



Deep Multi-class Adversarial Specularity Removal

John Lin^{1,2}(✉), Mohamed El Amine Seddik¹, Mohamed Tamaazousti¹,
Youssef Tamaazousti¹, and Adrien Bartoli²

¹ CEA, LIST, Point Courrier 184, 91191 Gif-sur-Yvette, France
{john.lin,mohamedelamine.seddik,mohamed.tamaazousti}@cea.fr,
youssef.tamaazousti@gmail.com

² Institut Pascal - UMR 6602 - CNRS/UCA/CHU, Clermont-Ferrand, France
adrien.bartoli@gmail.com

Abstract. We propose a novel learning approach, in the form of a fully-convolutional neural network (CNN), which automatically and consistently removes specular highlights from a single image by generating its diffuse component. To train the generative network, we define an adversarial loss on a discriminative network as in the GAN framework and combined it with a content loss. In contrast to existing GAN approaches, we implemented the discriminator to be a multi-class classifier instead of a binary one, to find more constraining features. This helps the network pinpoint the diffuse manifold by providing two more gradient terms. We also rendered a synthetic dataset designed to help the network generalize well. We show that our model performs well across various synthetic and real images and outperforms the state-of-the-art in consistency.

Keywords: Deep learning · GAN · Dichromatic reflection separation · Specular and diffuse components

1 Introduction

The appearance of an object depends on the way it reflects light. The materials constituting most objects have a dichromatic behaviour: they produce two types of reflection, namely diffuse and specular reflections. Shafer's dichromatic model [21] linearly combines these two terms for the image formation model. The fundamental difference between them is that the diffuse reflection does not change with the viewing direction, while the specular one does. The specular reflection is thus largely responsible for the tremendous difficulty of solving for the parameters of an image formation model. It is then appealing to assume the object's surface to be perfectly diffuse and rule out the specular reflection. This has been extensively used in various computer vision problems, including SLAM, image segmentation and object detection, to name but a few. The price to pay however is failure of the methods when the real object's reflection departs, sometimes even slightly, from diffusion.

Quite naturally, many approaches have been proposed to solve the problem of specular and diffuse separation, and to be applied as a preprocessing to many algorithms. The separation is also relevant in computer graphics since the specular component conveys precious information about the surface material and illumination [6, 14]. We can group model-based methods in two categories: multi-image and single-image approaches. Multi-image methods ingeniously use specular reflections’ physical properties to find and remove specularities from an image, such as their polarization properties [18, 28] or their dependance to the viewpoint to find matching specular and diffuse pixels from several images [9, 15, 16]. While obtaining good results, these methods are impractical to use because of the need for multiple images, special equipment (polarizer) or known object geometry. Single-image methods mostly rely on the Dichromatic Reflection Model and the fact that specularities retain the illumination’s color to do the separation [2, 11, 21, 22, 26]. However, the separation problem with a single image being ill-posed because of the ambiguity of the image formation process [1], they make strong assumptions about the scene such as a single illumination of known color, no saturated pixels and no nonlinearity of the capture device. This obviously hinders the generic applicability of the methods. Therefore, this is still a challenging and open problem.

In this paper, we propose a deep learning approach to overcome the limitations in applicability. The idea is that the network will work out the intricate relationships between an image and its diffuse part. Recently, a handful of learning-based methods have been proposed to solve the diffuse and specular separation problem [4, 17, 24]. Such data-driven approaches reduce the need to find hand-crafted features and priors, which might not even be relevant for the wide diversity of possible scenes [27]. An immediate challenge is to find a large scale real dataset, since it is extremely time-consuming to produce one. Therefore, we train our network on synthetic data. We specifically rendered the data to overcome some limitations known to the problem of separation, by including known causes of failure cases in hand-crafted methods. Another challenge of learning approaches is to generalize, all the more difficult when training with synthetic data. To overcome this limitation, we build our work on the fairly recent framework of Generative Adversarial Networks (GAN) [5], which we adapt to the separation problem. Just as in GANs, we have a generator network, which we call Specularity Removal Network, trained to generate the diffuse image, while the discriminator network is used only for training by determining whether specularities are well removed. The main difference with a classical GAN resides in the discriminator network, which is not a binary classifier but a categorical classifier. By increasing the number of classes, we help the discriminator pinpoint the desired manifold. This allows it to find more discriminative features for the task at hand. It also prevents an unwanted behaviour of the GAN on synthetic data *i.e.* to generate data that look synthetic. Our method takes a single RGB image as input and does not make any assumption about the scene. We show in the results Sect. 3 that our framework is more stable than existing methods [23, 25, 26] for a wide range of images, outperforming them qualitatively and quantitatively.

