# IntrinSeqNet: Learning to Estimate the Reflectance from Varying Illumination

Grégoire Nieto
Technicolor
gregoire.nieto@gmail.com

Mohammad Rouhani
InterDigital
mohammad.rouhani@interdigital.com

Philippe Robert
InterDigital
philippe.robert@interdigital.com

## Abstract

*Intrinsic image decomposition describes an image based on its reflectance and shading components. In this paper we tackle the problem of estimating the diffuse reflectance from a sequence of images captured from a fixed viewpoint under various illuminations. To this end we propose a deep learning approach to avoid heuristics and strong assumptions on the reflectance prior. We compare two network architectures: one classic 'U' shaped Convolutional Neural Network (CNN) and a Recurrent Neural Network (RNN) composed of Convolutional Gated Recurrent Units (CGRU). We train our networks on a new dataset specifically designed for the task of intrinsic decomposition from sequences. We test our networks on MIT and BigTime datasets and outperform state-of-the-art algorithms both qualitatively and quantitatively.*

## 1. Introduction

Intrinsic image decomposition describes an image based on its reflectance and shading components with many applications ranging from stabilization, re-colorization, relighting to texture and virtual object insertion to mention a few. Given an image $I$, this problem aims at disentangling the shading component $S$ from the albedo $R$, which is referred to as *reflectance* or *diffuse reflectance*: $I = R.S$. The problem of Single Image Intrinsic Decomposition (SIID) is an ill-posed problem as we have infinite number of potential solutions for a single image. To reduce the ambiguity of the decomposition, the given image can be accompanied by a sequence of images captured under different lighting condition. This prevents common failure cases of single image methods such as difficulties to handle hard cast shadows and bright specularities. We refer to it as Multiple Image Intrinsic Decomposition (MIID).

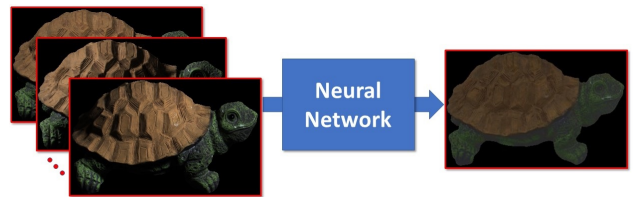Traditional approaches, regardless of using a single in-



Figure 1: At training time and inference time, our neural network processes a sequence of images $(I)_t$ captured under various (uncorrelated) illuminations, and produces the time-invariant reflectance image $R$.

put image or a sequence of images, estimate the reflectance through direct optimization. The quality of the result highly depends on heuristics modeling and hand-crafted priors on shading and reflectance. Classical optimization-based methods like [18] disentangle reflectance and shading by formulating strong assumptions on reflectance that are derived from prior knowledge. For instance, in [18], it is assumed that pixels with similar normalized intensity profile over time are likely to have the same shading. The result shown in Fig. 2(a) clearly depicts seams between clusters of pixels that are supposed to share the same shading over time. Increasing the number of clusters (Fig. 2(b)) does not improve the result, which implies that these methods strongly depend on the *assumption* over the reflectance priors. One of the main advantage of using deep learning over classic approaches is to get rid of explicit priors, and rather, to learn them implicitly.

Another flaw of prior-based approaches lies in *how* these priors are tuned. Fig. 2(c) displays what we obtain with different settings for $\alpha$ coefficient that modulates the regularization term. This term forces the solution to be close to the mean chromaticity image. Since the heuristics that are used to disambiguate the intrinsic decomposition are not directly derived from a mathematical formulation of the prob-
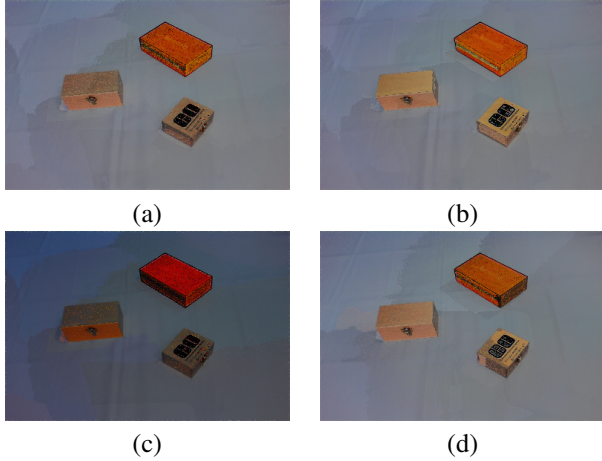
Figure 2: Results of [18]. (a) Reflectance obtained with 20 pixel clusters and $\alpha = 100.0$. (b) Effects of changing the number of pixel clusters (60 instead of 20). (c) Effects of different parameter tuning ($\alpha = 50.0$ instead of $\alpha = 100.0$). (d) Poor illumination invariance: reflectance estimated from a different set of shaded images (same parameters as in (a)).

lem, but rather from prior knowledge of how we expect the reflectance to be, finding the right coefficient to balance the cost function ends up to unhandy parameter tuning.

Moreover, intrinsic image decomposition should be invariant to illumination changes, which means that the decomposition should be consistent over a sequence of frames with varying illumination. This is the objective that [22] aims at, by training on sequences of various illuminations and forcing the predicted reflectance to be similar over the same sequence. However there is no guaranty that the network can learn this property at the inference time.

In this paper we propose a new intrinsic image decomposition method that estimates the reflectance image from a sequence of images captured from a fixed viewpoint under various illuminations (Fig. 1). To this end we propose a deep learning approach to avoid heuristics and strong assumptions on reflectance priors. We compare two network architectures: one classic 'U' shaped Convolutional Neural Network (CNN) and a Recurrent Neural Network (RNN) composed of Convolutional Gated Recurrent Units (CGRU). We train our network on a new dataset specifically designed for the task of intrinsic decomposition from sequences. We test our network on MIT [11] and Big-Time [22] datasets and outperform state-of-the-art algorithms both qualitatively and quantitatively.

