https://doi.org/10.2352/ISSN.2470-1173.2020.10.IPAS-182

OEC-cnn: a simple method for over-exposure correction in photographs

Zhao Gao[1,2], Eran Edirisinghe[2], Slava Chesnokov[1];
[1] Imaging & Vision Group, Arm Ltd, UK;
[2] Department of Computer Science, Loughborough University, Loughborough, UK

Abstract

Over-exposure happens often in daily-life photography due to the range of light far exceeding the capabilities of the limited dynamic range of current imaging sensors. Correcting over-exposure aims to recover the fine details from the input. Most of the existing methods are based on manual image pixel manipulation, and therefore are often tedious and time-consuming. In this paper, we present the first convolutional neural network (CNN) capable of inferring the photo-realistic natural image for the single over-exposed photograph. To achieve this, we propose a simple and lightweight Over-Exposure Correction CNN, namely OEC-cnn, and construct a synthesized dataset that covers various scenes and exposure rates to facilitate training. By doing so, we effectively replace the manual fixing operations with an end-to-end automatic correction process. Experiments on both synthesized and real-world datasets demonstrate that the proposed approach performs significantly better than existing methods and its simplicity and robustness make it a very useful tool for practical over-exposure correction. Our code and synthesized dataset will be made publicly available.

Introduction

Over-exposure artifacts occur when a limited dynamic range of digital imaging device records the bright scene of which luminance levels are higher than saturation levels of the sensor. In digital photography, over-exposure is also referred to as clipping, describing the data loss in a captured image. The clipped area of the image typically appears as a uniform area of the minimum or maximum brightness, which often loses rich details within the image. The captured low quality images will not only degrade the visual aesthetics of photos, but also affect the performance of further computer vision tasks such as object recognition. A well-exposed image is the one that feels just bright enough so that the highlights are as they feel the most natural and comfortable to look at. Theoretically, such a photograph contains no lost highlights, meaning all the details are clearly distinguishable and as close as possible to real life scenarios. Therefore, over-exposure is a crucial step for imaging processing and forms a particularly challenging problem. The over-exposure avoidance process can directly be performed by adjusting camera settings prior to capturing an image. For example, the photographer has ultimate control over how much light is actually allowed to hit the sensor when the shutter opens by changing the ISO settings, aperture settings and shutter speed. The tradeoff of preventing OE by this is that usually it will make the photo appear dimmer and suffer from sensor noise. Alternatively, over-exposure can be also alleviated by using high dynamic range (HDR) capture with tone mapping techniques. HDR imaging can be utilised to capture an HDR scene without over-exposure. With proper tone mapping, an HDR image can be compressed into a well-exposed low dynamic range (LDR) image. However, such techniques require expensive equipment and are not applicable for historical photographs.

In contrast, the OE correction can be performed as post-processing on the display side after the image is captured. In practice, many people use existing graphics editing software to salvage over-exposure by manually adding colour to the clipped area or adding a haze to the highlights. Unfortunately, this kind of user intervention is tedious and time-consuming since it cannot be automatically repeated when dealing with a large amount of over-exposed photographs. For humans who carry out these tasks manually, it is a very slow and painstaking process. Alternatively, since information in different colour channels of RGB images is known to be highly correlated, some automatic correction methods are based on this correlation to recover clipped information from unclipped channels as proposed in [21] [10] [1] [30]. Apart from the algorithmic complexity, pixel-manipulation based methods are not flexible enough to fully automate different cases of clipping such as colour clipping, where one or two channels are saturated and over-exposed, where all three channels are saturated. In this paper, the term over-exposure (OE) refers to both of the aforementioned cases. Recently, some deep learning based approaches on HDR reconstruction can be treated as image colour correction because their output HDR image can be tone mapped into the LDR image. However, existing tone mapping methods always fail to preserve the local details from the HDR domain when OE happens. Such methods take the indirect path to get the corrected image is suboptimal since the output of correction method is always expected to be in the same colour domain for more general display purpose (most of devices do not support HDR) as well as existing tone mapping methods always fail to preserve the local details from the HDR domain when OE happens.