In summary, our work addresses the aforementioned challenges and makes the following contributions:

- A new method of Single-image Specular-Diffuse Separation (SSDS), free from priors on the scene and capable of performing on a wide range of images.
- A new multi-class adversarial loss for the problem of SSDS.
- A new synthetic dataset, designed for the task of specular highlights removal.

2 Deep Specularity Removal

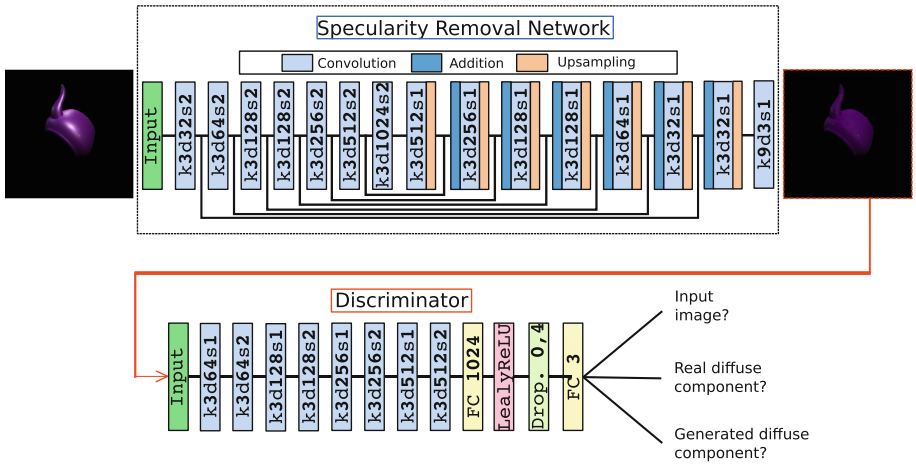


Fig. 1. Overview of our architecture. The Specularity Removal Network takes a single image as input and outputs its diffuse component. The discriminator network is used only for training. It takes the entire image as input and is trained to classify an image into three categories: the input image (I), the diffuse component (D) and the generated diffuse component (\hat{D}); while the generator is trained to fool the discriminator to classify \hat{D} as D .

Following the dichromatic model [21], when light hits an object, it is divided in two parts at the surface, at a ratio depending on the material’s refraction index. One part, called the specular reflection, is reflected at the surface, in the manner of a mirror reflection. The other part, called the diffuse reflection, penetrates the object and scatters before coming out and being reflected. These two bounced off parts of light then add up, and after integrating all lights coming from the upper hemisphere of the object surface, the image is formed. Integration being a linear operation, we can describe an image I as:

$$I = S + D, \quad (1)$$

where S is the specular image and D the diffuse image. The problem of Single image Specular and Diffuse Separation (SSDS) consists in estimating the specular S and/or the diffuse D component of a given image I , since a consistent estimation of one component suffices to retrieve the other by subtracting from I . Thus, we consider the problem of predicting the diffuse component D for which we have more visual cues since there are often purely diffuse pixels in images.