## 2. Related Work

**Intrinsic image decomposition from a single image**    The intrinsic image decomposition domain is mainly repre-

sented by methods that estimate the reflectance by using priors-based optimization, from a single image and without any additional data. For a long time they were largely dominated by the simple yet efficient Color Retinex algorithm [11]. It decomposes the image by assuming that a change in color is due to a change in shading when the chromaticity remains constant, and due to a change in reflectance otherwise. This idea derives from the prior knowledge that the reflectance is somehow closely related to the chromaticity. In [35], a non-local reflectance constraint is included to the Retinex formalism, by enforcing reflectance similarity between pixels that share similar chromaticities. Non-local reflectance priors were adopted by subsequent works [20, 34, 16, 26, 24, 9], as well as non local shading priors [20, 16] by adding valuable improvement to widely used local reflectance and shading priors. For example in [9], the dimensionality of the problem is reduced by adding a global sparsity prior, assuming that the final reflectance is made from a sparse set of basis colors. Given the ill-posed nature of the intrinsic decomposition problem, formulating novel hand-crafted priors has been a crucial issue. Deep learning is proposed as an alternative to bypass this approach by learning implicit priors *from the data itself*.

**Intrinsic image decomposition with deep learning**    Machine learning has been proposed as a solution to the problem in the seminal work on relative reflectance judgments [28], trained on human judgment dataset [3]. Likewise in [36], it is propose to learn relative reflectance judgments, but as a first step before recovering dense reflectance map like in [3]. Indeed machine learning is used as a way to estimate priors, not to estimate the final reflectance itself. Direct reflectance estimation by a CNN was an original idea by [27], improved by [30] that introduces the use of a U-net [29] for intrinsic decomposition. There has been many variations of these approaches [21, 2, 8, 15] but to our knowledge, only [22] suggests the use of sequences of illumination-varying images to train a neural network. However it is still a one-to-one process at inference time, thus well-adapted to classic CNN encoder-decoder architecture. On the other hand our method is a many-to-one process at training time *and* inference time.

**Intrinsic decomposition from image sequences**    Methods that benefit from temporal constraints by processing videos are numerous [20, 34, 16, 26, 24]. On the contrary, our approach does not require any temporal consistency between views in our sequences, nor extra information (such as depth [5, 20, 1] or user input [25]). Our problem formulated our way was initially addressed by [32]. They produce reasonable results by taking advantage of prior knowledge on natural images. However in [23] it is shown that such results may severely be altered by a biased illumina-

tion. More recent solutions were provided to the multi-view problem [19, 7] (movable viewpoint). In [19], the inferred 3D geometry imposes shading constraints for intrinsic decomposition. In [18], a state-of-the-art solution is proposed to solve the initial fixed viewpoint problem [32, 23]; they alleviate the lack of geometry by cleverly clustering pixels that share the same radiance profile over time. But like previous optimization-based work, their decomposition algorithm strongly depends on reflectance heuristics, handmade priors and coefficient tuning. As far as we know, we are the first to adopt neural network to solve this problem, although neural networks have already been proposed for other similar applications such as photometric stereo and general inverse rendering [31, 14].

**Recurrent Neural Network (RNN)**   RNNs are networks that loop over themselves, so the output is passed as an input at the next iteration. They are ideal to process sequences of correlated data, such as words in a sentence of frames in a video. The number of times the network loops over itself is not part of its architecture, which makes it flexible to any sequence length; this is analogous to how convolution is flexible for the spatial size of inputs. However during the training, the back-propagated loss gradient tends to vanish at every iteration: an issue known as *vanishing gradients* [12]. To overcome this problem, Long Short-Term Memory (LSTM) units have been proposed. They are composed of a memory cell that *remembers* information of previous sequences. The information flow over sequences is regulated by three *gates* inside the cell. Applied to images, convolutional LSTMs [33] replace the traditional dense matrix to vector operations by convolutions. In Gated Recurrent Units [6], the memory cell is also used as hidden layer, requiring fewer parameters while performing similarly.

## 3. Model

In this section we describe two architectures that we compare in section 5: a fully convolutional U-net and convolutional GRU, referred to as CGRU. To describe precisely the architectures, we refer to $T$ as the number of frames in a sequence, $C$ the number of input image channels, $F$ the number of output features (or output channels) and $(H, W)$ the spatial dimensions of input/output images. We feed the network with tensors of shape $T \times C \times H \times W$, and obtain a prediction tensor of size $F \times H \times W$. In practice we want to obtain RGB reflectance images, so $F = C = 3$. Images are resized so that their spatial dimensions are either $256 \times 384$, $384 \times 256$ or $256 \times 256$ depending on their original ratio.

### 3.1. U-net

We implemented a 'U' shaped architecture, that has already shown some potential in SIID [30, 15, 2, 22].

The original architecture used for medical segmentation [29, 13], is composed of an encoder and a mirrored decoder. This classic U-net architecture has been modified for adapting with our application. Like [22] it takes a sequence of images but stacks them along the channel dimension, so that the shape of the input tensor is $T.C \times H \times W$. It is composed of 6 convolutional layers for the encoder, 6 for the decoder and one central. Our U-net is lighter than [22]: each convolutional module consists of a single 2D convolution (with a stride of 2 for the encoder), a batch normalization and a *RELU*. Indeed due to the limited amount of data (every sequence is one single training example), we need to decrease the number of parameters to a reasonable level to avoid overfitting. In addition each convolutional layer of the decoder possesses a bilinear interpolation (that doubles the spatial dimension) and a concatenation of the input signal of corresponding spatial dimensions (skip-connection).

Since input images are stacked together in the channel dimension, the length of a sequence has to be fixed, otherwise the architecture changes (the number of convolutions of the first encoding layer). The flexible nature of RNNs makes them better-suited than CNNs to process sequences of images.

