

CGIntrinsics: Better Intrinsic Image Decomposition through Physically-Based Rendering

Zhengqi Li^[0000–0003–2929–8149] and Noah Snavely^[0000–0002–6921–6833]

Department of Computer Science & Cornell Tech, Cornell University

Abstract. Intrinsic image decomposition is a challenging, long-standing computer vision problem for which ground truth data is very difficult to acquire. We explore the use of synthetic data for training CNN-based intrinsic image decomposition models, then applying these learned models to real-world images. To that end, we present CGINTRINSICS, a new, large-scale dataset of physically-based rendered images of scenes with full ground truth decompositions. The rendering process we use is carefully designed to yield high-quality, realistic images, which we find to be crucial for this problem domain. We also propose a new end-to-end training method that learns better decompositions by leveraging CGINTRINSICS, and optionally IIW and SAW, two recent datasets of sparse annotations on real-world images. Surprisingly, we find that a decomposition network trained solely on our synthetic data outperforms the state-of-the-art on both IIW and SAW, and performance improves even further when IIW and SAW data is added during training. Our work demonstrates the surprising effectiveness of carefully-rendered synthetic data for the intrinsic images task.

1 Introduction

Intrinsic images is a classic vision problem involving decomposing an input image I into a product of reflectance (albedo) and shading images $R \cdot S$. Recent years have seen remarkable progress on this problem, but it remains challenging due to its ill-posedness. An attractive proposition has been to replace traditional hand-crafted priors with learned, CNN-based models. For such learning methods data is key, but collecting ground truth data for intrinsic images is extremely difficult, especially for images of real-world scenes.

One way to generate large amounts of training data for intrinsic images is to render synthetic scenes. However, existing synthetic datasets are limited to images of single objects [1, 2] (e.g., via ShapeNet [3]) or images of CG animation that utilize simplified, unrealistic illumination (e.g., via Sintel [4]). An alternative is to collect ground truth for real images using crowdsourcing, as in the Intrinsic Images in the Wild (IIW) and Shading Annotations in the Wild (SAW) datasets [5, 6]. However, the annotations in such datasets are sparse and difficult to collect accurately at scale.

Inspired by recent efforts to use synthetic images of scenes as training data for indoor and outdoor scene understanding [7–10], we present the first large-scale scene-level intrinsic images dataset based on high-quality physically-based rendering, which we call CGINTRINSICS (CGI). CGI consists of over 20,000 images of indoor scenes, based on the SUNCG dataset [11]. Our aim with CGI is to help drive significant progress towards solving the intrinsic images problem for Internet photos of real-world scenes. We find

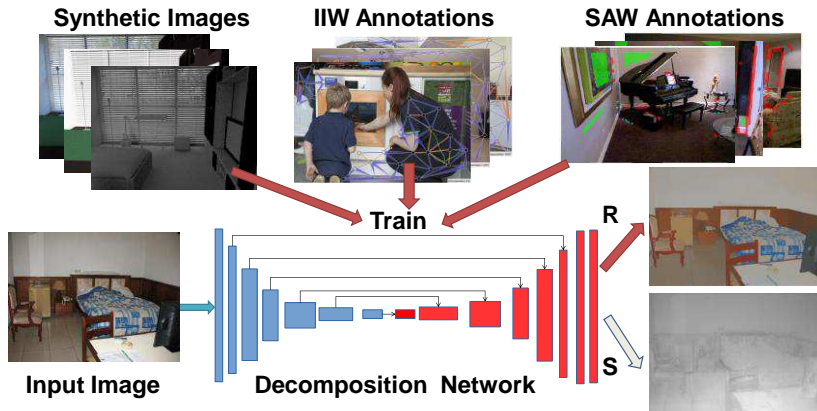


Fig. 1. Overview and network architecture. Our work integrates physically-based rendered images from our CGINTRINSICS dataset and reflectance/shading annotations from IIW and SAW in order to train a better intrinsic decomposition network.

that high-quality physically-based rendering is essential for our task. While SUNCG provides physically-based scene renderings [12], our experiments show that the details of how images are rendered are of critical importance, and certain choices can lead to massive improvements in how well CNNs trained for intrinsic images on synthetic data generalize to real data.

We also propose a new partially supervised learning method for training a CNN to directly predict reflectance and shading, by combining ground truth from CGI and sparse annotations from IIW/SAW. Through evaluations on IIW and SAW, we find that, surprisingly, decomposition networks trained solely on CGI can achieve state-of-the-art performance on both datasets. Combined training using both CGI and IIW/SAW leads to even better performance. Finally, we find that CGI generalizes better than existing datasets by evaluating on MIT Intrinsic Images, a very different, object-centric, dataset.

2 Related work

Optimization-based methods. The classical approach to intrinsic images is to integrate various priors (smoothness, reflectance sparseness, etc.) into an optimization framework [13–17, 5]. However, for images of real-world scenes, such hand-crafted prior assumptions are difficult to craft and are often violated. Several recent methods seek to improve decomposition quality by integrating surface normals or depths from RGB-D cameras [18–20] into the optimization process. However, these methods assume depth maps are available during optimization, preventing them from being used for a wide range of consumer photos.

Learning-based methods. Learning methods for intrinsic images have recently been explored as an alternative to models with hand-crafted priors, or a way to set the parameters of such models automatically. Barron and Malik [21] learn parameters of

a model that utilizes sophisticated priors on reflectance, shape and illumination. This approach works on images of objects (such as in the MIT dataset), but does not generalize to real world scenes. More recently, CNN-based methods have been deployed, including work that regresses directly to the output decomposition based on various training datasets, such as Sintel [22, 23], MIT intrinsics and ShapeNet [2, 1]. Shu *et al.* [24] also propose a CNN-based method specifically for the domain of facial images, where ground truth geometry can be obtained through model fitting. However, as we show in the evaluation section, the networks trained on such prior datasets perform poorly on images of real-world scenes.

Two recent datasets are based on images of real-world scenes. Intrinsic Images in the Wild (IIW) [5] and Shading Annotations in the Wild (SAW) [6] consist of sparse, crowd-sourced reflectance and shading annotations on real indoor images. Subsequently, several papers train CNN-based classifiers on these sparse annotations and use the classifier outputs as priors to guide decomposition [6, 25–27]. However, we find these annotations alone are insufficient to train a direct regression approach, likely because they are sparse and are derived from just a few thousand images. Finally, very recent work has explored the use of time-lapse imagery as training data for intrinsic images [28], although this provides a very indirect source of supervision.

