

Analysis of Coded Apertures for Defocus Deblurring of HDR Images

Luis Garcia, Lara Presa, Diego Gutierrez and Belen Masia

Universidad de Zaragoza

Abstract

In recent years, research on computational photography has reached important advances in the field of coded apertures for defocus deblurring. These advances are known to perform well for low dynamic range images (LDR), but nothing is written about the extension of these techniques to high dynamic range imaging (HDR).

In this paper, we focus on the analysis of how existing coded apertures techniques perform in defocus deblurring of HDR images. We present and analyse three different methods for recovering focused HDR radiances from an input of blurred LDR exposures and from a single blurred HDR radiance, and compare them in terms of the quality of their results, given by the perceptual metric HDR-VDP2. Our research includes the analysis of the employment of different statistical deconvolution priors, made both from HDR and LDR images, performing synthetic experiments as well as real ones.

Categories and Subject Descriptors (according to ACM CCS): I.4.3 [Image Processing and Computer Vision]: Enhancement—Sharpening and deblurring

1. Introduction

The field of computational photography has obtained impressive results in last years, improving conventional photography results. One well known problem that conventional cameras present is the limitation of the sensor to capture images with an extended dynamic range. In a conventional camera the dynamic range is limited and parts of the scene which present luminance out of the range would not be correctly represented. In this context HDR imaging (High Dynamic Range imaging) [RHD*10] is a strategy to capture and represent the extended luminance range present in real scenes.

Also in terms of defocus deblurring computational photography has reached important advances. Since image capture can be modelled as a convolution between the focused image and the blur kernel plus a noise function, recovering a sharp image is reduced to a deconvolution problem. However, traditional circular apertures present a very poor response in the frequency domain with multiple zero-crossings and attenuation in high frequencies. Thus, recovered images present poor quality. Coded apertures are designed to have an appropriate frequency response to solve this problem, placing them in the camera lens in order to code light before it reaches the camera sensor. The defocus blur is encoded

and high frequencies are better preserved in the original image, obtaining better deblurred images after deconvolution.

This work turns around both approaches, analysing the use of coded apertures for defocus deblurring techniques in HDR imaging. While it is well known that the use of coded apertures for defocus deblurring offers good performance with LDR images [ZN09], to our knowledge this is the first time that these techniques are extended to HDR imaging.

For this purpose, we rely on a coded aperture specifically designed for defocus deblurring of LDR images by Zhou et al. [ZN09] and use it to analyse this problem in HDR images. The pattern of this aperture can be seen in Figure 1 together with its power spectrum compared to that of a circular aperture. Note that this aperture offers a better frequency response for defocus deblurring than the circular aperture.

We propose and analyse three different processing models for recovering focused HDR images, one from a single blurred HDR radiance and two from an input of blurred LDR exposures, and analyse them first in a simulation environment and finally in real scenarios. We also analyse the use of deconvolution statistical priors, made both from HDR and from LDR images, taking into account the work of Pouli et al. [PCR10] and following the idea that, to solve HDR prob-

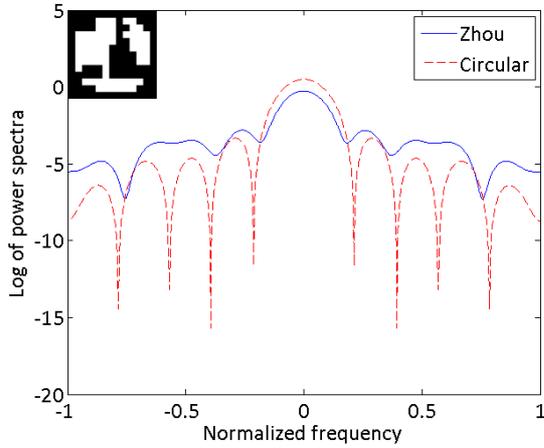


Figure 1: Power spectra of the coded aperture designed for defocus deblurring by Zhou et al. [ZN09] and a conventional circular aperture. Note how the coded aperture pattern offers better frequency response, as it avoids zero-crossings and reduces the attenuation in high frequencies.

lems, the use of HDR priors instead of LDR ones would lead to better results due to the existing statistical differences between both types of images.

2. Previous Work

Coded apertures have been traditionally used in astronomy since the 1960s to address SNR problems related to lensless imaging, coding the incoming high frequency x-rays and γ -rays. One well known pattern for this purpose are MURA patterns (Modified Uniformly Redundant Array) [GF89].

More recently, in the field of computational photography, Veeraghavan et al. [VRA*07] showed how coded apertures can be used to reconstruct 4D light fields from 2D sensor information. Coded apertures have also been used for solving the defocus deblurring problem. The main idea is to obtain coded apertures with better frequency response than the conventional circular aperture. Levin et al. [LDF07] designed a coded aperture optimized for depth recovery and a novel deconvolution method, in order to achieve all-in-focus images and a depthmap estimation simultaneously. Other techniques aimed at approximating depth and a focused image, although requiring multiple images, include that of Hiura and Matsuyama [HM98], who proposed a four-pinhole coded aperture, or the work by Liang et al. [LLW*08], who make use of multiple images captured with Hadamard-based aperture patterns. Yet another approach to recover focus and depth information of a scene was developed by Zhou et al. [ZLN09], in this case obtaining a pair of coded apertures through genetic algorithms and gradient descent search. On a separate work, Zhou et al. [ZN09] presented a metric that eval-

uates the goodness of a coded aperture for defocus deblurring based on the quality of the resulting deblurred image. Building on that work, Masia et al. studied the use of non-binary apertures for defocus deblurring [MCPG11]. More recently, Masia and colleagues [MPCG12] introduced perceptual metrics in the optimization process leading to an aperture design, and proved the benefits of these perceptually optimized coded apertures.