### 3.2. Convolutional Gated Recurrent Unit (CGRU)

Our second architecture is a Recurrent Neural Network. Using such RNNs has several advantages over CNNs:

- it needs fewer parameters (thousands instead of millions) because every image in the sequence is processed by units that share the same weights;

- since it loops over the sequence, it can process any number of frames;

- it requires less GPU memory because the recurrent loop can be run sequentially.

Following the idea that our model should have as few parameters as possible so it generalizes well to any kind of scene, we opt for Gated Recurrent Unit (GRU). Like [33] with LSTMs, we implemented a convolutional version of the GRU (CGRU), that replaces matrix-vector operations by convolutions, so that it processes information both *temporally* and *spatially*. Let $I_t$ be the current input frame and $h_t$ the hidden layer (or output) at time $t$, the next hidden layer $h_{t+1}$ or output of the unit is computed as followed:

$$z_t = \sigma(W_{zi} * I_t + W_{zh} * h_t + b_z) \tag{1}$$

$$r_t = \sigma(W_{ri} * I_t + W_{rh} * h_t + b_r) \tag{2}$$

$$\tilde{h}_t = \tanh(W_{hi} * I_t + W_{hh} * (r_t \circ h_t) + b_h) \tag{3}$$

$$h_{t+1} = (1 - z_t) \circ h_t + z_t \circ \tilde{h}_t \tag{4}$$

with $*$ denoting the convolution operation, $W$ and $b$ the weights and biases of the convolutions, $\circ$ the element-wise
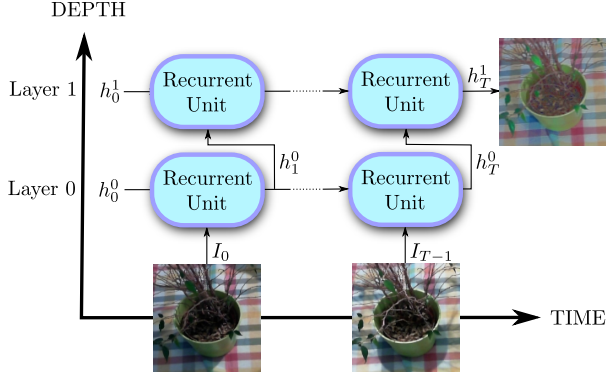
Figure 3: CGRUs can be stacked vertically to enable deep feature processing. Here two layers 0 and 1 are displayed. The sequence of output $(h^0)_t$ of the first layer 0 is provided to the second layer 1, that outputs a second sequence $(h^1)_t$. The final prediction $R$ is the output $h_T^1$ of the last layer.

product, $\sigma$ the sigmoid function and $\tanh$ the hyperbolic tangent. Hidden layers $h_t$ are tensors of shape $F \times H \times W$ and input frames $I_t$ are tensors of shape $C \times H \times W$.

The hidden layer $h$ behaves like the memory cell of a LSTM unit, forgetting and learning information from the successive frames that are fed to the unit. Intuitively we can think of it as containing all the illumination-invariant information of the sequence, which is expected to be the reflectance. At each iteration, a new input frame passes through the gates that select the information $h$ will forget with the reset gate image $r_t$, and update $h$ with the update gate image $z_t$. Initially the hidden layer is set to 0, which means we have no prior memory. Eventually when we have iterated over the whole sequence, the prediction $R$ is the final hidden layer $h_T$.

It is also possible to *stack* several layers of CGRUs. In Fig. 3 we represent two layers 0 and 1. Each recurrent unit output $h_t^l$ is not only passed *horizontally* to the next iteration to be jointly processed with $I_{t+1}$ by the same unit, but also *vertically* to another unit as a new input. Units from other layers do not share the same weights. Stacked layers can be seen as if the sequence was processed by successive convolutional layers, which enables deeper feature processing and structure extraction. The final prediction $R$ is the output $h_T^L$ of the last layer $L$ ($h_T^1$ in Fig. 3).

### 3.3. Achromatic Illumination Assumption

The assumption of the achromatic illumination is common in the SIID literature [3, 36, 8]. However it is not clear whether it improves or deteriorates the prediction. It assumes that the lighting intensity is constant over the sequence, equal to $(1,1,1)$, which is the case in most scenes but not always (multiple light sources, colored illumination,

dimming light). As a consequence the shading is greyscale; we can predict a greyscale reflectance $r$ and recover the tri-color reflectance $R$ with the median chromaticity of the sequence:

$$R = 3 \cdot r \cdot \operatorname{median}_t\left(\frac{I_t^c}{\sum_c I_t^c}\right) \text{ with } I_t^c \text{ the } c^{th} \text{ channel of } I_t.$$
(5)

The relevance of using such assumption is discussed in the section 5, where both models (achromatic that predicts $r$ with $F = 1$ and chromatic that predicts $R$ with $F = 3$) are confronted.

### 3.4. Loss

Our model fitting is supervised by comparing the prediction $R$ and the ground truth reflectance $R^*$ in the loss function. We force the prediction to be close to the ground truth not only in terms of color, but also in terms of gradient [27]. In [27] however, a L2-norm is used; we use instead an L1, known to be more robust to outliers. We do not apply the loss to the shading because our model implies that it is directly derived from the input sequence and the predicted reflectance. We tried to add the same shading smoothness constraint as in [18], but we obtained worse results.

In addition we apply a total variation [10] on the predicted reflectance. It comes from the observation that natural images are generally piecewise smooth. In the section 5 we study the influence of the total variation term on the quality of the images. The final loss function is

$$\mathcal{L}(R, R^*) = \|R^* - R\|_1 + \|\nabla R^* - \nabla R\|_1 + \|\nabla R\|_1. \text{ (6)}$$

### 4. Dataset

Traditionally, SIID methods train their models and evaluate their results on 4 datasets: MIT [11], MPI Sintel [4], IIW [3] and SAW [17]. Nevertheless only MIT is suited to the task of MIID, since the rest only contain single image examples. Recently a new dataset called BigTime [22] was introduced, containing sequences on images with time-varying illumination. However this dataset does not provide any sort of ground truth or reference data since it was designed specifically for unsupervised learning of SIID.

