UvA-DARE (Digital Academic Repository)

Invariant Descriptors for Intrinsic Reflectance Optimization

Baslamisli, A.S.; Gevers, T.

DOI

[10.1364/JOSAA.414682](https://doi.org/10.1364/JOSAA.414682)

Publication date

2021

Document Version

Final published version

Published in

Journal of the Optical Society of America. A, Optics, Image Science and Vision

License

Article 25fa Dutch Copyright Act

[Link to publication](#)

Citation for published version (APA):

Baslamisli, A. S., & Gevers, T. (2021). Invariant Descriptors for Intrinsic Reflectance Optimization. *Journal of the Optical Society of America. A, Optics, Image Science and Vision*, 38(6), 887-896. <https://doi.org/10.1364/JOSAA.414682>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Invariant descriptors for intrinsic reflectance optimization

ANIL S. BASLAMISLI^{1,*} AND THEO GEVERS^{1,2}

¹University of Amsterdam, Science Park 904, 1098XH Amsterdam, The Netherlands

²3DUniversum, Science Park 400, 1098XH Amsterdam, The Netherlands

*Corresponding author: a.s.baslamisli@uva.nl

Received 10 November 2020; revised 6 February 2021; accepted 11 May 2021; posted 12 May 2021 (Doc. ID 414682); published 1 June 2021

Intrinsic image decomposition aims to factorize an image into albedo (reflectance) and shading (illumination) sub-components. Being ill posed and under-constrained, it is a very challenging computer vision problem. There are infinite pairs of reflectance and shading images that can reconstruct the same input. To address the problem, Intrinsic Images in the Wild by Bell *et al.* provides an optimization framework based on a dense conditional random field (CRF) formulation that considers long-range material relations. We improve upon their model by introducing illumination invariant image descriptors: color ratios. The color ratios and the intrinsic reflectance are both invariant to illumination and thus are highly correlated. Through detailed experiments, we provide ways to inject the color ratios into the dense CRF optimization. Our approach is physics based and learning free and leads to more accurate and robust reflectance decompositions. © 2021 Optical Society of America

<https://doi.org/10.1364/JOSAA.414682>

1. INTRODUCTION

Intrinsic image decomposition aims to factorize an image into albedo (reflectance) and shading (illumination) sub-components [1]. The reflectance component represents pure material colors of a scene, independent of any illumination effect including geometry and camera viewpoint. The shading component is invariant to color (albedo) and encodes the illumination effects within a scene such as shadows, shading due to geometry, and ambient light. Due to their invariant properties, intrinsic images are favored by numerous computer vision and computational photography applications. For example, semantic segmentation algorithms may profit from the illumination invariant properties of the reflectance map for robust estimations [2]. In addition, intrinsic images can be used for physically plausible photo editing tasks such as recoloring [3], color transfer, and photo fusion [4].

In an ideal diffuse environment, element-wise multiplication of the reflectance R and shading S components properly approximates the observed image I :

$$I = R \times S. \quad (1)$$

Being ill posed and under-constrained, intrinsic image decomposition is a very challenging computer vision problem. Given an observed image, there exist twice as many unknowns, reflectance and shading, to predict. Hence, there are infinite pairs of reflectance and shading images that can reconstruct the same input. To address the problem, early methods adapt heuristics-based priors. For example, Land and McCann's

pioneering Retinex algorithm is based on the assumption that the reflectance images can directly be estimated by strong image gradients using a surface reconstruction algorithm [5]. Recent approaches use either an energy minimization (optimization) framework to constrain intrinsic image estimations based on handcrafted priors [6,7] or data driven deep-learning-based techniques [8,9].

One of the key advantages of optimization-based methods is that they do not require any labeled data, whereas deep-learning-based methods demand considerable amounts of training data to learn the model parameters. Optimization-based methods aim to regularize the intrinsic component computation to constrain their behavior based on handcrafted priors. Those priors are mostly based on observations, for example, global reflectance sparsity [10,11] and piece-wise constant reflectance [5,12]. However, physics-based invariant descriptors are mostly ignored for crafting priors. One example is the photometric invariant color ratios [13]. They eliminate shading information and thus represent an image featuring only albedo edges. In that sense, the color ratios and the intrinsic reflectance are both invariant to illumination and thus are highly correlated. However, this correlation has never been explored before. Thus, to our knowledge, this is the first work to compute dense intrinsic reflectance maps using color ratios.

Therefore, in this paper, we explore the use of photometric invariant descriptors in an optimization scheme to compute reflectance images. We demonstrate the effectiveness of such an approach by integrating the color ratios into

the dense conditional random field (CRF) optimization framework of Bell *et al.* [6]. Their model is publicly available (<https://github.com/seanbell/intrinsic>), and their CRF framework considers long-range relative material relations as in the case of color ratios. Our extensions to their code are straightforward, fully unsupervised, do not demand any additional data such as depth or image sequences, and do not require any domain calibration.

In particular, the main contribution of this work is not to chase the highest performance numbers, but to put forward a novel hypothesis that physics-based color ratios can be used to improve and robustify reflectance predictions. To sustain our hypothesis, (1) the ratios are integrated in the Color Retinex paradigm to provide complementary color cues. The results are particularly important, as almost all the optimization-based intrinsic image decomposition methods consider the Color Retinex approach in their framework. Thus, a more robust prior can be achieved. (2) The ratios are used to define an adaptive clustering initialization mechanism. Current work using the global reflectance sparsity prior treats the number of clusters as a hyper-parameter and decides it through a grid search. However, the process is costly, not effective, and dataset dependent. (3) The ratios are embedded into a pairwise term that uses handcrafted priors based on observations. Unlike handcrafted priors, the ratios are physics based and do not need domain calibration. (4) Finally, we demonstrate that these additions lead to reflectance predictions that are robust to outdoor natural shadow handling. Experiments on four different datasets demonstrate the merits of the proposed method.

