

HDR image reconstruction from a single exposure using deep CNNs

GABRIEL EILERTSEN, Linköping University, Sweden
JOEL KRONANDER, Linköping University, Sweden
GYORGY DENES, University of Cambridge, UK
RAFAŁ K. MANTIUK, University of Cambridge, UK
JONAS UNGER, Linköping University, Sweden

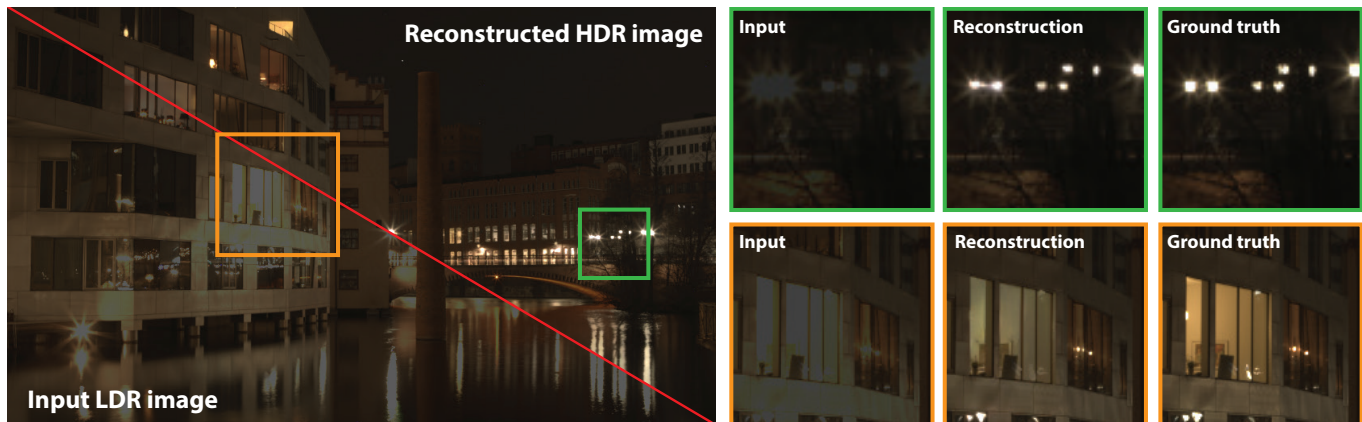


Fig. 1. The exposure of the input LDR image in the bottom left has been reduced by 3 stops, revealing loss of information in saturated image regions. Using the proposed CNN trained on HDR image data, we can reconstruct the highlight information realistically (top right). The insets show that the high luminance of the street lights can be recovered (top row), as well as colors and details of larger saturated areas (bottom row). The exposures of the insets have been reduced by 5 and 4 stops in the top and bottom rows, respectively, in order to facilitate comparisons. All images have been gamma corrected for display.

Camera sensors can only capture a limited range of luminance simultaneously, and in order to create high dynamic range (HDR) images a set of different exposures are typically combined. In this paper we address the problem of predicting information that have been lost in saturated image areas, in order to enable HDR reconstruction from a single exposure. We show that this problem is well-suited for deep learning algorithms, and propose a deep convolutional neural network (CNN) that is specifically designed taking into account the challenges in predicting HDR values. To train the CNN we gather a large dataset of HDR images, which we augment by simulating sensor saturation for a range of cameras. To further boost robustness, we pre-train the CNN on a simulated HDR dataset created from a subset of the MIT Places database. We demonstrate that our approach can reconstruct high-resolution visually convincing HDR results in a wide range of situations, and that it generalizes well to reconstruction of images captured with arbitrary and low-end cameras that use unknown camera response functions and post-processing. Furthermore, we compare to existing methods for HDR expansion, and show high quality results also for image based lighting. Finally, we evaluate the results in a subjective experiment performed on an HDR display. This shows that the reconstructed HDR images are visually convincing, with large improvements as compared to existing methods.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2017 Association for Computing Machinery.
0730-0301/2017/11-ART178 \$15.00
<https://doi.org/10.1145/3130800.3130816>

CCS Concepts: • **Computing methodologies** → *Image processing; Neural networks*;

Additional Key Words and Phrases: HDR reconstruction, inverse tone-mapping, deep learning, convolutional network

ACM Reference format:

Gabriel Eilertsen, Joel Kronander, Gyorgy Denes, Rafał K. Mantiuk, and Jonas Unger. 2017. HDR image reconstruction from a single exposure using deep CNNs. *ACM Trans. Graph.* 36, 6, Article 178 (November 2017), 15 pages. <https://doi.org/10.1145/3130800.3130816>

1 INTRODUCTION

High dynamic range (HDR) images can significantly improve the viewing experience – viewed on an HDR capable display or by means of tone-mapping. With the graphics community as an early adopter, HDR images are now routinely used in many applications including photo realistic image synthesis and a range of post-processing operations; for an overview see [Banterle et al. 2011; Dufaux et al. 2016; Reinhard et al. 2010]. The ongoing rapid development of HDR technologies and cameras has now made it possible to collect the data required to explore recent advances in deep learning for HDR imaging problems.

In this paper, we propose a novel method for reconstructing HDR images from low dynamic range (LDR) input images, by estimating missing information in bright image parts, such as highlights, lost due to saturation of the camera sensor. We base our approach on a

fully convolutional neural network (CNN) design in the form of a *hybrid dynamic range autoencoder*. Similarly to deep autoencoder architectures [Hinton and Salakhutdinov 2006; Vincent et al. 2008], the LDR input image is transformed by an encoder network to produce a compact feature representation of the spatial context of the image. The encoded image is then fed to an HDR decoder network, operating in the log domain, to reconstruct an HDR image. Furthermore, the network is equipped with skip-connections that transfer data between the LDR encoder and HDR decoder domains in order to make optimal use of high resolution image details in the reconstruction. For training, we first gather data from a large set of existing HDR image sources in order to create a training dataset. For each HDR image we then simulate a set of corresponding LDR exposures using a virtual camera model. The network weights are optimized over the dataset by minimizing a custom HDR loss function. As the amount of available HDR content is still limited we utilize transfer-learning, where the weights are pre-trained on a large set of simulated HDR images, created from a subset of the MIT Places database [Zhou et al. 2014].

Expansion of LDR images for HDR applications is commonly referred to as inverse tone-mapping (iTMO). Most existing inverse tone-mapping operators (iTMOs) are not very successful in reconstruction of saturated pixels. This has been shown in a number of studies [Akyüz et al. 2007; Masia et al. 2009], in which naïve methods or non-processed images were more preferred than the results of those operators. The existing operators focus on boosting the dynamic range to look plausible on an HDR display, or to produce rough estimates needed for image based lighting (IBL). The proposed method demonstrates a step improvement in the quality of reconstruction, in which the structures and shapes in the saturated regions are recovered. It offers a range of new applications, such as exposure correction, tone-mapping, or glare simulation.

The main contributions of the paper can be summarized as:

- (1) A deep learning system that can reconstruct a high quality HDR image from an arbitrary single exposed LDR image, provided that saturated areas are reasonably small.
- (2) A *hybrid dynamic range autoencoder* that is tailored to operate on LDR input data and output HDR images. It utilizes HDR specific transfer-learning, skip-connections, color space and loss function.
- (3) The quality of the HDR reconstructions is confirmed in a subjective evaluation on an HDR display, where predicted images are compared to HDR and LDR images as well as a representative iTMO using a random selection of test images in order to avoid bias in image selection.
- (4) The HDR reconstruction CNN together with trained parameters are made available online, enabling prediction from any LDR images: <https://github.com/gabrieleilertsen/hdrcnn>.

2 RELATED WORK

2.1 HDR reconstruction

In order to capture the entire range of luminance in a scene it is necessary to use some form of exposure multiplexing. While static scenes commonly are captured using multiplexing exposures in the time domain [Debevec and Malik 1997; Mann and Picard 1994;

Unger and Gustavson 2007], dynamic scenes can be challenging as robust exposure alignment is needed. This can be solved by techniques such as multi-sensor imaging [Kronander et al. 2014; Tocci et al. 2011] or by varying the per-pixel exposure [Nayar and Mitsunaga 2000] or gain [Hajisharif et al. 2015]. Furthermore, saturated regions can be encoded in glare patterns [Rouf et al. 2011] or with convolutional sparse coding [Serrano et al. 2016]. However, all these approaches introduce other limitations such as bulky and custom built systems, calibration problems, or decreased image resolution. Here, we instead tackle the problem by reconstructing visually convincing HDR images from single images that have been captured using standard cameras without any assumptions on the imaging system or camera calibration.

2.2 Inverse tone-mapping

Inverse tone-mapping is a general term used to describe methods that utilize LDR images for HDR image applications [Banterle et al. 2006]. The intent of different iTMOs may vary. If it is to display standard images on HDR capable devices, maximizing the subjective quality, there is some evidence that global pixel transformations may be preferred [Masia et al. 2009]. Given widely different input materials, such methods are less likely to introduce artifacts compared to more advanced strategies. The transformation could be a linear scaling [Akyüz et al. 2007] or some non-linear function [Masia et al. 2009, 2017]. These methods modify all pixels without reconstructing any of the lost information.