Because the MIT dataset alone is not sufficient neither for training or for evaluation, we crafted our own dataset, called Washington, composed of real and synthetic images. Our dataset combines the better of two worlds: sequences of real images with changing illumination and an image used as reflectance ground truth. Washington is made of indoor scenes of various objects on a table. Indoor scenes are chosen over outdoor so we are able to capture a large diversity of illuminations in a controlled environment, which makes it possible to estimate a pseudo-reflectance for supervised training. Indeed to obtain an acceptable ground truth
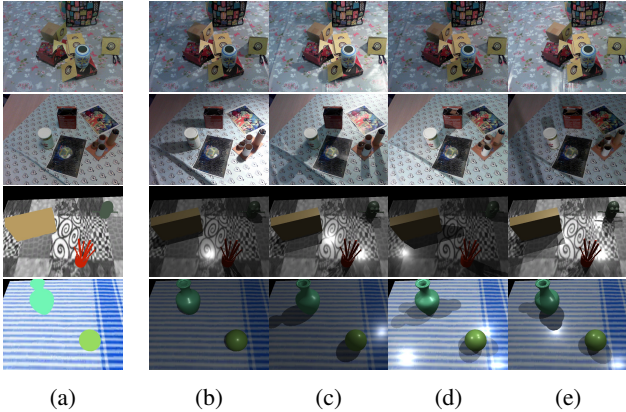
Figure 4: A sample of our training set. First two rows: real data. Last two rows: synthetic data. (a) Pseudo-reflectance (reference). (b-e) Some images from sample sequences.

and reference image, we capture an image under pseudo-ambient lighting, which makes the shading almost constant everywhere. Such image, called *pseudo-reflectance*, is not the proper reflectance since it still contains ambient occlusion (inter-shadowing of concave surfaces), but section 5 shows that it is sufficient for the supervision of MIID.

We create our real dataset (see the first two rows in Fig. 4) by capturing sequences of frames from a fixed viewpoint, and changing the illumination. We used 10 sets of objects; for each set we create ten various scenes, changing the position of the objects and the camera viewpoint, which makes a total of 100 scenes. From every of these 100 captured videos, we randomly extract 10 sequences of 30 uncorrelated frames (there is not continuous variation of shading over time). The complete dataset of 1000 sequences of real images is split into a training set, a validation set and a test set with a respective proportion of (80% / 10% / 10%). We make sure that the splits do not share any scenes or objects.

To cope with our issue of not having a ground truth reflectance (in particular for evaluation since a pseudo-reflectance image suits well to MIID supervision) we enhance our Washington dataset with 780 supplementary sequences of synthetic views (see the last two rows in Fig. 4). The synthetic scenes were created and rendered using Unity; they simulate a plane carrying various 3D objects from the daily life. Up to 3 point light sources of different intensities lit the scenes. In addition we use a large collection of textures and materials to render specularities. The dataset is split into a training set (700 sequences) and a test set (80 sequences).

# 5. Experiments

We perform several experiments to validate the performance of our models. First we study the relevance of the achromatic illumination assumption. Second we observe the influence of the number of input views on the predicted reflectance, and validate the recurrent model (CGRU) by decomposing the result frame by frame. Then we compare our models to state-of-the-art methods, in term of numerical, visual and runtime performance. Last but not least we address some limits of our models.

**Training** Our models (U-net and CGRU) are trained on our Washington training set (real images), unless otherwise stated. The 800 training sequences contain 15 images that are resized to $384 \times 256$. At each iteration, network is fed a single sequence, which means the batch size is one (huge due to memory requirement). We perform 40 epochs and the training lasts approximately 10 hours on an NVidia GTX 1080. We save the model for which the validation loss (computed on the real Washington validation set) is minimum. We use Adam optimizer with a learning rate of 0.0005.

**Evaluation metrics** Like previous experiments in the literature [27, 30, 22], we use the metrics MSE, LMSE and DSSIM to evaluate our results. MSE is a scale-invariant version of the Mean Square Error introduced by [11] to account for relative reflectance. LMSE is the Local Mean Square Error [11], which is the MSE computed on $16 \times 16$ overlapping patches (the size matters little according to the authors). DSSIM is the structural dissimilarity index. Predicted reflectance images are compared to reference image. In the case of real Washington dataset, for which we have no proper ground truth, we compare the prediction to the pseudo-reference.

## 5.1. Chromatic versus Achromatic Illumination Assumption

In this experiment we validate the use of the achromatic assumption that models the shading as being greyscale. Therefore, as detailed in the section 3, instead of predicting tricolor reflectance we predict a greyscale reflectance and recover the color thanks to the median chromaticity of the sequence (equation 5). This assumption is validated on the Washington real test set and on the MIT dataset. The model used for this experiment is CGRU with a depth (number of layers $L$) of 1. Although there is no significant numerical difference between the chromatic and the achromatic model on Washington dataset, we notice a strong difference on MIT dataset in favor of the achromatic model (LMSE = 0.0581 against 0.0866). The superiority of the achromatic model is qualitatively validated on Washington, as high-
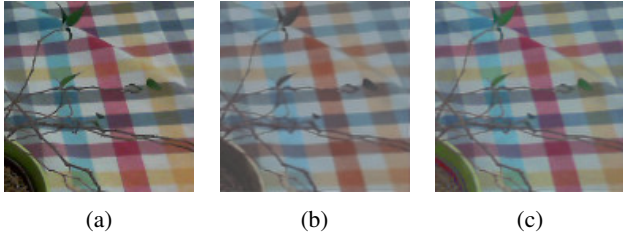
(a)                    (b)                    (c)