Synthetic datasets for real scenes. Synthetic data has recently been utilized to improve predictions on real-world images across a range of problems. For instance, [7, 10] created a large-scale dataset and benchmark based on video games for the purpose of autonomous driving, and [29, 30] use synthetic imagery to form small benchmarks for intrinsic images. SUNCG [12] is a recent, large-scale synthetic dataset for indoor scene understanding. However, many of the images in the PBRS database of physically-based renderings derived from SUNCG have low signal-to-noise ratio (SNR) and non-realistic sensor properties. We show that higher quality renderings yield much better training data for intrinsic images.

3 CGINTRINSICS Dataset

To create our CGINTRINSICS (CGI) dataset, we started from the SUNCG dataset [11], which contains over 45,000 3D models of indoor scenes. We first considered the PBRS dataset of physically-based renderings of scenes from SUNCG [12]. For each scene, PBRS samples cameras from good viewpoints, and uses the physically-based Mitsuba renderer [31] to generate realistic images under reasonably realistic lighting (including a mix of indoor and outdoor illumination sources), with global illumination. Using such an approach, we can also generate ground truth data for intrinsic images by rendering a standard RGB image I , then asking the renderer to produce a reflectance map R from the same viewpoint, and finally dividing to get the shading image $S = I/R$. Examples of such ground truth decompositions are shown in Figure 2. Note that we automatically mask out light sources (including illumination from windows looking outside) when creating the decomposition, and do not consider those pixels when training the network.

However, we found that the PBRS renderings are not ideal for use in training real-world intrinsic image decomposition networks. In fact, certain details in how images are rendered have a dramatic impact on learning performance:




Fig. 2. Visualization of ground truth from our CGINTRINSICS dataset. Top row: rendered RGB images. Middle: ground truth reflectance. Bottom: ground truth shading. Note that light sources are masked out when creating the ground truth decomposition.

Rendering quality. Mitsuba and other high-quality renderers support a range of rendering algorithms, including various flavors of path tracing methods that sample many light paths for each output pixel. In PBRS, the authors note that bidirectional path tracing works well but is very slow, and opt for Metropolis Light Transport (MLT) with a sample rate of 512 samples per pixel [12]. In contrast, for our purposes we found that bidirectional path tracing (BDPT) with very large numbers of samples per pixel was the only algorithm that gave consistently good results for rendering SUNCG images. Comparisons between selected renderings from PBRS and our new CGI images are shown in Figure 3. Note the significantly decreased noise in our renderings.

This extra quality comes at a cost. We find that using BDPT with 8,192 samples per pixel yields acceptable quality for most images. This increases the render time per image significantly, from a reported 31s [12], to approximately 30 minutes.¹ One reason for the need for large numbers of samples is that SUNCG scenes are often challenging from a rendering perspective—the illumination is often indirect, coming from open doorways or constrained in other ways by geometry. However, rendering is highly parallelizable, and over the course of about six months we rendered over ten thousand images on a cluster of about 10 machines.

Tone mapping from HDR to LDR. We found that another critical factor in image generation is how rendered images are tone mapped. Renderers like Mitsuba generally produce high dynamic range (HDR) outputs that encode raw, linear radiance estimates for each pixel. In contrast, real photos are usually low dynamic range. The process that takes an HDR input and produces an LDR output is called *tone mapping*, and in real

¹ While high, this is still a fair ways off of reported render times for animated films. For instance, each frame of Pixar’s *Monsters University* took a reported 29 hours to render [32].



Fig. 3. Visual comparisons between our CGI and the original SUNCG dataset. Top row: images from SUNCG/PBRs. Bottom row: images from our CGI dataset. The images in our dataset have higher SNR and are more realistic.

Dataset	Size	Setting	Rendered/Real	Illumination	GT type
MPI Sintel [34]	890	Animation	non-PB	spatial-varying	full
MIT Intrinsic [35]	110	Object	Real	single global	full
ShapeNet [2]	2M+	Object	PB	single global	full
IHW [5]	5230	Scene	Real	spatial-varying	sparse
SAW [6]	6677	Scene	Real	spatial-varying	sparse
CGINTRINSICS	20,000+	Scene	PB	spatial-varying	full

Table 1. Comparisons of existing intrinsic image datasets with our CGINTRINSICS dataset. PB indicates physically-based rendering and non-PB indicates non-physically-based rendering.

cameras the analogous operations are the auto-exposure, gamma correction, etc., that yield a well-exposed, high-contrast photograph. PBRs uses the tone mapping method of Reinhard *et al.* [33], which is inspired by photographers such as Ansel Adams, but which can produce images that are very different in character from those of consumer cameras. We find that a simpler tone mapping method produces more natural-looking results. Again, Figure 3 shows comparisons between PBRs renderings and our own. Note how the color and illumination features, such as shadows, are better captured in our renderings (we noticed that shadows often disappear with the Reinhard tone mapper).

In particular, to tone map a linear HDR radiance image I_{HDR} , we find the 90th percentile intensity value r_{90} , then compute the image $I_{\text{LDR}} = \alpha I_{\text{HDR}}^\gamma$, where $\gamma = \frac{1}{2.2}$ is a standard gamma correction factor, and α is computed such that r_{90} maps to the value 0.8. The final image is then clipped to the range $[0, 1]$. This mapping ensures that at most 10% of the image pixels (and usually many fewer) are saturated after tone mapping, and tends to result in natural-looking LDR images.

Using the above rendering approach, we re-rendered $\sim 20,000$ images from PBRs. We also integrated 152 realistic renderings from [30] into our dataset. Table 1 compares our CGI dataset to prior intrinsic image datasets. Sintel is a dataset created for an animated film, and does not utilize physical-based rendering. Other datasets, such as ShapeNet and MIT, are object-centered, whereas CGI focuses on images of indoor

scenes, which have more sophisticated structure and illumination (cast shadows, spatial-varying lighting, etc). Compared to IIW and SAW, which include images of real scenes, CGI has full ground truth and is much more easily collected at scale.

