Reflectance Adaptive Filtering Improves Intrinsic Image Estimation

Thomas Nestmeyer^{1,3} Peter V. Gehler^{1,2,3} ¹University of Würzburg, Germany ²Bernstein Center for Computational Neuroscience, Tübingen, Germany ³Max Planck Institute for Intelligent Systems, Tübingen, Germany

{tnestmeyer, pgehler}@tuebingen.mpg.de

Abstract

Separating an image into reflectance and shading layers poses a challenge for learning approaches because no large corpus of precise and realistic ground truth decompositions exists. The Intrinsic Images in the Wild (IIW) dataset provides a sparse set of relative human reflectance judgments, which serves as a standard benchmark for intrinsic images. A number of methods use IIW to learn statistical dependencies between the images and their reflectance layer. Although learning plays an important role for high performance, we show that a standard signal processing technique achieves performance on par with current state-of-the-art. We propose a loss function for CNN learning of dense reflectance predictions. Our results show a simple pixel-wise decision, without any context or prior knowledge, is sufficient to provide a strong baseline on IIW. This sets a competitive baseline which only two other approaches surpass. We then develop a joint bilateral filtering method that implements strong prior knowledge about reflectance constancy. This filtering operation can be applied to any intrinsic image algorithm and we improve several previous results achieving a new state-of-the-art on IIW. Our findings suggest that the effect of learning-based approaches may have been over-estimated so far. Explicit prior knowledge is still at least as important to obtain high performance in intrinsic image decompositions.

1. Introduction

Almost 40 years ago, the seminal paper of Barrow and Tenenbaum [3] conjectured that "A robust visual system should be organized around a noncognitive, nonpurposive level of processing that attempts to recover an intrinsic description of the scene". Their work motivates the task of decomposing an image into constituent layers such as surface reflectance, surface orientation, distance and incident illumination. Ever since, significant progress has been made on this problem, but the recovery of these physical properties of visual scenes or videos remains an open challenge. A successful model needs to resolve the ill-posedness of the problem and cope with the variety of image appearances.

A possible line of attack are supervised learning methods which have been used with great success for a wide range of computer vision applications. Standing out for superior performance combined with favorable runtimes is the class of Convolutional Neural Network (CNNs), a dominant contender for many vision problems. CNNs are mostly falling into the category of purposive models, guided by task specific goals such as image classification or recognition. The obvious question is whether CNNs will fare equally well on the problem of intrinsic image decompositions.

Several works have included CNN methods in systems that recover reflectance and shading layers [25, 26, 37, 38]. However, prior work uses CNNs mostly in combination with additional methods, such as Conditional Random Fields (CRFs), to achieve a dense image decomposition. An advantage of CRF models is their ability to encode prior information about the problem. In the pre-CNN time, intrinsic image methods were dominated by CRF models with carefully designed priors on reflectance, shading, and their combination. In this paper we attempt to answer the question whether prior terms are necessary when human annotation in the form of weak labels is available.

Acquiring accurate training data for intrinsic images is a challenge. The MIT intrinsic dataset [16] with 20 images and 10 (single color) light configurations was a first attempt to empirically validate intrinsic estimation techniques. It has served this purpose well, but lacks realism and diversity. Recently, [4] proposed an extension to multi illuminants, but without overcoming the limitations on extent. Another possible route to generate datasets is the use of computer graphics rendering engines. This has been explored by the authors of [5] who created a dataset of synthetic scenes rendered using the Blender open source rendering engine [19]. This led to a dataset of 32 single objects and 36 scene compositions which is still limited in terms of detail and diversity. The MPI-Sintel dataset has been created using the open

source movie Sintel [10] to serve as a benchmark for several problems such as optical flow estimation. While MPI-Sintel is more varied and complex, the type of scenes and visual appearance is still very different from real world data.

A significant attempt to overcome the lack of empirical data is the dataset of "Intrinsic Images in the Wild" (IIW) [6]. This dataset contains 5230 photos of mostly indoor scenes which have been annotated with a sparse set of relative reflectance judgments. From a small set of image locations, human judgments on pairs of neighboring locations have been collected, which provide whether one point is of darker or similar material reflectance. Although humans can be fooled with artificial setups [1], the perception of relative material reflectance is sufficient to provide mostly consistent label information for this large corpus of images (see [6] for an analysis). Along with the dataset, Bell et al. also formulate a performance metric (WHDR) that we will discuss in detail in Section 3.1. The IIW dataset allows to empirically validate intrinsic image estimation and also, the judgments have also been used to train models for intrinsic image decompositions [26, 37, 38].