A second category of iTMOs attempt to reconstruct saturated regions to mimic a true HDR image. These are expected to generate results that look more like a reference HDR, which was also indicated by the pair-wise comparison experiment on an HDR display performed by Banterle et al. [2009]. Meylan et al. [2006] used a linear transformation, but applied different scalings in highlight regions. Banterle et al. [2006] first linearized the input image, followed by boosting highlights using an expand map derived from the median cut algorithm. The method was extended for video processing, and with additional features such as automatic calibration and cross-bilateral filtering of the expand map [Banterle et al. 2008]. Rempel et al. [2007] also utilized an expand map, but computed this from Gaussian filtering in order to achieve real-time performance. Wang et al. [2007] applied inpainting techniques on the reflectance component of highlights. The method is limited to textured highlights, and requires some manual interaction. Another semi-manual method was proposed by Didyk et al. [2008], separating the image into diffuse, reflections and light sources. The reflections and light sources were enhanced, while the diffuse component was left unmodified. More recent methods includes the iTMO by Kovaleski and Oliveira [2014], that focus on achieving good results over a wide range of exposures, making use of a cross-bilateral expand map [Kovaleski and Oliveira 2009].

For an in-depth overview of inverse tone-mapping we refer to the survey by Banterle et al. [2009]. Compared to the existing iTMOs, our approach achieves significantly better results by learning from exploring a wide range of different HDR scenes. Furthermore, the reconstruction is completely automatic with no user parameters and runs within a second on modern hardware.

2.3 Bit-depth extension

A standard 8-bit LDR image is affected not only by clipping but also by quantization. If the contrast or exposure is significantly increased, quantization can be revealed as banding artifacts. Existing methods for decontouring, or bit-depth extension, include dithering methods that use noise in order to hide the banding artifacts [Daly and Feng 2003]. Decontouring can also be performed using low-pass filtering followed by quantization, in order to detect false contours [Daly and Feng 2004]. There are also a number of edge-preserving filters used for the same purpose. In this work we do not focus on decontouring, which is mostly a problem in under-exposed images. Also, since we treat the problem of predicting saturated image regions, the bit depth will be increased with the reconstructed information.

2.4 Convolutional neural networks

CNNs have recently been applied to a large range of computer vision tasks, significantly improving on the performance of classical supervised tasks such as image classification [Simonyan and Zisserman 2014], object detection [Ren et al. 2015] and semantic segmentation [Long et al. 2015], among others. Recently CNNs have also shown great promise for image reconstruction problems related to the challenges faced in inverse tone-mapping, such as compression artifact reduction [Svoboda et al. 2016], super-resolution [Ledig et al. 2016], and colorization [Iizuka et al. 2016]. Recent work on inpainting [Pathak et al. 2016; Yang et al. 2016] have also utilized variants of Generative Adversarial Networks (GANs) [Goodfellow et al. 2014] to produce visually convincing results. However, as these methods are based on adversarial training, results tend to be unpredictable and can vary widely from one training iteration to the next. To stabilize training, several tricks are used in practice, including restricting the output intensity, which is problematic for HDR generation. Furthermore, these methods are limited to a single image resolution, with results only shown so far for very low resolutions.

Recently, deep learning has also been successfully applied for improving classical HDR video reconstruction from multiple exposures captured over time [Kalantari and Ramamoorthi 2017]. In terms of reconstructing HDR from one single exposed LDR image, the recent work by Zhang and Lalonde [2017] is most similar to ours. They use an autoencoder [Hinton and Salakhutdinov 2006] in order to reconstruct HDR panoramas from single exposed LDR counterparts. However, the objective of this work is specifically to recover high intensities near the sun in order to use the prediction for IBL. Also, the method is only trained using rendered panoramas of outdoor environments where the sun is assumed to be in the same azimuthal position in all images. Given these restrictions, and that predictions are limited to 128×64 pixels, the results are only applicable for IBL of outdoor scenes. Compared to this work, we propose a solution to a very general problem specification without any such assumptions, and where any types of saturated regions are considered. We also introduce several key modifications to the standard autoencoder design [Hinton and Salakhutdinov 2006], and show that this significantly improves the performance.

Finally, it should be mentioned that the concurrent work by Endo et al. [2017] also treats inverse tone-mapping using deep learning algorithms, by using a different pipeline design. Given a single




Fig. 2. Zoom-in of an example of the components of the blending operation in Equation 1, compared to the ground truth HDR image. (a) is the input image, (b) is prediction, (c) is the blending mask, (d) is the blending of (a-b) using (c), and (e) is ground truth. Gamma correction has been applied to the images, for display purpose.

exposure input image, the method uses autoencoders in order to predict a set of LDR images with both shorter and longer exposures. These are subsequently combined using standard methods, in order to reconstruct the final HDR image.

3 HDR RECONSTRUCTION MODEL

3.1 Problem formulation and constraints

Our objective is to predict values of saturated pixels given an LDR image produced by any type of camera. In order to produce the final HDR image, the predicted pixels are combined with the linearized input image. The final HDR reconstructed pixel $\hat{H}_{i,c}$ with spatial index i and color channel c is computed using a pixel-wise blending with the blend value α_i ,

$$\hat{H}_{i,c} = (1 - \alpha_i) f^{-1}(D_{i,c}) + \alpha_i \exp(\hat{y}_{i,c}), \quad (1)$$

where $D_{i,c}$ is the input LDR image pixel and $\hat{y}_{i,c}$ is the CNN output (in the log domain). The inverse camera curve f^{-1} is used to transform the input to the linear domain. The blending is a linear ramp starting from pixel values at a threshold τ , and ending at the maximum pixel value,

$$\alpha_i = \frac{\max(0, \max_c(D_{i,c}) - \tau)}{1 - \tau}. \quad (2)$$

In all examples we use $\tau = 0.95$, where the input is defined to be in the range $[0, 1]$. The linear blending prevents banding artifacts between predicted highlights and their surroundings, as compared to a binary mask. It is also used to define the loss function in the training, as described in Section 3.4. For an illustration of the components of the blending, see Figure 2. Due to the blending predictions are focused on reconstructing around the saturated areas, and artifacts may appear in other image regions (Figure 2(b)).

The blending means that the input image is kept unmodified in the non-saturated regions, and linearization has to be made from either knowledge of the specific camera used or by assuming a

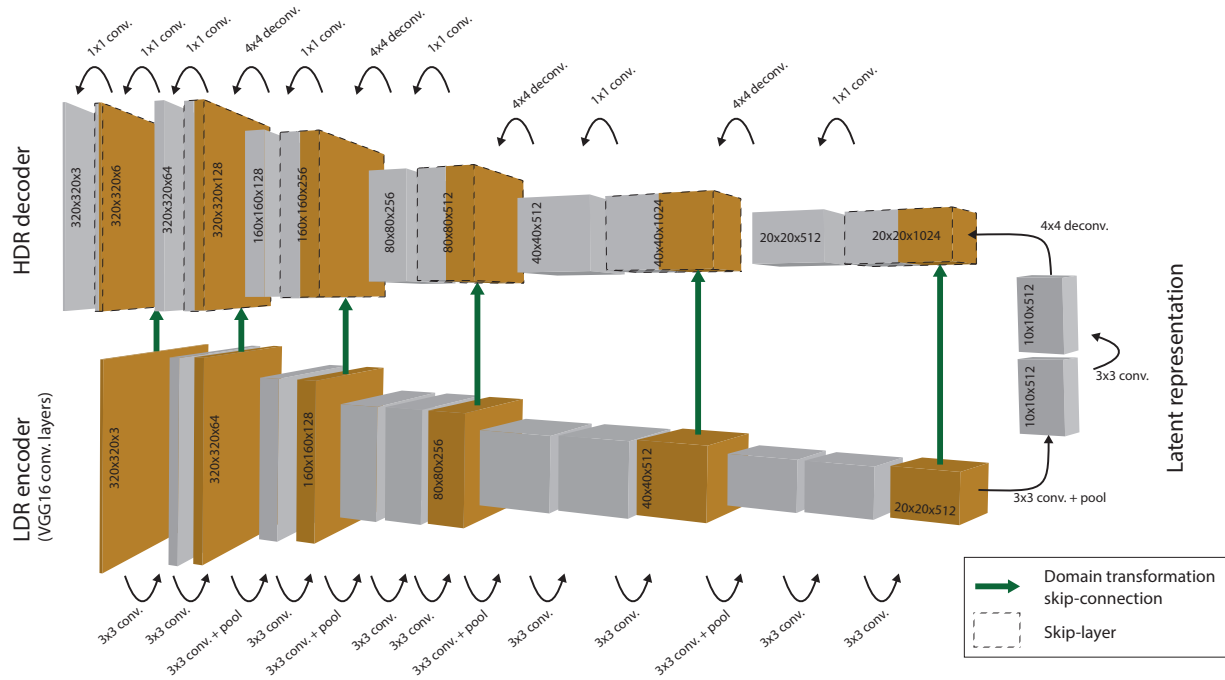


Fig. 3. Fully convolutional deep hybrid dynamic range autoencoder network, used for HDR reconstruction. The encoder converts an LDR input to a latent feature representation, and the decoder reconstructs this into an HDR image in the log domain. The skip-connections include a domain transformation from LDR display values to logarithmic HDR, and the fusion of the skip-layers is initialized to perform an addition. The network is pre-trained on a subset of the Places database, and deconvolutions are initialized to perform bilinear upsampling. While the specified spatial resolutions are given for a 320×320 pixels input image, which is used in the training, the network is not restricted to a fixed image size.