2.1 Overview

To solve the ill-posed problem of SSDS, we propose a generator \mathcal{G}_{θ_g} , parametrized by its weights θ_g . \mathcal{G}_{θ_g} is a feed-forward CNN, which takes I as input and generates D . To train \mathcal{G}_{θ_g} , we carefully rendered a realistic and diverse synthetic training set $\mathbb{T} = \{I_i, D_i\}_{i=1}^N$ of N images and their corresponding ground truth diffuse components. The dataset generation is discussed in Sect. 2.2. Formally, our method boils down to the following optimization problem:

$$\hat{\theta}_g = \arg \min_{\theta_g} \frac{1}{N} \sum_{i=1}^N \ell(\mathcal{G}_{\theta_g}(I_i), D_i), \quad (2)$$

where ℓ is our SSDS-specific loss function. ℓ is a weighted combination of two loss components. One of these components is a content-loss to drive the learning and the other is an adversarial loss to increase the accuracy of predicted diffuse components. The discriminative network is specifically designed for the task of SSDS, as shown in the results. It is trained to recognize whether or not specularities were well removed from an image, while the specular removal network is trained to fool them. The discriminator network and the loss are discussed in Sects. 2.4 and 2.5 respectively.

2.2 Synthetic Training Dataset

Our training set $\mathbb{T} = \{I_i, D_i\}_{i=1}^N$ consists of $N = 20000$ realistically rendered pairs of images. The generation is automated with a script using Blender and the Cycles engine. Each of the N training frames shows a single object from a set of 8 synthetic 3D models, rendered at the center of the image in a random orientation. We excluded 5 shapes of the training set to also create a test set of 1000 images used in the quantitative evaluation. The diffuse part is rendered with the Lambertian model [12] and the specular part with the Beckmann distribution of a microfacet model (Glossy BSDF in Blender). The specular roughness is randomly chosen in the range $[0.2, 0.5]$. We set a lower bound to this roughness so that the material does not become highly specular, almost mirror-like. To handle mirror-like surfaces, we would need a different dataset and maybe a different network all-together. Among these N frames, we have four sets of data, each one rendered to overcome a learning limitation.

Random Texture Objects. This first set has 10000 images rendered with four area lights, directed towards the object. The position of the lights as well as their intensity are randomized, to increase the diversity of the scenes. With the same intent, we set a random colored texture to each object. This simulates real life objects which do not always show uniform reflectance and also compels the network to render images of high quality by focusing on details. The lights' color is fixed to near white (two are slightly blue and the others yellow to imitate real lights).

White Objects. With only the first set, the generative network simply learns to remove white pixels from the image as it suffices for the training set. Therefore we add this second set with 4000 images of white objects, otherwise rendered in the same manner as the first set, to help the network grasp the difference between white objects and specularities.

Colored Lights. This third set contains 2000 images rendered as the first set but with a random color assigned to each light. This allows us to take into account the real life cases where the lights are not white, although it is not very common (which is why we render only 2000 images for this set). Having different light colors in the scene also ensures that our network does not overfit our dataset by only analyzing white pixels and their intensity.

Environment Maps. Finally, this last set was added to the corpus because in real life specular highlights sometimes spread over a large portion of the object instead of being small and localized. This is due to the inter-reflections that occur in real scenes, whereas our synthesized data only contain one object. To simulate this effect, we render this set with an environment map E_i , randomly sampled from a set of 6 High Dynamic Range (HDR) maps.

2.3 Specularity Removal Network

The specularity removal network is a CNN with skip connections, inspired by U-Net [19]. It takes in a single RGB image, decreases its resolution with strided convolutions before further processing and upsampling it to generate its diffuse counterpart. Figure 1 (top) depicts the architecture of our generator network, where k , d and s respectively mean the kernel size, the depth and stride of the convolutions. We use the ReLU activation function at each convolutional layer. The network was trained with images of size 256×256 , but can then be applied to images of any size as it is fully convolutional.