In this paper we develop two intrinsic image models: a CNN approach with appropriate loss function and a filtering technique to include strong prior knowledge about reflectance properties. We first design a CNN method that, in contrast to previous work, does not include prior information on shading smoothness [23], reflectance [27, 15, 31, 7], or combinations [2]. We design a loss function that enables end-to-end learning from the pairwise judgments. This leads to an interesting result: a simple multi-layer perceptron with no image context, just based on the pixels alone provides competitive performance, better or on par with current learning and non-learning models. We then develop a method from the other extreme, a dense filtering operation based on joint bilateral and guided filtering. This technique simplifies the processing pipeline of [7] and makes it possible to apply to any reflectance prediction. Our experiments show drastically improved state-of-the-art performance on IIW. Besides presenting the empirically best performing algorithm, our results reveal interesting observations about the current state of intrinsic image estimation. In summary, we believe that for intrinsic image estimation, it is the inclusion of prior knowledge through regularization, CRFs, or filtering that still drives the performance. To rely solely on learning approaches, the amount of available annotation may still be insufficient.

2. Related Work

Until recently there was a lack of empirical data to validate intrinsic images algorithms. Therefore, most of the literature revolved around the design of suitable priors. The recent work of [2] is a prominent example of a method that carefully trades the use of prior information with interesting representations that enable a detailed decomposition into several layers. Priors in [2] include terms on smoothness, parsimony and absolute values of reflectance, smoothness, surface isotropy and occluding contour priors on shape and a multivariate Gaussian fit to a spherical-harmonic illumination model. This lead to impressive results on the MIT intrinsic dataset [16], but the method is limited to single masked objects in a scene, and problems with complex illumination remain.

The work of [7] approaches the problem from a filtering perspective. After a filtering step followed by clustering, the pixels are grouped into regions of same reflectance, such that a simple shading term suffices to recover the full intrinsic decomposition. This method produces the best results on the IIW dataset but takes several minutes of processing time. In Section 5 we build on this work and propose a filtering technique that can be applied to any other intrinsic image estimation as well. This implements the idea of grouping pixels into sets of constant reflectance. Other works consider additional knowledge in order to recover reflectance and shading, as, *e.g.*, multiple images of the same scene with different lighting [34, 22], an interactive setting with user annotations [9, 8], or an additional depth layer as input [11].

The paper of [6] introduced the Intrinsic Images in the Wild dataset with human annotations giving relative reflectance judgments that served as the training and test set for different learning based methods. Using this data, the work of [6] was the first to compare different algorithms on a large corpus of real world scenes.

A first attempt to learn using the data from IIW was made by [26]. The authors used the relative judgment information in a multi-class setup and fine-tune an AlexNet CNN trained on ImageNet. Only the sparse annotation points that are required to compute the WHDR loss are predicted with this network and there is no step that turns them into a dense decomposition. The works of [37] and [38] are similar, both use a CNN to obtain pairwise judgment predictions, then followed by a step to turn the sparse information into a dense decomposition. Both methods achieve good results on IIW and take several seconds to process an image.

Similar to our work, in the sense that a dense intrinsic decomposition is predicted, is the work of [25]. A CNN is used to directly predict reflectance and shading with the objective function being the difference to ground truth decompositions. Since those are only available for the rendered dataset of MPI-Sintel, the authors report that the learned model does not generalize well to the real world images of IIW. An additional data term in the gradient domain is used by [24]. They also propose to use an adversary in order to remove typical generative CNN artifacts by discriminating between generated and ground truth decompositions.

Table 1. Overview of different intrinsic image estimation methods. For every method we note whether or not it uses a *CNN trained on IIW* and whether the *CNN decomposes densely* into intrinsic layers without an additional globalization step.

C	ĊŃŊ	CNN
Method	trained	decomposes
	on IIW	densely
Retinex [23]	×	×
Bell et al. 2014 [6]	×	×
Bi et al. 2015 [7]	×	×
Narihira et al. 2015a [26]	~	×
Zhou et al. 2015 [37]	~	×
Zoran et al. 2015 [38]	~	×
Narihira et al. 2015b [25]	×	~
Lettry et al. 2016 [24]	×	~
This paper	~	~

Therefore, this approach has the same limitation requiring dense ground truth decompositions and no results on IIW are available. To our knowledge, there is no CNN based method that predicts a dense intrinsic decomposition and works well for images from IIW.

In Table 1 we organized the related work along the dimensions that are relevant for the proposed method.

The work of [12] also trains a CNN from relative judgments with a ranking loss to predict pixel-wise labels, but for the application of recovering dense depth estimates. This involved the creation of a dataset with relative depth judgments in the spirit of IIW. However, in contrast to intrinsic images, it is possible to capture accurate ground truth depth for training and testing, making reflectance and shading estimation a more relevant target of learning from sparse pairwise comparisons.