In order to handle practical over exposure problems in photographs, an ideal solution is expected to have the following desirable properties: it is able to perform correction using a single model; it is efficient, effective and user-friendly; it can handle spatially variant exposure. To realize the ideals, we introduce the first end-to-end deep-learning based approach to automatically infer a colour corrected image from a single over-exposed input. We employ supervised learning using CNN directly to learn an end-to-end mapping from an 8-bit LDR over-exposed image to an 8-bit LDR colour corrected image. Our method differs fundamentally from existing pixel-manipulation methods, in that ours does not explicitly detect the clipped area or reconstruct HDR and these are implicitly achieved via hidden layers. In our method, the entire colour correction pipeline is fully obtained through learning. To obtain a sufficiently large amount of training data without utilizing the expensive camera devices, we propose a simulation of approximating the over exposure process. At the meanwhile, introducing the different degrees of over-exposure parameter greatly enriches the dataset and enables the model to handle the spatially variant exposure and arbitrary target for exposure. We name the proposed model over-exposure correction convolutional neural network (OEC-cnn) and the code will be made publicly available. Our main contributions are:

- A lightweight and robust colour correction network, namely OEC-cnn, is proposed for automatically recovering over-exposed images. To our best knowledge, this is the first end-to-end deep learning based approach to fix the exposure issues in photographs. A single OEC-cnn is able to deal with exposure on different levels, as well as spatially variant exposure.
- Currently, there is no public dataset dedicated to fix exposure issues. To address this, we provide a novel training data generation procedure to enable computer vision research on exposure correction, which effectively reformulates the conventional colour correction methods into a deep CNN.
- OEC-cnn exhibits perceptually appealing results on both synthetic over-exposed images and real-world over-exposed images, demonstrating its potential for practical image colour correction. Its simplicity and rubustness make it easy to be deployed into hardware or embedded into a built-in camera component, which provides a better alternative to the photo editing tools.

Related Work

Over-exposure Correction

Due to the limited dynamic range of common camera sensors, over-exposure appears commonly in daily-life photography, where saturated areas of the image will typically appear as a uniform brightness. The clipped area will commonly become completely white, though in the case that only one channel has clipped it may represent itself as an area of distorted colour. Interactive manual corrections could fix the issue, however pixel-level manipulation is an art form, which is both tedious and time-consuming. In colour images, such clipping may occur in any of the image's RGB colour channels independently, where just one- or two-colour channels of the output colour signals from the image sensor may become saturated. Since the information in different channels is known to be highly correlated, some automatic methods based on this knowledge have been proposed to address the over-exposure problem [21] [10] [1] [30]. There is a detailed review of existing methods based on the assumption of the partial over-exposure in [12]. By the nature of their hypothesis, the primary limitation of the most earlier works is that at least one of the colour channels cannot be saturated. When all three of the colour channels are clipped, those methods typically used to accomplish the task do not bring additional information to solve such ill-posed reconstruction problem. That causes the main issue of wrong colour propagations from adjacent regions.

Deep Learning based Approaches

Our idea is motivated by the outstanding performance on convolutional neural networks (CNN) for inverse problems in imaging [22]. As a powerful discriminative learning method, CNN has been successfully used in many low-level computer vision applications such as single image super-resolution(SISR) [4], image-demosaicking [8] and image denoising [31]. In those applications, pairs of degraded images and their high-quality counterparts can be easily generated. For example, SISR methods directly learns an end-to-end mapping between the low/high-resolution images; denoisers take advantages of noisy images and corresponding clear reference images. With those paired training datasets, CNN can be adopted to learn a mapping function between the degraded observations and their high-quality reference images. The success of deep neural networks depends critically on the data used for training. Deep learning approaches are data hungry and their performance is strongly correlated with the amount of available training data. As aforementioned, most available data in computer vision research are tailored to the problem of a specific research group and there is no publicly available over-exposed image datasets with corresponding ground truth images.