4 Learning Cross-Dataset Intrinsic

In this section, we describe how we use CGINTRINSICS to jointly train an intrinsic decomposition network end-to-end, incorporating additional sparse annotations from IIW and SAW. Our full training loss considers training data from each dataset:

$$\mathcal{L} = \mathcal{L}_{\text{CGI}} + \lambda_{\text{IIW}}\mathcal{L}_{\text{IIW}} + \lambda_{\text{SAW}}\mathcal{L}_{\text{SAW}}. \quad (1)$$

where \mathcal{L}_{CGI} , \mathcal{L}_{IIW} , and \mathcal{L}_{SAW} are the losses we use for training from the CGI, IIW, and SAW datasets respectively. The most direct way to train would be to simply incorporate supervision from each dataset. In the case of CGI, this supervision consists of full ground truth. For IIW and SAW, this supervision takes the form of sparse annotations for each image, as illustrated in Figure 1. However, in addition to supervision, we found that incorporating smoothness priors into the loss also improves performance. Our full loss functions thus incorporate a number of terms:

$$\mathcal{L}_{\text{CGI}} = \mathcal{L}_{\text{sup}} + \lambda_{\text{ord}}\mathcal{L}_{\text{ord}} + \lambda_{\text{rec}}\mathcal{L}_{\text{reconstruct}} \quad (2)$$

$$\mathcal{L}_{\text{IIW}} = \lambda_{\text{ord}}\mathcal{L}_{\text{ord}} + \lambda_{\text{rs}}\mathcal{L}_{\text{rsmooth}} + \lambda_{\text{ss}}\mathcal{L}_{\text{ssmooth}} + \mathcal{L}_{\text{reconstruct}} \quad (3)$$

$$\mathcal{L}_{\text{SAW}} = \lambda_{\text{S/NS}}\mathcal{L}_{\text{S/NS}} + \lambda_{\text{rs}}\mathcal{L}_{\text{rsmooth}} + \lambda_{\text{ss}}\mathcal{L}_{\text{ssmooth}} + \mathcal{L}_{\text{reconstruct}} \quad (4)$$

We now describe each term in detail.

4.1 Supervised losses

CGIntrinsic-supervised loss. Since the images in our CGI dataset are equipped with a full ground truth decomposition, the learning problem for this dataset can be formulated as a direct regression problem from input image I to output images R and S . However, because the decomposition is only up to an unknown scale factor, we use a scale-invariant supervised loss, $\mathcal{L}_{\text{siMSE}}$ (for ‘‘scale-invariant mean-squared-error’’). In addition, we add a gradient domain multi-scale matching term $\mathcal{L}_{\text{grad}}$. For each training image in CGI, our supervised loss is defined as $\mathcal{L}_{\text{sup}} = \mathcal{L}_{\text{siMSE}} + \mathcal{L}_{\text{grad}}$, where

$$\mathcal{L}_{\text{siMSE}} = \frac{1}{N} \sum_{i=1}^N (R_i^* - c_r R_i)^2 + (S_i^* - c_s S_i)^2 \quad (5)$$

$$\mathcal{L}_{\text{grad}} = \sum_{l=1}^L \frac{1}{N_l} \sum_{i=1}^{N_l} \|\nabla R_{l,i}^* - c_r \nabla R_{l,i}\|_1 + \|\nabla S_{l,i}^* - c_s \nabla S_{l,i}\|_1. \quad (6)$$

$R_{l,i}$ ($R_{l,i}^*$) and $S_{l,i}$ ($S_{l,i}^*$) denote reflectance prediction (resp. ground truth) and shading prediction (resp. ground truth) respectively, at pixel i and scale l of an image pyramid. N_l is the number of valid pixels at scale l and $N = N_1$ is the number of valid pixels at the original image scale. The scale factors c_r and c_s are computed via least squares.

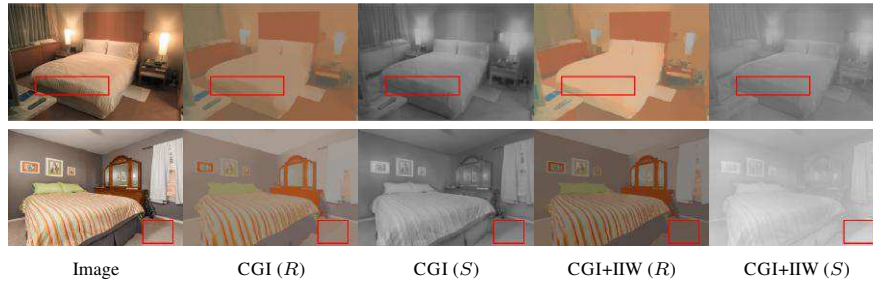


Fig. 4. Examples of predictions with and without IIW training data. Adding real IIW data can qualitatively improve reflectance and shading predictions. Note for instance how the quilt highlighted in first row has a more uniform reflectance after incorporating IIW data, and similarly for the floor highlighted in the second row.

In addition to the scale-invariance of $\mathcal{L}_{\text{simMSE}}$, another important aspect is that we compute the MSE in the linear intensity domain, as opposed to the all-pairs pixel comparisons in the log domain used in [22]. In the log domain, pairs of pixels with large absolute log-difference tend to dominate the loss. As we show in our evaluation, computing $\mathcal{L}_{\text{simMSE}}$ in the linear domain significantly improves performance.

Finally, the multi-scale gradient matching term $\mathcal{L}_{\text{grad}}$ encourages decompositions to be piecewise smooth with sharp discontinuities.

Ordinal reflectance loss. IIW provides sparse *ordinal* reflectance judgments between pairs of points (e.g., “point i has brighter reflectance than point j ”). We introduce a loss based on this ordinal supervision. For a given IIW training image and predicted reflectance R , we accumulate losses for each pair of annotated pixels (i, j) in that image: $\mathcal{L}_{\text{ord}}(R) = \sum_{(i,j)} e_{i,j}(R)$, where

$$e_{i,j}(R) = \begin{cases} w_{i,j}(\log R_i - \log R_j)^2, & r_{i,j} = 0 \\ w_{i,j}(\max(0, m - \log R_i + \log R_j))^2, & r_{i,j} = +1 \\ w_{i,j}(\max(0, m - \log R_j + \log R_i))^2, & r_{i,j} = -1 \end{cases} \quad (7)$$

and $r_{i,j}$ is the ordinal relation from IIW, indicating whether point i is darker (-1), j is darker (+1), or they have equal reflectance (0). $w_{i,j}$ is the confidence of the annotation, provided by IIW. Example predictions with and without IIW data are shown in Fig. 4.