3. Preliminaries

We work with linear RGB and the Lambertian reflectance assumption, which allows to separate every pixel in image $I \in [0,1]^{3 \times h \times w}$ into a product of reflectance Rand shading S, that is the pixel and channel-wise product I = RS. Further, we assume achromatic light which reduces the decomposition problem to a per-pixel scalar estimation problem. Namely, given a scalar $r_p \in [0,1]$ for each pixel p, we recover reflectance and shading as

$$R_p = \frac{r_p}{\frac{1}{3}\sum_c I_p^c} \cdot I_p, \qquad S_p = \frac{\frac{1}{3}\sum_c I_p^c}{r_p} \cdot \begin{pmatrix} 1\\1\\1 \end{pmatrix}, \quad (1)$$

where $c \in \{R, G, B\}$ denotes the color channel. Under these assumptions, the problem boils down to estimation of a single scalar per pixel $\mathbf{r} \in \mathbb{R}^{h \times w}$.

The same assumptions are commonplace in the literature and have been used, e.g., in [15]. We note that achromatic

light is often violated in the IIW dataset, especially in the presence of multiple light sources. As the proposed loss function WHDR only compares relative lightness and no color information, it is invariant to this choice.

3.1. A quantitative measure for intrinsic images

Accurate ground truth information in the form of image decompositions in reflectance and shading layers does not exist at scale. To empirically validate the quality of intrinsic image algorithms using the pairs of relative reflectance judgments alone, Bell et al. [6] introduced the WHDR metric (weighted human disagreement rate). We refer to their work for all details on the data annotation process, but will review the ingredients that we need for our development.

For every image, annotation is given in the form of pairs of image locations (i_1, i_2) for which a human reflectance judgment $J_i \in \{1, 2, E\}$ is provided. The judgment indicates whether point i_1 is darker than i_2 ($J_i = 1$), lighter ($J_i = 2$), or of equal reflectance ($J_i = E$). The confidence w_i of a judgment is defined via the CUBAM score of the two-decision model "points have the same reflectance" (and if not "does the darker point have darker reflectance" (see [6] for further details). The annotation set $\{(i_1, i_2, J_i, w_i)\}_{i=1,...,N_I}$ varies in size N_I for every image I in the range from 1 to 1181 with a median of 113.

Given a reflectance prediction R, first a relative classification for the set of annotated points is computed as

$$\hat{J}_{\delta}(R,i) = \begin{cases} 1 & \text{if } R_{i_2}/R_{i_1} > 1 + \delta \\ 2 & \text{if } R_{i_1}/R_{i_2} > 1 + \delta \\ E & \text{else}, \end{cases}$$
(2)

where $\delta \ge 0$ controls when two points are considered different. For large values of δ , two points would need to be farther apart to be judged as darker (resp. lighter).

Given these relative estimates, the WHDR loss is computed as the weighted average of how often the annotation and prediction disagree

WHDR_{$$\delta$$} $(J, R) = \frac{\sum_{i} w_{i} \cdot \mathbb{1}\left(J_{i} \neq \hat{J}_{\delta}(R, i)\right)}{\sum_{i} w_{i}}.$ (3)

Note that this loss does not evaluate the reflectance at all points in the image, but only at those for which labels are available. Therefore, it could also be evaluated on these points alone for an algorithm that does not provide a dense decomposition of the image.

The works of [26, 37, 38] use these relative annotations to train multi-class classifiers, predicting for every pair of patches its relative reflectance judgment $\{1, 2, E\}$ directly. Since this approach does not provide the actual values Rof the reflectance layer, further post-processing steps are required to produce a dense prediction. These post-processing



Figure 1. Visualization of the WHDR-Hinge loss dependent on the ratio R_{i_1}/R_{i_2} for $\delta = 0.1$ and $\xi = 0.05$. The value of δ controls where the decision boundary for darker/lighter or equal reflectance lightness is made. With the value ξ , a margin from this boundary is encouraged. For values $\xi > \delta$ the *E* class will always have a non-zero loss.

steps are separate from the classifiers and motivated by common intrinsic prior terms. We will circumvent any post-processing by directly predicting a dense reflectance map *R*.

4. Direct Reflectance Prediction with a CNN

We propose an objective function that makes direct use of the relative reflectance judgments by humans that the IIW dataset provides. This weak label information has been used in [37, 38] for CNN training already, treating it however as a multi-class classification problem. While a multi-class loss achieves good performance on pairs of points, this strategy requires an additional globalization step to propagate information to all pixels. Our aim is to directly decompose the entire image with a single forward pass of a CNN, avoiding any need for post-processing.

We will first discuss the loss function that we use and then describe the network architecture and training method.