2. RELATED WORK

The intrinsic image decomposition task is a very challenging computer vision problem, as it is ill posed and under-constrained. Given a single pixel value, it seeks to find the two unknown disentangled variables: reflectance and shading. As a result, different combinations of the reflectance and shading may achieve the same input. An example is illustrated in Fig. 1.

The pioneering work Retinex by Land and McCann, based on Mondrian patterns, which are composed of piece-wise distinctly colored patches, identifies strong image gradients assuming they correspond to true albedo changes [5]. Then, those selected gradients are re-integrated by a surface reconstruction algorithm such as Poisson. Successive research focuses on deriving priors, based on observations, that can explain the characteristics of the intrinsic images. Usually, an optimization

approach is taken to impose these constraints. For example, Gehler *et al.* and Shen *et al.* consider that the number of colors in a scene is limited; thus, they sample reflectances from a sparse set of colors [10,11]. Shen *et al.* and Zhao *et al.* observe that the distinct points with intensity-normalized texture structures tend to have the same reflectance values [14,15]. In a similar manner, Shen *et al.* assume that a local neighborhood of pixels with similar intensity values also has similar reflectance values [16]. Similarly, Garces *et al.* assume that changes in chromaticity usually correspond to changes in reflectance [17]. Jiang *et al.* make the observation that correlations between local mean luminance and local luminance in textured regions indicate illumination changes [18]. Inspired by the Retinex algorithm's assumption of piece-wise constant reflectances, other works impose piece-wise reflectance priors utilizing various regularization techniques [12,19–22]. Additional user guidance is also explored as a prior to associate reflectance values [16,23]. Moreover, when available, extra sensory information can be incorporated to obtain specific priors. For instance, Chen and Koltun, Jeon *et al.*, and Lee *et al.* explore supplementary depth and surface normal cues to constrain shading estimations [24–26]. More recently, near-infrared imagery-based priors have been explored by Cheng *et al.* [7]. Finally, image sequences are utilized to impose constant reflectance priors over time [27–29]. To conclude, single image intrinsic image decomposition based on optimizations usually derives handcrafted priors based on observation. In contrast, we propose to incorporate physics-based invariant descriptors that are not limited by the imaging conditions or the generalization quality of the handcrafted features.

Invariant descriptors are independent of a set of imaging conditions; thus, they represent simplified versions of the observed images. To derive them, physics-based reflection models are used. An example is the color ratios by Finlayson [13]. As illuminant invariants, they are used to construct histogram features for object recognition. Matas *et al.* integrate ratios into a graph framework for efficient object recognition [30]. Nayar and Bolle employ color ratios for robust pose estimation [31]. Barnard and Finlayson identify possible shadow segments by color ratios for color constancy and dynamic range compression [32]. Color ratio gradients are also employed for content-based image retrieval of non-uniform texture objects [33]. Later, Gevers and Smeulders extend the idea of color ratios to cross color ratios that are also invariant to the scene geometry [34]. Not only cross color ratios are independent of illumination and geometry, but also the intrinsic reflectance by definition. Thus, cross color ratios emphasizing true color variations are expected to provide

Fig. 1. Ill-posed nature of the problem. The left part is an incorrect intrinsic image decomposition, whereas the right part presents the ground-truth one. Both achieve the same input image through $R \times S = I$.

cues for intrinsic images. Nonetheless, this natural correlation between intrinsic images and cross color ratios has never been incorporated as prior knowledge before. To that end, we suggest the use of photometric invariant descriptors in optimization to compute intrinsic reflectance. Unlike handcrafted priors that are based on observation and require domain expertise, cross color ratios are based on physics-based reflections models. They are directly computed from the observed RGB image in a learning-free manner (unsupervised) and do not require additional sensory information or calibration.

Powerful deep-learning-based methods are also widely used for the task. Narihira *et al.* introduced the first work to predict reflectance and shading from a single RGB image in an end-to-end manner [8]. Shi *et al.* predict diffuse albedo, shading, and specular highlights from a single RGB image using millions of synthetic images of objects [35]. Baslamisli *et al.* apply a Retinex approach in a deep learning framework [9]. Li and Snavely utilize image sequences and a constant reflectance prior in a deep learning model [36]. Lettry *et al.* investigate adversarial learning [37]. Baslamisli *et al.* propose fine-grained shading decomposition [38]. We refer readers to the work of Sial *et al.*, which provides a comprehensive overview of deep-learning-based methods and large-scale datasets [39].