Training Dataset Generation

To solve the issue of lack of training data, we get motivated by the ideas that the dataset can be synthesized by simulating the camera rather than capturing the real images with expensive imaging devices. In [3], the authors train a single image contrast enhancer with the constructed training dataset. Multi-Exposure Image Fusion (MEF) and stack-based HDR methods are used to reconstruct the reference good-contrast image of a scene, while those under-exposure or over-exposure images of the scene can be naturally taken as low-contrast counterparts. DRHT [29] is a united framework consisting of two CNNs for HDR reconstruction and tone mapping, which uses the Adobe Photoshop software to empirically generate LDR images from collected HDR ground truth images with human supervision. Further, to adapt their model to the real images with high resolution, they also use the Physically based Rendering Technology (PBRT) [25] to generate the ground truth HDR scenes as well as the input and ground truth LDR images. Hdrcnn [5] aims to reconstruct the HDR image from an arbitrary single exposed LDR image. For training, data is gathered from a large set of existing HDR image sources in order to create a training dataset. For each HDR image, a set of corresponding LDR exposures are simulated using a virtual camera model. DrTMo [6] infers the HDR image from a single LDR input by synthesizing a set of LDR images and exposures from each HDR image. Both of Hdrcnn and DrTMo use Grossberg and Nayar's Database of Response Functions (DoRF) [9] to define camera response functions (CRFs). This database consists of 201 response curves for common brands of films, charge-coupled devices (CCDs), and digital cameras collected by the authors. Since using all the CRFs is redundant and unnecessarily increase the training time, DrTMo uses only representative five CRFs selected using k-means clustering while Hdrcnn uses a parametric sigmoid function to fit the mean of the collected camera curves by [9].

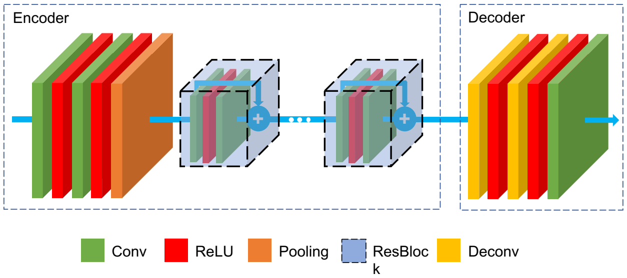

Figure 1: The architecture of the proposed over-exposure correction network (OEC).

Proposed Over-Exposure Correction CNN

Problem formulation

The aim of OE correction is to estimate a recovered image I^{REC} from an over-exposed input image I^{OE}: $I^{REC} = f(I^{OE})$, where $f : [0,255] \rightarrow [0,255]$. Both input and output data are LDR integer RGB values. Here, I^{OE} is obtained by the simulation of clipping from its counterpart ground truth image I^{GT}. The ground truth images are only available during training. In this context, $f(\cdot)$ can be considered as an ill-posed problem, which is to compensate lost information by over-exposure. The approach is similar to inferring a HDR image from a single LDR input: $I^{HDR} = f(I^{LDR})$, where $f : [0,255] \rightarrow \mathbb{R}^{+}$. The difference is the range of output values of those two tasks.

Our ultimate goal is to train an end-to-end mapping function f that estimates a given over-exposed input image from its corresponding ground truth counterpart. Learning f requires the estimation of network parameters $f(\Theta)$. This is achieved through optimising the average mean squared error (MSE) between the reconstructed images $f(I^{OE})$ and the corresponding desired images I^{GT}:

$$L(\Theta) = \frac{1}{N}\sum_{n=1}^{N} ||f(I_n^{OE};\Theta) - I_n^{GT}||^2, \tag{1}$$

where N is the number of training samples. For training images I_n^{GT}, $n = 1,...,N$ with corresponding I_n^{OE}, $n = 1,...,N$, we solve:

$$\hat{\Theta} = \text{argmin}\, \frac{1}{N}\sum_{n=1}^{N} L(f(I_n^{OE}), I_n^{GT}). \tag{2}$$

With this approach, our network can learn to create solutions that are highly similar to the ground truth images.