We also found that adding a similar ordinal term derived from CGI data can improve reflectance predictions. For each image in CGI, we over-segment it using superpixel segmentation [36]. Then in each training iteration, we randomly choose one pixel from every segmented region, and for each pair of chosen pixels, we evaluate \mathcal{L}_{ord} similar to Eq. 7, with $w_{i,j} = 1$ and the ordinal relation derived from the ground truth reflectance.

SAW shading loss. The SAW dataset provides images containing annotations of smooth (S) shading regions and non-smooth (NS) shading points, as depicted in Figure 1. These annotations can be further divided into three types: regions of constant shading, shadow boundaries, and depth/normal discontinuities.

We integrate all three types of annotations into our supervised SAW loss $\mathcal{L}_{\text{S/NS}}$. For each constant shading region (with N_c pixels), we compute a loss $\mathcal{L}_{\text{constant-shading}}$

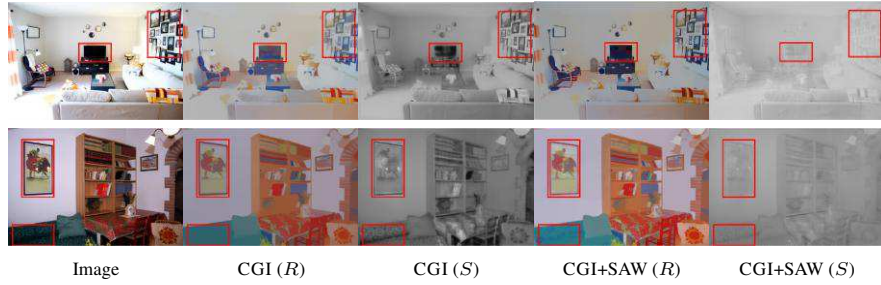


Fig. 5. Examples of predictions with and without SAW training data. Adding SAW training data can qualitatively improve reflectance and shading predictions. Note the pictures/TV highlighted in the decompositions in the first row, and the improved assignment of texture to the reflectance channel for the paintings and sofa in the second row.

encouraging the variance of the predicted shading in the region to be zero:

$$\mathcal{L}_{\text{constant-shading}} = \frac{1}{N_c} \sum_{i=1}^{N_c} (\log S_i)^2 - \frac{1}{N_c^2} \left(\sum_{i=1}^{N_c} \log S_i \right)^2. \quad (8)$$

SAW also provides individual point annotations at cast shadow boundaries. As noted in [6], these points are not localized precisely on shadow boundaries, and so we apply a morphological dilation with a radius of 5 pixels to the set of marked points before using them in training. This results in shadow boundary regions. We find that most shadow boundary annotations lie in regions of constant reflectance, which implies that for all pair of shading pixels within a small neighborhood, their log difference should be approximately equal to the log difference of the image intensity. This is equivalent to encouraging the variance of $\log S_i - \log I_i$ within this small region to be 0 [37]. Hence, we define the loss for each shadow boundary region (with N_{sd} pixels as:

$$\mathcal{L}_{\text{shadow}} = \frac{1}{N_{sd}} \sum_{i=1}^{N_{sd}} (\log S_i - \log I_i)^2 - \frac{1}{N_{sd}^2} \left(\sum_{i=1}^{N_{sd}} (\log S_i - \log I_i) \right)^2 \quad (9)$$

Finally, SAW provides depth/normal discontinuities, which are also usually shading discontinuities. However, since we cannot derive the actual shading change for such discontinuities, we simply mask out such regions in our shading smoothness term $\mathcal{L}_{\text{smooth}}$ (Eq. 11), i.e., we do not penalize shading changes in such regions. As above, we first dilate these annotated regions before use in training. Examples predictions before/after adding SAW data into our training are shown in Fig. 5.

4.2 Smoothness losses

To further constrain the decompositions for real images in IIW/SAW, following classical intrinsic image algorithms we add reflectance smoothness $\mathcal{L}_{\text{rsmooth}}$ and shading

smoothness $\mathcal{L}_{\text{smooth}}$ terms. For reflectance, we use a multi-scale ℓ_1 smoothness term to encourage reflectance predictions to be piecewise constant:

$$\mathcal{L}_{\text{rsmooth}} = \sum_{l=1}^L \frac{1}{N_l l} \sum_{i=1}^{N_l} \sum_{j \in \mathcal{N}(l,i)} v_{l,i,j} \|\log R_{l,i} - \log R_{l,j}\|_1 \quad (10)$$

where $\mathcal{N}(l, i)$ denotes the 8-connected neighborhood of the pixel at position i and scale l . The reflectance weight $v_{l,i,j} = \exp\left(-\frac{1}{2}(\mathbf{f}_{l,i} - \mathbf{f}_{l,j})^T \Sigma^{-1}(\mathbf{f}_{l,i} - \mathbf{f}_{l,j})\right)$, and the feature vector $\mathbf{f}_{l,i}$ is defined as $[\mathbf{p}_{l,i}, I_{l,i}, c_{l,i}^1, c_{l,i}^2]$, where $\mathbf{p}_{l,i}$ and $I_{l,i}$ are the spatial position and image intensity respectively, and $c_{l,i}^1$ and $c_{l,i}^2$ are the first two elements of chromaticity. Σ is a covariance matrix defining the distance between two feature vectors.

We also include a densely-connected ℓ_2 shading smoothness term, which can be evaluated in linear time in the number of pixels N using bilateral embeddings [38, 28]:

$$\mathcal{L}_{\text{ssmooth}} = \frac{1}{2N} \sum_i^N \sum_j^N \hat{W}_{i,j} (\log S_i - \log S_j)^2 \approx \frac{1}{N} \mathbf{s}^\top (I - N_b S_b^\top \bar{B}_b S_b N_b) \mathbf{s} \quad (11)$$

where \hat{W} is a bistochastic weight matrix derived from W and $W_{i,j} = \exp\left(-\frac{1}{2} \left\| \frac{\mathbf{p}_i - \mathbf{p}_j}{\sigma_p} \right\|_2^2\right)$. We refer readers to [38, 28] for a detailed derivation. As shown in our experiments, adding such smoothness terms to real data can yield better generalization.

