

# SILT: Self-supervised Lighting Transfer Using Implicit Image Decomposition

Nikolina Kubiak<sup>1</sup>

n.kubiak@surrey.ac.uk

Armin Mustafa<sup>1</sup>

armin.mustafa@surrey.ac.uk

Graeme Phillipson<sup>2</sup>

graeme.phillipson@bbc.co.uk

Stephen Jolly<sup>2</sup>

stephen.jolly@bbc.co.uk

Simon Hadfield<sup>1</sup>

s.hadfield@surrey.ac.uk

<sup>1</sup> University of Surrey

Centre for Vision, Speech and Signal  
Processing (CVSSP)  
Guildford, Surrey, UK

<sup>2</sup> BBC Research & Development

Salford, Greater Manchester, UK

---

## Abstract

We present *SILT*, a Self-supervised Implicit Lighting Transfer method. Unlike previous research on scene relighting, we do not seek to apply arbitrary new lighting configurations to a given scene. Instead, we wish to transfer the lighting style from a database of other scenes, to provide a uniform lighting style regardless of the input. The solution operates as a two-branch network that first aims to map input images of any arbitrary lighting style to a unified domain, with extra guidance achieved through implicit image decomposition. We then remap this unified input domain using a discriminator that is presented with the generated outputs and the style reference, *i.e.* images of the desired illumination conditions. Our method is shown to outperform supervised relighting solutions across two different datasets without requiring lighting supervision. The code and pre-trained models can be found [here](#).

## 1 Introduction

We propose the problem of lighting transfer where an input image under arbitrary illumination conditions is adapted to match the lighting of a style reference database. Previous approaches to a similar problem - arbitrary image relighting - either make simplifying assumptions about the scene (*e.g.* a single dominant object or a single light source) or require costly supervision where identical scenes must be captured under a large number of known lighting conditions. In contrast, our lighting transfer approach, *SILT*, is entirely self-supervised. *SILT* requires only a training dataset with multiple examples of the same scene and a style reference database. It is not necessary to know the ground truth lighting conditions for the training data, nor is it necessary for the target illumination to be present within the training dataset. This distinction between image relighting and lighting transfer is demonstrated in Fig. 1.

The proposed *SILT* method consists of a two-branch network. During training, the generator sees a number of input images of the same scene under different unknown illuminations. The model attempts to enforce similarity between the illumination conditions of the outputs

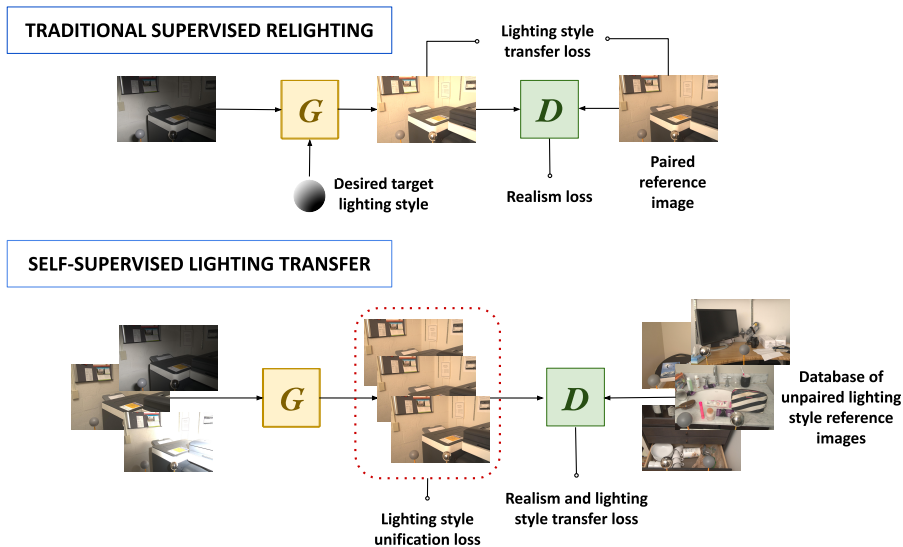


Figure 1: System overview: differences between relighting and lighting transfer.

while preserving the contents of the scene. The generated data is then presented to a discriminator that tries to distinguish the outputs from the examples of target illumination (*i.e.* the reference style images), thus enforcing the desired lighting conditions. At inference, SILT takes in a single input image and can adapt to previously unseen locations.