2.4 Discriminator Network

In classical GANs, the discriminator network is a binary classifier, trained to recognize real images from fake ones. Then, the generator is trained to fool it,

which results in high perceptual quality image generation, close to the data’s distribution. In our case, this framework does not fit for two reasons: (1) we do not want our generator to fit our synthetic data’s distribution (*e.g.* a single object in the center of the image), which still lacks diversity compared to real data in spite of our efforts; (2) visual quality is not the sole concern of our problem, the main one being to accurately remove specularities. Therefore we propose a new framework following the GAN paradigm of Goodfellow et al. [5], but differing in the discriminator network’s role.

Multi-class Adversarial Optimization. The discriminator network’s objective is to classify generated diffuse images from real ones but also from input images I , which show specularities. Instead of returning a single scalar value, it outputs a tensor of three values, standing for the probabilities of the image belonging to either one of the three classes and which add up to 1. For that we replaced the usual sigmoid activation function at the last layer by a softmax activation. We call \mathcal{D}_{θ_d} our discriminator network and $\hat{D} = \mathcal{G}_{\theta_g}(I)$ a diffuse image generated by \mathcal{G}_{θ_g} . The discriminator network is then optimized in an alternating manner with \mathcal{G}_{θ_g} to solve the adversarial min-max problem:

$$\min_{\theta_g} \max_{\theta_d} \sum_{i=1}^3 \mathbb{E}_{x_i \sim p_{x_i}} \left\{ \log \mathcal{D}_{\theta_d}^{(i)}(x_i) \right\} + \sum_{i=1, j=1, i \neq j}^3 \mathbb{E}_{x_j \sim p_{x_j}} \left\{ \log \left[1 - \mathcal{D}_{\theta_d}^{(i)}(x_j) \right] \right\}, \quad (3)$$

where:

- $\mathcal{D}_{\theta_d}^{(i)}$ denotes the i^{th} output of \mathcal{D}_{θ_d} and represents the probability of the image belonging to the class C_i .
- x_i is an image drawn from the distribution p_{x_i} which corresponds to C_i .
- $C_i \in \{C_1, C_2, C_3\} = \{C_I, C_D, C_{\hat{D}}\}$ is one of the three classes.

Note that Eq. (3) depends on θ_g when $j = 3$ and $x_j = \hat{D} = \mathcal{G}_{\theta_g}(I)$. The idea behind this multi-class discriminator is that recognizing D from \hat{D} will ensure visual quality as in a classical GAN, while the classification between I and \hat{D} will compel the discriminator to find features related to the sole difference between them *i.e.* specular highlights, thus ensuring accurate specular removal.

Architecture. An overview of the discriminator’s architecture can be seen in Fig. 1 (bottom). It takes as input the image of size 256×256 and decreases its resolution every other layer with convolutions of stride 2, going from 256×256 to 16×16 , while the number of kernel filters increases every other layer, going from 64 to 512. Every convolutional layer is followed by a LeakyReLU activation and a batch normalization except at the last layer. We then apply a Flatten layer before a Dense layer to form the 3-dimensional vector.

2.5 Loss Function

To stabilize the learning and at the same time train an efficient and accurate network, we define our loss to be a combination of a content loss and our SSDS-specific adversarial loss:

$$\ell(\mathcal{G}_{\theta_g}(I), D) = \ell_{content} + \lambda \times \ell_{SSDS}, \quad (4)$$

where λ is a regularization parameter to scale the two losses.

We use the Mean Square Error (MSE) as our content loss:

$$\ell_{content}(\mathcal{G}_{\theta_g}(I), D) = \frac{1}{WH} \sum_{x=1}^W \sum_{y=1}^H (D_{x,y} - \mathcal{G}_{\theta_g}(I_{x,y}))^2, \quad (5)$$

with W and H are respectively the width and the height of the input image. We explored other options, such as measuring the error on low layer feature maps of a discriminator network in order to extract relevant representations [13, 20], but it did not bring improvements on our rather simple data.