4.3 Reconstruction loss

Finally, for each training image in each dataset, we add a loss expressing the constraint that the reflectance and shading should reconstruct the original image:

$$\mathcal{L}_{\text{reconstruct}} = \frac{1}{N} \sum_{i=1}^N (I_i - R_i S_i)^2. \quad (12)$$

4.4 Network architecture

Our network architecture is illustrated in Figure 1. We use a variant of the ‘‘U-Net’’ architecture [28, 39]. Our network has one encoder and two decoders with skip connections. The two decoders output log reflectance and log shading, respectively. Each layer of the encoder mainly consists of a 4×4 stride-2 convolutional layer followed by batch normalization [40] and leaky ReLU [41]. For the two decoders, each layer is composed of a 4×4 deconvolutional layer followed by batch normalization and ReLU, and a 1×1 convolutional layer is appended to the final layer of each decoder.

5 Evaluation

We conduct experiments on two datasets of real world scenes, IIW [5] and SAW [6] (using test data unseen during training) and compare our method with several state-of-the-art intrinsic images algorithms. Additionally, we also evaluate the generalization of our CGI dataset by evaluating it on the MIT Intrinsic Images benchmark [35].

Method	Training set	WHDR	Method	Training set	WHDR
Retinex-Color [35]	-	26.9%	Ours (log, $\mathcal{L}_{\text{siMSE}}$)	CGI	22.7%
Garces <i>et al.</i> [17]	-	24.8%	Ours (w/o $\mathcal{L}_{\text{grad}}$)	CGI	19.7%
Zhao <i>et al.</i> [14]	-	23.8%	Ours (w/o \mathcal{L}_{ord})	CGI	19.9%
Bell <i>et al.</i> [5]	-	20.6%	Ours (w/o $\mathcal{L}_{\text{rsmooth}}$)	All	16.1%
Zhou <i>et al.</i> [25]	IIW	19.9%	Ours	SUNCG	26.1%
Bi <i>et al.</i> [44]	-	17.7%	Ours	CGI	17.8%
Nestmeyer <i>et al.</i> [45]	IIW	19.5%	Ours*	CGI	17.1%
Nestmeyer <i>et al.</i> [45]*	IIW	17.7%	Ours	CGI+IIW(O)	17.5%
DI [22]	Sintel	37.3%	Ours	CGI+IIW(A)	16.2%
Shi <i>et al.</i> [2]	ShapeNet	59.4%	Ours	All	15.5%
			Ours*	All	14.8%

Table 2. Numerical results on the IIW test set. Lower is better for WHDR. The table is split into two subtables for space (prior methods are shown in the left subtable, and our results are shown on the right). The “Training set” column specifies the training data used by each learning-based method: “-” indicates an optimization-based method. IIW(O) indicates original IIW annotations and IIW(A) indicates augmented IIW comparisons. “All” indicates CGI+IIW(A)+SAW. * indicates that CNN predictions are post-processed with a guided filter [45].

Network training details. We implement our method in PyTorch [42]. For all three datasets, we perform data augmentation through random flips, resizing, and crops. For all evaluations, we train our network from scratch using the Adam [43] optimizer, with initial learning rate 0.0005 and mini-batch size 16. We refer readers to the supplementary material for the detailed hyperparameter settings.

5.1 Evaluation on IIW

We follow the train/test split for IIW provided by [27], also used in [25]. We also conduct several ablation studies using different loss configurations. Quantitative comparisons of Weighted Human Disagreement Rate (WHDR) between our method and other optimization- and learning-based methods are shown in Table 2.

Comparing direct CNN predictions, our CGI-trained model is significantly better than the best learning-based method [45], and similar to [44], even though [45] was directly trained on IIW. Additionally, running the post-processing from [45] on the results of the CGI-trained model achieves a further performance boost. Table 2 also shows that models trained on SUNCG (i.e., PBRs), Sintel, MIT Intrinsic, or ShapeNet generalize poorly to IIW likely due to the lower quality of training data (SUNCG/PBRs), or the larger domain gap with respect to images of real-world scenes, compared to CGI. The comparison to SUNCG suggests the key importance of our rendering decisions.

We also evaluate networks trained jointly using CGI and real imagery from IIW. As in [25], we augment the pairwise IIW judgments by globally exploiting their transitivity and symmetry. The right part of Table 2 demonstrates that including IIW training data leads to further improvements in performance, as does also including SAW training data. Table 2 also shows various ablations on variants of our method, such as evaluating

Method	Training set	AP% (unweighted)	AP% (challenge)
Retinex-Color [35]	-	91.93	85.26
Garces <i>et al.</i> [17]	-	96.89	92.39
Zhao <i>et al.</i> [14]	-	97.11	89.72
Bell <i>et al.</i> [5]	-	97.37	92.18
Zhou <i>et al.</i> [25]	IIW	96.24	86.34
Nestmeyer <i>et al.</i> [45]	IIW	97.26	89.94
Nestmeyer <i>et al.</i> [45]*	IIW	96.85	88.64
DI [22]	Sintel+MIT	95.04	86.08
Shi <i>et al.</i> [2]	ShapeNet	86.62	81.30
Ours (log, $\mathcal{L}_{\text{silMSE}}$)	CGI	97.73	93.03
Ours (w/o $\mathcal{L}_{\text{grad}}$)	CGI	98.15	93.74
Ours (w/o $\mathcal{L}_{\text{ssmooth}}$)	CGI+IIW(A)+SAW	98.60	94.87
Ours	SUNCG	96.56	87.09
Ours	CGI	98.43	94.08
Ours	CGI+IIW(A)	98.56	94.69
Ours	CGI+IIW(A)+SAW	98.78	96.57

Table 3. Quantitative results on the SAW test set. Higher is better for AP%. The second column is described in Table 2. The third and fourth columns show performance on the unweighted SAW benchmark and our more challenging gradient-weighted benchmark, respectively.