We motivate our work by the fact that in complex, real-life scenarios there is rarely one *correct* way to describe the desired lighting conditions. Instead of trying to quantify a particular style (*e.g.* a TED talk), we can show the model examples of such an event. We do not need the re-styled materials to look like a specific, *best* talk - we just want the model to learn what such talks look like and copy the style. The same idea can be applied to other domains such as casual live music performances, filters for social media or domain adaptation for multi-view reconstruction. In summary, the contributions of our paper are threefold:

- We present the first formulation of the lighting transfer problem, with associated baselines and evaluation protocol.
- We demonstrate a flexible self-supervised approach to lighting transfer, performing on par with supervised relighting models.
- We show that jointly training a self-supervised image decomposition step allows us to better capture the complexities of real scenes, compared to a pre-trained model.

## 2 Literature review

As lighting transfer is a new field of inquiry, we will instead focus this discussion on the most relevant works from the field of image relighting.

**General-purpose relighting.** Last year the VIDIT dataset [6] was created and used for relighting challenges [7, 8]. It features indoor and outdoor scenes lit from 8 directions and captured under 5 temperature settings. The applications of VIDIT include [22] where Wang *et al.* tackle the task of relighting in stages. They simultaneously teach two network branches

to remove the effects of source illumination and to apply shadows according to target illumination, and then combine them to render the final result. In [9] Das *et al.* stack two U-Nets to achieve good quality reconstructions while keeping the computational cost very low. They expand on this idea in [5] and use 2 images captured from opposite directions as input, which allows for better reconstruction of dark regions. Gafton and Maraz [6] expand the pix2pixHD [23] architecture by adding light direction estimation layers, and train 8 network branches that can relight any image to 1 of 8 illumination directions covered by VIDIT. Puthussery *et al.* [19] perform relighting using wavelet decomposition and propose a new loss responsible for accurate shadow recasting. Their *gray loss* is calculated on blurred input and ground truth images, stripped from texture information to focus on illumination. As shown, in just a year VIDIT has found a number of applications. The dataset, however, is synthetic and unnaturally dark, making it difficult to apply models trained on it to real-life data.

Very few papers operate on real-life data captured in the wild. Zhou *et al.* [28] perform image decomposition and then relight scenes by swapping the shading maps of two images. Nevertheless, the approach is limited to relighting within a fixed location, and the lighting style captured by the low-level shading map cannot be applied to another scene. Murmann *et al.* [17] created a multi-illumination indoor dataset, published alongside an encoder-decoder relighting model.

**Reducing supervision requirements.** All of the previously described solutions operate in a supervised manner. Yet, such an approach makes training difficult as pairs of images of the same scene, captured under source and target illumination conditions, are rarely readily available outside of simulation. Therefore, a number of authors have recently begun to explore domain-specific mechanisms to reduce the supervision requirements of relighting.

Liu *et al.* [14] propose a Siamese autoencoder network for portrait relighting that decomposes the images into illumination and content embeddings. The process is supervised by augmenting the available data in a way that preserves its content while changing the light direction, thus creating training pairs. The relighting task is constrained via a simplified spherical harmonic lighting model. This, however, limits the relighting to light direction and intensity changes, and neglects the impact of colour on the white-balance of the outputs.

Liu *et al.* [13] and Yu *et al.* [26] proposed self-supervised image-based outdoor relighting models that first decompose the scene into constant and changing factors, and then synthesise new images by recombining the desired illumination conditions with the illumination-independent elements. Yu *et al.* show that during decomposition, phenomena such as shadows or specularities tend to be wrongly preserved in reflectance or normal maps, skewing the final renderings. To address this, they expand on their previous work [25] and add a channel that aims to capture this type of residual information. Inspired by [12], Liu *et al.* [13] learn from stacks of timelapse panoramas of the same location. They capture time-varying information using high-level features, allowing the system to mix and match illumination settings between locations. However, the method requires the target lighting conditions to be present in a panorama stack, which limits the model’s flexibility.