With respect to HDR imaging, for information about technical details we refer the reader to Reinhard and colleagues' book [RHD*10]. Another book authored by Banterle et al. [BAD11] was recently published, providing a different vision. Pouli et al. [PCR10] offer a useful analysis of the existing statistical differences between HDR and LDR images. There are also a series of works aimed at obtaining the optimal sequence of exposures needed to build HDR images [AR07, GN03, HDF10]. Finally, photographic hardware for HDR capture is another related line of research, where for instance the seminal work of Nayar et al. [NB03] significantly enhanced the dynamic range of a camera allowing to adapt the exposure of each pixel on the sensor.

3. Processing Models

The capture process of an image f is given by Equation 1:

$$f = f_0 * k + \eta \quad (1)$$

where f_0 is the focused scene, η is Gaussian white noise with standard deviation σ , and k is a convolution kernel determined by the aperture shape and the blur size.

In order to study the viability of the employment of coded apertures for defocus deblurring in HDR images, we simulate the capture process and attempt to recover a sharp image from the simulated blurred image.

Being f_0^{HDR} an HDR scene, we can use the approximation given by Equation 2 to simulate the capture of a High Dynamic Range radiance f^{HDR} only if we are able to capture it in one single shot.

$$f^{HDR} = f_0^{HDR} * k + \eta \quad (2)$$

Some existing cameras allow the capture of extended dynamic range, but in most cases HDR images are obtained by capturing series of LDR exposures and merging them later.

Then, being f_{0n}^{LDR} , ($n = 1, \dots, N$) a set of LDR exposures of the same focused HDR scene f_0^{HDR} , we can simulate the capture of the defocused HDR radiance by first simulating the capture of each exposure following Equation 3, and second merging them into a single HDR defocused radiance as expressed in Equation 4, g being the HDR merging operator.

$$f_n^{LDR} = f_{0n}^{LDR} * k + \eta \quad (3)$$

$$f^{HDR} = g(f_1^{LDR}, f_2^{LDR}, \dots, f_N^{LDR}) \quad (4)$$

Once f^{HDR} is obtained, we can recover the focused HDR

radiance \hat{f}_0^{HDR} by performing a single deconvolution. However, since we have the LDR defocused exposures, it is possible to deblur them separately with a set of N deconvolutions and merge them later to obtain \hat{f}_0^{HDR} , following Equation 5.

$$\hat{f}_0^{HDR} = g(\hat{f}_{01}^{LDR}, \hat{f}_{02}^{LDR}, \dots, \hat{f}_{0N}^{LDR}) \quad (5)$$

According to this, we present three different models for recovering focused HDR radiances:

1. **One-shot model:** Processing HDR radiance obtained with a single shot. Equation 2 is used to model the capture process and the focused radiance is recovered with a single deconvolution, as seen in Figure 2(a).
2. **HDR model:** Processing HDR radiance obtained by merging LDR exposures. Equations 3 and 4 are used and the focused HDR image is recovered with a single deconvolution. The pipeline of this processing is shown in Figure 2(b).
3. **LDR model:** Processing LDR exposures separately before merging. We follow Equation 3 to model the capture process of the N input images, and recover the focused LDR exposures with N deconvolutions, then merging them as in Equation 5 to obtain the HDR focused radiance. This pipeline can be seen in Figure 2(c).

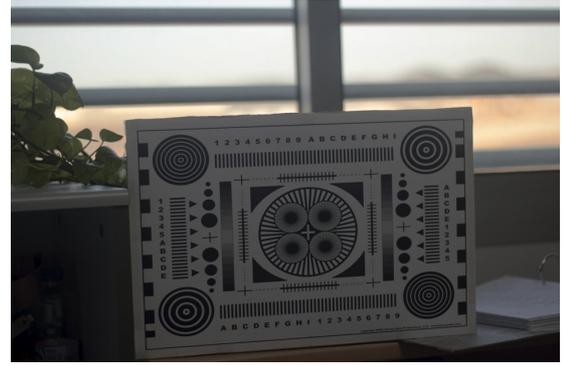
4. Simulation of Processing Models

First we analyse these three models by performing simulations in order to study their viability before proceeding to real experiments. To carry them out, we use one of the coded apertures developed by Zhou et al. [ZN09], which is shown in Figure 1. This aperture is known to work well for defocus deblurring LDR images.

For the simulations we use a set of seven HDR photographs with different dynamic ranges for the first model, and their three corresponding LDR exposures for the other two. One of them is shown in Figure 3. The main goal is to recover the focused HDR images with all three processing models. We use the perceptual metric *HDR-VDP2* [MKRH11] in order to assess the quality of the results. This metric works on luminance, comparing a reference HDR image with its distorted version, providing quality and visibility (probability of detection) measures based on a calibrated model of the human visual system. In this work we focus in obtaining the quality factor Q, a prediction of the quality degradation of the recovered HDR image with respect to the reference HDR image, expressed as a mean-opinion-score (with values between 0 and 100). This metric can not only work with HDR images, but also with their LDR counterparts.