Synthesizing Over-exposed Images

A key challenge for a learning-based OE correction is to obtain a sufficiently large amount of well-structured training data. Collecting the high-quality reference images is feasible, however it is very hard to capture the corresponding OE counterpart. Despite the sophisticated metering techniques have been requited, taking well-exposed photos with corresponding over-exposed images remains a challenge for normal users. To address this issue, we simulate the camera sensor to approximate the OE process. Considering each existing RGB ground truth as a real-world scene, we generate the OE image by passing the ground truth to the simplified camera pipeline: linearisation, exposure and gamma curve application. Towards this, we define the camera response curve using a parametrised sigmoid function:

$$f(I) = (1+\sigma)\frac{I^n}{I^n+\sigma}, \tag{3}$$

which is to represent the relation between the input and output of the camera sensor. $n = 0.9$ and $\sigma = 0.6$ are set to fit the mean of the database [9] of camera curves. The parameter settings and the camera curve fitting function have been proposed by Hdrcnn [5]. The inverse camera curve f^{-1} is adopted to perform linearization, followed by clipping occurs where the intensity falls outside the $[0,1]$ thresholds, then gamma correction is mimicked using function f for display purpose:

$$I^{linear} = f^{-1}(I^{GT}), \tag{4}$$

$$I^{clipped} = min(1, max(0, I^{linear} \times \tau)), \tag{5}$$

$$I^{OE} = f(I^{clipped}), \tag{6}$$

where τ is the degree of over-exposure. The higher the τ value is, the higher the saturation ratio of the pixels becomes. In the colour image, over-exposure area is seen as washed-out bright area that turns to completely white if all colour components are clipped. There is the other case that only one- or two channels are saturated, it may be represented as an area of distorted colour, such as an area of sky that appears more dominantly greener or yellower than it should be in reality. To enrich our training dataset, we generate different levels of over-exposure by changing the τ value.

Network Architecture

Fig. 1 shows the overall architecture of our network. The input to the network is a $W \times H \times c$ over-exposed LDR image where W, H and c are width, height, and the number of colour channels respectively. In the training stage, we use the exposed images from a range of exposure degrees (e.g., $\tau \in \{1.1, 1.5\}$) to train a single model. Given a test image whose exposure degree range, the learned single OEC-cnn model can utilzed to correct it without estimating its exposure level. Inspired by [13] [11] [16], we design our OEC-cnn as an encoder-decoder network with residual blocks. OEC-cnn structure is intentionally designed with simplicity in mind. Overall, it contains 5 types of layers which are shown in Fig. 1 with 5 different colours.

1. Conv: the convolutional layer with filter size of 9×9 and stride 1 for feature extraction;
2. ReLU: rectified linear unit [23] is utilised for the nonlinearity;
3. Pooling: the 3×3 convolutional layer with stride 2 is adopted for pooling, helps the compressed layer outputs to be invariant to small shifts and distortions of the input image;
4. ResBlock: each block contains two convolutional layers with 3×3 kernels and ReLU as the activation function. The reasons behind deprecation of batch normalization have been suggested already in [17].
5. Deconv: following the structure of deconvolution in [24], the deconvolutional layer has been applied to keep the size of the output unchanged.

Table 1: The Average SSIM, PSNR and NRMSE results of different numbers of residual blocks on CBSD68, Kodak24 and McMaster datasets.

	CBSD68			Kodak			McMaster			Overall		
#	SSIM	PSNR	NRMSE	SSIM	PSNR	NRMSE	SSIM	PSNR	NRMSE	SSIM	PSNR	NRMSE
4	0.986	30.527	0.076	0.984	30.615	0.079	0.961	30.328	0.083	0.977	30.492	0.079
8	0.986	29.905	0.080	0.985	29.992	0.083	0.960	30.118	0.084	0.977	30.015	0.082
12	0.986	31.011	0.073	0.985	31.156	0.075	0.962	30.895	0.077	0.978	31.023	0.075
16	0.987	31.216	0.069	0.989	31.678	0.063	0.977	32.044	0.065	0.984	31.654	0.065
20	0.987	31.350	0.066	0.985	29.680	0.075	0.971	31.399	0.069	0.981	30.813	0.069