## 3 Methodology

### 3.1 Image relighting posed as style transfer

Our model is an LSGAN [16] based on the pix2pix [23] U-Net [20] architecture. The multi-scale discriminator operates on patches and distinguishes between real and fake samples

at different spatial scales. Our generator’s architecture consists of 3 downsampling steps, 9 residual blocks and 3 upsampling steps, and does not require the LocalEnhancer add-on proposed in the pix2pix paper. To achieve self-supervision, the model structure is duplicated to create two branches; we refer to them as branches A and B.

The training of our generator is guided by a sum of three types of losses. Firstly, we use a generator GAN loss  $\mathcal{L}_g$  which measures the model’s ability to fool the discriminator. It has the opposite goal and acts as an adversary to the discriminator loss  $\mathcal{L}_d$  which ensures that the real and generated samples are classified as such. In our case, this distinction covers both image realism and the illumination style transfer accuracy. The losses can be expressed as

$$\mathcal{L}_g = \|D(G(I)) - 1\|_2 \quad \text{and} \quad \mathcal{L}_d = \|D(\tilde{I}) - 1\|_2 + \|D(G(I))\|_2, \quad (1)$$

where  $\tilde{I}$  and  $I$  correspond to reference and input samples, and  $D$  and  $G$  mark the discriminator and generator models.

Now, to enforce the characteristics of the target illumination, we need to first encourage the generator to translate all images, regardless of their initial lighting conditions, to a unified domain. To this end, we enforce similarity between the outputs  $\hat{I}$  of both branches using output similarity loss  $\mathcal{L}_{os}$  defined as

$$\mathcal{L}_{os} = \|\hat{I}_A - \hat{I}_B\|_1 + \|\nabla \hat{I}_A - \nabla \hat{I}_B\|_1 = \|G(I_A) - G(I_B)\|_1 + \|\nabla G(I_A) - \nabla G(I_B)\|_1. \quad (2)$$

In the equation,  $\nabla \hat{I}$  symbolises spatial gradients calculated over the generated images. We choose these two types of losses to ensure the correct matching of colour and brightness as well as edge information.

Finally, to prevent mode collapse, we need to make the output have the same visual content as the original image. This is controlled using the content preservation loss  $\mathcal{L}_{cp}$ , calculated between the input and output for each branch. This error can be described as

$$\mathcal{L}_{cp} = \mathcal{VGG}(I, \hat{I}) = \mathcal{VGG}(I, G(I)), \quad (3)$$

where  $\mathcal{VGG}$  represents the perceptual loss proposed in [9]. This is a weighted L1 loss calculated between the features extracted from input and output images using specific layers of a pre-trained VGG-19 network.

To summarise, the described model performs lighting style transfer, controlling the content and style of the output image. The former is preserved using a perceptual loss  $\mathcal{L}_{cp}$ . For the latter, we propose a new approach and split the task between the generator and discriminator. The former uses  $\mathcal{L}_{os}$  to ensure a unified style, and the latter controls the characteristics of said style by comparing generated data with style reference examples.

## 3.2 Decomposition-guided relighting

In addition to the simple style transfer approach, we propose to enhance the generator with a jointly learnt self-supervised image decomposition mechanism. The decomposition process, inspired by the Retinex theory [10], assumes that an image  $I$  can be expressed as a Hadamard product of its reflectance  $R$  and shading  $S$ , *i.e.*  $I = R \odot S$ . These correspond to the illumination-invariant vs the lighting-dependent factors, respectively.