We test four different noise levels ($\sigma = 0.0005, 0.001, 0.005$ and 0.05), and three different deconvolution variations based on Wiener deconvolution, whose formulation in frequency is given by Equation 6

$$\hat{F}_0 = \frac{F \cdot \bar{K}}{|K|^2 + |C|^2} \quad (6)$$



(a) Tone mapped HDR



(b) Overexposed (c) Medium exposed (d) Underexposed

Figure 3: Example of one of the HDR images used in simulation, with the three exposures merged to obtain it, with relative exposures of +2, 0 and -2 stops.

where \hat{F}_0 is the Fourier Transform of the recovered image, \bar{K} is the complex conjugate of K , $|K|^2 = K \cdot \bar{K}$ and $|C|^2 = |\sigma/F_0|^2$ is the Noise to Signal Ratio (NSR) matrix of the original image. From this deconvolution, we study these three different variations:

- **Wiener deconvolution without prior, with a constant NSR matrix.** Replacing $|C|^2$ in Equation 6 by a constant NSR matrix. We tested several values and found that there is a trade-off between noise and ringing in resulting images. We finally decided to set the NSR to 0.005, achieving good balance between both artifacts.
- **Wiener deconvolution using an HDR image prior.** Approximating $|F_0|^2$ in Equation 6 by a statistical prior matrix averaging power spectra of a series of 198 HDR images. We construct the prior employing *man-made (day and indoors)* HDR images from the database of Tania Pouli (<http://taniapouli.co.uk/research/statistics/>).
- **Wiener deconvolution using an LDR image prior.** Replacing $|F_0|^2$ as in the previous, using a prior of 198 *man-made (day and indoors)* LDR images instead, extracted from the database of Tania Pouli.

We explore the use of HDR priors in the one-shot and HDR models, given that we are deconvolving an HDR radiance, inspired by Pouli et al. [PCR10]. Note that we do not test the LDR model with an HDR prior since we are deconvolving LDR images in it. Since the aperture we are using is optimized for a noise level of $\sigma = 0.005$, we set this value as

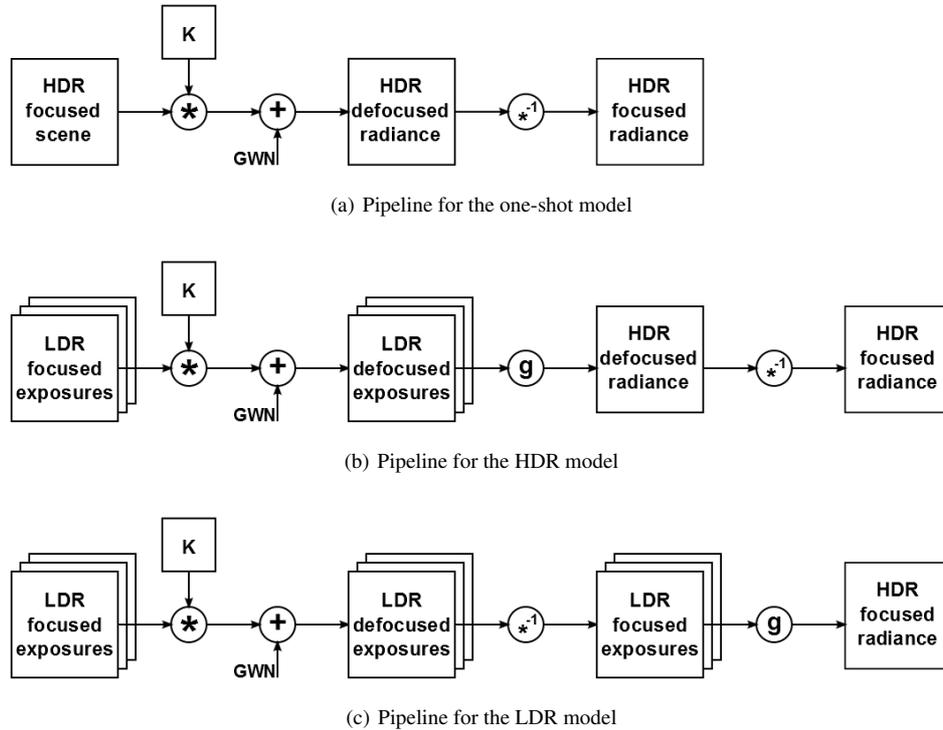


Figure 2: Pipelines for all different processing models, where k is the convolution kernel, \mathbf{GWN} is Gaussian White Noise, g is the HDR merging operator and $*$ is the convolution operator.

standard deviation of the Gaussian noise in our deconvolutions with priors.

5. Performance Comparison

Once all the simulations are finished, we compute the mean quality factor Q , given by the *HDR-VDP2* metric, of the seven images obtained with the three proposed processing models shown in Figure 2. For each model we analyse four different noise levels and the three different deconvolution variations explained in Section 4 (except for the LDR model, as explained). This information is collected in Figure 4.

We can see how the use of priors is strongly recommended for the one-shot model when image noise is very high. In this noisy scenario, an HDR prior offers better results than an LDR prior. However, when image noise decreases, all three different deconvolutions produce similar behaviours. As expected, using an HDR prior outperforms using an LDR prior in the HDR model, but we can see how the use of Wiener deconvolution with a constant NSR matrix seems to offer similar or even better quality all along the noise range. For the LDR model, the use of a constant NSR matrix in the deconvolution seems to offer better results than the LDR prior, although differences are not significant.

With regard to the comparison between all three process-

ing models, we can see how the one-shot model clearly derives in better results than the other two, and would be the ideal method, if the appropriate hardware becomes widely available. Meanwhile, the HDR model seems to perform worst. Note that the merging operation is a non-linear process, and therefore the deconvolution is performed over content which has been non-linearly transformed. Also, the added \mathbf{GWN} can be amplified during this process. It must be noted, however, that function g is approximately linear for a wide range of luminances. In the LDR model, three deconvolutions are performed, and it is well-known that deconvolution is a noisy process. However, in HDR images the relative difference between neighbour pixels is bigger than in LDR ones. This increases ringing significantly, and along with the amplified \mathbf{GWN} and non-linearity may be what causes the HDR model results to be the worst of all.

In terms of computational cost, the lowest is offered by the one-shot model, as it only requires one deconvolution, while the HDR model requires one deconvolution and one exposure fusion, and the LDR model requires one deconvolution for each exposure and one exposure fusion.