3. APPROACH

A. Image Formation Model

To describe an image I over the visible spectrum ω , we use the diffuse dichromatic reflection model for three color channels $c \in \{R, G, B\}$ as follows [40]:

$$I_c = m(\vec{n}, \vec{l}) \int_{\omega} f_c(\lambda) e(\lambda) s(\lambda) d\lambda. \quad (2)$$

In the equation, m is a function modeling the interaction between the surface normal $\vec{n}$ and the incoming light source direction $\vec{l}$ [e.g., Lambertian $m(\vec{n}, \vec{l}) = (\vec{n} \cdot \vec{l})$]. λ represents the wavelength, f indicates the camera spectral sensitivity, e describes the spectral power distribution of the light source, and s denotes the surface reflectance, i.e., the albedo. Then, assuming a linear sensor response and narrowband filters ($f_c(\lambda_c)$) [41], the equation can be simplified as follows:

$$I_c = m(\vec{n}, \vec{l}) e(\lambda_c) s(\lambda_c) = m(\vec{n}, \vec{l}) e_c s_c. \quad (3)$$

The equation models an image I by the multiplication of surface geometry and light interaction $m(\vec{n}, \vec{l})$, reflectance s_c , and light source properties e_c per pixel. Using Eq. (3), intrinsic images are defined (per pixel) as follows:

$$\begin{aligned} I_c &= S_c \times R_c, \\ S_c &= m(\vec{n}, \vec{l}) e_c, \\ R_c &= s_c, \end{aligned} \quad (4)$$

where an image I_c is modeled by the element-wise multiplication of its shading S_c and reflectance R_c components (intrinsic). Most works assume white light, but if the light source e is colored, then that color information is embedded in the shading (illumination) component.

B. Color Ratios

For two neighboring (adjacent) pixels x_1 and x_2 , locally constant illumination can be assumed $e_c^{x_1} = e_c^{x_2}$ [5]. Based on this property, color ratios F are computed for $c \in \{R, G, B\}$ at two neighboring pixels x_1 and x_2 as follows:

$$F_1 = \frac{R^{x_1}}{R^{x_2}}, \quad F_2 = \frac{G^{x_1}}{G^{x_2}}, \quad F_3 = \frac{B^{x_1}}{B^{x_2}}. \quad (5)$$

We illustrate the invariant properties of (single) ratios by plugging Eq. (3) into Eq. (5) for F_1 as follows (holds also for F_2 and F_3):

$$F_1 = \frac{(m(\vec{n}, \vec{l}))^{x_1} e_{R^{x_1}} s_{R^{x_1}}}{(m(\vec{n}, \vec{l}))^{x_2} e_{R^{x_2}} s_{R^{x_2}}}. \quad (6)$$

Since two neighboring pixels x_1 and x_2 share the same locally constant illumination, $e_{R^{x_1}}$ and $e_{R^{x_2}}$ are cancelled out in the equation. If we assume that neighboring pixels share the same geometry $(m(\vec{n}, \vec{l}))^{x_1} = (m(\vec{n}, \vec{l}))^{x_2}$, i.e., locally smooth surfaces, then the ratios are invariant to photometric effects.

C. Cross Color Ratios

To overcome the geometry constraint of Eq. (6), cross color ratios are preferred, which are not only independent of the illumination conditions but also invariant to the object geometry. For neighboring pixels x_1 and x_2 , cross color ratios are defined as

$$M_1 = \frac{R^{x_1} G^{x_2}}{R^{x_2} G^{x_1}}, \quad M_2 = \frac{R^{x_1} B^{x_2}}{R^{x_2} B^{x_1}}, \quad M_3 = \frac{G^{x_1} B^{x_2}}{G^{x_2} B^{x_1}}. \quad (7)$$

As a result, the assumption that neighboring points share the same geometry, i.e., locally smooth surfaces, is no longer required for the cross color ratio that is invariant to photometric effects. We illustrate the invariant properties of the cross ratios by plugging Eq. (3) into Eq. (7) for M_1 as follows (holds also for M_2 and M_3):

$$M_1 = \frac{(m(\vec{n}, \vec{l}))^{x_1} e_{R^{x_1}} s_{R^{x_1}} (m(\vec{n}, \vec{l}))^{x_2} e_{G^{x_2}} s_{G^{x_2}}}{(m(\vec{n}, \vec{l}))^{x_2} e_{R^{x_2}} s_{R^{x_2}} (m(\vec{n}, \vec{l}))^{x_1} e_{G^{x_1}} s_{G^{x_1}}} = \frac{s_{R^{x_1}} s_{G^{x_2}}}{s_{R^{x_2}} s_{G^{x_1}}}. \quad (8)$$

Since two neighboring pixels x_1 and x_2 share the same locally constant illumination, $e_{R^{x_1}}$ and $e_{R^{x_2}}$, and $e_{G^{x_1}}$ and $e_{G^{x_2}}$ are cancelled out. Moreover, since R^{x_1} and G^{x_1} and R^{x_2} and G^{x_2} share the same local geometry, $(m(\vec{n}, \vec{l}))$ components are also cancelled out. Therefore, Eq. (7) can be used to compute a measure for photometric invariance without any additional constraint. The ratio encodes a relative reflectance relation between two pixels. If there is no intrinsic color change (reflectance between x_1 and x_2), then the ratio is one. Sensor artifacts or noise may slightly deviate the value from one. On the other hand, when the ratio deviates significantly from one, it manifests an intrinsic color change. A threshold can be defined or learned to differentiate significant deviations, identifying true color changes. This property can also be used to measure the number of distinct color edges in an image, which is explained in Section 4.B.

D. Intrinsic Image Optimization

An overview is given for the optimization framework of [6]. The problem is formulated by finding the decomposition (reflectance R^* and shading S^*) that most likely matches the RGB image I based on a number of priors under probability distribution p as follows:

$$R^*, S^* = \operatorname{argmax}_{R, S} p(R, S | I). \quad (9)$$

A number of priors are utilized such that the framework assigns a high probability to the decompositions that are consistent with these priors. In particular, the priors are: (i) pixels that are nearby and have similar chromaticity or intensity should also have similar reflectance, (ii) reflectances are piecewise-constant (Retinex), (iii) reflectances are sampled from a sparse set (global reflectance sparsity), (iv) certain shading values are *a priori* more likely than others, (v) neighboring pixels have similar shading (shading smoothness), and (vi) shading is grayscale, or the same color as the light source (color constancy). Then, the framework incorporates the priors in a global sense using a dense CRF.