Figure 5: Qualitative validation of the achromatic illumination assumption of a sample of the Washington validation set. (a) Reference. (b) Chromatic model. (c) Achromatic model.
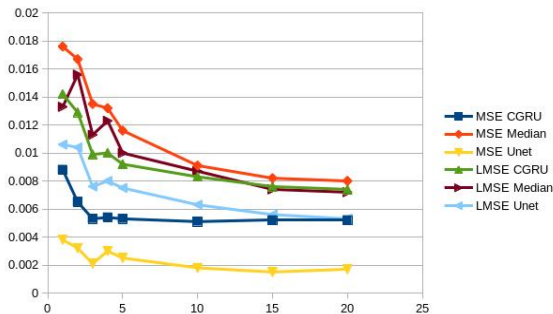


Figure 6: Performance of U-net and CGRU on Washington real test set with an increasing number of input frames. The median image of the sequence is used as baseline.

lighted in Fig. 5. The achromatic model tends to produce images with colors that are more faithful to the original. In addition we observe that the obtained tricolor reflectance from the chromatic model becomes more and more greyish as we increase the depth of the recurrent network ($L > 1$). For all these reasons we performed the next experiments with the use of the achromatic model.

## 5.2. The Influence of the Number of Images in a Sequence

We analyze the influence of the number of images on the predicted reflectance by testing the U-net and CGRU ($L = 1$) models on our Washington real test set. The Fig. 6 illustrates the results obtained with a number of input images $T$ that varies between 1 and 20. The figure clearly shows the superiority of U-net against CGRU on this test set. As expected, increasing the number of input frames increases the performance. However CGRU behaves differently: it quickly converges towards its maximum performance after a few frames, contrarily to U-net whose performance w.r.t. the number of input frames appears more linear.

We show the non-linear behavior of the CGRU in Fig.

7, that displays two sequences of the validation set and the predicted output over time (b-i). Notice that the shadow formed by the flower pot on the top sequence quickly disappears from the prediction. Although the convergence is fast, the output is not disturbed by new challenging input frames that contain hard cast shadows (Fig. 7 (f)) nor over/under-exposed frames (Fig. 7 (h-i)). The bottom sequence is another example of a robust prediction: when the light direction abruptly changes from Fig. 7 (h) to Fig. 7 (i), no shadow is added at the right of the green hose, but the left shadow of the hose and the coffin are largely attenuated. This is a case where the prediction overperforms the reference. It means that the recurrent model, by its simplicity, easily generalizes to produce reflectance images that are even better than the pseudo-reflectance images.

A clear disadvantage of U-net is that it needs to be retrained each time we change the number of input frames, because of its rigid architecture. On the contrary the prediction of CGRU ($L = 1$) is improved only by feeding it another image, because it processes the input images sequentially. The number of images CGRU ($L = 1$) can process without the need for another training is virtually infinite. Note that it is the case for most optimization-based methods like [32, 23, 18]. We also tested [18] with different numbers of input images and it is always outperformed by our models (*e.g.* LMSE = 0.0198 for $T = 5$ or LMSE = 0.017 for $T = 20$). In the next subsection we will see further comparisons with state-of-the-art methods.

## 5.3. Comparison with State-Of-the-Art Methods

We compare our models to state-of-the-art methods [32, 23, 18, 22]. In [32, 23, 18], the same assumptions as ours are made: the viewpoint is fixed and we process a sequence of 15 images. For [18], the $\alpha$ parameter value that produces best results in our dataset is 100; we set the number of clusters to 20. We choose to compare to [32] because despite its age it is one of the few methods whose results still compete with [18]. For [23], we keep the same parameters as in the original paper. We also compare our models to the recent single image method [22]: we obtain the results by computing the temporal median of the predicted reflectance over the sequence. Note that we use the temporal median of the input raw sequence as baseline.

**Numerical results** Numerical results are presented in the table 1. It is shown that our models largely outperform all other methods. U-net seems to perform the best, followed by CGRU with a single depth layer (CGRU-1). Surprisingly deep recurrent networks ($L > 1$) do not outperform the single layer one; networks with 3 and 4 layers are not shown because they produce worse results, only CGRU-2 is showed ($L = 1$). Visually we observe that the deeper the CGRU, the blurrier the predictions tend to be. Training on

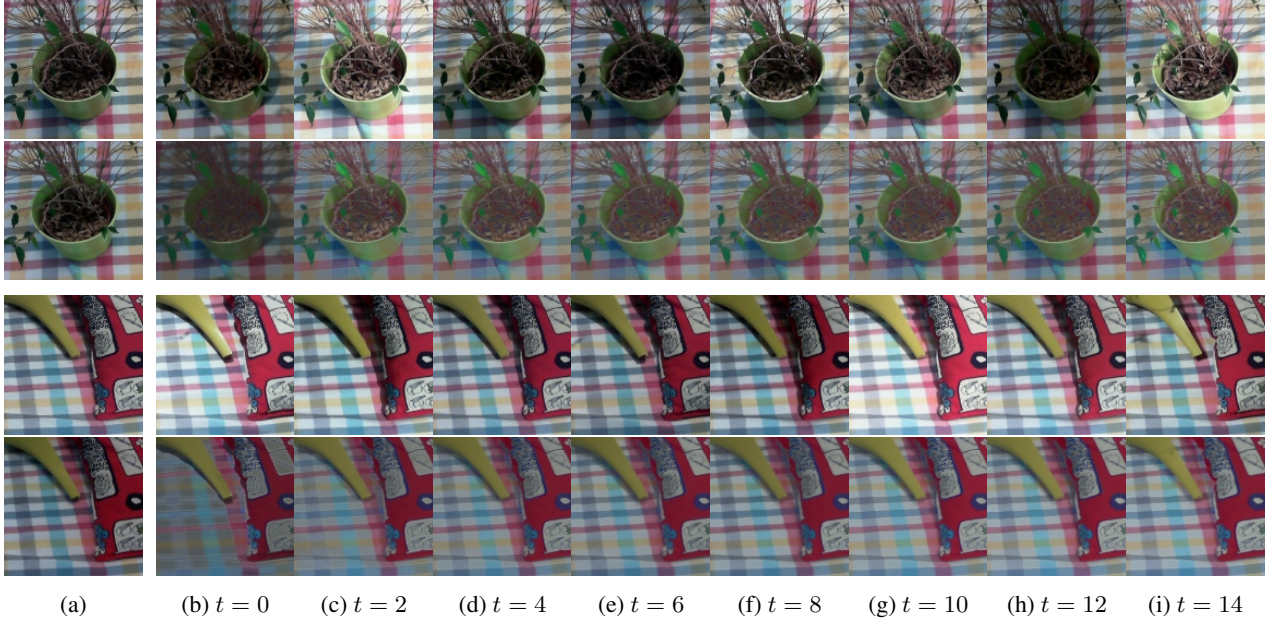|  | (a) | (b) $t = 0$ | (c) $t = 2$ | (d) $t = 4$ | (e) $t = 6$ | (f) $t = 8$ | (g) $t = 10$ | (h) $t = 12$ | (i) $t = 14$ |