In Figure 5 we show the result of one of the noisy simulations ($\sigma = 0.05$) using the one-shot model, with both priors. We can see how the use of an HDR prior slightly reduces the

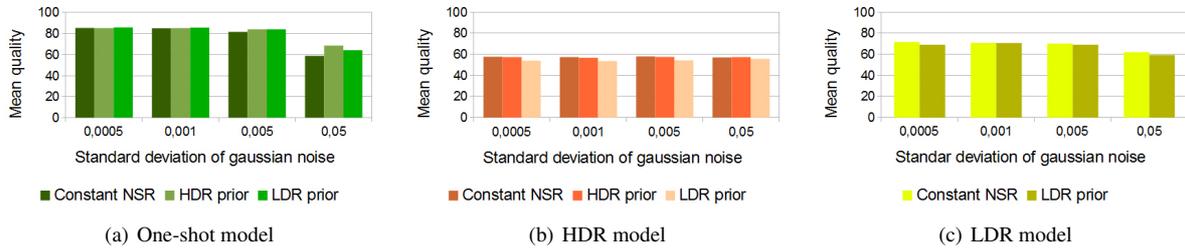


Figure 4: Mean Q obtained with the HDR-VDP2 metric for each processing model, with all different combinations of noise level and deconvolution prior.

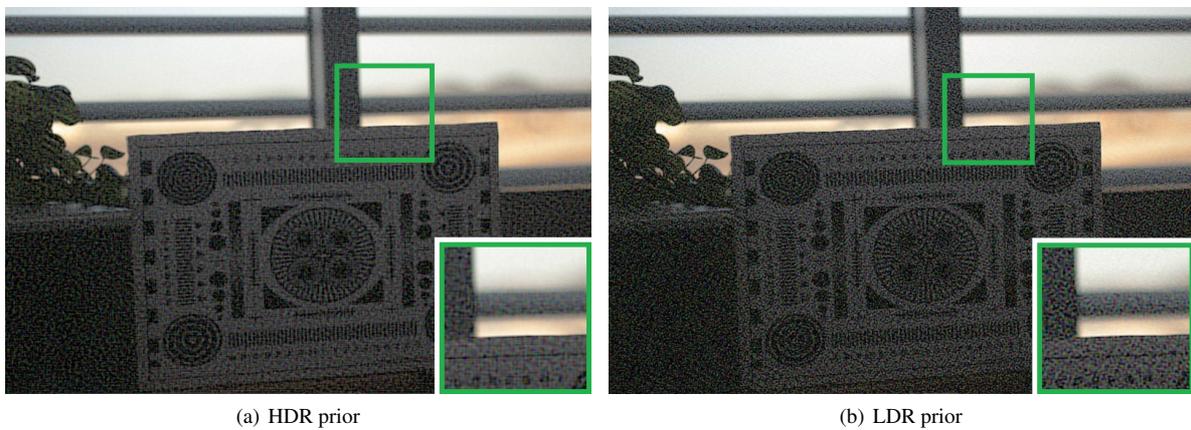


Figure 5: Comparison between images recovered after simulation of the one-shot model, with HDR and LDR priors and $\sigma = 0.05$. Note how the use of the HDR prior instead of the LDR one slightly reduces image noise.

recovered image noise. In Figure 6 we show an example of the same HDR scene recovered with the HDR model, with both priors, this time with $\sigma = 0.0005$. In this low noise scenario we can appreciate how the use of an HDR prior instead of an LDR one results in a reduction of ringing artifacts.

6. Validation in Real Scenarios

After performing the simulations we proceed to validate the same processes in real scenarios. We cannot validate the one-shot model in real scenes because of the lack of the required equipment: an HDR camera that allows to capture an HDR image with a single shot. For this reason, physical validation is restricted to the HDR and LDR models. We use a DSLR camera Canon EOS 500D with an EF 50mm f/1.8 II lens for all the tests. The same coded aperture used in simulation (Figure 1) is printed and inserted into the camera lens.

6.1. Image capture process

We construct a scene with a large luminance range and capture three images using the multi-bracketing camera option

set to relative exposures of +2, 0 and -2 stops. For these captures we fix the ISO setting value at 100 and aperture size at F2.0, leading to exposure times of 1/5, 1/20 and 1/80 seconds.

We place the scene 180 cm away from the camera, and set the focus plane at 120 cm, leading to a defocus distance of 60 cm. We also take three exposures of the well focused scene to obtain a ground truth HDR image that allows comparison, using the same capture parameters described above.

All images are taken in RAW format, with a size of 4752x3168 pixels. To reduce computational time and cost we resize images by a factor of 0.2, reducing them to 951x634 pixels.

6.2. System calibration

In order to recover the focused HDR image of the scene we need to know the PSF of the capture system (i.e. the response to an impulse) to use it as the kernel in the deconvolution process. To calibrate the PSF at the depth of interest (60 cm) we use a LED mounted on a pierced thick black cardboard in