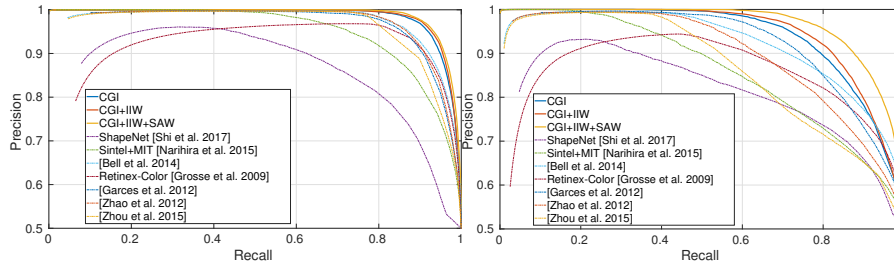


Fig. 6. Precision-Recall (PR) curve for shading images on the SAW test set. Left: PR curves generated using the unweighted SAW error metric of [28]. Right: curves generated using our more challenging gradient-weighted metric.

losses in the log domain and removing terms from the loss functions. Finally, we test a network trained on *only* IIW/SAW data (and not CGI), or trained on CGI and fine-tuned on IIW/SAW. Although such a network achieves $\sim 19\%$ WHDR, we find that the decompositions are qualitatively unsatisfactory. The sparsity of the training data causes these networks to produce degenerate decompositions, especially for shading images.

5.2 Evaluation on SAW

To evaluate our shading predictions, we test our models on the SAW [6] test set, utilizing the error metric introduced in [28]. We also propose a new, more challenging

error metric for SAW evaluation. In particular, we found that many of the constant-shading regions annotated in SAW also have smooth image intensity (e.g., textureless walls), making their shading easy to predict. Our proposed metric downweights such regions as follows. For each annotated region of constant shading, we compute the average image gradient magnitude over the region. During evaluation, when we add the pixels belonging to a region of constant shading into the confusion matrices, we multiply the number of pixels by this average gradient. This proposed metric leads to more distinguishable performance differences between methods, because regions with rich textures will contribute more to the error compared to the unweighted metric.

Figure 6 and Table 3 show precision-recall (PR) curves and average precision (AP) on the SAW test set with both unweighted [28] and our proposed challenge error metrics. As with IIW, networks trained solely on our CGI data can achieve state-of-the-art performance, even *without* using SAW training data. Adding real IIW data improves the AP in term of both error metrics. Finally, the last column of Table 3 shows that integrating SAW training data can significantly improve the performance on shading predictions, suggesting the effectiveness of our proposed losses for SAW sparse annotations.

Note that the previous state-of-the-art algorithms on IIW (e.g., Zhou *et al.* [25] and Nestmeyer *et al.* [45]) tend to overfit to reflectance, hurting the accuracy of shading predictions. This is especially evident in terms of our proposed challenge error metric. In contrast, our method achieves state-of-the-art results on *both* reflectance and shading predictions, in terms of all error metrics. Note that models trained on the original SUNCG, Sintel, MIT intrinsics or ShapeNet datasets perform poorly on the SAW test set, indicating the much improved generalization to real scenes of our CGI dataset.

Qualitative results on IIW/SAW. Figure 7 shows qualitative comparisons between our network trained on all three datasets, and two other state-of-the-art intrinsic images algorithms (Bell *et al.* [5] and Zhou *et al.* [25]), on images from the IIW/SAW test sets. In general, our decompositions show significant improvements. In particular, our network is better at avoiding attributing surface texture to the shading channel (for instance, the checkerboard patterns evident in the first two rows, and the complex textures in the last four rows) while still predicting accurate reflectance (such as the mini-sofa in the images of third row). In contrast, the other two methods often fail to handle such difficult settings. In particular, [25] tends to overfit to reflectance predictions, and their shading estimates strongly resemble the original image intensity. However, our method still makes mistakes, such as the non-uniform reflectance prediction for the chair in the fifth row, as well as residual textures and shadows in the shading and reflectance channels.

5.3 Evaluation on MIT intrinsic images

For the sake of completeness, we also test the ability of our CGI-trained networks to generalize to the MIT Intrinsic Images dataset [35]. In contrast to IIW/SAW, the MIT dataset contains 20 real objects with 11 different illumination conditions. We follow the same train/test split as Barron *et al.* [21], and, as in the work of Shi *et al.* [2], we directly apply our CGI trained networks to MIT testset, and additionally test fine-tuning them on the MIT training set.

We compare our models with several state-of-the-art learning-based methods using the same error metrics as [2]. Table 4 shows quantitative comparisons and Figure 8 shows