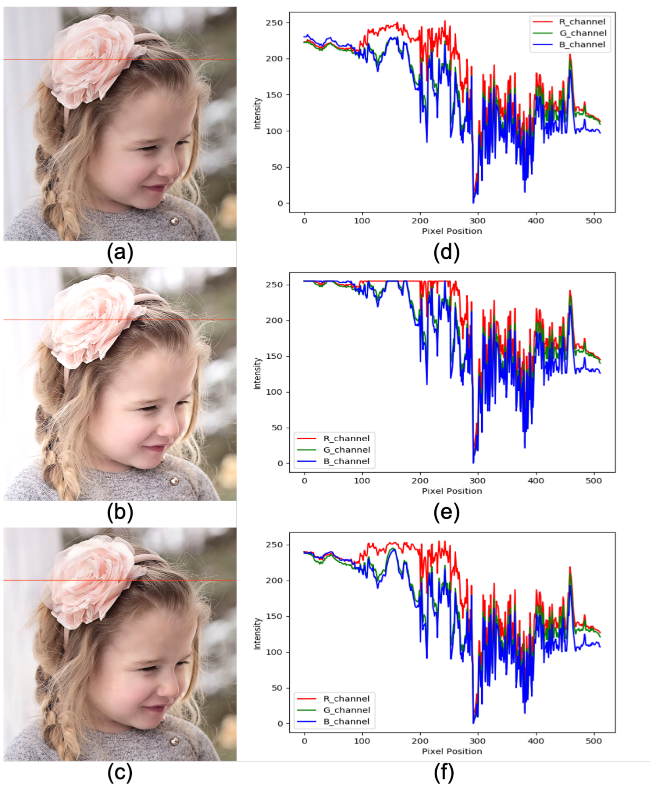

Figure 2: Example of colour distortion due to the clipping. (a) A correct exposure image as ground truth, (b) the corresponding simulated over-exposed image, (c) the constructed result. (d), (e) and (f) are the plots of the RGB intensities versus the pixel index along the highlighted horizontal lines of images (a), (b) and (c) respectively.

Experimental Results

Experimental settings

We implement our model using PyTorch package, and utilise DIV2K [2] and CelebA [18] as ground truth images and generate corresponding over-exposed images from them as training dataset. Images are cropped into 178×178 size patches. We train with a batch size of 32 for 200 epoches using Adam [14] with a learning rate of 1×10^{-4} without weight decay or dropout. In the experiment, we use a Nvidia Tesla P100 for training our OEC. We build the model using PyTorch version 0.3.1. The operating system of the server is Ubuntu 16.04.1 with CUDA 9.1 installed.

Investigation on the depth of network

We investigate the proposed OEC-cnn with different number of building blocks and test them on the synthesized over-exposed image datasets from three publicly available datasets: CBSD68 [20], Kodak [7] and McMaster [32]. Those datasets haven't been seen in the training stage. Due to the nature of over-exposure, it is clear that high exposure rate leads to more detailed information lost in the images. Therefore, It is expected that the saturation rate is in direct ratio to the exposure rate in the figure. Though under same exposure value, different datasets demonstrate different saturation rate because of the various contents. The average saturation rate of CBSD68, which contains lots of landscape scenes, is already 0.915% due to the existence of over-exposure, while other two datasets are relatively well-exposed and their initial saturation rate is close to 0%.

Standard techniques for quality assessment, including Structural Similarity Index (SSIM), Peak Signal to Noise Ratio (PSNR) and normalized root-mean-square error (NRMSE), have been adopted. The evaluation results are documented in Table 1. In our study, PSNR and NRMSE are more sensitive to the number of residual blocks while SSIM keeps stable. SSIM measures how "similar" various metrics of the images are, while PSNR and NRMSE are essentially the measurements of "error" between the two images and both of them are based on the mean-square error. Table 1 shows that OEC with more residual blocks gives better performance.

Comparison to the ground truth

Fig. 2 shows an example of dealing with colour image. Within an image, clipping can occur in different cases: one or two channels clipped with some information left, three channels all clipped without any information left. In this example, we demonstrate that exposure area detection is sub-optimal and adds computational complexity. In our method, we are capable of dealing with all variations of clipping in the same manner.

We compare the proposed method to two deep learning based HDR prediction methods: Hdrcnn [5] and DrTMO [6], see Fig. 3 . These two approaches can be treated as over-exposure correction methods because their output HDR image can be tone mapped into LDR domain. As illustrated in Fig. 4, differences were quantitatively assessed using HDR-VDP-2 [19]. This metric calculates visibility differences based on human perception and the visualized differences increases from blue to red.