There are many general purpose solutions that use inverse rendering techniques to achieve this kind of decomposition [11, 12, 13, 14, 15, 16, 17]. Of particular interest to us is the state-of-the-art decomposition model by Liu *et al.* [18], combining indoor scenes and unsupervised

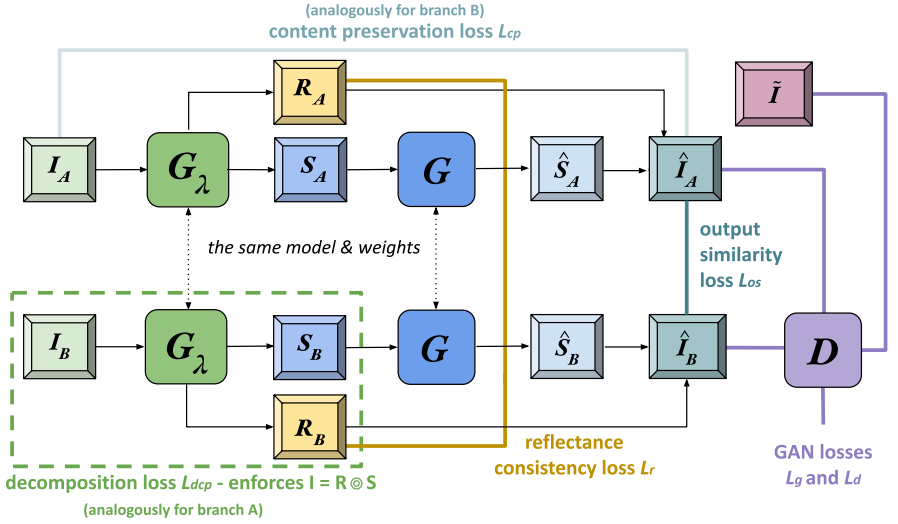


Figure 2: All parts and losses of SILT, the implicit decomposition model.

learning. The authors propose a scheme in which the style of reflectance and shading images is learnt based on unpaired samples. At the same time, their generator  $G_\lambda$  learns to decompose images into underlying factors that match their corresponding learnt distributions, *i.e.*  $\{R, S\} = G_\lambda(I)$ . In our use case, we do not require a set of reference reflectance and shading maps. We simply apply self-supervision asking the network to extract two features from the image that, when multiplied, produce an accurate reconstruction of the original input. As such they should serve roles similar to reflectance and shading yet this is never defined explicitly. Nevertheless, in the interest of clarity, we will use the reflectance-shading naming convention throughout the paper.

During training, our network must learn to decompose images and to optimise this procedure to support the downstream lighting transfer task. We guide the process by minimising the residual between the input image and the product of its generated reflectance and shading maps. This is done using the decomposition loss  $\mathcal{L}_{dcp}$ , expressed as

$$\mathcal{L}_{dcp} = \|I - R \odot S\|_1. \quad (4)$$

Since we only want to perform the style transfer on the shading component  $S$ , we feed its corresponding map into the generator  $G$  instead of using the original image. The reflectance maps should remain unchanged so these are passed via a skip connection and multiplied with the newly generated shading image  $\hat{S}$  to get the output image, as shown in Fig. 2. The output image formulation can now be described as  $\hat{I} = R \odot \hat{S} = R \odot G(S)$ .

Additionally, we want to enforce consistency between reflectance maps generated for both branches as they should be identical regardless of the original illumination conditions. We control this using the reflectance loss  $\mathcal{L}_r$  expressed as

$$\mathcal{L}_r = \|R_A - R_B\|_2, \quad (5)$$

where  $R_A$  and  $R_B$  are the reflectances obtained during decomposition for each input pair.

## 4 Experiments

### 4.1 Ablation study

During these experiments all models were trained on the Multi-Illumination dataset [17]. Its data is split into train and test sets, with 985 and 30 scenes respectively. Each scene is photographed under 25 fixed lighting directions located in the upper hemisphere relative to the camera. For this paper we picked nine illumination conditions, consistent across all scenes of the dataset, and chose one unpaired image setting to act as target lighting domain. During training, we paired the nine images up and each unique pair was presented to the generator once per epoch. At test time, the models were fed single images of previously unseen scenes, captured under the same nine illumination conditions; no additional reference data was needed. For our experiments the images were resized to  $768 \times 512$  pixels.

Our method was trained for approx. 180k iterations and used the same default training hyperparameters as the pix2pix GAN; this holds true for all discussed model versions. We report additional information regarding the model’s complexity in Appendix C. In the following sections we measure the performance of the tested models using SSIM and PSNR metrics (higher is better) and report the perceptual (VGG) loss scores (lower is better). For results reported in tables, the top score is shown in bold and the second best is underlined.