First, a set of possible labels (colors) is selected from the observed RGB image as the initial reflectance predictions using clustering. The clustering is based on the chromaticity computed from the RGB image, which separates reflectances under ideal conditions. Then, each reflectance pixel is labeled with a chromaticity value chosen from the clusters such that $p(R, S | I)$ is maximized. Second, the reflectances are tuned by minimizing the discontinuities in the shading estimations. That is achieved by the optimization process aiming to minimize the energy function $E(x)$, composed of pairwise $E_p(x)$ and unary $[E_s(x) \text{ and } E_l(x)]$ costs, by imposing the prior constraints:

$$E(x) = \omega_p E_p(x) + \omega_s E_s(x) + \omega_l E_l(x), \quad (10)$$

where $E_p(x)$ is the pairwise reflectance term that forces pixels that are nearby in position, chromaticity, and intensity to be assigned to the same surface reflectance, $E_s(x)$ enforces the shading smoothness, $E_l(x)$ penalizes extreme values of shading for too many pixels, and x denotes the optimal labeling (chromaticity value) for a pixel.

4. EXPERIMENTS

We show the potential merits of using color ratios by conducting the following experiments: (1) an ablation study of the cross color ratios on the Color Retinex performance, (2) the effect on the cluster initialization scheme, (3) the influence of the pairwise term, and (4) the performance on real world shadow removal. Experiments are provided on the real world object level MIT intrinsics dataset [42] and real world scene level Intrinsic Images in the Wild (IIW) of complex indoor scenes for reflectance predictions [6]. We report the mean, median, and trimean values to measure the central tendency of the error metrics to evaluate the robustness gains. Following the common practice, the local mean-squared error (LMSE) with a window size of 20 for the MIT dataset and the weighted human disagreement rate (WHDR) for IIW, as provided by the authors, are reported as error metrics. Lower values are better for both cases. Finally, we demonstrate the merits of our approach for the shadow removal

application on the Image Shadow Triplets Dataset (ISTD) [43] and the Shadow Removal Dataset (SRD) [44] of real world outdoor scenes with a variety of qualitative evaluations.

A. Combination with Color Retinex

As an ablation study, we integrate cross color ratios, Eq. (7), into the Color Retinex algorithm. The Color Retinex algorithm identifies reflectance changes by two threshold values, one for strong brightness changes and the other for strong chromaticity changes. If a pixel satisfies both thresholds, then it is classified as a reflectance change. Then, using a surface reconstruction algorithm, those pixels are re-integrated to achieve the final reflectance estimation. Although powerful, the algorithm is based on the Mondrian world hypothesis assuming piece-wise distinctly colored patches for the gradient separation. However, that is not always the case for real world images, as they might contain weak color transitions or very strong photometric effects (e.g., shadow casts) that also cause strong gradient changes. Therefore, to show the merits, we integrate the ratios into the Color Retinex paradigm to provide complementarity making the algorithm more robust. For the experiments, we use the Color Retinex implementation of [42].

For the integration, we first compute the cross color ratios by Eq. (7) in RGB space. Before computation, Gaussian smoothing is applied to the image. To merge the ratios and achieve a unified map, we combine the three ratios by their geometric mean: $\sqrt[3]{M1 \times M2 \times M3}$. It is applied in log space to avoid numerical instabilities. Note that the arithmetic mean can also be preferred in RGB space. This way, all color information is properly utilized. One might also use a single color ratio such as $M1$ (red-green), yet given an image dominated by blue color, the ratio will not be sufficiently responsive. Finally, a small threshold value of 0.02 is set to eliminate possible sensor artifacts and noise.

The combination is achieved such that the response maps of Color Retinex and cross color ratios are fused using the *logical or* operation. Therefore, they become complementary to each other. If a pixel is labeled not reflectance by both, then it is discarded, whereas if a pixel is labeled as reflectance at least by one of the algorithms, it is marked as true reflectance change. We demonstrate our contribution on the MIT intrinsics dataset of real world objects and IIW real world indoor scenes for the reflectance predictions in Table 1.

The results show that the cross color ratios further improve the Color Retinex algorithm, and performance gains are observed. All three metrics—mean, median, and trimean—are improved. The improvements are more significant for median

Table 1. Combination of Color Retinex and Cross Color Ratios (CCR)^a

	Color Retinex	CCR-Assisted Color Retinex
Mean (LMSE)	0.0315	0.0313
Median (LMSE)	0.0303	0.0295
Trimean (LMSE)	0.0314	0.0309
Mean (WHDR)	33.88	33.86
Median (WHDR)	33.10	33.00
Trimean (WHDR)	33.26	33.14

^aThe ratios further improve Color Retinex on all metrics and make it more robust.

and trimean metrics such that Color Retinex becomes more robust. As a result, simply using an elementary setup, we manage to improve the reconstruction quality of the Color Retinex's reflectance estimations for both object and scene level cases. Note that the computation of the ratios does not add any overhead, is completely learning free, and realized in real time. They are computed directly from the input RGB image (unsupervised). The results are particularly important, as a large number of works consider the Color Retinex approach in their optimizations [7,12,20–22,45]. The results suggest that future work intending to use the Color Retinex gradient constraint should also consider the cross color ratios for improved and more robust estimations.