Results on real-world images

As shown in Fig. 5, we collected some photos from internet, which are known as bad exposure examples, and then test our algorithms on those images. For the "bridesmaid" photo, our method can differentiate the skin from dresses and recover the details of dresses well. The skin recovered is close to the realistic case. In the photo of "bride and groom", our method effectively recovers the details of the body even background and people are

Figure 3: Over-exposure correction results by different methods. From left to right: over-exposed input, Hdrcnn, DrTMO, ours and reference.

both burned out. In the “flower” case, the edges are reserved and colour is more vivid.

Conclusion

In this paper, we proposed a CNN model, namely OEC-cnn, for over-exposure correction. Besides, we construct a synthesized over-exposed training dataset which effectively reformulates the conventional colour correction methods into a deep CNN. OEC-cnn exhibits perceptually appealing results on both synthetic over-exposed images and real-world over-exposed images. Taking its performance and robustness into consideration, OEC-cnn is very competitive for practical applications.

Acknowledgments

This research was funded by ARM Ltd and as part of the EPSRC Centre for Doctoral Training in Embedded Intelligence grant number EP/L014998/1.

References

[1] M. A. Abebe, A. Booth, J. Kervec, T. Pouli, and M.-C. Larabi. Towards an automatic correction of over-exposure in photographs: Application to tone-mapping. *Computer Vision and Image Understanding*, 168:3 – 20, 2018. Special Issue on Vision and Computational Photography and Graphics.

[2] E. Agustsson and R. Timofte. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, July 2017.

[3] J. Cai, S. Gu, and L. Zhang. Learning a deep single image contrast enhancer from multi-exposure images. *IEEE Transactions on Image Processing*, 27(4):2049–2062, 2018.

[4] C. Dong, C. C. Loy, K. He, and X. Tang. Image super-resolution using deep convolutional networks. *IEEE transactions on pattern analysis and machine intelligence*, 38(2):295–307, 2016.

[5] G. Eilertsen, J. Kronander, G. Denes, R. K. Mantiuk, and J. Unger. Hdr image reconstruction from a single exposure using deep cnns. *ACM Transactions on Graphics (TOG)*, 36(6):178, 2017.

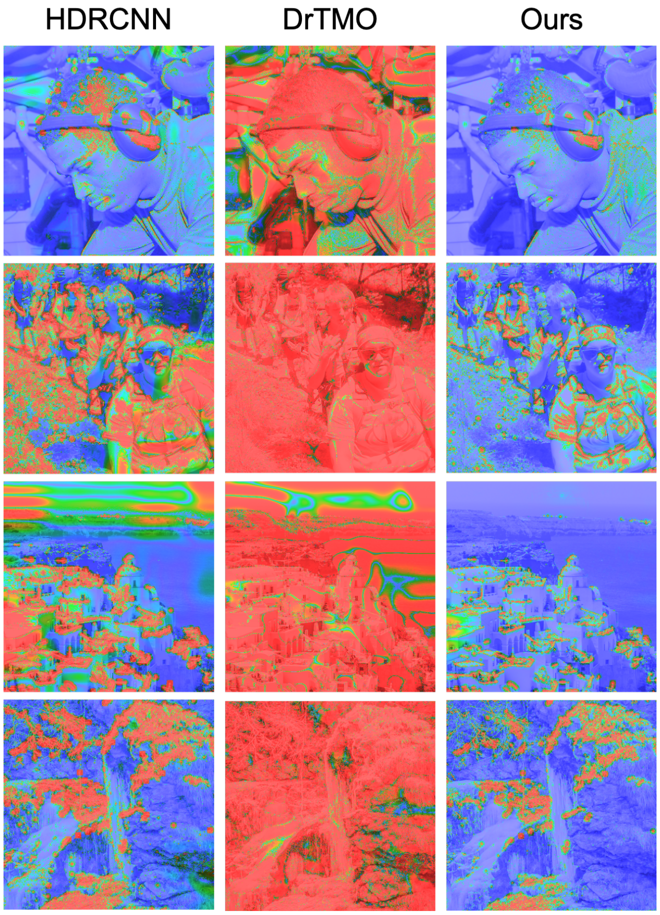

Figure 4: Comparisons with the corrected images generated by each method and the ground truth images using the HDR-VDP-2 metric. From left to right: Hdrcnn, DrTMO, ours.