**Decomposition.** First, we explore the benefits of the jointly-learnt self-supervised decomposition approach. We compare our basic model with no decomposition against our network using the fixed pre-trained decomposition approach of Liu *et al.* [15] and our full model with a jointly-learnt decomposition step based on Liu *et al.*’s architecture. For the pre-trained model, we used the generated reflectances as-is but also tried averaging them across the scene to support reflectance consistency. Below we show results for the latter, more successful approach. For the  $R$  and  $S$  maps generated during the jointly-learnt decomposition step, please see Appendix A.

The outcomes of the study are measured quantitatively and qualitatively, with results in Table 1 and Fig. 3. The basic version and the full implicit decomposition versions of our self-supervised model achieve similar performance across all metrics. The generated images, however, show that the latter network, SILT, provides a closer colour match and makes a better attempt at specular reconstruction (bottom row). Even though the decomposition step of Liu *et al.* was pretrained on indoor scenes, it cannot produce results on par with our implicit decomposition model. Its generated images are too dark and contain more prominent artefacts.

Decomposition	SSIM $\uparrow$	PSNR $\uparrow$	VGG $\downarrow$
None	<u>0.800</u>	<u>17.995</u>	<u>0.323</u>
Liu <i>et al.</i> [15]	0.689	15.874	0.423
Learnt	<b>0.811</b>	<b>18.667</b>	<b>0.320</b>

Table 1: Decomposition architecture study

**Losses.** Next, we review the losses used by our best self-supervised solution, *i.e.* the implicit decomposition model. We start with just the GAN losses -  $\mathcal{L}_g$  and  $\mathcal{L}_d$  - and the decomposition loss  $\mathcal{L}_{dep}$  as these are responsible for the core functionality of our network. Then, we gradually add the other losses. The changes linked to each component are presented in Table 2.



Figure 3: Visual results across different decomposition settings.

The results show that the most meaningful improvements in the SSIM and VGG metrics stem from the content preservation loss  $\mathcal{L}_{cp}$ . Interestingly, using just the output similarity loss  $\mathcal{L}_{os}$  leads to a performance decrease. However, when this loss is linked with  $\mathcal{L}_{cp}$  (line #4), the performance across all metrics improves visibly. This synergy suggests that simply trying to equate two outputs without constraining their content creates an ill-posed task and more guidance is needed to generate meaningful renders. Finally, we can observe a slight increase in SSIM and larger improvements in PSNR when  $\mathcal{L}_r$  is added. This proves that enforcing reflectance consistency between the branches leads to better  $R$  and  $S$  decomposition and, consequently, more accurate lighting style transfer achieved by the generator  $G$ . In the future experiments we will use the model using all losses, corresponding to line #5 in the table.

#	core	$\mathcal{L}_{os}$	$\mathcal{L}_{cp}$	$\mathcal{L}_r$	SSIM $\uparrow$	PSNR $\uparrow$	VGG $\downarrow$
1	✓				0.730	17.075	0.428
2	✓	✓			0.682	17.142	0.495
3	✓		✓		0.798	17.575	0.319
4	✓	✓	✓		<u>0.805</u>	<u>17.712</u>	<b>0.317</b>
5	✓	✓	✓	✓	<b>0.811</b>	<b>18.667</b>	0.320

Table 2: Loss ablation study

## 4.2 State-of-the-art comparisons on the *Multi-Illumination* dataset

Next, we compare the performance of the best version of our self-supervised SILT model with the state-of-the-art supervised method by Murmann *et al.* [17] and the pix2pix baseline [23]. Murmann *et al.*'s solution originally focused on relighting from a ‘base’ image to any of the desired lighting directions; we refer to this as *one2any* translation. To adapt their approach to our target application, we retrained it in *any2one* fashion, *i.e.* so that it can convert any lighting condition seen during training to the reference illumination. The same approach is followed for pix2pix. We train the baseline with the default model settings and, since we do not train with HD data, only the global generator is used. It is worth noting that both of the comparisons are trained in a supervised manner. They still use the same 9 input light directions as our model, but these are paired with a target image of the same scene that our generator never sees.



