

MODELING REALISTIC DEGRADATIONS IN NON-BLIND DECONVOLUTION

Jérémy Anger[†], Mauricio Delbracio[§], and Gabriele Facciolo[†]

[†]CMLA, ENS Cachan, CNRS, Université Paris-Saclay, 94235 Cachan, France

[§]IIE, Universidad de la República, Uruguay

ABSTRACT

Most image deblurring methods assume an over-simplistic image formation model and as a result are sensitive to more realistic image degradations. We propose a novel variational framework, that explicitly handles pixel saturation, noise, quantization, as well as non-linear camera response function due to e.g., gamma correction. We show that accurately modeling a more realistic image acquisition pipeline leads to significant improvements, both in terms of image quality and PSNR. Furthermore, we show that incorporating the non-linear response in both the data and the regularization terms of the proposed energy leads to a more detailed restoration than a naive inversion of the non-linear curve. The minimization of the proposed energy is performed using stochastic optimization. A dataset consisting of realistically degraded images is created in order to evaluate the method.

Index Terms— Non-blind deconvolution, image deblurring, saturation, quantization, gamma correction

1. INTRODUCTION

One of the major sources of image blur is due to camera motion during the sensor integration time. This phenomenon is most visible in low light conditions, when the integration time has to be long enough to capture a minimum amount of photons. In this situation, any strong light source present in the scene will certainly lead to pixel saturation, since the dynamic range to capture will be too large for the sensor.

Most motion deblurring strategies consist in estimating a blur kernel (which represents the effect of the camera motion in the image plane), and then deconvolving the blurred image with the estimated kernel. In this paper, we propose a non-blind deconvolution algorithm, which assumes that the kernel is known. The simplest image acquisition model is

$$v = u * k + n, \quad (1)$$

We thank Jean-Michel Morel for fruitful comments and discussions. Work partly financed by Agencia Nacional de Investigación e Innovación (ANII, Uruguay) grant FCE_1_2017_135458; Office of Naval research grant N00014-17-1-2552, Programme ECOS Sud – UdelaR - Paris Descartes U17E04, DGA Astrid project « filmer la Terre » n° ANR-17-ASTR-0013-01, MENRT; DGA PhD scholarship jointly supported with FMJH.



(a) Degraded image. (b) $D_{\gamma\text{inv}}\text{-TV}$. (c) $D_{\text{full}}\text{-TV}^{\gamma}$.

Fig. 1: Image deconvolution can be significantly improved by defining a data fitting term that considers the whole image pipeline (quantization, noise, saturation, gamma correction) as shown in (c). Details best seen in the electronic version.

where u represents the sharp noiseless ideal image, $*$ denotes the convolution operator, k is a known blurring kernel which we assume stationary, v is the observed blurry image, and n is a realization of white Gaussian or Poisson noise, depending on the formulation. The inverse problem defined in (1) is linear, but significantly ill-posed. There is a huge amount of work seeking to restore images under this formulation [1]. Most methods are casted as minimizations of an energy of the form

$$E(u) = D(u; v) + \lambda R(u), \quad (2)$$

where $D(u; v)$ (denoted $D(u)$ from now on) is a data fitting term that enforces the image formation model (1) and $R(u)$ is a regularizer that imposes prior knowledge on the solution. The total variation penalization [2] is often used:

$$R(u) = \text{TV}(u) = \int |\nabla u(\mathbf{x})| d\mathbf{x}. \quad (3)$$

Although interesting, the model (1) is over-simplistic. In more realistic scenarios due to the physical image acquisition and the complex processing pipeline, non invertible non-linear degradations, such as, quantization and compression, can occur. Under these circumstances, a more accurate forward model is needed:

$$v = Q_q(S_c(u * k + n)^{\frac{1}{\gamma}}), \quad (4)$$

where $S_c(u) = \min(c, u)$ is the pixel saturation operator, $Q_q(u) = q \cdot \text{round}(\frac{u}{q})$ is the pixel quantization of step q and γ is a gamma correction coefficient, generally introduced by the camera manufacturer (usually $q = \frac{1}{256}$, as $u(\mathbf{x}) \in [0, 1]$).