4.1. WHDR hinge loss

We construct a proxy loss for the WHDR that can be used for supervised training. The formulation is an adaption of the ε -insensitive loss for regression [33] for this problem setup. We define

$$\begin{split} \ell_{\delta,\xi}\left(J,R,i\right) &= \\ \begin{cases} \max\left(0, \frac{R_{i_1}}{R_{i_2}} - \frac{1}{1+\delta+\xi}\right) & \text{if } J_i = 1 \\ \max\left(0, \left\{\frac{1}{1+\delta-\xi} - \frac{R_{i_1}}{R_{i_2}}, \\ \frac{R_{i_1}}{R_{i_2}} - (1+\delta-\xi) \right) & \text{if } J_i = E \\ \max\left(0, 1+\delta+\xi - \frac{R_{i_1}}{R_{i_2}}\right) & \text{if } J_i = 2, \end{split}$$

which is visualized in Fig. 1. The scalar δ is the threshold of the WDHR_{δ} and we introduce the hyper-parameter ξ , which is the margin between the neighbouring classes 1, *E* and 2, *E*.

The pipeline of supervised training is simple. A network produces a dense decomposition R, which is then used to arrive at relative judgments for two pixel locations based on the ratio of the predicted R values. The loss in Eq. (4) is then weighted and summed over all annotated pixel pairs, similar to Eq. (3), and the error is propagated backwards to compute the gradients of the network parameters.

As with the standard hinge-loss commonly used for binary SVM training, the sub-gradients of the WHDR-hinge loss can be easily computed.

4.2. Train and test data set

The IIW dataset does not come with a pre-defined train, validation and test split. We adopt the split suggested by [26] into 80% training and 20% test images, putting the first of every five images sorted by file name in the test set. In order to properly evaluate different models, we additionally split the data into a separate validation set, with the ratios of 70% training, 10% validation and 20% test. We keep the test set of [26], and use from every series of 10 images the seventh in the validation set.

4.3. Network architecture of the CNN

We take the linearized RGB images in the range [0, 1] as input, evaluate a series of n convolutional layers with f filters each, acting on a kernel of size k, with a ReLU as nonlinear activation function in between. The padding in the convolutions is chosen based on k, so as to not change the resolution. The output of all nonlinearities is concatenated and convolved with a 1×1 filter to fuse the information of skipped layers. A last sigmoidal activation function bounds the single channel output r, on which the WHDR-Hinge loss, as given in Section 4.1, operates during training.

One final layer recovers RGB reflectance R and shading S from the scalar reflectance intensity r, as given in Eq. (1), to output the final dense intrinsic image decomposition.

Resolving light intensity The last nonlinearity in the network acting on r is included since ambiguity about the light intensity in an image cannot be solved. It is only possible to determine reflectance and shading up to a constant $\alpha \in (0, \infty)$, since $I = RS = (\alpha R) \left(\frac{1}{\alpha}S\right)$. Therefore, to keep the reflectance values bounded, we employ a sigmoidal activation function to limit the scalar reflectance intensity to be in the range [0, 1].

4.4. Experiments

For all experiments in this paper we use the open source deep learning framework caffe [20] utilizing the ADAM solver [21] with a learning rate of 0.001, momentum of $\beta_1 = 0.9$ and momentum-2 of $\beta_2 = 0.999$. All training images are resized to a fixed 256×256 pixel resolution to be able to process them in batches. More details about data augmentation and label analysis are included in the supplementary.



Figure 2. (a) WHDR for different network depths n and number of filters f. Missing data is the result of memory limit on our graphics card. (b) WHDR for different thresholds δ and margins ξ .

Network Hyper-Parameters For the network layout described above we performed an extensive parameter sweep over a varying number of kernel widths $k \in \{1, 3, 5, 7\}$, layers $n \in \{1, ..., 9\}$, and filters $f \in \{2^1, ..., 2^9\}$. The results on the validation set for k = 1 are shown in Fig. 2(a). The number of layers n has only small influence on the performance above n > 2, similar with $f > 2^4$. An unexpected finding was that the kernel size k has little to no effect and 1×1 convolutions work just as well as those with bigger kernels. This means that the network only learns a pixel-wise lookup table, but at the same time, the network performs already better in WHDR than most state-of-the-art methods in the literature. This amounts to a re-scaling of the reflectance intensity at every pixel separately, no context needed. We will discuss this further in Section 4.5. From this analysis we chose a network of $n = 5, f = 2^5$, and k = 1 as the basis for all future experiments.

In addition to this basic setup we also played with different network layouts, *e.g.*, without skip connections and with a U-net like architecture [30], tried PReLU [18] nonlinearities in between and dilated convolutions [35] to widen the receptive field, but did not find better results. In general we found that simpler networks perform better, what we believe is the outcome of the amount of weakly labeled training data.