certain camera curve f . We do not attempt to perform linearization or color correction with the CNN. Furthermore, information lost due to quantization is not recovered. We consider these problems separate for the following reasons:

- (1) **Linearization:** The most general approach would be to linearize either within the network or by learning the weights of a parametric camera curve. We experimented with both these approaches, but found them to be too problematic given any input image. Many images contain too little information in order to evaluate an accurate camera curve, resulting in high variance in the estimation. On average a carefully chosen assumed transformation performs better.
- (2) **Color correction:** The same reasoning applies to color correction. Also, this would require all training data to be properly color graded, which is not the case. This means that given a certain white balancing transformation of the input, the saturated regions are predicted within this transformed color space.
- (3) **Quantization recovery:** Information lost due to quantization can potentially be reconstructed from a CNN. However, this problem is more closely related to super-resolution and compression artifact reduction, for which deep learning techniques have been successfully applied [Dong et al. 2015; Ledig et al. 2016; Svoboda et al. 2016]. Furthermore, a number of filtering techniques can reduce banding artifacts due to quantization [Bhagavathy et al. 2007; Daly and Feng 2004].

Although we only consider the problem of reconstructing saturated image regions, we argue that this is the far most important part when transforming LDR images to HDR, and that it can be used to cover a wide range of situations. Typical camera sensors can capture between 8 and 12 stops of dynamic range, which is often sufficient to register all textured areas. However, many scenes contain a small number of pixels that are very bright and thus saturated. These can be reconstructed with the proposed method, instead of capturing multiple exposures or using dedicated HDR cameras. Our method is not intended to recover the lower end of the dynamic range, which is below the noise floor of a sensor. Instead, the problem of underexposed areas is best addressed by increasing exposure time or gain (ISO). This will result in more saturated pixels, which then can be recovered using our approach.

3.2 Hybrid dynamic range autoencoder

Autoencoder architectures transform the input to a low-dimensional latent representation, and a decoder is trained to reconstruct the full-dimensional data [Hinton and Salakhutdinov 2006]. A denoising autoencoder is trained with a corrupted input, with the objective of reconstructing the original uncorrupted data [Vincent et al. 2008]. This is achieved by mapping to a higher level representation that is invariant to the specific corruption. We use the same concept for reconstruction of HDR images. In this case the corruption is clipped highlights, and the encoder maps the LDR to a representation that can be used by the decoder for HDR reconstruction. This means that

the encoder and decoder work in different domains of pixel values, and we design them to optimally account for this. Since our objective is to reconstruct larger images than is practical to use in training, the latent representation is not a fully connected layer, but a low-resolution multi-channel image. Such a fully convolutional network (FCN) enables predictions at any resolution that is a multiple of the autoencoder downscaling factor.

The complete autoencoder design is depicted in Figure 3. Convolutional layers followed by max-pooling encodes the input LDR in a $\frac{W}{32} \times \frac{H}{32} \times 512$ latent image representation, where W and H are the image width and height, respectively. The encoder layers correspond to the well-known VGG16 network [Simonyan and Zisserman 2014], but without the fully connected layers.

While the encoder operates directly on the LDR input image, the decoder is responsible for producing HDR data. For this reason the decoder operates in the log domain. This is accomplished using a loss function that compares the output of the network to the log of the ground truth HDR image, as explained in Section 3.4. For the image upsampling, we use deconvolutional layers with a spatial resolution of 4×4 initialized to perform bilinear upsampling [Long et al. 2015]. While nearest neighbor up-sampling followed by convolution has been shown to alleviate artifacts that are common in decoder deconvolutions [Odena et al. 2016], we have not experienced such problems, and instead use the more general deconvolutional layers. All layers of the network use ReLU activation functions, and after each layer of the decoder a batch normalization layer [Ioffe and Szegedy 2015] is used.

3.3 Domain transformation and skip-connections

The encoding of the input image means that much of the high resolution information in earlier layers of the encoder are lost. The information could potentially be used by the decoder to aid reconstruction of high frequency details in saturated regions. Thus, we introduce skip-connections that transfer data between both high and low level features in the encoder and decoder.

Skip-connections have been shown to be useful for constructing deeper network architectures which improve performance in a variety of tasks [He et al. 2016]. For autoencoders, where layers have different spatial resolution, a separate residual stream can be maintained in full resolution, with connections to each layer within the autoencoder [Pohlen et al. 2017]. Alternatively, skip-connections between layers of equal resolution in encoder and decoder have also been shown to boost performance in a variety of imaging tasks using autoencoders [Ronneberger et al. 2015; Zhang and Lalonde 2017].

Our autoencoder uses skip-connections to transfer each level of the encoder to the corresponding level on the decoder side. Since the encoder and decoder process different types of data (see Section 3.2), the connections include a domain transformation described by an inverse camera curve and a log transformation, mapping LDR display values to a logarithmic HDR representation. Since the camera curve is unknown, we have to assume its shape. Although a sigmoid function fits well with camera curves in general [Grossberg and Nayar 2003], its inverse is not continuous over \mathbb{R}^+ . The linearization of the skip-connections is therefore done using a gamma function $f^{-1}(x) = x^\gamma$, where $\gamma = 2$.

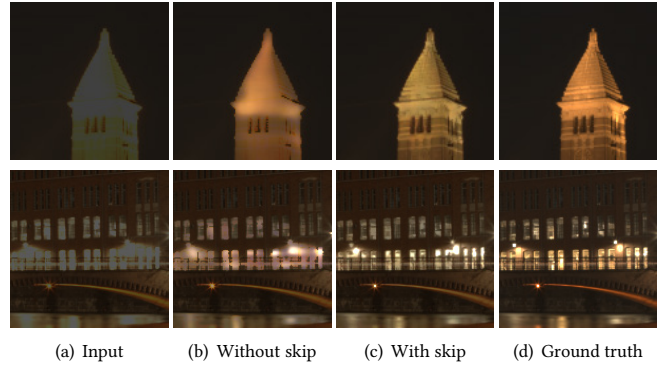


Fig. 4. Zoom-ins of reconstruction without (b) and with (c) the domain transformation skip-connections. The plain autoencoder architecture can reconstruct high luminance, but without skip-connections the detail information around saturated regions cannot be fully exploited.

A skip-connected layer is typically added to the output of the layer at the connection point. However, to allow for additional freedom, we concatenate the two layers along the feature dimension. That is, given two $W \times H \times K$ dimensional layers, the concatenated layer is $W \times H \times 2K$. The decoder then makes a linear combination of these, that reduces the number of features back to K . This is equivalent to using a convolutional layer with a filter size of 1×1 , where the number of input and output channels are $2K$ and K , respectively, as depicted in Figure 3. More specifically, the complete LDR to HDR skip connection is defined as

$$\tilde{h}_i^D = \sigma \left(\mathbf{W} \left[\log \left(f^{-1} \left(\mathbf{h}_i^E \right) + \epsilon \right) \right] + \mathbf{b} \right). \quad (3)$$

The vectors \mathbf{h}_i^E and \mathbf{h}_i^D denote the slices across all the feature channels $k \in \{1, \dots, K\}$ of the encoder and decoder layer tensors $\mathbf{y}^E, \mathbf{y}^D \in \mathbb{R}^{W \times H \times K}$, for one specific pixel i . Furthermore, \tilde{h}_i^D is the decoder feature vector with information fused from the skip-connected vector \mathbf{h}_i^E . \mathbf{b} is the bias of the feature fusion, and σ is the activation function, in our case the rectified linear unit (ReLU). A small constant ϵ is used in the domain transformation in order to avoid zero values in the log transform. Given K features, \mathbf{h}^E and \mathbf{h}^D are $1 \times K$ vectors, and \mathbf{W} is a $2K \times K$ weight matrix, which maps the $2K$ concatenated features to K dimensions. This is initialized to perform an addition of encoder and decoder features, setting the weights as

$$\mathbf{W}_0 = \begin{bmatrix} 1 & 0 & \dots & 0 & 1 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 & 0 & 1 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 1 & 0 & 0 & \dots & 1 \end{bmatrix}, \quad \mathbf{b}_0 = \mathbf{0}. \quad (4)$$

During training, these weights can be optimized to improve the performance of the skip-connection. Since the linear combination of features is performed in the log domain, it corresponds to multiplications of linear HDR data. This is an important characteristic of the domain transformation skip-connections as compared to existing skip architectures.




Fig. 5. Zoom-ins of reconstructions with different relative weight of illuminance and reflectance, λ in Equation 7. A higher weight of illuminance will in general better predict high intensity regions (b), while a higher weight of reflectance is better at deducing local colors and details (c).

An example of the impact of the described skip-connection architecture is shown in Figure 4. The autoencoder design is able to reconstruct HDR information from an encoded LDR image. However, all information needed by the decoder has to travel through the intermediate encoded representation. Adding the skip-connections enables a more optimal use of existing details.