To avoid modeling the effects of the non-linear processing on the noise, as done by White *et al.* [3], we approximate the forward model (4) by

$$v = Q_q(S_c(u * k)^{\frac{1}{\gamma}}) + n, \quad (5)$$

where n is assumed white and Gaussian. While this is an approximation, the effect of the gamma correction on shot noise (which follows a Poisson distribution) can be assimilated to a variance-stabilizing transform [4].

Due to model mismatch, traditional approaches, that assume the simple linear model (1), need to impose a strong image prior to overcome these degradations. In this work, we propose to adapt the data fitting term to explicitly account for these typical degradations. This yields better results as is illustrated in Fig. 1, where a naive restoration model (1) using total variation regularization is compared to the proposed model (5) with a gamma corrected TV regularization.

The paper is organized as follows. In Section 2 we review state-of-the-art methods that consider realistic acquisition models. In section 3 we present a deconvolution method that works under real practical degradations, such as saturation or quantization, while considering gamma correction in a rigorous way. Finally, in Section 4 we demonstrate the effectiveness of our approach on a new dataset of degraded images and conclude in Section 5.

2. RELATED WORK

Although image deconvolution has received significant attention in the past decades [5], there are only few works that address the deconvolution problem under a realistic image pipeline (saturated pixels, quantization, gamma correction).

Cho *et al.* [6] proposed a robust method that explicitly model outliers in the degradation process. However, this method is only effective when outliers are sparse in well-localized areas (e.g., saturated regions). Gregson *et al.* [7] proposed a variational stochastic deconvolution framework, inspired on stochastic tomography reconstruction [8], that works with different image priors. The method can handle saturation by discarding saturated pixels and uses a prior in non-linear space. The method is later extended to blind deconvolution [9], where a two step reconstruction that improved the saturation handling is introduced. The first step reconstructs the latent image by discarding unreliable blurred pixels, and the second one works on the regions that were masked out in the first phase. Our method is simpler and does not need to distinguish reliable from unreliable pixels.

Whyte *et al.* [3] claim that while saturation can be handled by discarding saturated pixels, a better solution is obtained by modifying the data term to handle saturation explicitly. The saturation operator (clipping) is approximated with a smooth function allowing to compute its derivative. In this work, we present a similar approach, but it does not require to approximate the non-smooth saturation operator. The authors also

proposed a split update between reliable and unreliable pixels. Although it effectively reduces ringing, it introduces blur as we show in the experimental section.

Camera response functions, including the gamma curve, are typically invertible functions. As such, some methods [6] directly invert the non-linear curve before deconvolving the image. However, if the image was quantized in the non-linear space, inverting the response curve results in non uniform quantization in the linear pixel space.

3. METHOD

In this section, we present different formulations that incorporate data fitting terms to handle the following degradations: saturation, quantization, and gamma correction. Each problem is formulated as an energy of the form (2), that is minimized using the Stochastic Deconvolution framework [10]. This framework is based on a coordinate descent algorithm with a Metropolis-Hastings strategy guiding the pixel sampling. The method is derivative-free and can be applied to any energy minimization problem, although its convergence is not guaranteed for non-smooth functions.

At each step of Stochastic Deconvolution, a pixel is drawn either by selecting a pixel nearby the previous sampled one or by randomly choosing a new position. Given a pixel position, the method evaluates the difference of energy that a small increase or decrease of the given pixel value would produce. If the energy decreases, the new value is kept and the algorithm is more likely to chose a nearby pixel in the next iteration. Since this process affects a single pixel of the solution at a time, the energy change due to the data and regularization terms can be computed by evaluating a small number of pixels surrounding the sampled one [10]. Convolution boundaries are handled by padding the image and considering the data fitting term only on valid pixels. In what follows, unless otherwise specified, we use a total variation penalization (3).

Saturation. Pixel saturation occurs when the scene dynamic range is larger than the one captured by the camera sensor. In this case, high intensities are clipped to the maximum sensor capacity, resulting in information loss. The saturation model that we use is very simple, yet leads to competitive results. If c is the sensor saturation limit, the considered data term is

$$D_S(u) = \|\min(c, u * k) - v\|^2. \quad (6)$$

Instead of discarding saturated pixels, this formulation expresses the fact that the estimated sharp image convolved by the motion blur kernel has to be saturated in the same pixels as the observed image, even if the exact intensity values are lost. Note that the data fitting term is equal to zero in the saturated pixels that match. This implies that these regions are regularized by the prior, independently of the regularization strength (controlled by λ).