B. Adaptive Cluster Initialization

The optimization of Eq. (9) is initialized by k-means clustering of pixel features. RGB pixels are transformed into pixel intensity, red chromaticity, and green chromaticity as features, because chromaticity information perfectly separates reflectances under ideal conditions. Then, each pixel is labeled with a reflectance by the cluster centers. That imposes the global reflectance sparsity constraint [6,10].

Mostly, the number of clusters is treated as a hyper-parameter and decided through a grid search. However, the process is costly, not effective, and dataset dependent. It needs to be adjusted per domain. The task is an active area of research (e.g., [46,47]). Here, we take a different approach and use the cross color ratios to adaptively determine the number of clusters. In a homogeneously (single) colored surface, the ratios are constant (i.e., one). They significantly deviate from one when there is a color change. Thus, the number of distinct color ratio values correspond to the number of distinct colors in an image. Therefore, the number of distinct color ratio values can be used

Table 2. Effect of k for the K-Means Algorithm for the MIT Dataset^a

	Mean (LMSE)	Median (LMSE)	Trimean (LMSE)
Fixed $k = 10$	0.0484	0.0400	0.0438
Fixed $k = 20$ (default)	0.0597	0.0434	0.0506
Fixed $k = 30$	0.0499	0.0397	0.0437
Adaptive k by ratios	0.0443	0.0356	0.0388
Adaptive k by ratios w/ratio features	0.0431	0.0353	0.0366

^aAdaptive setting of the k value significantly improves the reflectance estimations.

to determine the number of (true color) clusters per image. To take advantage of this property, we represent each RGB pixel x as its cross color ratios ($M1(x)$, $M2(x)$, $M3(x)$) computed by Eq. (7). When at least one of the three features significantly deviates from one, we count it as a unique color. The same ratios are counted only once. We define the significance by the *round* function. As a result, the number of clusters can be adaptively set for each image, unlike other methods that use the same (fixed) number of clusters for all images. Finally, we include cross color ratios as features in the clustering. The ratios are first normalized by the maximum value, then weighted by 0.5 for MIT and 10 for IIW experiments. Our contribution is presented in Table 2 and demonstrated in Fig. 2 for the MIT dataset and in Table 3 for IIW. An advanced Retinex model [48] is also included as a comparison.

For the MIT dataset, Fig. 2 shows that the default approach produces extra clusters for shadows and strong shadings. On the other hand, strong shadings are decently handled, and the region under the strong shadow is visible now in the *raccoon* object due to the adaptive setting of the k value. In addition, Table 2 shows that the ratio driven adaptive k leads to more

Fig. 2. Effect of the ratio driven clustering. Default model extra clusters shadows and strong shadings as reflectance. Adaptive setting k by the ratios makes the model more accurate and robust to photometric effects such as strong shading and shadows.

Table 3. Effect of k for the K-Means Algorithm for the IIW Dataset^{a, b}

	Mean (WHDR)	Median (WHDR)	Trimean (WHDR)
STAR [48]	31.88	32.24	32.05
Fixed $k = 10$	21.23	18.59	19.58
Fixed $k = 20$ (default)	20.08	17.91	18.26
Fixed $k = 70$	19.80	17.52	18.50
Adaptive k by ratios w/ratio features	19.50	17.06	17.55
Adaptive k by ratios w/ratio features*	17.51	15.33	15.87

^aAdaptive setting of the k value further improves the reflectance estimations for indoor scene level images.

^bFurther improvements are also achieved with a guided filter post-processing (*) [49].

robust and accurate results improving all three metrics. Finally, additionally including the cross color ratios as features in the clustering scheme further boosts the reconstruction performance. For the IIW dataset, Table 3 demonstrates that the ratio driven adaptive k leads to more robust and accurate results also for scene level images. The improvements are more significant for median and trimean metrics such that the ratios make the algorithm more robust. Further improvements can also be achieved with a guided filter post-processing [49].

Calculation of the features ($M1(x)$, $M2(x)$, $M3(x)$) is realized in real time and does not introduce any overhead. The experiments suggest that a global reflectance sparsity prior is beneficial considering the ratio driven adaptive setting of the number of clusters. Therefore, future work intending to use the global reflectance sparsity prior should also consider the adaptive setting of the number of clusters by the ratios for improved and more robust estimations.

C. Pairwise Reflectance

Finally, we inject the color ratios into the pairwise reflectance term $E_p(x)$ of Eq. (10). The term forces pixels that are nearby in position, chromaticity, and intensity to have the same reflectance value. Under white illumination, two pixels having the same reflectance should also have the same chromaticity. However, the white illumination assumption is not realistic for real world scenes. Thus, it is limited by real world applications. Therefore, we modify the function such that it also considers cross color ratios along with nearby pixels in position, pixel chromaticity, and intensity. Thus, the term is further constrained by physics-based descriptors that are not limited in any assumption such as white light. The ratio feature is defined as the geometric mean of the three cross color ratios, the same as in Section 4.A.

Fig. 3. Reflectance evaluations for IIW dataset. Problematic parts are marked with green bounding boxes. The final model further handles discontinuities in the reflectance. It becomes more robust to direct light effects and also to specular highlights.

We further weight it with a Gaussian to be suitable for CRF. The results are provided in Fig. 3 for the scene level IIW estimations. It can be observed that the ratios further handle discontinuities in the reflectance map estimations, making them more robust to direct light effects and also to specular highlights.

In addition, our final model contribution employing also adaptive k-means is presented in Fig. 4 and in Table 4 for the MIT dataset. An advanced Retinex model [48] is included as a comparison. The quantitative results demonstrate that the final model achieves significantly better results on all metrics due to the cross color ratio injections. The visual results demonstrate a significant degree of photometric invariance, capable of handling even strong shadow casts.