3.4 HDR loss function

A cost function formulated directly on linear HDR values will be heavily influenced by high luminance values, leading to underestimation of important differences in the lower range of luminances. The few existing deep learning systems that predict HDR have treated this problem by defining the objective function in terms of tone-mapped luminance [Kalantari and Ramamoorthi 2017; Zhang and Lalonde 2017]. In our system the HDR decoder is instead designed to operate in the log domain. Thus, the loss is formulated directly on logarithmic HDR values, given the predicted log HDR image $\hat{\mathbf{y}}$ and the linear ground truth \mathbf{H} ,

$$\mathcal{L}(\hat{\mathbf{y}}, \mathbf{H}) = \frac{1}{3N} \sum_{i,c} |\alpha_i (\hat{y}_{i,c} - \log(H_{i,c} + \epsilon))|^2, \quad (5)$$

where N is the number of pixels. Since $H_{i,c} \in \mathbb{R}^+$, the small constant ϵ removes the singularity at zero pixel values. The cost formulation is perceptually motivated by the close to logarithmic response of the human visual system (HVS) in large areas of the luminance range, according to the Weber-Fechner law [Fechner 1965]. The law implies a logarithmic relationship between physical luminance and the perceived brightness. Thus, a loss formulated in the log domain makes perceived errors spread approximately uniformly across the luminance range.

As described in Section 3.1, we use only the information from the predicted HDR image $\hat{\mathbf{y}}$ around saturated areas. This is also reflected by the loss function in Equation 5 where the blend map α from Equation 2 is used to spatially weight the error.

Treating the illuminance and reflectance components separately makes sense from a perceptual standpoint, as the visual system may indirectly perform such separation when inferring reflectance or




Fig. 6. Zoom-in of a reconstruction with different loss functions. The input (a) is exposure corrected and clipped to have a large amount of information lost. The direct pixel loss (b) is more prone to generating artifacts as compared to the illuminance + reflectance loss (c).

discounting illumination [Gilchrist and Jacobsen 1984]. We therefore also propose another, more flexible loss function that treats illuminance and reflectance separately. The illuminance component I describes the global variations, and is responsible for the high dynamic range. The reflectance R stores information about details and colors. This is of lower dynamic range and modulates the illuminance to create the final HDR image, $H_{i,c} = I_i R_{i,c}$. We approximate the log illuminance by means of a Gaussian low-pass filter G_σ on the log luminance $L\hat{\mathbf{y}}$,

$$\begin{aligned} \log(I_i^{\hat{\mathbf{y}}}) &= (G_\sigma * L\hat{\mathbf{y}})_i, \\ \log(R_{i,c}^{\hat{\mathbf{y}}}) &= \hat{y}_{i,c} - \log(I_i^{\hat{\mathbf{y}}}). \end{aligned} \quad (6)$$

Since the estimation is performed in the log domain, the log reflectance is the difference between $\hat{\mathbf{y}}$ and log illuminance. $L\hat{\mathbf{y}}$ is a linear combination of the color channels, $L_i^{\hat{\mathbf{y}}} = \log(\sum_c w_c \exp(\hat{y}_{i,c}))$, where $w = \{0.213, 0.715, 0.072\}$. The standard deviation of the Gaussian filter is set to $\sigma = 2$. The resulting loss function using I and R is defined as

$$\begin{aligned} \mathcal{L}_{IR}(\hat{\mathbf{y}}, \mathbf{H}) &= \frac{\lambda}{N} \sum_i |\alpha_i (\log(I_i^{\hat{\mathbf{y}}}) - \log(I_i^{\mathbf{y}}))|^2 \\ &+ \frac{1-\lambda}{3N} \sum_{i,c} |\alpha_i (\log(R_{i,c}^{\hat{\mathbf{y}}}) - \log(R_{i,c}^{\mathbf{y}}))|^2, \end{aligned} \quad (7)$$

where $\mathbf{y} = \log(\mathbf{H} + \epsilon)$ to simplify notation. The user-specified parameter λ can be tuned for assigning different importance to the illuminance and reflectance components. If not stated otherwise, we use the illuminance + reflectance (I/R) loss with $\lambda = 0.5$ for all results in this paper. This puts more importance on the illuminance since the error in this component generally is larger. Figure 5 shows examples of predictions where the optimization has been performed with different values of λ . With more relative weight on illuminance, high intensity areas are in general better predicted, which *e.g.* could benefit IBL applications. If the reflectance is given more importance, local colors and details can be recovered with higher robustness, for better quality *e.g.* in post-processing applications.

The visual improvements from using the I/R loss compared to the direct loss in Equation 5 are subtle. However, in general it tends to produce less artifacts in large saturated areas, as exemplified in Figure 6. One possible explanation is that the Gaussian low-pass filter in the loss function could have a regularizing effect, since it makes the loss in a pixel influenced by its neighborhood. This




Fig. 7. Histograms over two LDR datasets and the pre-processed HDR dataset. The creation of the HDR data is described in Section 4. For the LDR data, the probabilities show a large increase close to 1, indicating saturated information. The HDR dataset contains such information, represented by the tail of decreasing frequency.

observation is further supported by the comparison in Table 1, where the I/R loss lowers the final error in Equation 5 by more than 5%, demonstrating better generalization to the test data.

4 HDR IMAGE DATASET

A key challenge for a learning-based HDR reconstruction is to obtain a sufficiently large set of well structured training data. However, an increasing amount of HDR content has become available, in particular through recent HDR video datasets [Azimi et al. 2014; Boitard et al. 2014; Froehlich et al. 2014; Kronander et al. 2014]. We were able to gather a total of 1121 HDR images and 67 HDR video sequences. The sources of the data are specified in the supplementary document. 4 video sequences and 95 images are separated from these to create the test set used throughout this paper, and the rest are used for training. Since consecutive frames from a video sequence are expected to be very similar, we use every 10th frame. Together with the static images, this results in a total of ~ 3700 HDR images. Using a virtual camera, a carefully designed set of data augmentation operations are then applied in order to improve robustness of the predictions.

Considering each HDR image a real-world scene, we set up a virtual camera that captures a number of random regions of the scene using a randomly selected camera calibration. This provides us with an augmented set of LDR and corresponding HDR images that are used as input and ground truth for training, respectively. The regions are selected as image crops with random size and position, followed by random flipping and resampling to 320×320 pixels. The camera calibration incorporates parameters for exposure, camera curve, white balance and noise level. These are randomly selected, with the camera curve defined as a parametric function fitted to the database of camera curves collected by Grossberg and Nayar [2003]. For details on the augmentation, we refer to Appendix A.

In total we capture $\sim 125K$ training samples from the HDR dataset using the virtual camera. This augmentation is responsible for creating a final trained model that generalizes well to a wide range of images captured with different cameras.



Fig. 8. A 150×150 pixels zoom-in of a reconstruction. Using pre-training the CNN is in general more consistent and can reconstruct smaller highlights better.

4.1 Image statistics

It is important to note that the dynamic range statistics of LDR and HDR images differ considerably. Figure 7 shows averaged histograms over two typical LDR datasets, as well as our HDR dataset of 125K images. The LDR data are composed of around 2.5M and 200K images for Places [Zhou et al. 2014] and Flickr, respectively. Inspecting the LDR histograms, they show a relatively uniform distribution of pixel values, except for distinct peaks near the maximum value representing information lost due to saturation. In the HDR histogram on the other hand, pixels are not saturated, and are instead represented by an exponentially decaying long tail. Although there are not many pixels with extreme intensities, these are very important to properly learn a plausible HDR image.

5 TRAINING

To initialize the weights in the network we use different strategies for different parts of the network. As we use the convolutional layers from the well-known VGG16 network, we can use pre-trained weights available for large scale image classification on the Places database [Zhou et al. 2014] to initialize the encoder. The decoder deconvolutions are initiated for bilinear upsampling, and the fusions of skip-connection layers are initiated to perform addition of features (Equation 4). For the convolutions within the latent image representation (right-most of Figure 3) and the final feature reduction (top-left in Figure 3) we use Xavier initialization [Glorot and Bengio 2010].

Minimization is performed with the ADAM optimizer [Kingma and Ba 2014], with a learning rate of 5×10^{-5} , on the loss function in Equation 7. In total 800K steps of back-propagation are performed, with a mini-batch size of 8, taking approximately 6 days on an Nvidia Titan X GPU.

5.1 Pre-training on simulated HDR data

As we have a limited amount of HDR data at hand, we use transfer learning by pre-training the entire network on a large simulated HDR dataset. To this end, we select a subset of the images in the Places database [Zhou et al. 2014], requiring that the images should not contain saturated image regions. Given the set of all Places images \mathbb{P} , this subset $\mathbb{S} \subset \mathbb{P}$ is defined as

$$\mathbb{S} = \{D \mid D \in \mathbb{P}, p_D(255) < \xi\}, \quad (8)$$

where p_D is the image histogram. For the threshold we use $\xi = 50/256^2$. Thus, if less than 50 pixels (0.076% of the 256^2 pixels of an image) have maximum value, we use this in the training set. For



Fig. 9. The input images (a) have been exposure corrected, followed by camera transformation, quantization and clipping. 5% of the pixels are saturated and contain no information. Visually convincing reconstructions (b) can be made in a wide range of situations. The reconstructions correspond well to the ground truth HDR images (c). The exposures of the images have been reduced to show the differences, and all images have been gamma corrected for display.

the Places database this gives $\sim 600K$ images of the total set size of $\sim 2.5M$. The averaged histogram over the subset \mathbb{S} , plotted in Figure 7, does not show the peak of saturated pixels as the original set \mathbb{P} . By linearizing the images $D \in \mathbb{S}$ using the inverse of the camera curve f in Equation 10 and increasing the exposure, $H = sf^{-1}(D)$, we create a simulated HDR training dataset.