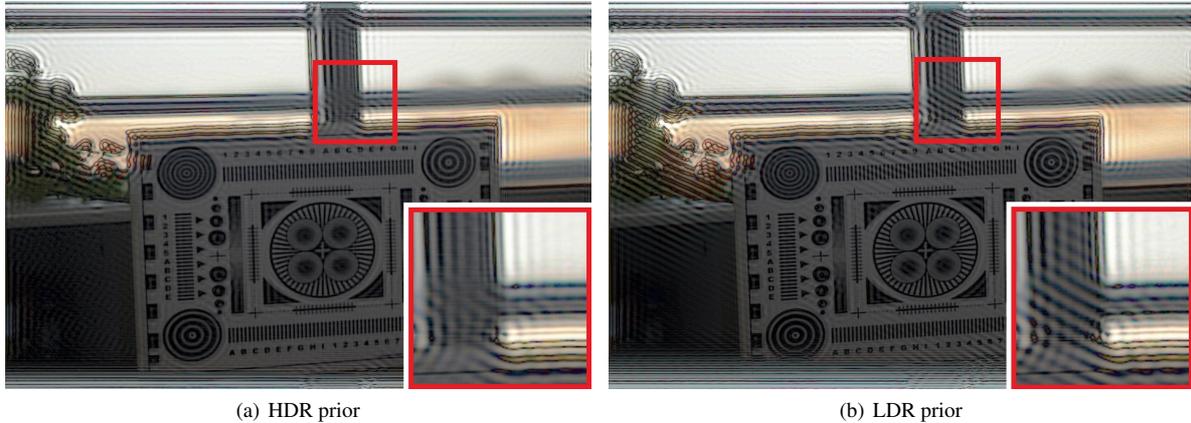


Figure 6: Comparison between images recovered after simulation of the HDR model, with HDR and LDR priors and $\sigma = 0.0005$. Note how using the HDR prior instead of the LDR one seems to reduce image ringing.

order to make a point light source. We lock the focus plane at 120 cm and place the cardboard with the LED at 180 cm. In order to be coherent with image capture, we obtain three images, one for each exposure value, with the same capture parameters used to capture the scene. The central detail of these images is shown in Figure 7. We also obtain an HDR image of the montage to obtain the PSF that we will use in the deconvolution in the HDR model.

The cropped greyscaled image of the LED serves us as PSF, after thresholding it in order to eliminate residual light, and normalizing it to preserve energy in the deconvolution process. Note that the threshold value changes for each PSF, increasing with the exposure value: 0.39 for underexposed, 0.5 for well-exposed and 0.8 for overexposed. For the PSF used in the HDR model the threshold value is 0.2. The resulting PSFs are shown in Figure 8. After resizing the kernel, its size is 14×14 pixels.

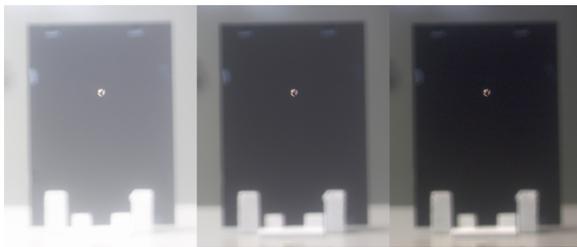


Figure 7: Central detail of the three different exposures used to recover the PSFs.

6.3. Deblurred image recovery

Once we obtain the PSFs we recover the sharp images following the HDR and LDR models. For the HDR one we

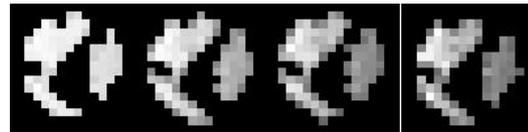


Figure 8: PSFs obtained for deconvolution. From left to right: PSF for the high, central and low exposure, used in the LDR model, and PSF obtained by merging the three exposures used in the HDR model.

merge the defocused exposures into a defocused HDR radiance and obtain the deblurred HDR image performing a single deconvolution using the HDR kernel, as in Figure 2(b). For the LDR model we perform one deconvolution for each defocused exposure, using the corresponding PSF for each one, and then we merge the resulting recovered exposures into the focused HDR image, as in Figure 2(c). In each case, we carry out the same Wiener deconvolution variations described in Section 4, excluding again the use of an HDR prior for the LDR model.

7. Results and Discussion

Once we perform all the experiments, we compare the results of both models. We compute the quality factor Q given by the *HDR-VDP2* metric of the HDR recovered images and show the results in two different scenes. We also check the effect of the use of the different deconvolution variations, specially those which employ deconvolution priors.

7.1. Model comparison

For our first scene, in Figure 9 we show the quality factor Q , given by the *HDR-VDP2* metric, of the HDR images re-

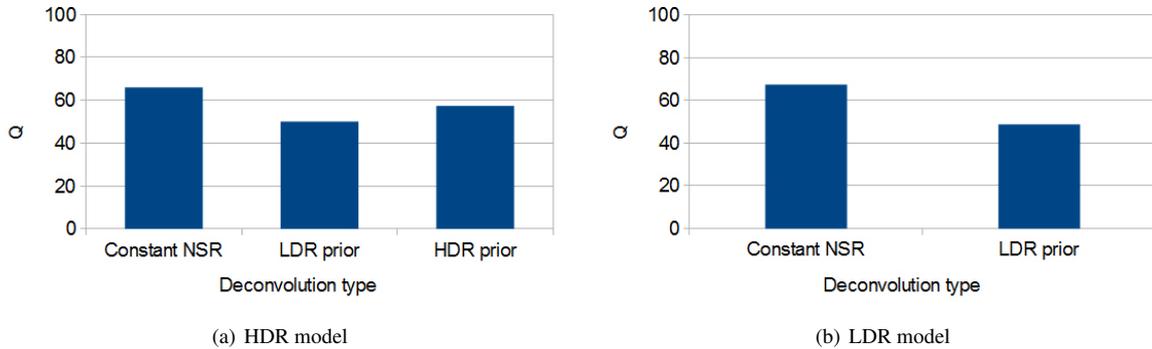


Figure 9: Quality factor Q obtained with the HDR-VDP2 metric for our first real scene, for each processing model and deconvolution prior. We can observe how in the HDR model the HDR prior outperforms the LDR one, and how both LDR and HDR models using constant NSR offer similar quality.

covered with each processing model. These results indicate that, while simulation results suggested that the LDR model offered better results than the HDR model (see Figure 4), real experiments point out that both models offer very similar qualities. Note also that, according to the metrics, the use of priors results in worse performance. We explore this fact further in Subsection 7.2.