Hyperparameters of the WHDR-Hinge loss To minimize the WHDR rate consistent with $\delta = 0.1$ from [6] we optimized the loss hyper-parameter δ, ξ on the validation set. The influence is shown in Fig. 2(b) and the final parameters used for training are $\delta = 0.12$ and $\xi = 0.08$.

4.5. Discussion of the results

Many methods build on the Retinex assumption [23], which states that strong image gradients are reflectance edges and small gradients are explained by shading. Under the assumption of smooth shading, local gradient estimation would only require a small receptive field, but there is no possibility that a method can resolve shading from a single pixel alone, e.g., see the famous illusion of [1] for a

Table 2. Comparison of some intrinsic image approaches on IIW. An extended comparison is in Fig. 5.

Method	Retinex	[6]	[37]	ours	[38]	[7]
WHDR	26.9	20.6	19.9	19.5	17.9	17.7



Figure 3. Lookup table in HSV space, generated by our direct prediction network, for varying hue and saturation and a constant value/brightness of 255. Left: The input image I, Right: The single channel reflectance intensity r predicted by the CNN.

counter example. Still, in terms of WHDR, this method performs better than most methods [31, 14, 36, 6, 37] on IIW, see the table Table 2 for an empirical comparison to a few approaches. Since there is a direct pixel-to-reflectance relationship, we can visualize a "lookup-table" mapping RGB pixels to reflectance, see Fig. 3. A big portion of colors is judged to have more or less the same reflectance intensity as white. Blue is mostly judged being darker than green, which is biologically plausible. Green light contributes the most to the intensity perceived by humans, and blue light the least [29]. There is a portion of very light reflectance for fully saturated green, even brighter than from white pixels. This may be a result of the Helmholtz-Kohlrausch effect [13], humans perceive colored light brighter than white light. This may lead to wrong human reflectance judgments under the circumstances of bright saturated colors.

A simple re-scaling of an image into [a, 1] with a = 0.55 (see Fig. 4 for an example), then using the image as a reflectance estimate results to a WHDR score of 25.7 on the test set. This low score is due to an in-balance of relative judgments, 2/3 of which are equal judgments. A rescaling to [0.55, 1] makes most equal judgments correct and compromises the unequal judgments. The output of the CNN is on average in the range [0.48, 0.96] but its nonlinearity accounts for small variations in color and therefore makes more un-equal judgments correct. Remember that the CNN predicts a dense reflectance map on the test set, unaware of the point pairs performance will be evaluated on. The CNN implements no explicit prior knowledge, we will show next that encouraging piecewise constant reflectance improves the result.

5. Reflectance Filtering

A common prior of intrinsic image estimation is to have only a sparse set of reflectances present in a scene [27, 15,



Figure 4. (a) Mean WHDR (in %) on training and validation set. As reflectance image we take the input image after rescaling it into the range [a, 1]. (b) An example input image. (c) The same image scaled to have a lower bound of a = 0.55, where the mean WHDR is minimal.

31]. We describe a new technique to include this prior knowledge using an image filter. This allows an easy integration into existing techniques and we found that filtering reflectance estimates always improves performance.

5.1. Image-aware filtering

A general linear translation-variant filtering process is defined as

$$q_i = \sum_j W_{ij}(I)p_j,\tag{5}$$

where the input image p is smoothed under the guidance of an image I to the filtered output image q. Here i, j denote pixels and the sum runs over the entire image. Two examples are the joint bilateral filter and the guided filter, whose weights for Eq. (5) we summarize next.

The (joint) bilateral filter The joint bilateral filter [28] is an extension to the bilateral filter [32] which uses feature difference in a (potentially different) guidance image to spatially smooth pixels in the input image. It defines the weights as

$$W_{ij}(I) = \frac{1}{K_i} \exp\left(-\frac{|x_i - x_j|^2}{\sigma_s^2} - \frac{|I_i - I_j|^2}{\sigma_r^2}\right), \quad (6)$$

with x_i being pixel coordinates. This means that pixels that are both close spatially and in intensity in the guidance image will be smoothed more. The normalization K_i is chosen to ensure $\sum_i W_{ij} = 1$.

The guided filter The guided filter [17] is a fast alternative to the joint bilateral filter, it is also edge-preserving, and has better behavior near edges. It is based on a locally linear model $\forall i \in \omega_k : q_i = a_k I_i + b_k$, where a_k, b_k are linear coefficients assumed to be constant in the square window ω_k centered at pixel k of size r. The linearity guarantees that q has an edge only if I has an edge, since $\nabla q = a \nabla I$. Solving

for the coefficients that minimize the difference between q and p leads to the weights

$$W_{ij}(I) = \frac{1}{|\omega|^2} \sum_{k:(i,j)\in\omega_k} \left(1 + \frac{(I_i - \mu_k)(I_j - \mu_k)}{\sigma_k^2 + \varepsilon}\right)$$
(7)

where μ_k and σ_k^2 are the mean and variance of I in ω_k , $|\omega_k|$ is the number of pixels in ω_k and ε a constant parameter similar to the range variance σ_r^2 in the bilateral filter. Especially for larger spatial scales, the guided filter benefits from not having the quadratic dependency on the filtering kernel size. We refer to [17] for a more thorough discussion.