The discriminative part of the loss is defined on the probabilities of the discriminator and updates the weights of the generator via the gradient of:

$$\ell_{SSDS} = \log \left\{ D_{\theta_d}^{(3)}(\mathcal{G}_{\theta_g}(I)) \right\} - \log \left\{ D_{\theta_d}^{(1)}(\mathcal{G}_{\theta_g}(I)) \right\} - \log \left\{ D_{\theta_d}^{(2)}(\mathcal{G}_{\theta_g}(I)) \right\}, \quad (6)$$

where, as a reminder, 3, 2 and 1 correspond to \hat{D} , D and I respectively. Compared to the adversarial loss of the GAN as formulated by Goodfellow et al. [5], our multi-class adversarial loss has two more terms, which translate into more gradients for the back-propagation.

2.6 Training Details

We trained our networks from scratch simultaneously on an NVIDIA GeForce GTX using the dataset described in Sect. 2.2. We trained the generator and the discriminator in an alternating manner with batches of size 16. We scaled the range of the input images to $[0, 1]$. Our final model was trained for 30,000 iterations at a learning rate of $2 \cdot 10^{-4}$ and a decay of 0 using the ADAM optimizer [10]. The generator and the discriminator were updated at each iteration to solve the adversarial min-max (3). In addition, the generator was updated to solve the optimization problem (2) via the gradient of the loss (4), with a regularization parameter set to $\lambda = 10^{-3}$. We implemented our models in Keras [3].

3 Experiments

In this section, we evaluate our Specularity Removal Network quantitatively on synthetic data, for which we have ground-truth, and qualitatively on real data. We compare the performance of our method with three states-of-the-art separation methods: Tan et al. [26] and Shen et al. [23], which are model-based

methods, and Shi et al. [24] which is learning-based. The code for the hand-crafted methods are available on the authors’ webpages and we downloaded the model of [24] on the author’s GitHub¹. We also show the contribution of the multi-class adversarial loss in Sect. 3.3 and discuss the limits in Sect. 3.4.

Table 1. (Left) Results of the **diffuse** component estimation task of the different methods on five synthetic test images. Best view in PDF. (Right) Quantitative results of the diffuse component estimation task in terms of the different metrics. The values represent the mean distance between the generated diffuse component and ground-truth over 1000 randomly sampled synthetic test images. Baseline (AE) is a simple autoencoder without GAN training and Baseline refers to our method with binary classification (classical GAN framework).

Methods	L2	DSSIM	NET	LIN-NET
Tan et al. [26]	0.020	0.052	0.408	0.054
Shen et al. [23]	<u>0.016</u>	0.048	<u>0.380</u>	0.052
Shi et al. [24]	0.222	0.410	2.363	0.313
Baseline (AE)	0.033	<u>0.046</u>	<u>0.336</u>	<u>0.044</u>
Baseline (GAN)	0.017	0.059	0.474	0.064
Ours	0.014	0.035	0.271	0.034

3.1 Evaluation on Synthetic Data

We provide qualitative and quantitative results regarding the estimation of the diffuse component of synthetic images in Table 1. Table 1 (left) shows the results of the diffuse component recovery with the different baseline methods and our method, on five images out of our test set of 1000 images. As can be noticed, our method’s outcome is the perceptually closest to the sought diffuse components. Note that, the learning-based approach of Shi et al. [24] shows reconstruction artifacts on the edges of the image. Their method actually has to work with a mask to segment the foreground object, which is limited as the mask is not provided in most scenarios.

Table 1 (right) provides quantitative results of the different methods and two Baseline methods (AE for autoencoder and GAN for classical GAN) averaged on the test set. We use two standard metrics, namely L2 and DSSIM [7], and also consider recently developed perceptual metrics based on computing the similarity of two images in the feature space of a neural network. Indeed, Zhang et al. [29] showed that these metrics provide an embedding of images which agrees surprisingly well with human judgment. In particular, we consider the

¹ <https://github.com/shi-jian/shapenet-intrinsics>.

metric corresponding to the Squeeze network [8] (denoted NET) and its linear version as proposed by [29] (denoted LIN-NET). This quantitative comparison agrees with the perceptual analysis as our method outperforms all the methods by a large margin on all the metrics. Specifically, it outperforms the Baseline method which consists in considering a binary classification for the discriminator rather than the proposed multi-class classification.