We show the result of both models, using constant NSR, in Figure 10, in order to offer a visual comparison of how both models perform. We also show the original (blurred) HDR radiance and the ground truth ideal HDR radiance. We can see how visual appearance is consistent with the results yielded by the metric. The image recovered with the HDR model shows more ringing due possibly to the biggest relative difference between neighbour pixels (see also Section 5). Furthermore, attending to the highlighted details and comparing recovered and original images we see how both models are able to recover the well-focused HDR radiance (see e.g. book titles or text in the lens box in the images). These images prove that the employment of coded apertures for defocus deblurring of HDR images is viable and presents a good performance.

We test again our approximations performing the experiments in a new scene, in order to check if results correlate with the first ones. In Figure 11 we show the quality factor Q given by the HDR-VDP2 metric for this second scene.

Again, the use of priors derives in worse results than the use of a constant NSR, for both processing models. In Figure 12 we show the HDR images of this scene recovered with the HDR and LDR models with constant NSR. As we can see, both models offer good results when recovering the focused image, and again the HDR model exhibits slightly more ringing than its LDR counterpart.

In Section 5 we have already pointed out possible causes

for one model performing better than the other in simulation. When incorporating results in real scenarios, the Q metric seems to indicate similar results for both models, although it would be advisable to perform more tests with more data. Also, the HDR-VDP2 metric works only with luminance values, not taking into account color, and while it has been specifically tested for some types of distortions, such as white noise or Gaussian blur, it has not been designed nor tested for e.g. ringing artifacts. Finally, modelling noise as GWN is another source of inaccuracy, an approximation, since image noise does not follow a Gaussian distribution.

7.2. Effects of using a prior

As shown in Figures 9 and 11, in real experiments we see that both HDR and LDR models perform much better when no deconvolution prior is used. We inspect the images recovered with both priors in order to know why this happens. If we carefully observe these images we can appreciate a grid shaped distortion, as seen in Figure 13. This distortion clearly reduces the visual quality of the images recovered with deconvolution prior. Further, we notice again that HDR prior outperforms LDR prior in the HDR model, as it minimizes, but not completely removes, this distortion.

We explore the variation of σ in the deconvolution process and see the impact of this alteration in the described distortion. This variation corresponds to a higher weight for the deconvolution prior. In Figure 14 we see some of the images obtained with different σ in the deconvolution process for the LDR model. We see how increasing this value we obtain a better reduction of prior distortion and ringing. In exchange, we find that this increase leads to less sharp results, resulting in a trade-off between both effects.

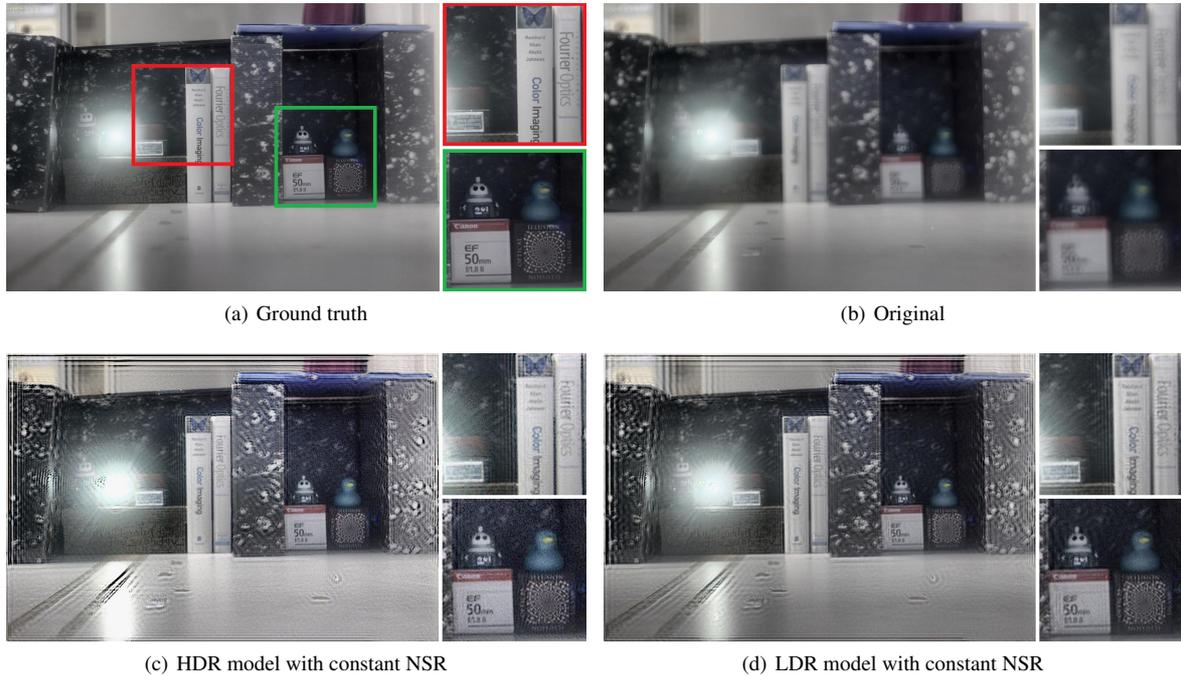


Figure 10: HDR results obtained for our first real scene with the best processing models in terms of Q (c,d), compared to the ground truth and original images, all of them tonemapped. Here we see how both models offer good and similar results.