Figure 7: Our neural network performing on two sequences of the validation dataset. (a) Top: median image. Bottom: reference (pseudo-reflectance). (b-i) Input images $I_t$ (top) and output images $h_t$ (bottom) at times $t = 0, 2, 4, 6, 8, 10, 12$ and $14$. Best viewed numerically.

| Metric | Washington Virtual Dataset | | | Washington Real Dataset | | | | MIT Dataset | | |
|---|---|---|---|---|---|---|---|---|---|---|
|  | MSE | LMSE | DSSIM | Variance | MSE | LMSE | DSSIM | MSE | LMSE | DSSIM |
| Median | 0.1498 | 0.0586 | 0.3613 | 125 | 0.0082 | 0.0074 | 0.2139 | 0.0573 | 0.1511 | 0.1400 |
| [32] | 0.1334 | 0.0689 | 0.3556 | 111 | 0.0087 | 0.0105 | 0.3039 | 0.0497 | 0.1364 | 0.1470 |
| [23] | 0.1378 | 0.0748 | 0.3433 | 103 | 0.0077 | 0.0076 | 0.2001 | 0.0448 | 0.1089 | 0.1289 |
| [22] (median) | 0.0770 | 0.1062 | 0.3138 | **2** | 0.0095 | 0.0303 | 0.2363 | 0.0292 | 0.0892 | 0.1498 |
| [19] | 0.0718 | 0.0885 | 0.3002 | 79 | 0.0107 | 0.0173 | 0.2541 | 0.0381 | 0.0920 | 0.1426 |
| CGRU-2-TV | **0.0661** | 0.0598 | 0.3179 | 15 | 0.0049 | 0.0097 | 0.2206 | 0.0377 | 0.0770 | 0.1462 |
| CGRU-2-L1* | 0.0718 | 0.0596 | 0.2914 | 27 | 0.0046 | 0.0093 | 0.2394 | 0.0266 | 0.0662 | 0.1335 |
| CGRU-2-L1 | 0.1017 | 0.0611 | 0.3701 | 46 | 0.0037 | 0.0066 | 0.2060 | 0.0365 | 0.0698 | 0.1374 |
| CGRU-1-TV | 0.0820 | 0.0506 | 0.3149 | 69 | 0.0061 | 0.0067 | 0.2198 | 0.0264 | 0.0592 | 0.1354 |
| CGRU-1-L1* | 0.0712 | 0.0596 | **0.2780** | 40 | 0.0064 | 0.0096 | 0.2365 | 0.0263 | 0.0632 | 0.1315 |
| CGRU-1-L1 | 0.0715 | 0.0533 | 0.3216 | 58 | 0.0052 | 0.0076 | 0.2192 | 0.0246 | 0.0581 | 0.1342 |
| U-net-TV | 0.0836 | **0.0439** | 0.3222 | 35 | 0.0023 | 0.0057 | 0.2071 | 0.0283 | 0.0755 | 0.1418 |
| U-net-L1* | 0.0824 | 0.0536 | 0.3305 | 32 | 0.0019 | 0.0058 | 0.2046 | 0.0407 | 0.0862 | **0.1271** |
| U-net-L1 | 0.0847 | 0.0509 | 0.3233 | 23 | **0.0015** | **0.0056** | **0.1993** | **0.0236** | **0.0578** | 0.1309 |

Table 1: Numerical results on our test dataset. The lower the better (in bold). Variance is $10^5 \times$ the original value.

the virtual train set in addition to the real train set (models marked with a * in the table 1) only improves results on the virtual test set. However the performance of deep CGRU $L > 1$ are improved when additional synthetic data is used, from which we can infer that their mediocre performance compared to simpler models is due to overfitting.

**Illumination invariance criterion** In addition to previously presented metrics, we add an illumination invariance criterion. It is the variance of the reflectances from several *sequences* of images of the same scene captured under different illuminations. Since in our Washington dataset we have 10 different sequences of illuminations per scene, we compute the variance over the 10 predicted reflectances; the