Over-exposed Image Proposed Method

Figure 5: Qualitative results on real over-exposed images.

[6] Y. Endo, Y. Kanamori, and J. Mitani. Deep reverse tone mapping. *ACM Trans. Graph*, 36(6), 2017.

[7] R. Franzen. Kodak lossless true color image suite: Photocd pcd0992, 2002.

[8] M. Gharbi, G. Chaurasia, S. Paris, and F. Durand. Deep joint demosaicking and denoising. *ACM Transactions on Graphics (TOG)*, 35(6):191, 2016.

[9] M. D. Grossberg and S. K. Nayar. Determining the camera response from images: What is knowable? *IEEE Transactions on Pattern Analysis & Machine Intelligence*, (11):1455–1467, 2003.

[10] D. Guo, Y. Cheng, S. Zhuo, and T. Sim. Correcting over-exposure in photographs. *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 515–521, 2010.

[11] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.

[12] L. Hou, H. Ji, and Z. Shen. Recovering over-/underexposed regions in photographs. *SIAM journal on imaging sciences*, 6(4):2213–2235, 2013.

[13] J. Kim, J. Kwon Lee, and K. Mu Lee. Accurate image super-resolution using very deep convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1646–1654, 2016.

[14] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[15] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, pages 1097–1105. Curran Associates, Inc., 2012.

[16] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. P. Aitken, A. Tejani, J. Totz, Z. Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *CVPR*, volume 2, page 4, 2017.

[17] B. Lim, S. Son, H. Kim, S. Nah, and K. M. Lee. Enhanced deep residual networks for single image super-resolution. In *The IEEE conference on computer vision and pattern recognition (CVPR) workshops*, volume 1, page 4, 2017.

[18] Z. Liu, P. Luo, X. Wang, and X. Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, 2015.

[19] R. Mantiuk, K. J. Kim, A. G. Rempel, and W. Heidrich. Hdr-vdp-2: a calibrated visual metric for visibility and quality predictions in all luminance conditions. In *ACM Transactions on graphics (TOG)*, volume 30, page 40. ACM, 2011.

[20] D. Martin, C. Fowlkes, D. Tal, and J. Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Proc. 8th Int'l Conf. Computer Vision*, volume 2, pages 416–423, July 2001.

[21] S. Z. Masood, J. Zhu, and M. F. Tappen. Automatic correction of saturated regions in photographs using cross-channel correlation. In *Computer Graphics Forum*, volume 28, pages 1861–1869. Wiley Online Library, 2009.

[22] M. T. McCann, K. H. Jin, and M. Unser. A review of convolutional neural networks for inverse problems in imaging. *arXiv preprint arXiv:1710.04011*, 2017.

[23] V. Nair and G. E. Hinton. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th international conference on machine learning (ICML-10)*, pages 807–814, 2010.

[24] A. Odena, V. Dumoulin, and C. Olah. Deconvolution and checker-

board artifacts. *Distill*, 2016.
[25] M. Pharr, W. Jakob, and G. Humphreys. *Physically based rendering: From theory to implementation*. Morgan Kaufmann, 2016.
[26] P. Sermanet, D. Eigen, X. Zhang, M. Mathieu, R. Fergus, and Y. LeCun. Overfeat: Integrated recognition, localization and detection using convolutional networks. *arXiv preprint arXiv:1312.6229*, 2013.
[27] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
[28] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015.
[29] X. Yang, K. Xu, Y. Song, Q. Zhang, X. Wei, and R. W. Lau. Image correction via deep reciprocating hdr transformation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1798–1807, 2018.
[30] L. Yuan and J. Sun. Automatic exposure correction of consumer photographs. In *European Conference on Computer Vision*, pages 771–785. Springer, 2012.
[31] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang. Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE Transactions on Image Processing*, 26(7):3142–3155, 2017.
[32] L. Zhang, X. Wu, A. Buades, and X. Li. Color demosaicking by local directional interpolation and nonlocal adaptive thresholding. *Journal of Electronic imaging*, 20(2):023016, 2011.