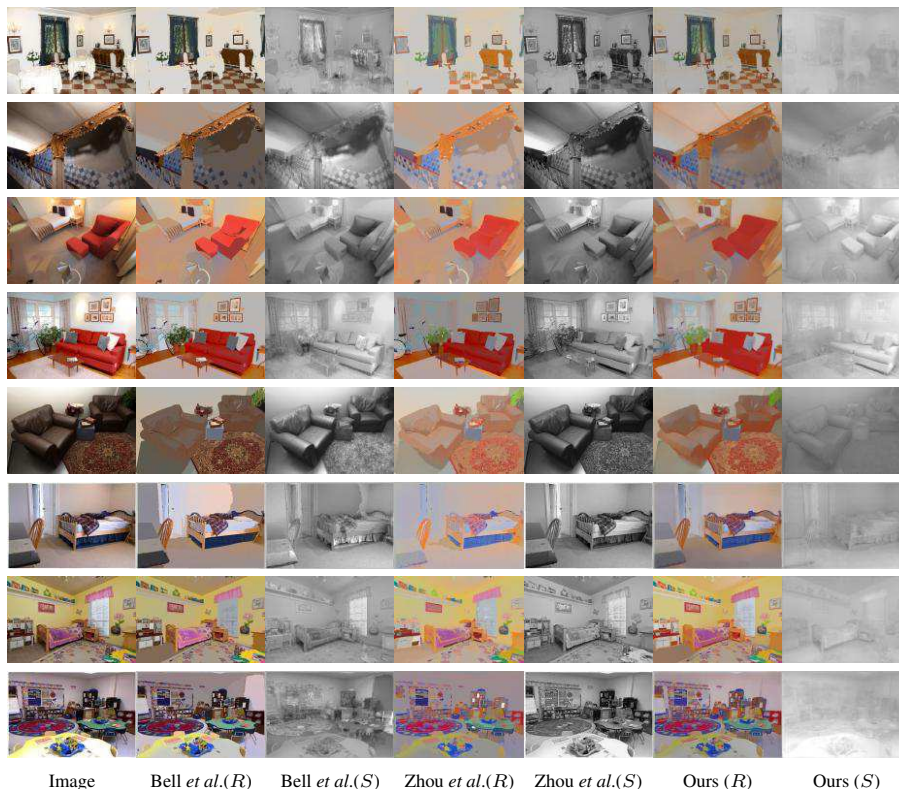


Fig. 7. Qualitative comparisons on the IIW/SAW test sets. Our predictions show significant improvements compared to state-of-the-art algorithms (Bell *et al.* [5] and Zhou *et al.* [25]). In particular, our predicted shading channels include significantly less surface texture in several challenging settings.

qualitative results. Both show that our CGI-trained model yields better performance compared to ShapeNet-trained networks both qualitatively and quantitatively, even though like MIT, ShapeNet consists of images of rendered objects, while our dataset contains images of scenes. Moreover, our CGI-pretrained model also performs better than networks pretrained on ShapeNet and Sintel. These results further demonstrate the improved generalization ability of our CGI dataset compared to existing datasets. Note that SIRFS still achieves the best results, but as described in [22, 2], their methods are designed specifically for single objects and generalize poorly to real scenes.

6 Conclusion

We presented a new synthetic dataset for learning intrinsic images, and an end-to-end learning approach that learns better intrinsic image decompositions by leveraging datasets with different types of labels. Our evaluations illustrate the surprising effectiveness of

Method	Training set	MSE		LMSE		DSSIM	
		refl.	shading	refl.	shading	refl.	shading
SIRFS [21]	MIT	0.0147	0.0083	0.0416	0.0168	0.1238	0.0985
DI [22]	Sintel+MIT	0.0277	0.0154	0.0585	0.0295	0.1526	0.1328
Shi <i>et al.</i> [2]	ShapeNet	0.0468	0.0194	0.0752	0.0318	0.1825	0.1667
Shi <i>et al.</i> [2]*	ShapeNet+MIT	0.0278	0.0126	0.0503	0.0240	0.1465	0.1200
Ours	CGI	0.0221	0.0186	0.0349	0.0259	0.1739	0.1652
Ours*	CGI +MIT	0.0167	0.0127	0.0319	0.0211	0.1287	0.1376

Table 4. Quantitative Results on MIT intrinsics testset. For all error metrics, lower is better. The second column shows the dataset used for training. * indicates models fine-tuned on MIT.




Fig. 8. Qualitative comparisons on MIT intrinsics testset. Odd rows: reflectance predictions. Even rows: shading predictions. * are the predictions fine-tuned on MIT.

our synthetic dataset on Internet photos of real-world scenes. We find that the details of rendering matter, and hypothesize that improved physically-based rendering may benefit other vision tasks, such as normal prediction and semantic segmentation [12].

Acknowledgments. We thank Jingguang Zhou for his help with data generation. This work was funded by the National Science Foundation through grant IIS-1149393, and by a grant from Schmidt Sciences.

References

1. Janner, M., Wu, J., Kulkarni, T., Yildirim, I., Tenenbaum, J.B.: Self-Supervised Intrinsic Image Decomposition. In: *Neural Information Processing Systems*. (2017)
2. Shi, J., Dong, Y., Su, H., Yu, S.X.: Learning non-Lambertian object intrinsics across ShapeNet categories. In: *Proc. Computer Vision and Pattern Recognition (CVPR)*. (2017) 5844–5853
3. Chang, A.X., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., et al.: ShapeNet: An information-rich 3D model repository. *arXiv preprint arXiv:1512.03012* (2015)
4. Butler, D.J., Wulff, J., Stanley, G.B., Black, M.J.: A naturalistic open source movie for optical flow evaluation. In: *Proc. European Conf. on Computer Vision (ECCV)*. (2012) 611–625
5. Bell, S., Bala, K., Snavely, N.: Intrinsic images in the wild. *ACM Trans. Graphics* **33**(4) (2014) 159
6. Kovacs, B., Bell, S., Snavely, N., Bala, K.: Shading annotations in the wild. In: *Proc. Computer Vision and Pattern Recognition (CVPR)*. (2017) 850–859
7. Richter, S.R., Vineet, V., Roth, S., Koltun, V.: Playing for data: Ground truth from computer games. In: *Proc. European Conf. on Computer Vision (ECCV)*. (2016) 102–118
8. Ros, G., Sellart, L., Materzynska, J., Vazquez, D., Lopez, A.M.: The SYNTHIA Dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In: *Proc. Computer Vision and Pattern Recognition (CVPR)*. (2016) 3234–3243
9. Gaidon, A., Wang, Q., Cabon, Y., Vig, E.: Virtual worlds as proxy for multi-object tracking analysis. In: *Proc. Computer Vision and Pattern Recognition (CVPR)*. (2016) 4340–4349
10. Richter, S.R., Hayder, Z., Koltun, V.: Playing for benchmarks. In: *Proc. Int. Conf. on Computer Vision (ICCV)*. (2017) 2232–2241
11. Song, S., Yu, F., Zeng, A., Chang, A.X., Savva, M., Funkhouser, T.: Semantic scene completion from a single depth image. In: *Proc. Computer Vision and Pattern Recognition (CVPR)*. (2017) 190–198
12. Zhang, Y., Song, S., Yumer, E., Savva, M., Lee, J.Y., Jin, H., Funkhouser, T.: Physically-based rendering for indoor scene understanding using convolutional neural networks. In: *Proc. Computer Vision and Pattern Recognition (CVPR)*. (2017) 5057–5065
13. Land, E.H., McCann, J.J.: Lightness and retinex theory. *Josa* **61**(1) (1971) 1–11
14. Zhao, Q., Tan, P., Dai, Q., Shen, L., Wu, E., Lin, S.: A closed-form solution to retinex with nonlocal texture constraints. *Trans. on Pattern Analysis and Machine Intelligence* **34**(7) (2012) 1437–1444
15. Rother, C., Kiefel, M., Zhang, L., Schölkopf, B., Gehler, P.V.: Recovering intrinsic images with a global sparsity prior on reflectance. In: *Neural Information Processing Systems*. (2011) 765–773
16. Shen, L., Yeo, C.: Intrinsic images decomposition using a local and global sparse representation of reflectance. In: *Proc. Computer Vision and Pattern Recognition (CVPR)*. (2011) 697–704
17. Garces, E., Munoz, A., Lopez-Moreno, J., Gutierrez, D.: Intrinsic images by clustering. *Computer Graphics Forum (Proc. EGSR 2012)* **31**(4) (2012)
18. Chen, Q., Koltun, V.: A simple model for intrinsic image decomposition with depth cues. In: *Proc. Computer Vision and Pattern Recognition (CVPR)*. (2013) 241–248
19. Barron, J.T., Malik, J.: Intrinsic scene properties from a single RGB-D image. In: *Proc. Computer Vision and Pattern Recognition (CVPR)*. (2013) 17–24
20. Jeon, J., Cho, S., Tong, X., Lee, S.: Intrinsic image decomposition using structure-texture separation and surface normals. In: *Proc. European Conf. on Computer Vision (ECCV)*. (2014)