The simulated HDR dataset is prepared in the same manner as in Section 4, but at 224×224 pixels resolution and without resampling. The CNN is trained using the ADAM optimizer with learning rate 2×10^{-5} for 3.2M steps with a batch size of 4.

The result of this pre-training on synthetic data leads to a significant improvement in performance. Small highlights, which sometimes are underestimated, are better recovered, as illustrated in Figure 8, and less artifacts are introduced in larger saturated regions. Table 1 shows that the error is reduced by more than 10%.

6 RESULTS

In this section we present a number of examples, verifying the quality of the HDR reconstructions. Additional visual examples can be found in the supplementary material and video. Furthermore, for prediction using any LDR image the CNN together with trained parameters can be downloaded from: <https://github.com/gabrieleilertsen/hdrcnn>.

6.1 Test errors

To justify the different model and training strategies explained in Section 3 and 5, we evaluate the success of the optimization in Table 1. The errors of the different configurations have been averaged over the test set of 95 HDR images, reconstructed at 1024×768 pixels resolution. The LDR images used for reconstruction use virtual

exposures and clipping such that 5% of the pixels in each image are saturated.

Table 1 shows how different errors are affected by the different training strategies. The CNN without skip-connections can drastically reduce the MSE of the input. However, adding the skip-connections reduces error by an additional 24%, and creates images with substantial improvements in details (Figure 4). Comparing the two different loss functions in Equation 5 and 7, the latter I/R loss shows a lower error both in terms of I/R and direct MSE, with a reduction of 5.8%. Finally, with pre-training and the I/R loss the best training performance is accomplished, lowering the error by 10.7% as compared to no pre-training. All the trainings/pre-trainings that do not use our pre-trained parameters have been initialized

Table 1. Different MSEs evaluated over the test set. Rows show different training strategies, while columns evaluate with different errors. The direct MSE is from Equation 5, while the I/R, I and R MSEs use Equation 7 with $\lambda = 0.5, 1$ and 0, respectively. The reference is the input image without reconstruction. Adding skip-connections improves the result to a large extent. The illuminance + reflectance loss lowers both direct MSE and in terms of illuminance and reflectance, as compared to optimizing for only the direct loss. Pre-training has a significant impact on the result.

	Direct	I/R	I	R
Reference	0.999	0.890	0.712	0.178
Without skip-conn.	0.249	0.204	0.102	0.102
Direct loss (eq. 5)	0.189	0.159	0.090	0.069
I/R loss (eq. 7)	0.178	0.150	0.081	0.068
Pre-train. + I/R loss	0.159	0.134	0.069	0.066




Fig. 10. The learned residual of the image in Figure 1 (top left), together with the ground truth residual (top right). The plots show relative luminance values across the two scanlines marked in the images. Complex regions are predicted convincingly (bottom left), but for the very intense spotlight the signal is underestimated (bottom right). The underestimation is less pronounced when training with higher illuminance weight (λ in Equation 7).

with VGG16 encoder weights trained for classification of the Places dataset [Zhou et al. 2014].

6.2 Comparisons to ground truth

Figure 9 demonstrates a set of predictions on HDR images from the test set that have been transformed to LDR by the virtual camera described in Section 4. The examples demonstrate successful HDR reconstruction in a variety of situations. In night scenes, colors and intensities of street lights and illuminated facades can be restored with very convincing quality. The same goes for specular reflections and other smaller highlights in day scenes. Furthermore, in situations where there is some small amount of information left in any of the color channels, details and colors of larger areas can be reconstructed to be very close to the ground truth HDR image. For example, in the right-most column the large area light source does not contain any visual details when inspecting the input image. However, some small amount of invisible information enables recovery of details. Also, while all channels are saturated in the sky in the third column, the spatial context can be utilized to infer a blue color.

In order to visualize the information that is reconstructed by the CNN, Figure 10 shows the residual, $\hat{r} = \max(0, \hat{H} - 1)$. That is, only the information in highlights is shown, which is not present in the input image $D \in [0, 1]$. The information corresponds well with the ground truth residual, $r = \max(0, H - 1)$. The complete input and output signals are also plotted for two different scanlines across



Fig. 11. Reconstruction from Canon 5DS R camera JPEG (top). The examples are the same as in Figure 9, but with the camera set to store JPEG, applying unknown image transformations and compression. Reconstruction can also be performed with hand-held low-end cameras, such as the iPhone 6S (bottom), where the input image is more degraded.

the images. Complex lighting of street lights and windows can be recovered convincingly (bottom left). However, in some situations of very intense light sources, the luminance is underestimated (bottom right). We elaborate on such limitations in Section 8.

6.3 Reconstruction with real-world cameras

In order to show that the HDR reconstruction model generalizes to real-world cameras, Figure 11 shows two of the scenes from Figure 9, captured using a Canon 5DS R camera's JPEG mode (top row) and with an iPhone 6S camera (bottom row). Both these cameras provide more realistic scenarios as compared to the virtual camera. Nevertheless, reconstructions of equal quality can be done from camera JPEGs. The iPhone images are more degraded, shot in dark conditions without a tripod, but the reconstructed information comply well with the image characteristics. To further explore the possibilities in reconstructing everyday images, Figure 12 displays a set of iPhone images, taken in a variety of situations. The examples not only demonstrate the method's ability to generalize to a different camera, but also to a wide range of situations, such as skin tones, fire and caustics. The most limiting factor when reconstructing from iPhone images is the hard JPEG compression that



Fig. 12. Predictions on iPhone camera images. Plausible reconstructions can be made of skin tones, caustics and fire (row 1-3). Large saturated areas can be recovered if there still is information in one of the color channels (bottom row).

has been applied, which results in small blocking artifacts close to highlights, affecting the final reconstruction. In order to compensate for these, the brightness of the images has been increased by a small factor, followed by clipping. This removes the artifacts caused by harder compression in the brightest pixels, which improves the final reconstruction quality.

6.4 Changing clipping point

In order to demonstrate the behavior of the reconstruction with varying amount of information loss, Figure 13 shows predictions using different virtual exposure times. As the training of the CNN uses a virtual camera with different exposure settings, part of the objective is to minimize the difference between these, apart from a scaling factor. However, since more information is available in highlights with shorter exposure, in most cases there will be visible differences in the reconstruction quality, as exemplified by the figure.

6.5 Comparison to iTMOs

Figure 14 shows the HDR reconstruction compared to three existing methods for inverse tone-mapping. These are examples of local methods that apply different processing in saturated areas in order to boost the dynamic range of an LDR image. The results can successfully convey an impression of HDR when viewed on an HDR

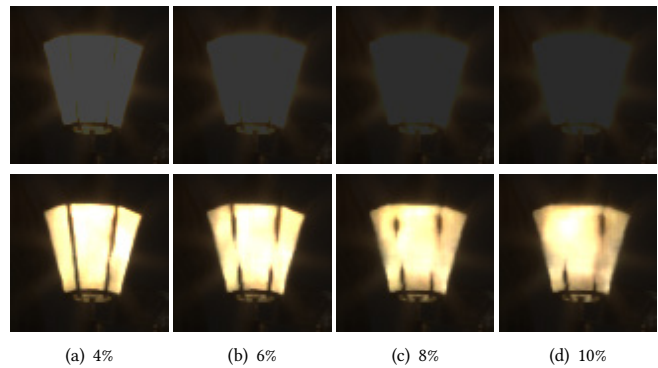


Fig. 13. Zoom-ins of reconstructions (bottom row) with different exposure settings of the input (top row). The numbers indicate how large fraction of the total number of pixels are saturated in the input. The images have then been scaled to have the same exposure after clipping has been applied. Although illuminance is predicted at approximately the same level, more details are available in reconstruction of the shorter exposure images.




Fig. 14. Comparison to some existing iTMOs. Since the iTMO results usually are calibrated for an HDR display, they have been scaled to the same range for comparison. Although they can deliver an impression of increased dynamic range by boosting highlights, when inspecting the saturated image regions little information have actually been reconstructed. The CNN we use can make a prediction that is significantly closer to the true HDR image.

capable display. However, when inspecting the highlights, local information is not recovered by a naïve scaling of image highlights. With our CNN we are able to predict saturated regions based on a high level understanding of the context around the area, which makes it possible to actually recover convincing colors and structural content.