Figure 4: Visual results on the Multi-Illumination dataset.

The results of the experiments are shown in Table 3 and Fig. 4. SILT comes first across all three performance metrics. Our relit images have the best colour balance and compare favourably against Murmann *et al.*'s colder tones and pix2pix's warm tint (top row). When it comes to specularities (bottom row), the model of Murmann *et al.* appears to smudge them. The pix2pix baseline does not make any changes. Our solution, on the other hand, shifts the input reflection to the left, as in the reference image. Since specularities are particularly hard to re-render, we believe this is a fair reconstruction attempt.

Model	SSIM $\uparrow$	PSNR $\uparrow$	VGG $\downarrow$
Murmann <i>et al.</i> [17]	0.763	13.951	0.381
pix2pix	<u>0.799</u>	<u>18.534</u>	<u>0.333</u>
SILT (ours)	<b>0.811</b>	<b>18.667</b>	<b>0.320</b>

Table 3: Results on the Multi-Illumination dataset

### 4.3 State-of-the-art comparisons on the VIDIT dataset

To make comparisons against VIDIT-based models, we retrain SILT on the synthetic VIDIT dataset, and compare our work against recent supervised methods. For this purpose, we choose two supervised solutions trained in *any2one* fashion that share quantitative results in their corresponding papers and run evaluations on the publicly available sets of VIDIT. Additionally, we train pix2pix on VIDIT and use it as a baseline.

The chosen methods differ slightly in their approach to the task. Gafton and Maraz [9] split the VIDIT training set 90:10 into train:test sets, and translate the images within one temperature setting (4500K). They separately measure the ability of their model to relight to all of the 8 directions. For our comparisons in this paper we pick one of them - East (E). Wang *et al.* [22] train their model to convert all input directions and temperatures of the training set to 4500K-E, and then during inference convert from 6500K-N to 4500K-E (as in the AIM2020 challenge [1]). In the interest of fairness, we train two versions of our self-supervised model and the baseline to follow the same rules as mentioned above.

The results of the experiments are presented in Table 4 and Fig. 5. The supervised solution of Gafton and Maraz outperforms our model and the baseline in terms of PSNR. This is likely due to their supervised light direction classifier which is trained on VIDIT to pro-



vide the GAN with information regarding the input light direction. Apart from this addition, their generator follows the original pix2pix design. SILT also focuses on the target illumination but does not attempt to discern that of the source. This suggests that gaining more understanding of not only the reference but also the input illumination may lead to meaningful performance improvements. When it comes to other metrics, undisclosed by Gafton & Maraz, we beat the baseline in terms of SSIM and the perceptual (VGG) score.

The second set of experiments, trained on all temperatures and directions, tackles a more complex task as the networks now have to simultaneously deal with light colour and direction changes. Nevertheless, our solution outperforms the supervised model of Wang *et al.* in terms of SSIM and achieves a similar PSNR score. We also beat the baseline across all metrics. Visually, all methods struggle to reconstruct the dark region. Wang *et al.* and the pix2pix baseline attempt to lighten the area yet their reconstructions are very blurry and, while lighter, not similar to the ground truth. Our self-supervised SILT method manages to translate the colour setting but does not sufficiently lighten up the dark region.

Model	SSIM $\uparrow$	PSNR $\uparrow$	VGG $\downarrow$
Gafton and Maraz [9]	n/a	<b>21.45</b>	n/a
pix2pix	<u>0.619</u>	17.80	<u>0.350</u>
SILT (ours)	<b>0.654</b>	<u>17.88</u>	<b>0.341</b>
Wang <i>et al.</i> [17]	<u>0.596</u>	<b>17.59</b>	n/a
pix2pix	0.447	15.01	<u>0.419</u>
SILT (ours)	<b>0.606</b>	<u>17.00</u>	<b>0.371</b>

Table 4: Results on the VIDIT dataset. Some metrics were not provided by the authors and are marked as not available (n/a).