21. Barron, J.T., Malik, J.: Shape, illumination, and reflectance from shading. *Trans. on Pattern Analysis and Machine Intelligence* **37**(8) (2015) 1670–1687
22. Narihira, T., Maire, M., Yu, S.X.: Direct intrinsics: Learning albedo-shading decomposition by convolutional regression. In: *Proc. Int. Conf. on Computer Vision (ICCV)*. (2015) 2992–2992
23. Kim, S., Park, K., Sohn, K., Lin, S.: Unified depth prediction and intrinsic image decomposition from a single image via joint convolutional neural fields. In: *Proc. European Conf. on Computer Vision (ECCV)*. (2016) 143–159
24. Shu, Z., Yumer, E., Hadap, S., Sunkavalli, K., Shechtman, E., Samaras, D.: Neural face editing with intrinsic image disentangling. In: *Proc. Computer Vision and Pattern Recognition (CVPR)*. (2017) 5444–5453
25. Zhou, T., Krahenbuhl, P., Efros, A.A.: Learning data-driven reflectance priors for intrinsic image decomposition. In: *Proc. Int. Conf. on Computer Vision (ICCV)*. (2015) 3469–3477
26. Zoran, D., Isola, P., Krishnan, D., Freeman, W.T.: Learning ordinal relationships for mid-level vision. In: *Proc. Int. Conf. on Computer Vision (ICCV)*. (2015) 388–396
27. Narihira, T., Maire, M., Yu, S.X.: Learning lightness from human judgement on relative reflectance. In: *Proc. Computer Vision and Pattern Recognition (CVPR)*. (2015) 2965–2973
28. Li, Z., Snavely, N.: Learning intrinsic image decomposition from watching the world. In: *Proc. Computer Vision and Pattern Recognition (CVPR)*. (2018)
29. Beigpour, S., Serra, M., van de Weijer, J., Benavente, R., Vanrell, M., Penacchio, O., Samaras, D.: Intrinsic image evaluation on synthetic complex scenes. *Int. Conf. on Image Processing* (2013)
30. Bonneel, N., Kovacs, B., Paris, S., Bala, K.: Intrinsic decompositions for image editing. *Computer Graphics Forum (Eurographics State of the Art Reports 2017)* **36**(2) (2017)
31. Jakob, W.: Mitsuba renderer (2010) <http://www.mitsuba-renderer.org>.
32. Takahashi, D.: How Pixar made Monsters University, its latest technological marvel. <https://venturebeat.com/2013/04/24/the-making-of-pixars-latest-technological-marvel-monsters-university/> (2013)
33. Reinhard, E., Stark, M., Shirley, P., Ferwerda, J.: Photographic tone reproduction for digital images. In: *ACM Trans. Graphics (SIGGRAPH)*. (2002)
34. Butler, D.J., Wulff, J., Stanley, G.B., Black, M.J.: A naturalistic open source movie for optical flow evaluation. In: *Proc. European Conf. on Computer Vision (ECCV)*. (2012)
35. Grosse, R., Johnson, M.K., Adelson, E.H., Freeman, W.T.: Ground truth dataset and baseline evaluations for intrinsic image algorithms. In: *Proc. Int. Conf. on Computer Vision (ICCV)*. (2009)
36. Achanta, R., Shaji, A., Smith, K., Lucchi, A., Fua, P., Süsstrunk, S.: SLIC superpixels compared to state-of-the-art superpixel methods. *Trans. on Pattern Analysis and Machine Intelligence* **34**(11) (2012)
37. Eigen, D., Puhrsch, C., Fergus, R.: Depth map prediction from a single image using a multi-scale deep network. In: *Neural Information Processing Systems*. (2014) 2366–2374
38. Barron, J.T., Adams, A., Shih, Y., Hernández, C.: Fast bilateral-space stereo for synthetic defocus. In: *Proc. Computer Vision and Pattern Recognition (CVPR)*. (2015) 4466–4474
39. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: *Proc. Computer Vision and Pattern Recognition (CVPR)*. (2017) 6967–5976
40. Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: *Proc. Int. Conf. on Machine Learning*. (2015) 448–456
41. He, K., Zhang, X., Ren, S., Sun, J.: Delving deep into rectifiers: Surpassing human-level performance on ImageNet classification. In: *Proc. Int. Conf. on Computer Vision (ICCV)*. (2015)
42. : Pytorch (2016) <http://pytorch.org>.

43. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. CoRR **abs/1412.6980** (2014)
44. Bi, S., Han, X., Yu, Y.: An l_1 image transform for edge-preserving smoothing and scene-level intrinsic decomposition. ACM Trans. Graph. **34** (2015) 78:1–78:12
45. Nestmeyer, T., Gehler, P.V.: Reflectance adaptive filtering improves intrinsic image estimation. In: Proc. Computer Vision and Pattern Recognition (CVPR). (2017)