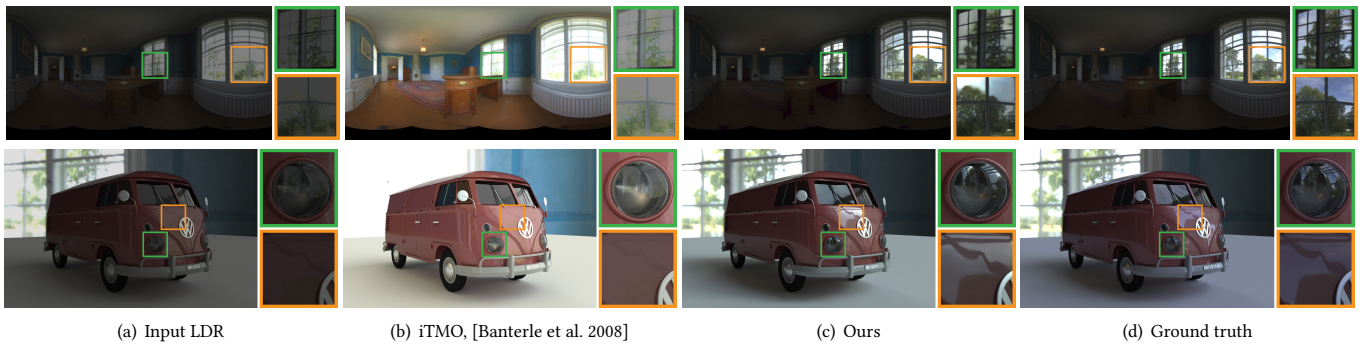


Fig. 15. IBL using reconstructed highlights. The top row shows the panoramas that are used for IBL in the bottom row. Rendering with the LDR input gives a dark and undynamic result. The iTMO boosts brightness to alleviate the problems. With our reconstruction, although all details cannot be recovered in the large saturated areas of the windows, the estimated luminance enables a visually convincing rendering that is much closer to the ground truth.

6.6 Image based lighting

An important application of HDR imaging is image based lighting (IBL), where the omnidirectional lighting of a particular scene is captured in a panorama and used for re-lighting of a computer graphics model. For this task it is of major importance that the entire range of luminances are present within the image, in order to convey a faithful and dynamic rendering. Figure 15 shows a panorama from an indoor scene, where the majority of illuminance is due to two windows. In the LDR image the windows are saturated, and the result when used for IBL is overall dark and of low contrast. The iTMO by Banterle *et al.* [2008] can accomplish a more dynamic rendering by boosting image highlights, although the result is far from the ground truth. With our learning-based method, the reconstructed panorama shows some loss of details in the large saturated areas of the windows, but the illuminance is estimated convincingly. This makes it possible to render a result that is very close to the ground truth.

7 EVALUATION

In order to assess the perceived visual quality of our novel HDR reconstruction, we performed a subjective pairwise comparison experiment. 15 participants took part in the experiment, aged 19 – 40 with normal or corrected-to-normal full color vision.

7.1 Setup

We used a calibrated projector-based HDR display, similar to the one described by Seetzen *et al.* [2004], at a viewing distance of 80 *cm*. The display consisted of an Acer P1276 1024 × 768 DLP projector with removed color wheel, and a 9.7" 2048 × 1536 iPad Retina display panel with removed backlight. The maximal and minimal luminance of the display was 5000 cd/m^2 and 0.1 cd/m^2 , yielding a maximal contrast range of 50 000:1.

7.2 Stimuli

Each stimulus consisted of a pair of images identical in terms of content, but reproducing one of the four processing or display scenarios: (1) LDR image with its dynamic range clamped to that of a typical DSLR camera (10.5 stops); (2) ground truth HDR image; (3) iTMO technique by Banterle *et al.* [2008], shown to perform the best in the

evaluation [Banterle *et al.* 2009]; (4) output from our CNN-based HDR reconstruction (Pre-train + I/R loss). To prevent overall image brightness from affecting the outcome of the experiment, we fixed the luminance of the 90th percentile pixel value for all methods to 180 cd/m^2 . To avoid bias in the selection of images, we used a randomly selected sample of 25 images from a pool of the 95 images in the test set.

7.3 Task

We used a two-alternative forced choice experiment, where in each trial the participants were shown a pair of images side-by-side. The task was to select the image that looked more *natural*. Our definition of natural involved “the depth and vividness that bears most resemblance to a view that you could experience looking through a window”. Participants were given unlimited time to look at the images and to make their decisions. Before each session, participants were briefed about their task both verbally and in writing, followed by a short training session to gain familiarity with the display equipment and the task. We used a full pairwise comparison design, in which all pairs were compared. Each observer compared each pair three times, resulting in 450 trials per observer. The order of the stimuli as well as their placement on the screen were randomized.

7.4 Results

The result of the pairwise comparison experiment is scaled in just-objectionable-differences (JODs), which quantify the relative quality differences. The scaling is performed using publicly available software¹, which formulates the problem as a Bayesian inference under the Thurstone Case V assumptions, and uses a maximum-likelihood-estimator to find relative JOD values. Unlike standard scaling procedures, the Bayesian approach robustly scales pairs of conditions for which there is unanimous agreement. Since JOD values are relative, we fix the starting point of the JOD scale at 0 for the LDR images. When two points are a single unit apart on the JOD space, approximately 75% of the population are expected to perceive an objectionable quality difference.

The results of the subjective quality experiment are shown in Figure 16. Unexpectedly, the iTMO technique by Banterle *et al.* was

¹Pairwise comparison scaling software for Matlab: <https://github.com/mantiuk/pwcmp>




Fig. 16. Results of the subjective quality experiment. The error bars represent 95% confidence intervals computed by bootstrapping. All values on the JOD scale are relative to the LDR images. Negative scores indicate a lower perceived naturalness for the iTMO technique when compared with LDR images. The output of our CNN-based HDR reconstruction method surpasses LDR and is comparable to the original HDR images in most cases.

judged to be the least natural, even less so than the LDR images. This can be explained by the operator’s attempt to inverse the camera response curve, which often results in reduced contrast and inaccurate colors. Figure 17 row (a) shows the 19th image from the evaluation, where the effect of over-boosted colors can be easily observed. As expected, LDR images were rated worse than the original HDR images in almost all cases. Most participants were mainly accustomed to standard display monitors, which, as reflected upon by some subjects during the unstructured post-experiment interview, might have affected their perception of “natural”. With more exposure to HDR displays we expect the perceived quality of LDR images to drop in the future. According to the data, our CNN-based images are very likely to perform better than their original LDR counterparts. The number of times our CNN images were picked is slightly less but comparable to the original HDR images. Figure 17 rows (b) and (c) illustrate two scenes, where the algorithm succeeds in estimating the spatial texture of high luminance areas, producing plausible results. The performance of the CNN is scene-dependent, and can sometimes introduce artifacts in the highlights that affect perceived naturalness. One example is depicted in Figure 17 row (d), where the artificial texture of the sun is immediately obvious. Overall, the CNN-based reconstruction improves the subjective quality as compared to the input LDR images, which is not always the case for the state-of-the-art iTMO techniques. The reconstructed images are in most cases comparable to ground truth HDR images, as evidenced by the quality differences of less than 1 JOD.

8 CONCLUSION AND FUTURE WORK

HDR reconstruction from an arbitrary single exposed LDR image is a challenging task. To robustly solve this problem, we have presented a hybrid dynamic range autoencoder. This is designed and trained taking into account the characteristics of HDR images in the model architecture, training data and optimization procedure. The quality and versatility of the HDR reconstruction have been demonstrated through a number of examples, as well as in a subjective experiment.

8.1 Limitations

There is a content-dependent limit on how much missing information the network can handle which is generally hard to quantify. Figure 18 shows two examples of difficult scenes that are hard to reconstruct. The first row has a large region with saturation in all color channels, so that structures and details cannot be inferred. However, illuminance may still be estimated well enough in order to allow for high quality IBL, as demonstrated in Figure 15. The second row of Figure 18 shows a situation where besides a similar loss of spatial structures, extreme intensities are also underestimated. This is also demonstrated in Figure 10 with the intense spotlight. The plot also shows that the problem can be alleviated by altering the illuminance weight λ in Equation 7. However, underestimation of highlights is also an inherent problem of the training data. Some of the HDR images used to create the training data show saturated pixels in high intensity regions. For example, the sun in Figure 18 is saturated in the ground truth HDR image.

There is also a limitation on how much compression artifacts that can be present in the input image. If there are blocking artifacts around highlights, these will impair the reconstruction performance to some extent.

8.2 Future work

Recovering saturated image regions is only one of the problems in reconstructing HDR images. Another, less prominent issue is quantization artifacts, which can be alleviated using existing methods [Daly and Feng 2003]. However, we believe that bit-depth extension also can be achieved by means of deep learning. For this purpose existing architectures for compression artifact reduction [Svoboda et al. 2016] or super resolution [Ledig et al. 2016] are probably better suited.

The complementary problem of reconstructing saturated pixels, is the recovery of dark regions of an image, that have been lost due to quantization and noise. This problem is also significantly different




Fig. 17. Examples of typical images in the subjective experiment. a) *image 19* demonstrates how the inaccurate colors of iTMO reduces perceived realism. b) *image 21* with a successful HDR reconstruction. c) *image 4* shows how even an inaccurate estimate of the highlight luminance still produces plausible results. d) *image 8* is an example of unsuccessful reconstruction, producing artifacts that heavily affect the perceived naturalness.

from ours in that noise will be a main issue when increasing the exposure of an image.