3.2 Evaluation on Real Data

Our network being trained on synthetic images, the fact that it outperforms the other methods on synthetic data can seem natural. However, in this section we evaluate our method on real images and show that it performs consistently. The results can be seen in Fig. 2. Our network performs well on a wide range of images, from images resembling our training data with a black background to complex scenes such as the wooden objects and the earth balloon. This attests of a better generalization from our Specularity Removal Network, while artifacts on the edges and on the object are still visible in the results of Shi et al. [24]. Our discriminative network constrains the generator to understand the distribution of specularities and help it remove them from the image by training themselves to differentiate I from D . Visually, our method consistently outperforms [24].

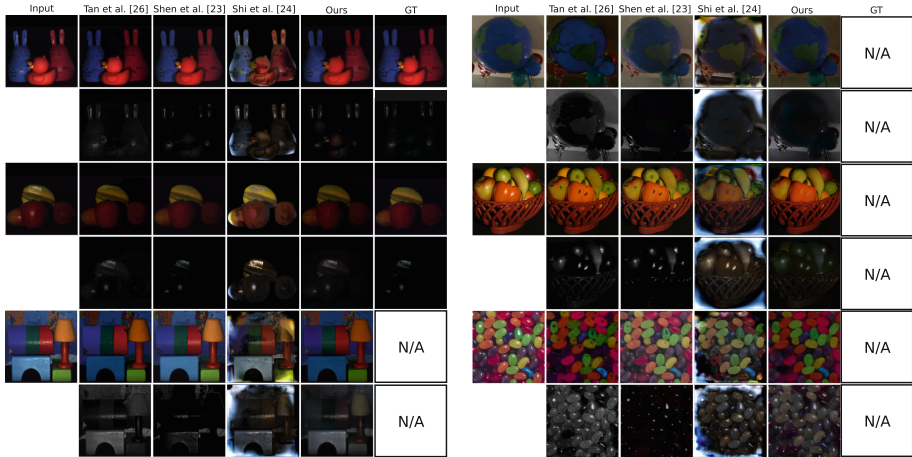


Fig. 2. Results and comparison of our method on real images. The input image is on the left and for each image, the top row is the diffuse component and the bottom row is the specular component. Ground-truths are provided on the right when available. Our baselines include Tan et al. [26], Shen et al. [23] and Shi et al. [24]. Our specular component is obtained by subtracting our estimated diffuse component to the input image. Best view in PDF.

In visual comparison to hand-crafted methods, we perform slightly worse than Shen et al. [23] on the first and the second images (animals and fruits)

and slightly worse than Tan et al. [26] on the first and third examples (animals and wooden objects). This can be explained by the fact that these images were taken in laboratory conditions and fit perfectly the hypothesis of the Dichromatic Reflection Model [21], on which hand-crafted models are built. However, albeit subjective since there is no metric to evaluate specularity removal, our method performs best on the other examples, showing its consistency. On the fifth image (the fruit basket), we can clearly see that purely specular (saturated) pixels lose too much energy with hand-crafted methods because they don't show any diffuse color underneath, while our method consistently begins to inpaint them. The same goes for the last image (beans) where our method does not leave holes like the others. In summary, our method separates the reflection components with the best consistency compared to the state-of-the-art.

3.3 Contribution of Our Multi-class GAN

Figure 3 shows learning curves for our method and for a classical binary GAN with the exact same parameters. The amplitudes of the oscillations of both curves easily tell us that the multi-class adversarial loss allows for a more stable learning, while instability is common to adversarial frameworks. It also shows that convergence comes faster and is more accurate, which is visible in the images generated by the two methods (right of Fig. 3).