#### 4.4 Cross-dataset results

To verify the generality of our method, we test the performance of SILT on data from outside of its training dataset. For these experiments we choose the model trained on the Multi-Illumination dataset as we believe it represents real-life scenarios better than the dark VIDIT

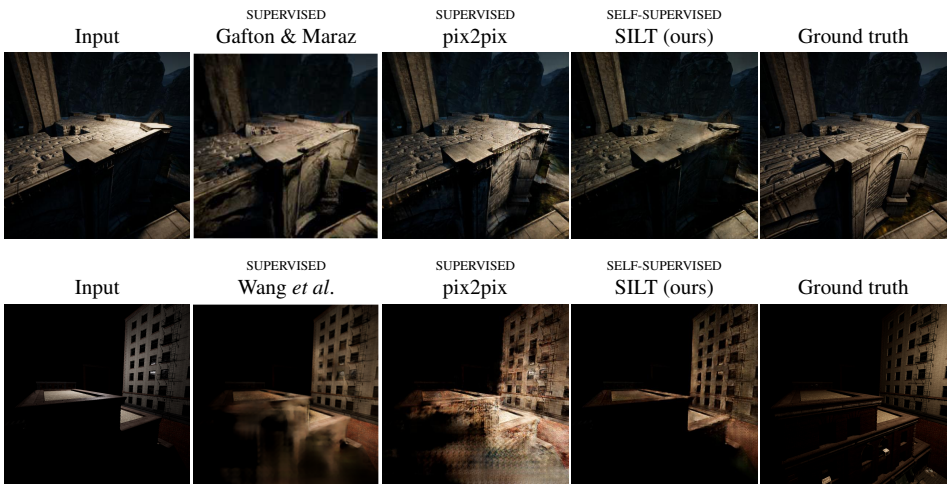


Figure 5: Visual results on the VIDIT dataset. Top: just 4500K, bottom: all temperatures.




Figure 6: Visual results across different data using SILT trained on the Multi-Illumination dataset. The top row shows input images and the bottom row - their corresponding outputs.

data. We show a few output images re-styled using SILT in Fig. 6. For examples showing results obtained by running SILT on higher definition data, please see Appendix B.

The images used for demonstrations in Fig. 6 come from VIDIT and Places [27] datasets. Since the style reference for the Multi-Illumination dataset is quite light and warm-toned, we can see such changes in the re-styled outputs. In VIDIT image a) more texture details are recovered, in comparison with the VIDIT-trained model (see: Fig. 5, top row). The model is similarly successful in re-styling the dark parking image from Places a). In the remaining images we can primarily see changes in the white balance.

## 5 Conclusions

In this paper we proposed SILT, a self-supervised solution for general-purpose lighting transfer, a previously uncharted domain. We also demonstrated that using self-supervised image decomposition can lead to better physical accuracy in the re-styled images. To reap these gains, one can approach this process implicitly, and let the network learn what features to extract to optimise the downstream lighting style transfer task.

Unsurprisingly, our solution has some limitations. When we train the model on VIDIT and do not show it ground truth images (due to our self-supervised approach), SILT struggles to fill in the dark regions and instead only adjusts the colour/white balance (see Fig. 5). While the supervised methods are better at inserting textured content in such areas, their solutions are rarely fully accurate, and low-light region enhancement remains an open problem.

Our future goals are to improve the lighting transfer accuracy and to expand the discussed solution to include people and/or more objects in our scenes. While advances in domain-specific relighting are being made, no existing approaches can cover real-life scenarios where full-bodied humans exist in complex environments. Therefore, we wish to fill this niche and move a step closer towards comprehensive lighting enhancement.

**Acknowledgements.** This work was partially supported by the British Broadcasting Corporation (BBC) and the Engineering and Physical Sciences Research Council’s (EPSRC) industrial CASE project “Computational lighting in video” (voucher number 19000034).