This model works well on small saturated regions. For larger regions, we observe an over-smoothed restoration and missing pixel values cannot be properly restored (see Fig. 4e).

Quantization. A direct way to handle quantization is to explicitly introduce it as a constrained minimization problem

$$\arg \min_u R(u) \text{ s.t. } Q(u * k) = v, \quad (7)$$

where Q is the quantization operator. Note that this problem considers a noiseless observation. The data term in the associated Lagrangian relaxation is

$$D_{Q\text{fw}}(u) = \|Q(u * k) - v\|^2. \quad (8)$$

Let us denote Equation (8) as the *forward quantization energy*. This data fitting term is piecewise constant; in general, a small perturbation in u does not introduce any change on the energy, making the energy difficult to optimize. More importantly, this model does not exploit the nature of the problem: given two different images, u_1 and u_2 such that $Q(u_1 * k) = Q(u_2 * k) \neq v$, the data fitting term (8) leads to the same cost, whereas one image could be closer to the true latent sharp one. Hence the cost should favor one over the other.

Thus, we propose to replace the constraint $Q(u * k) = v$ by $(u * k)(\mathbf{x}) \in Q^{-1}(v(\mathbf{x}))$ where $Q^{-1}(s) = [s - \frac{q}{2}, s + \frac{q}{2}]$ is the quantization error interval centered at s . This yields

$$\arg \min_u R(u) \text{ s.t. } (u * k)(\mathbf{x}) \in Q^{-1}(v(\mathbf{x})), \forall \mathbf{x}. \quad (9)$$

In this formulation, if the estimate is not within the quantization error, we can compute a distance to the interval, namely

$$D_{Q\text{cx}}(u) = \left\| \left(|u * k - v| - \frac{q}{2} \right)_+ \right\|^2, \quad (10)$$

where $(\cdot)_+ = \max(\cdot, 0)$. Let us denote (10) as the *convexified quantization energy*. With this formulation the penalization is zero when the current residual is within the quantization error, while it is quadratic when the estimated sharp image u is far from the solution. Overfitting the observed image further than the quantization is thus avoided and model can successfully restore the image even with a low regularization weight.

Gamma correction. Since images are stored in a non-linear color space, through the use of gamma correction, the deconvolution cannot be performed directly. Indeed, if not properly handled, the gamma correction produces ringing around strong edges during deconvolution [11].

The usual way to deal with gamma correction is to apply the inverse function directly on the observed image [6], leading to

$$D_{\gamma\text{inv}}(u) = \|u * k - v^\gamma\|^2. \quad (11)$$

In this case, the model is fitted in linear space.

In this work, we argue that the data fitting should be computed directly in the non-linear color space,

$$D_\gamma(u) = \|(u * k)^{\frac{1}{\gamma}} - v\|^2, \quad (12)$$

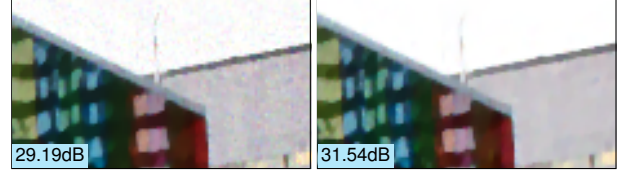


Fig. 2: Effect of a gamma corrected data fitting term. On the left, the deconvolution is performed in linear space; on the right, it is performed in gamma-corrected space.

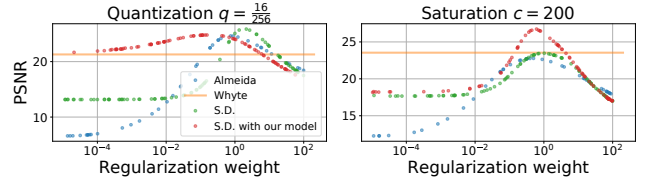


Fig. 3: Numerical results on individual degradations. Both plots indicates the behavior of the methods for various regularization weights under a large degradation.

where v is the observed image in non-linear space, and u is the restored sharp image in linear space. Fitting the model in the non-linear space reduces the importance of bright regions and improves the restoration of dark regions, which are more sensitive to the eyes. The effects of fitting the data in the gamma corrected space are easily visible to the human eye, as shown in Fig. 2.