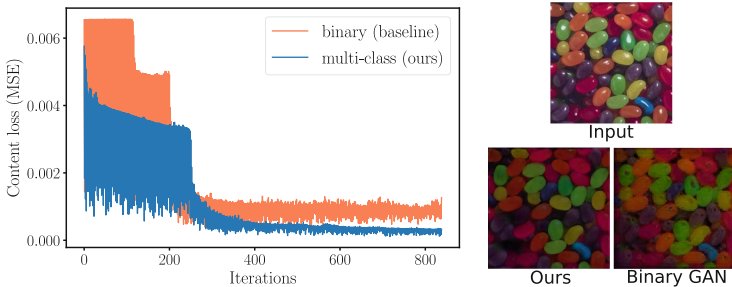


Fig. 3. Learning curves of a classical GAN and our multi-class GAN. (Right) Generated diffuse components.

3.4 Limits

Our network is of course not perfect and can have failure cases in a real life application. First, as mentioned before it does not handle mirror surfaces, which would require a completely different definition of the problem (different data and priors). Furthermore, despite our efforts (white objects and multi-class discriminators), the network still tends to darken the images. This is visible in the earth balloon example of Fig. 2 and shows our network still misses a step of generalization. This can be explained by the simplicity of our data, especially the lack of a

background which would normally provide context to work on. Not any random background however, which would have been easy to add, but a true background containing illumination information. We tried to add a white background to help the network discern white material from white specular highlights but did not see any noticeable improvement. We also tried the patchGAN formulation of the GAN framework without noticeable changes.

4 Conclusion

We have proposed a new method to separate diffuse and specular reflections based on a learning approach, which can better infer the complex relations between the object, the lighting condition and the image than existing approaches. Our method takes advantage of synthetic data generation which allows us to easily obtain a large amount of labeled data. We generated our own training set and augmented our data in such a way to account for difficult cases encountered in real scenarios, in order to help the network better generalize to these situations. We also trained our model in a GAN framework adapted to reflection component separation. For that, we defined a new multi-class adversarial loss, which helps the training process by providing more gradients and more precise features. This results in a model that removes specularity from a single image, without any assumption made about the scene. We evaluated our model on both synthetic and real data. Our method outperforms the state-of-the-art in consistency across various scenes.

In future work, we would like to investigate the temporal coherence for live applications in a continuous video stream. Our Specularity Removal Network is not perfect and might not be entirely consistent from one frame to the next. Our network would also greatly benefit from data with more complex scenes in the training set.

References

1. Adelson, E.H., Pentland, A.P.: The perception of shading and reflectance. In: Perception as Bayesian Inference, pp. 409–423 (1996)
2. An, D., Suo, J., Ji, X., Wang, H., Dai, Q.: Fast and high quality highlight removal from a single image. arXiv preprint [arXiv:1512.00237](https://arxiv.org/abs/1512.00237) (2015)
3. Chollet, F., et al.: Keras (2015)
4. Funke, I., Bodenstedt, S., Riediger, C., Weitz, J., Speidel, S.: Generative adversarial networks for specular highlight removal in endoscopic images. In: Medical Imaging 2018: Image-Guided Procedures, Robotic Interventions, and Modeling, vol. 10576, p. 1057604. International Society for Optics and Photonics (2018)
5. Goodfellow, I.: NIPS 2016 tutorial: generative adversarial networks. arXiv preprint [arXiv:1701.00160](https://arxiv.org/abs/1701.00160) (2016)
6. Hara, K., Nishino, K., Ikeuchi, K.: Determining reflectance and light position from a single image without distant illumination assumption, p. 560. IEEE (2003)
7. Hore, A., Ziou, D.: Image quality metrics: PSNR vs. SSIM. In: 2010 20th International Conference on Pattern Recognition (ICPR), pp. 2366–2369. IEEE (2010)