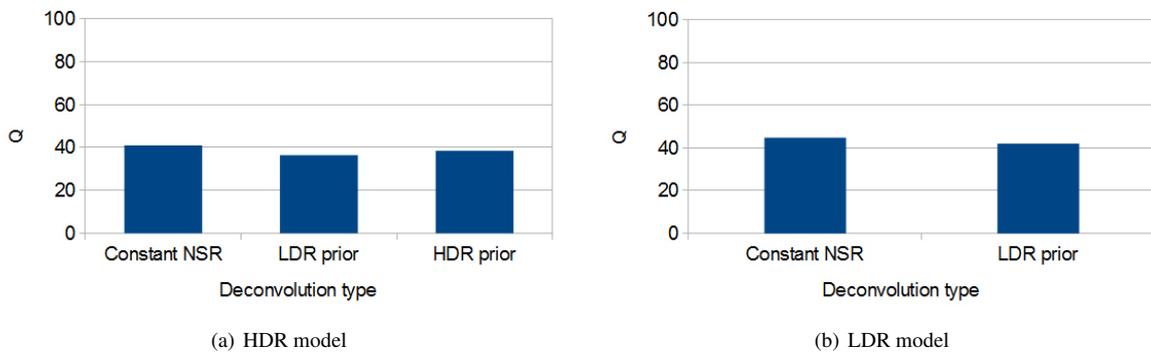


Figure 11: Quality factor Q obtained with the HDR-VDP2 metric for our second real scene, for each processing model and deconvolution prior. Note that in the HDR model the HDR prior outperforms the LDR one, and that in both models the use of a constant NSR offers the best results.

8. Conclusions and Future Work

In this paper we explore for the first time, to our knowledge, the use of coded apertures for defocus deblurring of HDR images, showing that these techniques, which used to be employed in LDR images, can be extended for HDR imaging. We implement three different processing models, either responding to an input of an HDR defocused radiance or a series of LDR defocused exposures of the same scene.

The one-shot model offers the best results in simulation, but due to the limited dynamic range of our camera, we are not able to capture HDR images with a single shot. Thus, we could not test this model as properly as we would like to, so the first future work that follows this paper is to perform more experiments in this way, employing more advanced cameras that allow the capture of extended dynamic range images with just one shot. As said, this would be the ideal

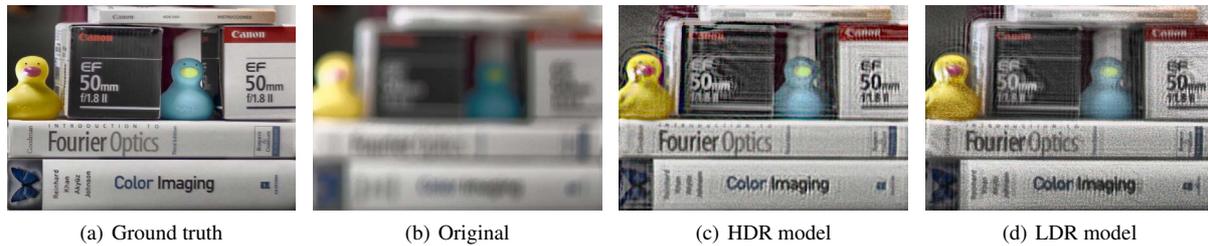


Figure 12: HDR results obtained for our second real scene with the best processing models in terms of Q , compared to the ground truth and original images, all of them tonemapped. We can see how both models are able to recover sharp details such as the book titles.

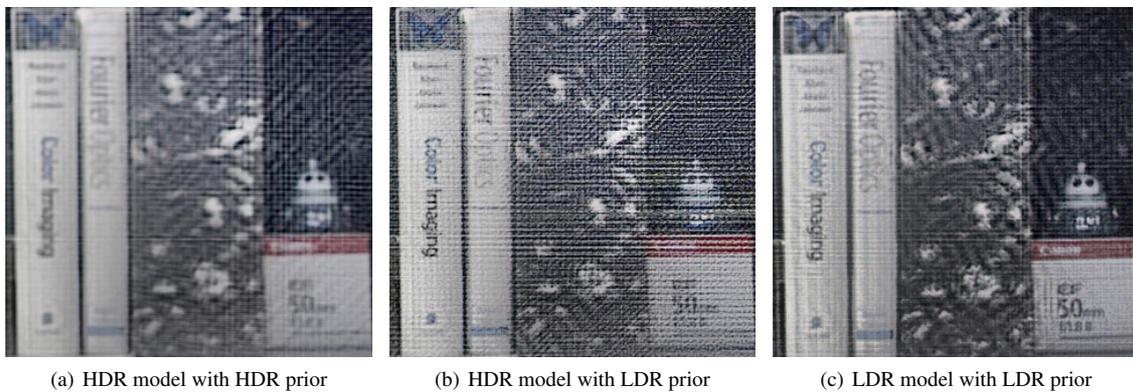


Figure 13: Detail of our recovered images of the first real scene using priors, where we can appreciate a clear grid shape distortion. Note that, in the HDR model, using an HDR prior instead of an LDR one reduces this effect. All the images are tonemapped.

processing model if the appropriate hardware became widely available to the general public.

The two other processing models are validated with real experiments, finding that both of them are viable and allow the recovery of a focused image. We show that the proposed HDR model seems to perform as good as the LDR one in practice, despite the fact that simulations indicate otherwise, and reduces the computational cost as it only requires one deconvolution.

We conclude that the use of deconvolution priors made of HDR images instead of conventional LDR priors leads to better performance. However, maybe due to the fact that the prior we are employing is far from optimal, the best results come when no prior is employed in the process. From this, and relying on the work of Pouli et al. [PCR10], we believe that more research related to HDR priors is needed. Since many optimization problems benefit from the use of statistical regularities of the images, and taking into account the advances on HDR imaging, the construction of good HDR priors is another avenue of future work.

One of the immediate applications of these new priors,

which is highly related to our work, is the design of optimal aperture patterns for defocus deblurring of HDR images. As the aperture we have employed [ZN09] is obtained by means of a genetic algorithm which uses prior information of LDR images, we believe that it is possible to obtain new coded apertures optimized specifically for HDR images.