5.2. Filtering for piecewise constancy

We need to define a guidance image to fully specify the filtering operation. An ideal guidance image would group pixels into regions of constant reflectance. We will refer the filtered image with BF(*method*, guidance) for the bilateral filter and GF(*method*, guidance) for the guided filter, respectively.

Using a flattened image as guidance The method of [7] formulates an optimization problem to group pixels into regions of similar reflectance. This provides a good candidate for a suitable guidance image. The piecewise flattened image is found by minimizing $E = E_l + \alpha E_g + \beta E_a$, with the *local flattening energy*

$$E_{l} = \sum_{i} \sum_{j \in N_{h}(i)} \exp\left(-\frac{\|f_{i} - f_{j}\|_{2}^{2}}{2\sigma^{2}}\right) \|q_{i} - q_{j}\|_{1}, \quad (8)$$

where $N_h(i)$ is the $h \times h$ neighborhood of the *i*-th pixel, q_i is the output RGB vector, $f_i = [\kappa \cdot l_i, a_i, b_i]$, with $[l_i, a_i, b_i]$ being the input vector in CIELab color space and κ, σ are hyper-parameters. A global sparsity energy is defined as

$$E_g = \sum_{i \in S_r} \sum_{j \in S_r} w_{ij} \| q_i - q_j \|_1,$$
(9)

with the same affinity weights w_{ij} as in Eq. (8) and S_r being the set of representative pixels which are closest to the average color in their superpixels. To avoid the trivial solution, a *data term for image approximation* is added:

$$E_a = \|q - p\|_2^2. \tag{10}$$

See [7] for how to solve the resulting optimization problem. We will refer to the result of this L_1 -flattening optimization from now on also simply as *flat*.

Filtering Results Using "flat" as guidance, we again found the best hyper-parameters on the validation set ($\sigma_r = 15, \sigma_s = 28$ for the joint bilateral and $r = 45, \varepsilon = 3$ for the guided filter). Using the CNN predictions as input to

Table 3. Comparison of filtering performance for intrinsic image estimation methods under varying guidance images. We report the improvement in mean WHDR over the images in the test split from Narihira et al. [26], and for [38] results on their respective test set (marked with an asterisk) before and relatively to it after one filtering operation.

Method	F(CNN, CNN)	F(CNN, flat)	F([38], flat)*
unfiltered	19.49	19.49	17.85
BF	-0.6	-1.38	-1.47
GF	-0.25	-1.8	-1.98

the filter, we find the result of BF(CNN,flat) to improve to 18.1 and GF(CNN,flat) further to 17.7. This is on par with the current state-of-the art (17.67) which is the full pipeline of [7]: the flat image is clustered, followed by a CRF and another energy minimization step. We note that using the L_1 flattened result directly as reflectance image has 20.9 WHDR, which shows that there is complementary information in the CNN output and the guidance image.

This use of the flattened image as a guidance in a filter, extends [7] and allows application to other intrinsic image decompositions. We apply filtering to the second best method [38] on IIW. Their work proposes to create a sparse representation of the image by using the centers of a superpixelization. Patches around those centers are extracted and a CNN is used to provide an ordinal relationship via the three-way classification into "darker", "equal", and "lighter". This sparse result is then again densified by solving a constrained quadratic optimization problem to produce a full reflectance image.

The application of GF([38], flat) improves WHDR from 17.85 to 16.38. Repeated application of the filter further improves the output down to 15.78 after three applications of the guided filter. This result represents the new state-of-the-art by a large margin.

Filtering with the image itself A conceptually easier choice is to filter using the input image itself as guidance. We applied this to the Direct CNN predictions and searched for the hyper-parameters (see supplementary) to find $\sigma_s = 22$, $\sigma_r = 20$. Bilateral filtering with the input already improved test performance from 19.5 to 18.9. Guided filtering also improved a bit to 19.2 with r = 7 and $\varepsilon = 52$.

5.3. Discussion

We found throughout that guided filtering improves performance, see Table 3. There are some cases where the joint bilateral filter outperforms the guided filter, but in general the latter has better performance. Also the guided filter is magnitudes faster. We summarize a comparison with recent methods and state-of-the-art in Fig. 5.