Fig. 4. Additional effect of the cross color ratios as a pairwise term. The full model has a significant degree of photometric invariance, capable of handling strong shadows.

Fig. 5. Default model completely fails handling shadow casts. The final proposed model driven by the photometric invariant color ratios is more robust to natural outdoor real world shadow handling. It can now differentiate drastic changes in pixel values and attribute them to the related intrinsics.

Table 4. Effect of the Additional Pairwise Term for the MIT Dataset^a

	Mean (LMSE)	Median (LMSE)	Trimean (LMSE)
STAR [48]	0.0478	0.0438	0.0451
Default model	0.0597	0.0434	0.0506
Final model	0.0424	0.0325	0.0349

^aIt further improves the reflectance estimations on all metrics and makes the model more robust.

D. Scene Level Real World Shadow Removal Application

In this section, we evaluate the proposed algorithm on a real world application. Reflectance images being photometric invariants should be shadow free by definition. Therefore, we demonstrate the significant benefits that the ratios bring to the reflectance estimations' shadow handling performance. To this end, we provide qualitative results on the ISTD [43] and the SRD [44]. The results are presented in Fig. 5. It can be observed that the photometric invariant ratios further robustify the model. The final model can handle strong shadow casts in real world natural scenes significantly better. We can handle even shadows on monochromatic surfaces, as illustrated in the second row. The default model fails and produces strong shadow artifacts polluting the reflection estimations in all cases.

E. Comparison with Prior State of the Art

Finally, in this section, we provide a deep qualitative comparison on ISTD and SRD with state-of-the-art methods: (1) Krebs *et al.* recover the parameters of the dichromatic reflection model using quadratic programming [50]; (2) Liu *et al.* introduce an

unsupervised intrinsic image decomposition model that learns the latent features of reflectance and shading from unsupervised and uncorrelated data [51]; (3) Bi *et al.* use edge-preserving smoothing, Dirichlet process Gaussian mixture model for adaptive clustering, and a superpixel-based CRF optimization [45]; and (4) STAR is an advanced structure and texture aware Retinex model [48]. Further, we provide two supervised deep learning models: (5) CGIntrinsics combines four different scene level real and synthetic datasets for training [52]; and (6) ShapeNet boosts the intrinsic correlations with special decoder links and is trained on 2.5M synthetic images [35]. Thus, the evaluation is profoundly diverse. The results are provided in Fig. 6 for ISTD and Fig. 7 for SRD.

The evaluations further prove the quality of the proposed invariant descriptors for natural outdoor shadow handling. Figure 6 shows that Krebs *et al.* tend to produce smooth reflectance maps, but cannot handle shadow casts. Liu *et al.* fail to generate proper colors. The images appear way too dull, and in the first two columns, the scenes have lost their intrinsic color. It also tends to generate bright yellowish artifacts. Bi *et al.* using CRF optimization and adaptive clustering appear to be the closest work to ours. Yet, their method also fails to handle shadow casts. Their adaptive clustering even generates additional shadow-like artifacts on the first image. Thus, with the contribution of color ratios, it is clear that our model does not just choose the cluster number adaptively, but the descriptors further constrain the reflectance predictions. STAR generates too bright images, so that most of the structures and colors are not visible anymore. Similarly, supervised deep learning models CGIntrinsics and ShapeNet fail to generate shadow-free

Fig. 6. State-of-the-art comparisons on shadow cast handling on ISTD [43]. Both unsupervised and supervised models fail to handle shadows. Our method driven by the photometric invariant color ratios is more robust, containing significantly fewer shadow cues and fewer shadow artifacts; the colors are more vivid and realistic, and the structures are well preserved. Images are best viewed in color and on the electronic version.

albedo images. ShapeNet produces further color artifacts and distortions.

Similar behavior is also observed in Fig. 7 for SRD. Krebs *et al.* tend to produce smooth reflectance maps, but fail to handle shadow casts. Liu *et al.* again fail to generate proper colors. For example, the yellow color of the bench in the last row is completely washed away. Moreover, their model tends to generate bright yellowish artifacts contaminating the reflectance maps. Finally, Bi *et al.* results are similar to Krebs *et al.*, producing smooth reflectance maps, but they cannot properly handle shadow casts.

On the other hand, our method driven by photometric invariant color ratios is more robust in handling natural outdoor real world shadow casts for two different domains of ISTD and SRD. Compared with others, our reflectance predictions contain significantly fewer shadow cues and fewer shadow artifacts. Furthermore, the colors are more vivid and realistic, and the structures are well preserved and not polluted with undesired color artifacts. The results provided on two distinct dataset domains further demonstrate the generalization capability of the photometric invariant color ratios. In addition, the results

Fig. 7. State-of-the-art comparisons on shadow cast handling on SRD [44]. Compared with others, our albedo estimations contain very little or almost no shadow artifacts, the colors are more vivid and realistic, and the structures are well preserved. Images are best viewed in color and on the electronic version.

reveal a surprising fact that none of the state-of-the-art models is able to handle natural outdoor shadow casts.

5. CONCLUSION AND FUTURE WORK

We demonstrated the concept of physics-based invariant descriptors used in a CRF optimization framework to improve reflectance predictions by making them more accurate and more robust to shading variations. We provided three particular ways to combine the ratios into the optimization, and a final combination is presented. Experiments were provided on four datasets: a real world object level, a real world indoor scene level, and two different outdoor scene level image sets with various shadow casts. Quantitative improvements were achieved on both object level and scene level indoor scenes. On the other hand, the significance of the proposed method is better observed by the qualitative results. The visual results showed that the default model has no sense of shadows, as it cannot differentiate whether the changes in pixel intensities are caused by a color change or photometric effects. Incorporating the color ratios further enhances the model, making it aware of the photometric cues including strong shadings and shadow casts. Thus, the model can now differentiate drastic changes in pixel values and attribute them to the related intrinsics.