Another direction for future work is to investigate how to improve reconstruction of images that are degraded by compression artifacts. The most straightforward solution would be to augment the training data with compression artifacts. However, this also runs the risk of lowering the quality for reconstructions of images without any compression applied.

Finally, although recent development in generative adversarial networks (GAN) [Goodfellow et al. 2014; Radford et al. 2015] shows promising result in a number of imaging tasks [Ledig et al. 2016; Pathak et al. 2016], they have several limitations. An important challenge for future work is to overcome these, in order to enable high-resolution and robust estimation.

ACKNOWLEDGMENTS

The authors would like to thank Francesco Banterle for the invaluable discussions and help with inverse tone-mapping operators, and the anonymous reviewers for helping in improving the manuscript. This work was funded through Linköping University Center for Industrial Information Technology (CENIIT), the Swedish Science Council through Grant 2015-05180, and the Wallenberg Autonomous Systems Program (WASP).

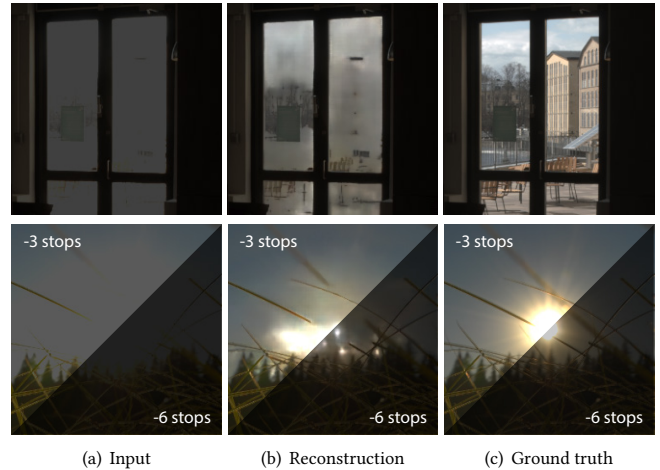


Fig. 18. Zoom-ins of areas where reconstruction fails. (Top) A large fraction of the pixels are saturated, and the structures cannot be recovered properly. (Bottom) the intensity and shape of the sun are not estimated correctly.

APPENDIX A: DATA AUGMENTATION

In this appendix we specify the details of the virtual camera used in Section 4 for augmentation of HDR images.

A.1 Random cropping

For each mega-pixel of HDR data, N sub-images are selected at random positions and sizes. For the trainings we perform, we choose $N = 10$, which results in a final training set of $\sim 125K$ images. The sizes are drawn uniformly from the range [20%, 60%] of the size of an input image, followed by bilinear resampling to 320×320 pixels. Since 320 pixels corresponds to 20%–60% of the original images, the training is optimized to account for images that are between $320/0.6 = 533$ and $320/0.2 = 1600$ pixels.

A.2 Exposure

The exposure of each cropped image is set so that clipping removes a fraction v of the image information. The fraction is uniformly drawn in the range $v \in [0.05, 0.15]$. To accomplish this, v is used to define an exposure scaling s ,

$$s = \frac{1}{H_{1-v}}, \text{ s.t. } \sum_{i=H_{min}}^{H_{th}} p_H(i) = 1 - v, \quad (9)$$

where p_H is the histogram of the HDR image H . Thus, H_{1-v} is the $1 - v$ percentile, and this defines the scaling s which is applied to the image in order to remove 5 – 15% of the information when clipping is performed (see Equation 11).

A.3 Camera curve

To approximate different camera curves we use a parametric function, in form of a sigmoid,

$$f(H_{i,c}) = (1 + \sigma) \frac{H_{i,c}^n}{H_{i,c}^n + \sigma}. \quad (10)$$

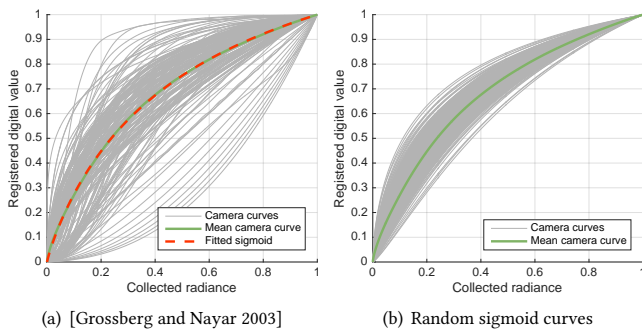


Fig. 19. A sigmoid function fits well to the mean of the collected camera curves by Grossberg and Nayar [2003] (a). We simulate camera curves from a normal distribution centered at the fitted parameters (b).

The scaling $1 + \sigma$ is used to normalize the output so that $f(1) = 1$. We fit the parameters n and σ to the mean of the database of camera curves collected by Grossberg and Nayar [2003], where $n = 0.9$ and $\sigma = 0.6$ gives a good fit as shown in Figure 19. For random selection of camera curves in the training data preparation, we draw the parameters from normal distributions around the fitted values, $n \sim \mathcal{N}(0.9, 0.1)$ and $\sigma \sim \mathcal{N}(0.6, 0.1)$. As demonstrated in Figure 19 this creates a continuous range that do not include extreme functions such as gamma curves with $\gamma > 1$.

A.4 Other

Augmentation in terms of colors is accomplished in the HSV color space, where we modify the hue and saturation channels. The hue is altered by means of adding a random perturbation $\tilde{h} \sim \mathcal{N}(0, 7)$. The same is done for the saturation, but with a narrower distribution, $\tilde{s} \sim \mathcal{N}(0, 0.1)$.

Finally, a small amount of additive Gaussian noise is injected, with a standard deviation randomly selected in the range $\sigma \in [0, 0.01]$. Also, images are flipped horizontally with a probability of 0.5. The processed linear images H represent the ground truth data, while the inputs for training are clipped at 1 and quantized,

