IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI) 30 (2008) 562–576
- Litvinov, A., Schechner, Y.Y.: Addressing radiometric nonidealities: A unified framework. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2005)
- Mann, S., Mann, R.: Quantigraphic imaging: Estimating the camera response and exposures from differently exposed images. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2001)
- Park, J., Tai, Y.W., Sinha, S.N., So Kweon, I.: Efficient and robust color consistency for community photo collections. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2016)
- Díaz, M., Sturm, P.: Radiometric calibration using photo collections. In: IEEE International Conference on Computational Photography (ICCP). (2011)
- Li, C., Lin, S., Zhou, K., Ikeuchi, K.: Radiometric calibration from faces in images. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2017)
- Kao, Y.L., Chen, Y.S., Ouhyoung, M.: Progressive-crf-net: single image radiometric calibration using stacked cnns. In: ACM SIGGRAPH 2018 Posters. (2018) 1–2
- Kim, S.J., Lin, H.T., Lu, Z., Süsstrunk, S., Lin, S., Brown, M.S.: A new in-camera imaging model for color computer vision and its application. IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI) 34 (2012) 2289–2302
- Cheng, D., Prasad, D.K., Brown, M.S.: Illuminant estimation for color constancy: why spatial-domain methods work and the role of the color distribution. J. Opt. Soc. Am. A **31** (2014) 1049–1058
- 30. Chen, C., Chen, Q., Xu, J., Koltun, V.: Learning to see in the dark. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2018)
- Guo, X., Li, Y., Ling, H.: Lime: Low-light image enhancement via illumination map estimation. IEEE Transactions on Image Processing 26 (2016) 982–993