## References

- [1] Jonathan T. Barron and Jitendra Malik. Shape, illumination, and reflectance from shading. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2015.
- [2] Anil S. Baslamisli, Yang Liu, Sezer Karaoglu, and Theo Gevers. Physics-based Shading Reconstruction for Intrinsic Image Decomposition. *arXiv preprint 2009.01540*, 2020.
- [3] Qifeng Chen and Vladlen Koltun. Photographic image synthesis with cascaded refinement networks. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [4] Sourya Dipta Das, Nisarg A. Shah, and Saikat Dutta. MSR-Net: Multi-Scale Relighting Network for One-to-One Relighting. *Workshop on Differentiable Vision, Graphics and Physics in Machine Learning at NeurIPS*, 2020.
- [5] Sourya Dipta Das, Nisarg A. Shah, Saikat Dutta, and Himanshu Kumar. DSRN: an Efficient Deep Network for Image Relighting. *arXiv preprint 2102.09242*, 2021.
- [6] Majed El Helou, Ruofan Zhou, Johan Barthas, and Sabine Süsstrunk. VIDIT: Virtual image dataset for illumination transfer. *arXiv preprint arXiv:2005.05460*, 2020.
- [7] Majed El Helou, Ruofan Zhou, Sabine Süsstrunk, Radu Timofte, et al. AIM 2020: Scene relighting and illumination estimation challenge. In *Proceedings of the European Conference on Computer Vision Workshops (ECCVW)*, 2020.
- [8] Majed El Helou, Ruofan Zhou, Sabine Süsstrunk, and Radu Timofte. NTIRE 2021 depth guided image relighting challenge. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2021.
- [9] Paul Gafton and Erick Maraz. 2d image relighting with image-to-image translation. *arXiv preprint 2006.07816*, 2020.
- [10] Edwin Land. The Retinex Theory of Color Vision. *Scientific American*, 237(6), 1977.
- [11] Zhengqi Li and Noah Snavely. CGIntrinsics: Better Intrinsic Image Decomposition Through Physically-Based Rendering. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 2018.
- [12] Zhengqi Li and Noah Snavely. Learning intrinsic image decomposition from watching the world. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [13] Andrew Liu, Shiry Ginosar, Tinghui Zhou, Alexei A. Efros, and Noah Snavely. Learning to factorize and relight a city. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020.
- [14] Yang Liu, Alexandros Neophytou, Sunando Sengupta, and Eric Sommerlade. Relighting images in the wild with a self-supervised siamese auto-encoder. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2021.

- [15] Yunfei Liu, Yu Li, Shaodi You, and Feng Lu. Unsupervised Learning for Intrinsic Image Decomposition From a Single Image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [16] Xudong Mao, Qing Li, Haoran Xie, Raymond Y.K. Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [17] Lukas Murmann, Michael Gharbi, Miika Aittala, and Fredo Durand. A multi-illumination dataset of indoor object appearance. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2019.
- [18] Thomas Nestmeyer and Peter V. Gehler. Reflectance Adaptive Filtering Improves Intrinsic Image Estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [19] Densen Puthussery, Hrishikesh P. S., Melvin Kuriakose, and Jiji C. V. WDRN : A Wavelet Decomposed RelightNet for Image Relighting. In *Proceedings of the European Conference on Computer Vision Workshops (ECCVW)*, 2020.
- [20] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, 2015.
- [21] Soumyadip Sengupta, Jinwei Gu, Kihwan Kim, Guilin Liu, David Jacobs, and Jan Kautz. Neural inverse rendering of an indoor scene from a single image. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2019.
- [22] Li-Wen Wang, Wan-Chi Siu, Zhi-Song Liu, Chu-Tak Li, and Daniel P. K. Lun. Deep Relighting Networks for Image Light Source Manipulation. In *Proceedings of the European Conference on Computer Vision Workshops (ECCVW)*, 2020.
- [23] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [24] Chen Wei, Wenjing Wang, Wenhan Yang, and Jiaying Liu. Deep Retinex Decomposition for Low-Light Enhancement. In *British Machine Vision Conference (BMVC)*, 2018.
- [25] Ye Yu and William A. P. Smith. Inverserendernet: Learning single image inverse rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [26] Ye Yu, Abhimitra Meka, Mohamed Elgharib, Hans-Peter Seidel, Christian Theobalt, and William A. P. Smith. Self-supervised Outdoor Scene Relighting. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020.
- [27] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10 million image database for scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017.

- 
- [28] Tinghui Zhou, Philipp Krähenbühl, and Alexei A. Efros. Learning data-driven reflectance priors for intrinsic image decomposition. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2015.