As expected, the quantitative performance increases also



Figure 5. Comparison of performance to different intrinsic image estimation methods. Over all images in the test split from Narihira et al. [26], we report the statistics of the individual WHDRs on the images. The red line represents the median, the black line the mean. Results of Zoran et al. 2015 [38], are based on reflectance predictions provided by the authors which are generated on a different test split. All methods that are evaluated on this different test set are marked with an asterisk as they are not directly comparable.

the qualitative results. For an impression to assess the qualitative performance, we refer to Fig. 6. Results on a larger number of sample images and in comparison to more related work can be found in the supplementary material.

Runtime analysis In Fig. 7 we show the runtime of different algorithms against their WHDR. All methods of this paper are colored in green. We collected the timing estimates from the respective statements in the corresponding publications. By construction, our direct prediction CNN with only a few filters is fast at test time (180 fps on GPU) but it requires further filtering for better results. The bilateral filter adds around 2 s per image on CPU and the guided filter less than 0.1 s. The bottleneck of the filtering approach is the computation of the L_1 flattened guidance image.

6. Conclusion

In this paper we have proposed methods that are on opposing ends of employed prior knowledge. This led to both the best results on IIW and valuable insights into the current state of intrinsic image estimation. We presented the first end-to-end CNN method, trained on the WHDR-



Figure 6. Sample decompositions of the (a) input image with the IIW ID 71341 into (b) reflectance and (c) shading by the method of [38]. Filtering it using the flat guidance image in (d) results in the intrinsic layers (e) and (f) of our final model. (g)-(l) are the same as (a)-(f) for IIW ID 58346.



Figure 7. WHDR against runtime for related work. Mean WHDR of competing methods is evaluated on the Narihira test split [26] on the decompositions provided by [6] and the project pages of [7, 37]. Methods with an appended asterisk were evaluated on the test split given in [38]. For methods which are evaluated in [6], we used the reported runtimes on the corresponding project page. Methods developed in this work are plotted in green, previous results are plotted in blue.

Hinge loss, that predicts a dense result without any postprocessing step. Our finding is that a context-free per-pixel judgment is sufficient for competitive results. We believe that this should set a new lower bar for learning methods on IIW. While this observation may be attributed to an inherent bias in IIW, we have no qualified reason to believe so. We still conjecture that good results correlate with low WHDR numbers, and note that human performance sets a high bar with a median WHDR of only 7.5% [6]. This has not been attained by any automatic method so far. We further develop a filtering technique to implement the assumption of piecewise constant and sparse reflectance. This extends the work of [7] and makes it possible to apply their reflectance grouping to other decompositions. We find that a filtered CNN output is on par with the best published learning based methods and further improve the initial result of [38] to 15.78% on its testset, which is the lowest WHDR performance for a dense decomposition.

In summary, the findings of this paper suggest that it is still the use of strong prior knowledge in intrinsic estimation algorithms that drives empirical performance. More research will be necessary to build combined models and enable learning from sparse pairwise judgments. A future direction is to replace the expensive optimization of [7] with CNN inference, which would enable fast algorithms for high quality intrinsic video decompositions.

All code, models, and results are available at https://ps.is.tue.mpg.de/research_projects/reflectance-filtering.