Future work considering the (1) Color Retinex gradient constraint, (2) global reflectance sparsity prior, (3) pair-wise constant reflectance prior, or (4) natural shadow cast handling should also consider incorporating the photometric invariant descriptors to make their frameworks more accurate and robust. Unlike handcrafted priors, which are based on observations and mostly computationally intensive, the color ratios are physics based, calculated directly from the RGB images in an unsupervised manner, and realized in real time. The process does not involve any specialized filters or labeled data, and it does not add any additional complexity to the system. Thus, further improvements are attainable without any overhead. Our future work will

focus on steering deep learning models by photometric invariant descriptors.

Funding. Horizon 2020 Framework Programme (688007).

Acknowledgment. This project was funded by the EU Horizon 2020 (TrimBot2020). We thank Partha Das for his contribution to the experiments.

Disclosures. The authors declare no conflicts of interest.

Data Availability. Data underlying the results presented in this paper are not publicly available at this time but may be obtained from the authors upon reasonable request.

REFERENCES

1. H. G. Barrow and J. M. Tenenbaum, "Recovering intrinsic scene characteristics from images," in *Computer Vision Systems* (1978), pp. 3–26.
2. A. S. Baslamisli, T. T. Groenestege, P. Das, H. A. Le, S. Karaoglu, and T. Gevers, "Joint learning of intrinsic images and semantic segmentation," in *European Conference on Computer Vision* (2018).
3. C. Xu, Y. Han, G. Baci, and M. Li, "Fabric image recolorization based on intrinsic image decomposition," *Text. Res. J.* **89**, 3617–3631 (2019).
4. S. Beigpour and J. van de Weijer, "Object recoloring based on intrinsic image estimation," in *IEEE International Conference on Computer Vision* (2011).
5. E. H. Land and J. J. McCann, "Lightness and retinex theory," *J. Opt. Soc. Am.* **61**, 1–11 (1971).
6. S. Bell, K. Bala, and N. Snavely, "Intrinsic images in the wild," *ACM Trans. Graph.* **33**, 159 (2014).
7. Z. Cheng, Y. Zheng, S. You, and I. Sato, "Non-local intrinsic decomposition with near-infrared priors," in *IEEE International Conference on Computer Vision* (2019).
8. T. Narihira, M. Maire, and S. X. Yu, "Direct intrinsics: learning albedo-shading decomposition by convolutional regression," in *IEEE International Conference on Computer Vision* (2015).
9. A. S. Baslamisli, H. A. Le, and T. Gevers, "CNN based learning using reflection and retinex models for intrinsic image decomposition," in *IEEE Conference on Computer Vision and Pattern Recognition* (2018).