9. Acknowledgements

We would like to thank the reviewers for their valuable comments. This research has been funded by the European Commission, Seventh Framework Programme, through the projects GOLEM (Marie Curie IAPP, grant agreement no.: 251415) and VERVE (Information and Communication Technologies, grant agreement no.: 288914), and by the Spanish Ministry of Science and Technology (TIN2010-21543). Belen Masia is supported by an FPU grant from the Spanish Ministry of Education and by a NVIDIA Graduate Fellowship.

References

[AR07] AKYÜZ A., REINHARD E.: Noise reduction in high dynamic range imaging. *Journal of Visual Communication and Im-*

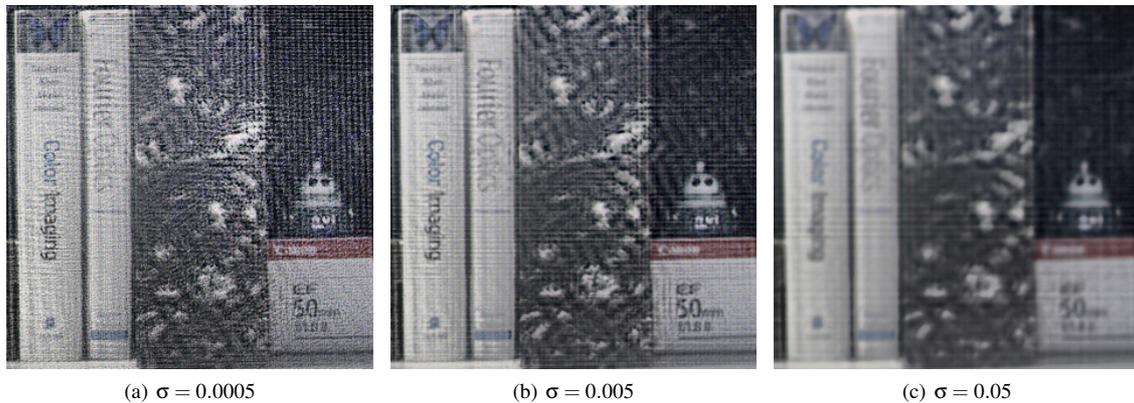


Figure 14: Effect of the variation of σ in the deconvolution for the LDR model. We can see a trade-off between the grid shape distortion and image sharpness. All the images are tonemapped.

- age Representation (JVCIR) (2007). 2
- [BADC11] BANTERLE F., ARTUSI A., DEBATTISTA K., CHALMERS A.: *Advanced High Dynamic Range Imaging: Theory and Practice*. AK Peters (CRC Press), 2011. 2
- [GF89] GOTTESMAN S., FENIMORE E.: New family of binary arrays for coded aperture imaging. *Applied Optics*, 20 (1989), 4344–4352. 2
- [GN03] GROSSBERG M. D., NAYAR S. K.: High dynamic range from multiple images: Which exposures to combine. In *ICCV Workshop CPMCV* (2003). 2
- [HDF10] HASINOFF S., DURAND F., FREEMAN W.: Noise-optimal capture for high dynamic range photography. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2010). 2
- [HM98] HIURA S., MATSUYAMA T.: Depth measurement by the multi-focus camera. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (Washington DC, USA, 1998), IEEE Computer Society. 2
- [LDF07] LEVIN A., FERGUS R., DURAND F., FREEMAN W.: Image and depth from a conventional camera with a coded aperture. *ACM Transactions on Graphics* 26, 3 (2007). 2
- [LLW*08] LIANG C., LIN T., WONG B., LIU C., , CHEN H.: Programmable aperture photography: multiplexed light field acquisition. *ACM Transactions on Graphics* 27, 3 (2008). 2
- [MCPG11] MASIA B., CORRALES A., PRESA L., GUTIERREZ D.: Coded apertures for defocus deblurring. In *Symposium Iberoamericano de Computacion Grafica* (Faro, Portugal, 2011). 2
- [MKRH11] MANTIUK R., KIM K. J., REMPEL A. G., HEIDRICH W.: HDR-VDP-2: A calibrated visual metric for visibility and quality predictions in all luminance conditions. *ACM Transactions on Graphics* 30, 40 (2011). 3
- [MPCG12] MASIA B., PRESA L., CORRALES A., GUTIERREZ D.: Perceptually optimized coded apertures for defocus deblurring. *Computer Graphics Forum* (2012). To appear. 2
- [NB03] NAYAR S., BRANZOI V.: Adaptive dynamic range imaging: Optical control of pixel exposures over space and time. In *IEEE International Conference on Computer Vision (ICCV)* (2003). 2
- [PCR10] POULI T., CUNNINGHAM D., REINHARD E.: Statistical regularities in low and high dynamic range images. *ACM Symposium on Applied Perception in Graphics and Visualization (APGV)* (July 2010). 1, 2, 3, 9
- [RHD*10] REINHARD E., HEIDRICH W., DEBEVEC P., PATANAIK S., WARD G., MYSZKOWSKI K.: *High Dynamic Range Imaging: Acquisition, Display and Image-Based Lighting*. Morgan Kaufmann Publishers, 2010. 1, 2
- [VRA*07] VEERARAGHAVAN A., RASKAR R., AGRAWAL A., MOHAN A., TUMBLIN J.: Dappled photography: mask enhanced cameras for heterodyned light fields and coded aperture refocusing. *ACM Transactions on Graphics* 26 (July 2007). 2
- [ZLN09] ZHOU C., LIN S., NAYAR S.: Coded aperture pairs for depth from defocus. In *IEEE International Conference on Computer Vision (ICCV)* (Kyoto, Japan, 2009). 2
- [ZN09] ZHOU C., NAYAR S. K.: What are Good Apertures for Defocus Deblurring? In *IEEE International Conference on Computational Photography* (San Francisco, CA, USA, 2009). 1, 2, 3, 9