References

- E. H. Adelson. Lightness perception and lightness illusions. New Cogn. Neurosci, 339, 2000.
- [2] J. T. Barron and J. Malik. Shape, illumination, and reflectance from shading. *TPAMI*, 2015.
- [3] H. Barrow and J. Tenenbaum. Recovering intrinsic scene characteristics from images. *Computer vision systems*, pages 3 – 26, 1978.
- [4] S. Beigpour, A. Kolb, and S. Kunz. A comprehensive multiilluminant dataset for benchmarking of the intrinsic image algorithms. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 172–180, 2015.
- [5] S. Beigpour, M. Serra, J. van de Weijer, R. Benavente, M. Vanrell, O. Penacchio, and D. Samaras. Intrinsic image evaluation on synthetic complex scenes. In *ICIP*, pages 285– 289, 2013.
- [6] S. Bell, K. Bala, and N. Snavely. Intrinsic images in the wild. ACM Transactions on Graphics (TOG), 33(4):159, 2014.
- [7] S. Bi, X. Han, and Y. Yu. An L1 image transform for edgepreserving smoothing and scene-level intrinsic decomposition. ACM Trans. Graph., 34(4):78:1–78:12, July 2015.
- [8] N. Bonneel, K. Sunkavalli, J. Tompkin, D. Sun, S. Paris, and H. Pfister. Interactive Intrinsic Video Editing. ACM Transactions on Graphics (SIGGRAPH Asia 2014), 33(6), 2014.
- [9] A. Bousseau, S. Paris, and F. Durand. User-assisted intrinsic images. In ACM Transactions on Graphics (TOG), volume 28, page 130. ACM, 2009.
- [10] D. J. Butler, J. Wulff, G. B. Stanley, and M. J. Black. A naturalistic open source movie for optical flow evaluation. In *Computer Vision–ECCV 2012*, pages 611–625. Springer, 2012.
- [11] Q. Chen and V. Koltun. A simple model for intrinsic image decomposition with depth cues. In *Computer Vision* (*ICCV*), 2013 IEEE International Conference on, pages 241– 248. IEEE, 2013.
- [12] W. Chen, Z. Fu, D. Yang, and J. Deng. Single-image depth perception in the wild. *Neural Information Processing Systems (NIPS)*, 2016.
- [13] D. Corney, J.-D. Haynes, G. Rees, and R. B. Lotto. The brightness of colour. *PloS one*, 4(3):e5091, 2009.
- [14] E. Garces, A. Munoz, J. Lopez-Moreno, and D. Gutierrez. Intrinsic images by clustering. In *Computer Graphics Forum*, volume 31, pages 1415–1424. Wiley Online Library, 2012.
- [15] 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 (NIPS), pages 765–773, 2011.
- [16] R. Grosse, M. K. Johnson, E. H. Adelson, and W. T. Freeman. Ground truth dataset and baseline evaluations for intrinsic image algorithms. In *Computer Vision, 2009 IEEE 12th International Conference on*, pages 2335–2342. IEEE, 2009.
- [17] K. He, J. Sun, and X. Tang. Guided image filtering. In Computer Vision–ECCV 2010, pages 1–14. Springer, 2010.

- [18] K. He, X. Zhang, S. Ren, and J. Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. arXiv preprint arXiv:1502.01852, 2015.
- [19] http://www.blender.org/.
- [20] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. In *Proceedings of the ACM International Conference on Multimedia*, pages 675–678. ACM, 2014.
- [21] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *The International Conference on Learning Representations (ICLR)*, volume 3, 2015.
- [22] P.-Y. Laffont and J.-C. Bazin. Intrinsic decomposition of image sequences from local temporal variations. In *ICCV*, 2015.
- [23] E. H. Land and J. McCann. Lightness and retinex theory. *JOSA*, 61(1):1–11, 1971.
- [24] L. Lettry, K. Vanhoey, and L. Van Gool. Darn: a deep adversial residual network for intrinsic image decomposition. *arXiv preprint arXiv:1612.07899*, 2016.
- [25] T. Narihira, M. Maire, and S. X. Yu. Direct intrinsics: Learning albedo-shading decomposition by convolutional regression. In *International Conference on Computer Vision* (*ICCV*), 2015.
- [26] T. Narihira, M. Maire, and S. X. Yu. Learning lightness from human judgement on relative reflectance. In *Proceed*ings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 2965–2973, 2015.
- [27] I. Omer and M. Werman. Color lines: Image specific color representation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2004.
- [28] G. Petschnigg, R. Szeliski, M. Agrawala, M. Cohen, H. Hoppe, and K. Toyama. Digital photography with flash and no-flash image pairs. ACM transactions on graphics (TOG), 23(3):664–672, 2004.
- [29] C. Poynton. Digital video and HD: Algorithms and Interfaces. Elsevier, 2012.
- [30] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *MIC-CAI*, 2015.
- [31] L. Shen and C. Yeo. Intrinsic images decomposition using a local and global sparse representation of reflectance. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 697–704. IEEE, 2011.
- [32] C. Tomasi and R. Manduchi. Bilateral filtering for gray and color images. In *Computer Vision*, 1998. Sixth International Conference on, pages 839–846. IEEE, 1998.
- [33] V. Vapnik. The Nature of Statistical Learning Theory. Springer, 1993.
- [34] Y. Weiss. Deriving intrinsic images from image sequences. In Computer Vision, 2001. ICCV 2001. Proceedings. Eighth IEEE International Conference on, volume 2, pages 68–75. IEEE, 2001.
- [35] F. Yu and V. Koltun. Multi-scale context aggregation by dilated convolutions. arXiv preprint arXiv:1511.07122, 2015.
- [36] Q. Zhao, P. Tan, Q. Dai, L. Shen, E. Wu, and S. Lin. A closed-form solution to retinex with nonlocal texture con-

straints. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 34(7):1437–1444, 2012.

- [37] T. Zhou, P. Krähenbühl, and A. A. Efros. Learning datadriven reflectance priors for intrinsic image decomposition. In *ICCV*, 2015.
- [38] D. Zoran, P. Isola, D. Krishnan, and W. T. Freeman. Learning ordinal relationships for mid-level vision. In *Proceedings* of the IEEE International Conference on Computer Vision, pages 388–396, 2015.