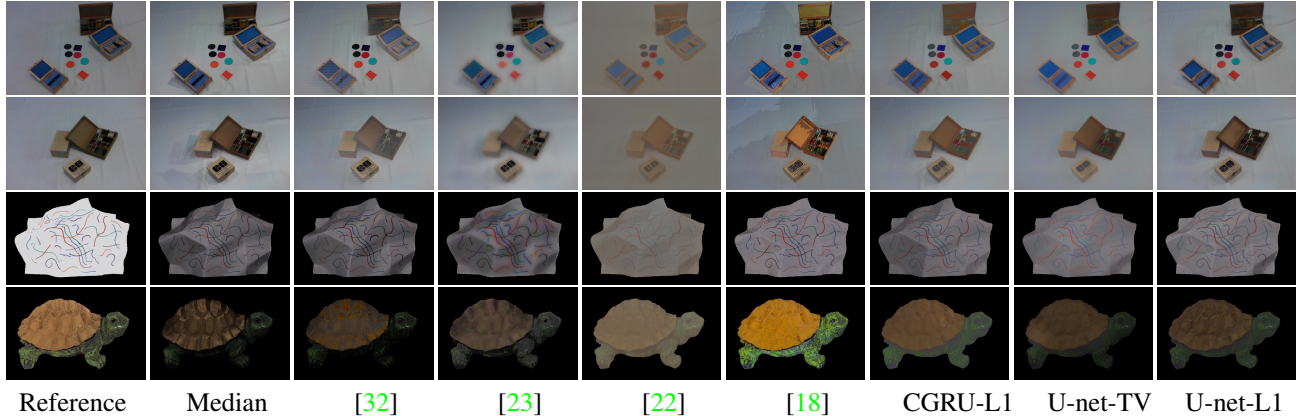| Reference | Median | [32] | [23] | [22] | [18] | CGRU-L1 | U-net-TV | U-net-L1 |

Figure 8: Comparison with state-of-the-art method on real Washington (first two rows) and MIT (last two rows) datasets. The median image is used as baseline for comparison. "TV" means that the total variation was used in the loss; we write "L1" otherwise.

lower variance the better. The SIID method [22] outperforms the others by far. We believe the way the model was trained enforced the illumination invariance of their output (which was a key contribution of their work). Nevertheless the quality of the predicted reflectance is inferior to ours.

**Loss function** We compare the performance our models trained with either the L1 norm only (written "L1" on the table 1), or the additional total variation term (written "TV"). We observe that the TV term penalizes the network in general. However its resulting prediction is visually interesting. In Fig. 8 we can state that the TV network performs better in the case of a simple uniform reflectance (*e.g.* for the paper on the third row). In contrast it performs poorly in the presence of highly detailed images such as the turtle (last row), smoothing out the details on the shell. The reason why it generally underperforms the simple yet generic L1 loss, is probably because it oversimplifies the predicted reflectance, making it piece-wise constant where texture should be left untouched. Nevertheless even trained with the TV loss, the U-net model still outperforms state-of-the-art approaches.

**Runtime performance** Like in [24] (where they use a geometric proxy and user scribbles), our implementation is real-time. In contrast, [34] processes a single frame in a minute, and [16] in ten minutes. The faster network is U-net (3ms to process 15 frames of size $384 \times 256$), then CGRU-1 (7ms) and finally CGRU-2 (17ms). BigTime [22] is also quite fast even when applied to the whole sequence (4ms). Other methods [32, 23, 18] are iterative and take approximately 30s. Note that we could obtain the reflectance with [32] faster if we computed the pseudo-inverse directly instead of solving the problem iteratively as we did.

### 5.4. Limitations

Our model is not exempt of limits, especially when it comes to decompose images that have a complex shading, such as the panther in the MIT dataset. Although cast shadows are removed or at least well attenuated, shading in the form of ambient occlusion is still clearly visible. We suspect that this incapability to completely remove self-shadowing on non-convex parts of the object may be due to the fact that the model training is supervised by incorrect reflectance images. Since the pseudo-ambient lighting used to capture so called ground truth reflectance does not prevent ambient occlusion and inter-shadowing, it is not surprising that it fails at inference stage.

## 6. Conclusion

We have presented an end-to-end method to estimate the reflectance from a sequence of images that are captured from the same view under various illuminations. Contrary to state-of-the-art approaches, we do not rely on any prior knowledge on reflectance, nor hand-crafted priors, but rather learn from the data itself. As a consequence no parameter tuning is needed at inference time. Two different models are proposed to solve this problem, based on U-net and recurrent network (CGRU), while the earlier has the advantage of being fast and provides better results, the latter is more flexible and requires less memory storage as it processes any number of views sequentially. Both networks process sequences in real-time and outperform state-of-the-art methods. Moreover a new dataset has been provided, including sequences of images and their ground truth reflectances. We hope this will encourage people to train and evaluate their networks for the tasks of single and multiple image intrinsic decomposition.

# References

[1] A. Alperovich and B. Goldluecke. A Variational Model for Intrinsic Light Field Decomposition. In *SpringerLink*, pages 66–82. Springer, Cham, Nov. 2016. 2

[2] A. S. Baslamisli, H.-A. Le, and T. Gevers. CNN Based Learning Using Reflection and Retinex Models for Intrinsic Image Decomposition. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 2, 3

[3] S. Bell, K. Bala, and N. Snavely. Intrinsic Images in the Wild. *ACM Trans. on Graphics (SIGGRAPH)*, 33(4), 2014. 2, 4

[4] D. Butler, J. Wulff, G. Stanley, and M. Black. MPI-Sintel optical flow benchmark: Supplemental material. Technical report, MPI-IS-TR-006, MPI for Intelligent Systems, 2012. 4

[5] Q. Chen and V. Koltun. A Simple Model for Intrinsic Image Decomposition with Depth Cues. In *2013 IEEE International Conference on Computer Vision*, pages 241–248, Dec. 2013. 2

[6] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio. Learning phrase representations using RNN encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014. 3

[7] S. Duchêne, C. Riant, G. Chaurasia, J. Lopez-Moreno, P.-Y. Laffont, S. Popov, A. Bousseau, and G. Drettakis. Multi-View Intrinsic Images of Outdoors Scenes with an Application to Relighting. *ACM Transactions on Graphics*, 2015. 3

[8] Q. Fan, J. Yang, G. Hua, B. Chen, and D. Wipf. Revisiting Deep Intrinsic Image Decompositions. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 2, 4