10. P. V. Gehler, C. Rother, M. Kiefel, L. Zhang, and B. Schölkopf, "Recovering intrinsic images with a global sparsity prior on reflectance," in *Advances in Neural Information Processing Systems* (2011).
11. L. Shen and C. Yeo, "Intrinsic images decomposition using a local and global sparse representation of reflectance," in *IEEE Conference on Computer Vision and Pattern Recognition* (2011).
12. J. T. Barron and J. Malik, "Shape, illumination, and reflectance from shading," *IEEE Trans. Pattern Anal. Mach. Intell.* **37**, 1670–1687 (2015).
13. G. D. Finlayson, "Colour object recognition," master's thesis (Simon Fraser University, 1992).
14. L. Shen, P. Tan, and S. Lin, "Intrinsic image decomposition with non-local texture cues," in *IEEE Conference on Computer Vision and Pattern Recognition* (2008).
15. Q. Zhao, P. Tan, Q. Dai, L. Shen, E. Wu, and S. Lin, "A closed-form solution to retinex with non-local texture constraints," *IEEE Trans. Pattern Anal. Mach. Intell.* **34**, 1437–1444 (2012).
16. J. Shen, X. Yang, X. Li, and Y. Jia, "Intrinsic image decomposition using optimization and user scribbles," *IEEE Trans. Cyber.* **43**, 425–436 (2013).
17. E. Garces, A. Munoz, J. Lopez-Moreno, and D. Gutierrez, "Intrinsic images by clustering," in *Computer Graphics Forum* (2012).
18. X. Jiang, A. J. Schofield, and J. L. Wyatt, "Correlation-based intrinsic image extraction from a single image," in *European Conference on Computer Vision* (2010).
19. M. F. Tappen, W. T. Freeman, and E. H. Adelson, "Recovering intrinsic images from a single image," *IEEE Trans. Pattern Anal. Mach. Intell.* **27**, 1459–1472 (2005).
20. S. Ding, B. Sheng, Z. Xie, and L. Ma, "Intrinsic image estimation using near-10 sparse optimization," *Visual Comput.* **33**, 355–369 (2017).
21. Y. Li and M. S. Brown, "Single image layer separation using relative smoothness," in *IEEE Conference on Computer Vision and Pattern Recognition* (2014).
22. B. Sheng, P. Li, Y. Jin, P. Tan, and T. Y. Lee, "Intrinsic image decomposition with step and drift shading separation," *IEEE Trans. Vis. Comput. Graph.* **26**, 1332–1346 (2020).
23. A. Bousseau, S. Paris, and F. Durand, "User-assisted intrinsic images," in *ACM SIGGRAPH Asia 2009* (2009), paper 130.
24. Q. Chen and V. Koltun, "A simple model for intrinsic image decomposition with depth cues," in *IEEE International Conference on Computer Vision* (2013).
25. J. Jeon, S. Cho, X. Tong, and S. Lee, "Intrinsic image decomposition using structure-texture separation and surface normals," in *European Conference on Computer Vision* (2016).
26. K. J. Lee, Q. Zhao, X. Tong, M. Gong, S. Izadi, S. U. Lee, P. Tan, and S. Lin, "Estimation of intrinsic image sequences from image+depth video," in *European Conference on Computer Vision* (2012).
27. Y. Weiss, "Deriving intrinsic images from image sequences," in *IEEE International Conference on Computer Vision* (2001).
28. W. Gong, W. Xu, L. Wu, X. Xie, and Z. Cheng, "Intrinsic image sequence decomposition using low-rank sparse model," *IEEE Access* **7**, 4024–4030 (2019).
29. P. Y. Laffont and J. C. Bazin, "Intrinsic decomposition of image sequences from local temporal variations," in *IEEE International Conference on Computer Vision* (2015).
30. J. Matas, R. Marik, and J. Kittler, "On representation and matching of multi-coloured objects," in *IEEE International Conference on Computer Vision* (1995).
31. S. K. Nayar and R. M. Bolle, "Reflectance based object recognition," *Int. J. Comput. Vis.* **17**, 219–240 (1996).
32. K. Barnard and G. D. Finlayson, "Shadow identification using colour ratios," in *Color and Imaging Conference* (2000).
33. T. Gevers and A. Smeulders, "Color constant ratio gradients for image segmentation and similarity of texture objects," in *IEEE Conference on Computer Vision and Pattern Recognition* (2001).
34. T. Gevers and A. Smeulders, "Object recognition based on photometric color invariants," in *Scandinavian Conference on Image Analysis* (1997).
35. J. Shi, Y. Dong, H. Su, and S. X. Yu, "Learning non-Lambertian object intrinsics across shapenet categories," in *IEEE Conference on Computer Vision and Pattern Recognition* (2017).
36. Z. Li and N. Snavely, "Learning intrinsic image decomposition from watching the world," in *IEEE Conference on Computer Vision and Pattern Recognition* (2018).
37. L. Lettry, K. Vanhoey, and L. van Gool, "DARN: a deep adversarial residual network for intrinsic image decomposition," in *IEEE Winter Conference on Applications of Computer Vision* (2018).
38. A. S. Baslamisli, P. Das, H. A. Le, S. Karaoglu, and T. Gevers, "ShadingNet: image intrinsics by fine-grained shading decomposition," arXiv:1912.04023 (2019).
39. H. A. Sial, R. Baldrich, and M. Vanrell, "Deep intrinsic decomposition trained on surreal scenes yet with realistic light effects," *J. Opt. Soc. Am. A* **37**, 1–15 (2020).
40. S. Shafer, "Using color to separate reflection components," *Color Res. Appl.* **10**, 210–218 (1985).
41. G. D. Finlayson, M. S. Drew, and B. V. Funt, "Color constancy: generalized diagonal transforms suffice," *J. Opt. Soc. Am. A* **11**, 3011–3019 (1994).
42. R. Grosse, M. K. Johnson, E. H. Adelson, and W. T. Freeman, "Ground truth dataset and baseline evaluations for intrinsic image algorithms," in *IEEE International Conference on Computer Vision* (2009).
43. J. Wang, X. Li, L. Hui, and J. Yang, "Stacked conditional generative adversarial networks for jointly learning shadow detection and shadow removal," in *IEEE Conference on Computer Vision and Pattern Recognition* (2018).
44. L. Qu, J. Tian, S. He, Y. Tang, and R. W. H. Lau, "Deshadownet: a multi-context embedding deep network for shadow removal," in *IEEE Conference on Computer Vision and Pattern Recognition* (2017).
45. S. Bi, X. Han, and Y. Yu, "An L1 image transform for edge-preserving smoothing and scene-level intrinsic decomposition," *ACM Trans. Graph.* **34**, 1 (2015).
46. M. Yan, "Methods of determining the number of clusters in a data set and a new clustering criterion," Ph.D. thesis (Virginia Tech, 2005).
47. X. Li, W. Liang, X. Zhang, S. Qing, and P. C. Chang, "A cluster validity evaluation method for dynamically determining the near-optimal number of clusters," *Soft Comput.* **24**, 9227–9241 (2020).
48. J. Xu, Y. Hou, D. Ren, L. Liu, F. Zhu, M. Yu, H. Wang, and L. Shao, "Star: a structure and texture aware retinex model," *IEEE Trans. Image Process.* **29**, 5022–5037 (2020).
49. T. Nestmeyer and P. V. Gehler, "Reflectance adaptive filtering improves intrinsic image estimation," in *IEEE Conference on Computer Vision and Pattern Recognition* (2017).
50. A. Krebs, Y. Benezeth, and F. Marzani, "Intrinsic RGB and multispectral images recovery by independent quadratic programming," *PeerJ Comput. Sci.* **6**, e256 (2020).
51. Y. Liu, Y. Li, S. You, and F. Lu, "Unsupervised learning for intrinsic image decomposition from a single image," in *IEEE Conference on Computer Vision and Pattern Recognition* (2020).
52. Z. Li and N. Snavely, "CGIntrinsics: better intrinsic image decomposition through physically-based rendering," in *European Conference on Computer Vision* (2018).