8. Iandola, F.N., Han, S., Moskewicz, M.W., Ashraf, K., Dally, W.J., Keutzer, K.: SqueezeNet: AlexNet-level accuracy with 50x fewer parameters and <0.5 MB model size. arXiv preprint [arXiv:1602.07360](https://arxiv.org/abs/1602.07360) (2016)
9. Jachnik, J., Newcombe, R.A., Davison, A.J.: Real-time surface light-field capture for augmentation of planar specular surfaces. In: 2012 IEEE International Symposium on Mixed and Augmented Reality (ISMAR), pp. 91–97. IEEE (2012)
10. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. arXiv preprint [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) (2014)
11. Klinker, G.J., Shafer, S.A., Kanade, T.: The measurement of highlights in color images. *Int. J. Comput. Vis.* **2**(1), 7–32 (1988)
12. Lambert, J.H.: *Photometria sive de mensura et gradibus luminis, colorum et umbrae*. Klett (1760)
13. Ledig, C., et al.: Photo-realistic single image super-resolution using a generative adversarial network. arXiv preprint (2016)
14. Lin, S., Lee, S.W.: Estimation of diffuse and specular appearance. In: The Proceedings of the Seventh IEEE International Conference on Computer Vision, vol. 2, pp. 855–860. IEEE (1999)
15. Lin, S., Li, Y., Kang, S.B., Tong, X., Shum, H.-Y.: Diffuse-specular separation and depth recovery from image sequences. In: Heyden, A., Sparr, G., Nielsen, M., Johansen, P. (eds.) ECCV 2002. LNCS, vol. 2352, pp. 210–224. Springer, Heidelberg (2002). https://doi.org/10.1007/3-540-47977-5_14
16. Lin, S., Shum, H.Y.: Separation of diffuse and specular reflection in color images. In: Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, CVPR 2001, vol. 1, p. I. IEEE (2001)
17. Meka, A., Maximov, M., Zollhoefer, M., Chatterjee, A., Richardt, C., Theobalt, C.: Live intrinsic material estimation. arXiv preprint [arXiv:1801.01075](https://arxiv.org/abs/1801.01075) (2018)
18. Nayar, S.K., Fang, X.S., Boulton, T.: Separation of reflection components using color and polarization. *Int. J. Comput. Vis.* **21**(3), 163–186 (1997)
19. Ronneberger, O., Fischer, P., Brox, T.: U-Net: convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) MICCAI 2015. LNCS, vol. 9351, pp. 234–241. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24574-4_28
20. Seddik, M.E.A., Tamaazousti, M., Lin, J.: Generative collaborative networks for single image super-resolution. [arXiv:1902.10467](https://arxiv.org/abs/1902.10467) (2019)
21. Shafer, S.A.: Using color to separate reflection components. *Color Res. Appl.* **10**(4), 210–218 (1985)
22. Shen, H.L., Cai, Q.Y.: Simple and efficient method for specular removal in an image. *Appl. Opt.* **48**(14), 2711–2719 (2009)
23. Shen, H.L., Zhang, H.G., Shao, S.J., Xin, J.H.: Chromaticity-based separation of reflection components in a single image. *Pattern Recogn.* **41**(8), 2461–2469 (2008)
24. Shi, J., Dong, Y., Su, H., Stella, X.Y.: Learning non-Lambertian object intrinsics across ShapeNet categories. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 5844–5853. IEEE (2017)
25. Shi, W., et al.: Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1874–1883 (2016)
26. Tan, R.T., Ikeuchi, K.: Separating reflection components of textured surfaces using a single image. *IEEE Trans. Pattern Anal. Mach. Intell.* **27**(2), 178–193 (2005)

27. Weiss, Y.: Deriving intrinsic images from image sequences. In: 2001 Proceedings of the Eighth IEEE International Conference on Computer Vision, ICCV 2001, vol. 2, pp. 68–75. IEEE (2001)
28. Wolff, L.B., Boulton, T.E.: Constraining object features using a polarization reflectance model. *IEEE Trans. Pattern Anal. Mach. Intell.* **13**(7), 635–657 (1991)
29. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. *arXiv preprint* (2018)