[9] P. V. Gehler, C. Rother, M. Kiefel, L. Zhang, and B. Schölkopf. Recovering Intrinsic Images with a Global Sparsity Prior on Reflectance. In *Proceedings of the 24th International Conference on Neural Information Processing Systems*, NIPS'11, pages 765–773, USA, 2011. Curran Associates Inc. 2

[10] B. Goldluecke and D. Cremers. An approach to vectorial total variation based on geometric measure theory. In *2010 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 327–333, June 2010. 4

[11] R. Grosse, M. K. Johnson, E. H. Adelson, and W. T. Freeman. Ground truth dataset and baseline evaluations for intrinsic image algorithms. In *2009 IEEE 12th International Conference on Computer Vision*, pages 2335–2342, Sept. 2009. 2, 4, 5

[12] S. Hochreiter, Y. Bengio, and P. Frasconi. Gradient Flow in Recurrent Nets: The Difficulty of Learning Long-Term Dependencies. In J. Kolen and S. Kremer, editors, *Field Guide to Dynamical Recurrent Networks*. IEEE Press, 2001. 3

[13] V. Iglovikov and A. Shvets. TernausNet: U-Net with VGG11 Encoder Pre-Trained on ImageNet for Image Segmentation. *ArXiv e-prints*, 2018. 3

[14] S. Ikehata. CNN-PS: CNN-based Photometric Stereo for General Non-Convex Surfaces. In *The European Conference on Computer Vision (ECCV)*, Sept. 2018. 3

[15] M. Janner, J. Wu, T. D. Kulkarni, I. Yildirim, and J. Tenenbaum. Self-supervised intrinsic image decomposition. In *Advances in Neural Information Processing Systems*, pages 5938–5948, 2017. 2, 3

[16] N. Kong, P. V. Gehler, and M. J. Black. Intrinsic Video. In *SpringerLink*, pages 360–375. Springer, Cham, Sept. 2014. 2, 8

[17] B. Kovacs, S. Bell, N. Snavely, and K. Bala. Shading Annotations in the Wild. *Computer Vision and Pattern Recognition (CVPR)*, 2017. 4

[18] P. Laffont and J. Bazin. Intrinsic Decomposition of Image Sequences from Local Temporal Variations. In *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 433–441, Dec. 2015. 1, 2, 3, 4, 6, 8

[19] P.-Y. Laffont. *Intrinsic Image Decomposition from Multiple Photographs*. PhD thesis, Inria / University of Nice Sophia-Antipolis, Oct. 2012. 3, 7

[20] 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 *Proceedings of the 12th European Conference on Computer Vision - Volume Part VI*, ECCV'12, pages 327–340, Berlin, Heidelberg, 2012. Springer-Verlag. 2

[21] L. Lettry, K. Vanhoey, and L. V. Gool. DARN: A Deep Adversarial Residual Network for Intrinsic Image Decomposition. *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1359–1367, 2018. 2

[22] Z. Li and N. Snavely. Learning Intrinsic Image Decomposition From Watching the World. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 2, 3, 4, 5, 6, 7, 8

[23] Y. Matsushita, S. Lin, S. B. Kang, and H.-Y. Shum. Estimating Intrinsic Images from Image Sequences with Biased Illumination. In *Computer Vision - ECCV 2004*, pages 274–286. Springer, Berlin, Heidelberg, May 2004. 2, 3, 6, 7, 8

[24] A. Meka, G. Fox, M. Zollhöfer, C. Richardt, and C. Theobalt. Live User-Guided Intrinsic Video for Static Scenes. *IEEE Transactions on Visualization and Computer Graphics*, 23(11):2447–2454, Nov. 2017. 2, 8

[25] A. Meka, M. Maximov, M. Zollhöfer, A. Chatterjee, H.-P. Seidel, C. Richardt, and C. Theobalt. LIME: Live Intrinsic Material Estimation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 2

[26] A. Meka, M. Zollhöfer, C. Richardt, and C. Theobalt. Live Intrinsic Video. *ACM Trans. Graph.*, 35(4):109:1–109:14, July 2016. 2

[27] T. Narihira, M. Maire, and S. X. Yu. Direct Intrinsics: Learning Albedo-Shading Decomposition by Convolutional Regression. In *Proceedings of the 2015 IEEE International Conference on Computer Vision (ICCV)*, ICCV '15, pages 2992–2992, Washington, DC, USA, 2015. IEEE Computer Society. 2, 4, 5

[28] 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. 2

[29] O. Ronneberger, P. Fischer, and T. Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, pages 234–241. Springer, Cham, Oct. 2015. 2, 3

[30] J. Shi, Y. Dong, H. Su, and X. Y. Stella. Learning non-lambertian object intrinsics across shapenet categories. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5844–5853. IEEE, 2017. 2, 3, 5

[31] T. Taniai and T. Maehara. Neural inverse rendering for general reflectance photometric stereo. In *International Conference on Machine Learning*, pages 4864–4873, 2018. 3

[32] 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. 2, 3, 6, 7, 8

[33] S. Xingjian, Z. Chen, H. Wang, D.-Y. Yeung, W.-K. Wong, and W.-c. Woo. Convolutional LSTM network: A machine learning approach for precipitation nowcasting. In *Advances in Neural Information Processing Systems*, pages 802–810, 2015. 3

[34] G. Ye, E. Garces, Y. Liu, Q. Dai, and D. Gutierrez. Intrinsic Video and Applications. *ACM Trans. Graph.*, 33(4):80:1–80:11, July 2014. 2, 8

[35] Q. Zhao, P. Tan, Q. Dai, L. Shen, E. Wu, and S. Lin. A Closed-Form Solution to Retinex with Nonlocal Texture Constraints. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(7):1437–1444, July 2012. 2

[36] T. Zhou, P. Krähenbühl, and A. A. Efros. Learning Data-Driven Reflectance Priors for Intrinsic Image Decomposition. In *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 3469–3477, Dec. 2015. 2, 4