Author Biography

Zhao Gao is a PhD candidate in computer science at Loughborough University, funded by ARM and EPSRC. From scientific research to commercial applications, she has applied deep learning approaches into various computer vision tasks. She has adopted convolutional neural networks on ill-posed image reconstructions problems such as: image super-resolution, image demosaicking, image denoising and over-exposure correction. She has also investigated deep learning based computer vision on virtual reality and augmented reality, such as depth reconstruction, gazing tracking, eye contact on video conference and so on.

Prof. Eran Edirisinghe completed his MSc and PhD at Loughborough University in 1996 and 1999, respectively. He was appointed as the Head, Department of Computer Science, Loughborough University in August 2011 and was promoted to a Chair in July 2012. His research has been funded by the EPSRC, TSB and industry. He has also acted as a reviewer for many esteemed journals, international grant awarding bodies and on various advisory roles to industry and public sector organisations.

Dr. Viacheslav Chesnokov received his B.Sc and M.Sc degrees in physics and applied mathematics from the Moscow Institute of Physics and Technology, Russia, in 1991, and PhD in Computer Science (Computational Neuroscience) from De Montfort University, UK, in 2001. He was a co-founder and CTO of Apical Limited, a UK company specializing in advanced imaging technologies. His current research interests include Image Processing /CV (he is an inventor of widely used in industry IridixTM and Assertive DisplayTM) and Neural Networks.

JOIN US AT THE NEXT EI!

IS&T International Symposium on

Electronic Imaging

SCIENCE AND TECHNOLOGY

Imaging across applications . . . Where industry and academia meet!

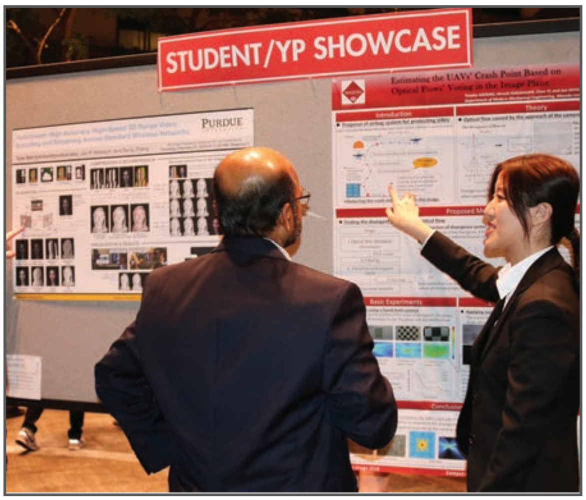

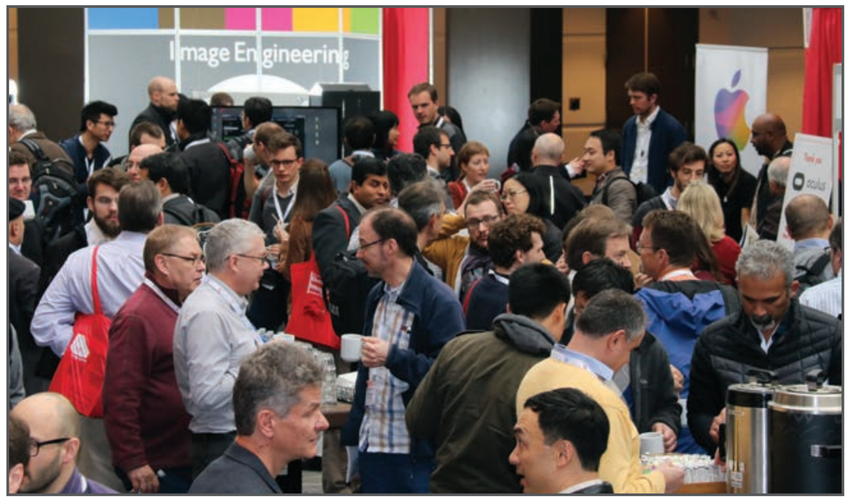

• SHORT COURSES • EXHIBITS • DEMONSTRATION SESSION • PLENARY TALKS •
• INTERACTIVE PAPER SESSION • SPECIAL EVENTS • TECHNICAL SESSIONS •

www.electronicimaging.org

IS&T
imaging.org