$$D_{i,c} = \lfloor 255 \min(1, f(H_{i,c})) + 0.5 \rfloor / 255. \quad (11)$$

REFERENCES

- A. O. Akçüz, R. Fleming, B. E. Riecke, E. Reinhard, and H. H. Bühlhoff. 2007. Do HDR Displays Support LDR Content?: A Psychophysical Evaluation. *ACM Trans. Graph.* 26, 3, Article 38 (2007).
- M. Azimi, A. Banitalebi-Dehkordi, Y. Dong, M. T. Pourazad, and P. Nasiopoulos. 2014. Evaluating the Performance of Existing Full-Reference Quality Metrics on High Dynamic Range (HDR) Video Content. In *International Conference on Multimedia Signal Processing (ICMSP '14)*, Vol. 1. 789.
- F. Banterle, A. Artusi, K. Debattista, and A. Chalmers. 2011. *Advanced High Dynamic Range Imaging: Theory and Practice*. AK Peters (CRC Press).
- F. Banterle, K. Debattista, A. Artusi, S. Pattanaik, K. Myszkowski, P. Ledda, and A. Chalmers. 2009. High Dynamic Range Imaging and Low Dynamic Range Expansion for Generating HDR Content. *Computer Graphics Forum* 28, 8 (2009), 2343–2367.
- F. Banterle, P. Ledda, K. Debattista, M. Bloj, A. Artusi, and A. Chalmers. 2009. A Psychophysical Evaluation of Inverse Tone Mapping Techniques. *Computer Graphics Forum* 28, 1 (2009), 13–25.
- F. Banterle, P. Ledda, K. Debattista, and A. Chalmers. 2006. Inverse Tone Mapping. In *Proceedings of the 4th International Conference on Computer Graphics and Interactive Techniques in Australasia and Southeast Asia (GRAPHITE '06)*. ACM, 349–356.
- F. Banterle, P. Ledda, K. Debattista, and A. Chalmers. 2008. Expanding Low Dynamic Range Videos for High Dynamic Range Applications. In *Proceedings of the 24th Spring Conference on Computer Graphics (SCCG '08)*. ACM, 33–41.
- S. Bhagavathy, J. Llach, and J. f. Zhai. 2007. Multi-Scale Probabilistic Dithering for Suppressing Banding Artifacts in Digital Images. In *The IEEE International Conference on Image Processing (ICIP '07)*, Vol. 4. IV – 397–IV – 400.
- R. Boitard, R. Cozot, D. Thoreau, and K. Bouatouch. 2014. Survey of Temporal Brightness Artifacts in Video Tone Mapping. In *Second International Conference and SME Workshop on HDR imaging (HDI2014)*.
- S. J. Daly and X. Feng. 2003. Bit-depth extension using spatiotemporal microdither based on models of the equivalent input noise of the visual system. In *Proceedings of SPIE*, Vol. 5008. 455–466.
- S. J. Daly and X. Feng. 2004. Decontouring: prevention and removal of false contour artifacts. In *Proceedings of SPIE*, Vol. 5292. 130–149.
- P. E. Debevec and J. Malik. 1997. Recovering High Dynamic Range Radiance Maps from Photographs. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '97)*. 369–378.
- P. Didyk, R. Mantiuk, M. Hein, and H.P. Seidel. 2008. Enhancement of Bright Video Features for HDR Displays. *Computer Graphics Forum* 27, 4 (2008), 1265–1274.
- C. Dong, Y. Deng, C. Change L., and X. Tang. 2015. Compression Artifacts Reduction by a Deep Convolutional Network. In *The IEEE International Conference on Computer Vision (ICCV '15)*.
- F. Dufaux, P. Callet, R.K. Mantiuk, and M. Mrak (Eds.). 2016. *High Dynamic Range Video: From Acquisition, to Display and Applications*. Vol. 1. Academic Press.
- Y. Endo, Y. Kanamori, and J. Mitani. 2017. Deep Reverse Tone Mapping. *ACM Trans. Graph.* 36, 6, Article 177 (2017).
- G. Fechner. 1965. *Elements of psychophysics*. Holt, Rinehart & Winston.
- J. Froehlich, S. Grandinetti, B. Eberhardt, S. Walter, A. Schilling, and H. Brendel. 2014. Creating Cinematic Wide Gamut HDR-Video for the Evaluation of Tone Mapping Operators and HDR-Displays. In *Proceedings of SPIE 9023, Digital Photography X. 90230X–90230X–10*.
- A. Gilchrist and A. Jacobsen. 1984. Perception of lightness and illumination in a world of one reflectance. *Perception* 13, 1 (1984), 5–19.
- X. Glorot and Y. Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics (AISTATS '10)*, Vol. 9. 249–256.
- I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. 2014. Generative Adversarial Nets. In *Advances in Neural Information Processing Systems 27*. 2672–2680.
- M. D. Grossberg and S. K. Nayar. 2003. What is the space of camera response functions?. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR '03)*, Vol. 2. II–602–9.
- S. Hajisharif, J. Kronander, and J. Unger. 2015. Adaptive dualISO HDR reconstruction. *EURASIP Journal on Image and Video Processing* 2015, 1 (2015), 41.
- K. He, X. Zhang, S. Ren, and J. Sun. 2016. Deep Residual Learning for Image Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR '16)*.
- G. E. Hinton and R. R. Salakhutdinov. 2006. Reducing the dimensionality of data with neural networks. *Science* 313, 5786 (2006), 504–507.
- S. Iizuka, E. Simo-Serra, and H. Ishikawa. 2016. Let There Be Color!: Joint End-to-end Learning of Global and Local Image Priors for Automatic Image Colorization with Simultaneous Classification. *ACM Trans. Graph.* 35, 4 (2016), 110:1–110:11.
- S. Ioffe and C. Szegedy. 2015. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. In *Proceedings of the 32nd International Conference on Machine Learning (PMLR '15)*, Vol. 37. 448–456.
- N. K. Kalantari and R. Ramamoorthi. 2017. Deep High Dynamic Range Imaging of Dynamic Scenes. *ACM Trans. Graph.* 36, 4 (2017), 144:1–144:12.
- D. P. Kingma and J. Ba. 2014. Adam: A Method for Stochastic Optimization. *CoRR* abs/1412.6980 (2014). <http://arxiv.org/abs/1412.6980>
- R. P. Kovaleski and M. M. Oliveira. 2009. High-quality brightness enhancement functions for real-time reverse tone mapping. *The Visual Computer* 25, 5 (2009), 539–547.
- R. P. Kovaleski and M. M. Oliveira. 2014. High-Quality Reverse Tone Mapping for a Wide Range of Exposures. In *27th Conference on Graphics, Patterns and Images (SIBGRAPI '14)*. 49–56.
- J. Kronander, S. Gustavson, G. Bonnet, A. Ynnerman, and J. Unger. 2014. A Unified Framework for Multi-Sensor HDR Video Reconstruction. *Signal Processing: Image Communications* 29, 2 (2014), 203 – 215.
- C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, et al. 2016. Photo-realistic single image super-resolution using a generative adversarial network. *arXiv preprint arXiv:1609.04802* (2016).
- J. Long, E. Shelhamer, and T. Darrell. 2015. Fully Convolutional Networks for Semantic Segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR '15)*.
- S. Mann and R.W. Picard. 1994. *Being 'undigital' with cameras: Extending Dynamic Range by Combining Differently Exposed Pictures*. Technical Report 323. M.I.T. Media Lab Perceptual Computing Section. 422–428 pages.

- B. Masia, S. Agustin, R. W. Fleming, O. Sorkine, and D. Gutierrez. 2009. Evaluation of Reverse Tone Mapping Through Varying Exposure Conditions. *ACM Trans. Graph.* 28, 5 (2009), 160:1–160:8.
- B. Masia, A. Serrano, and D. Gutierrez. 2017. Dynamic range expansion based on image statistics. *Multimedia Tools and Applications* 76, 1 (2017), 631–648.
- L. Meylan, S. Daly, and S. Süsstrunk. 2006. The Reproduction of Specular Highlights on High Dynamic Range Displays. *Color and Imaging Conference* 2006, 1 (2006), 333–338.
- S. K. Nayar and T. Mitsunaga. 2000. High dynamic range imaging: spatially varying pixel exposures. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR '00)*, Vol. 1. 472–479.
- A. Odena, V. Dumoulin, and C. Olah. 2016. Deconvolution and Checkerboard Artifacts. *Distill* (2016). <http://distill.pub/2016/deconv-checkerboard>.
- D. Pathak, P. Krähenbühl, J. Donahue, T. Darrell, and A. Efros. 2016. Context Encoders: Feature Learning by Inpainting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR '16)*. 2536–2544.
- T. Pohlen, A. Hermans, M. Mathias, and B. Leibe. 2017. Full-Resolution Residual Networks for Semantic Segmentation in Street Scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR '17)*.
- A. Radford, L. Metz, and S. Chintala. 2015. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434* (2015).
- E. Reinhard, G. Ward, S. N. Pattanaik, P. E. Debevec, W. Heidrich, and K. Myszkowski. 2010. *High dynamic range imaging: acquisition, display, and image-based lighting (2nd ed.)*. Morgan Kaufmann.
- A. G. Rempel, M. Trentacoste, H. Seetzen, H. D. Young, W. Heidrich, L. Whitehead, and G. Ward. 2007. Ldr2Hdr: On-the-fly Reverse Tone Mapping of Legacy Video and Photographs. *ACM Trans. Graph.* 26, 3, Article 39 (2007).
- S. Ren, K. He, R. Girshick, and J. Sun. 2015. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. In *Advances in Neural Information Processing Systems* 28. 91–99.
- O. Ronneberger, P. Fischer, and T. Brox. 2015. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *Proceedings of Medical Image Computing and Computer-Assisted Intervention (MICCAI '15)*. 234–241.
- M. Rouf, R. Mantiuk, W. Heidrich, M. Trentacoste, and C. Lau. 2011. Glare encoding of high dynamic range images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR '11)*. 289–296.
- H. Seetzen, W. Heidrich, W. Stuerzlinger, G. Ward, L. Whitehead, M. Trentacoste, A. Ghosh, and A. Vorozcovs. 2004. High Dynamic Range Display Systems. *ACM Trans. Graph.* 23, 3 (2004), 760–768.
- A. Serrano, F. Heide, D. Gutierrez, G. Wetzstein, and B. Masia. 2016. Convolutional Sparse Coding for High Dynamic Range Imaging. *Computer Graphics Forum* 35, 2 (2016), 153–163.
- K. Simonyan and A. Zisserman. 2014. Very Deep Convolutional Networks for Large-Scale Image Recognition. *CoRR* abs/1409.1556 (2014).
- P. Svoboda, M. Hradis, D. Barina, and P. Zemcik. 2016. Compression Artifacts Removal Using Convolutional Neural Networks. In *Journal of WSCG*, Vol. 24. 63–72.
- M. D. Tocci, C. Kiser, N. Tocci, and P. Sen. 2011. A Versatile HDR Video Production System. *ACM Trans. Graphics* 30, 4 (2011), 41:1–41:10.
- J. Unger and S. Gustavson. 2007. High-dynamic-range video for photometric measurement of illumination. In *Proceedings of SPIE*, Vol. 6501.
- P. Vincent, H. Larochelle, Y. Bengio, and P.-A. Manzagol. 2008. Extracting and Composing Robust Features with Denoising Autoencoders. In *Proceedings of the 25th International Conference on Machine Learning (ICML '08)*. ACM, 1096–1103.
- L. Wang, L.-Y. Wei, K. Zhou, B. Guo, and H.-Y. Shum. 2007. High Dynamic Range Image Hallucination. In *Proceedings of the 18th Eurographics Conference on Rendering Techniques (EGSR'07)*. 321–326.
- C. Yang, X. Lu, Z. Lin, E. Shechtman, O. Wang, and H. Li. 2016. High-Resolution Image Inpainting using Multi-Scale Neural Patch Synthesis. *arXiv preprint arXiv:1611.09969* (2016).
- J. Zhang and J.-F. Lalonde. 2017. Learning High Dynamic Range from Outdoor Panoramas. *arXiv preprint arXiv:1703.10200* (2017).
- B. Zhou, A. Lapedriza, J. Xiao, A. Torralba, and A. Oliva. 2014. Learning Deep Features for Scene Recognition using Places Database. In *Advances in Neural Information Processing Systems* 27. 487–495.