Furthermore, Gregson *et al.* [10] proposed to adapt the TV regularization in order to account for the non-linearity of the eye sensitivity. This is done by defining a new regularizer using a 3×3 neighborhood that computes absolute differences in the non-linear space. As such, the noise, amplified in dark region due to the gamma correction, is better taken into account. For our experiments, we use a similar regularization, expressed as $\text{TV}^\gamma(u) = \text{TV}(u^{\frac{1}{\gamma}})$, which is the total variation of the gamma-corrected image.

Model composition. We have seen how to independently address deconvolution under saturation, quantization and gamma correction. We now propose a straightforward data term that combines these degradations in a single one,

$$D_{\text{full}}(u) = \left\| \left(\left(\min(c, u * k) \right)^{\frac{1}{\gamma}} - v \right)_+ \right\|^2. \quad (13)$$

While saturation is independent from the others, quantization and gamma correction interact with each other. Quantization in non-linear space leads to non-uniform quantization in the linear space. In dark regions, where the human eye is sensitive, the quantization is less than one graylevel and the methods are usually unaffected. However, in bright regions, the gamma correction compresses the dynamic and the quantization leading to larger errors.

For this reason, a simple model such as the naive gamma

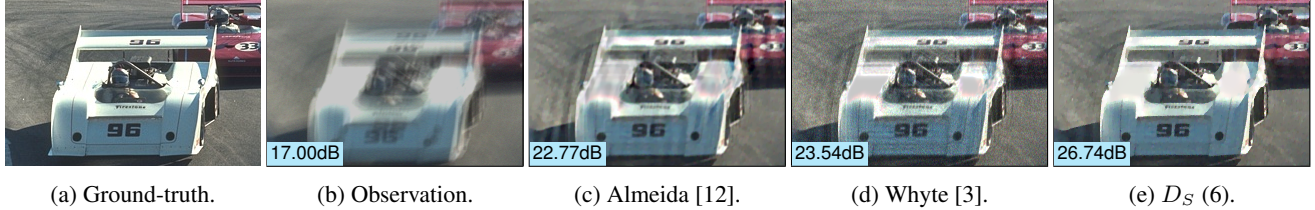


Fig. 4: Large saturation results (image clipped at intensity 200). The proposed model present less artifacts than [12] and [13].

inversion (11) that considers the gamma correction invertible even though the image is quantized, produces artifacts especially around bright regions. Our model effectively handles the interactions between all degradations.

4. EXPERIMENTS

First, we study the effectiveness of the different models individually. Our results are compared to two state-of-the-art methods. We compared our results with the methods of Almeida *et al.* [12] and Whyte *et al.* [3]. The first one uses the TV prior, which is representative of the literature, and handles accurately the boundary conditions. The second one is based on the Richardson-Lucy deconvolution algorithm [14], which is more robust to ringing, and in this version, also handles boundary conditions as well as saturation. Then, we present qualitative and quantitative results on a synthetic but realistic dataset, and show that modeling the complete degradation pipeline significantly improves the results.

Individual degradations. We evaluate two of our degradation models: saturation and quantization. Gamma correction is not evaluated individually as it can be directly inverted if no other degradation is present. For each modality, we apply the forward model to images of BSDS300 [15], vary the strength of the degradation and record the best PSNR obtained by optimizing the regularization weight. Fig. 3 shows the PSNR obtained by varying the regularization weight for the different methods under quantization with 16 levels and saturation at intensity 200. We note that, while the PSNR of our method for quantization is slightly lower than the others, it is more stable when sweeping the regularization weight. For saturation, our model clearly outperforms both [12] and [3], this is confirmed by the qualitative evaluation shown in Fig. 4.

A realistic model. To assess the gain of our individual models over a traditional model that does not consider the degradations, we created a realistic dataset. The dataset was created from eight sharp natural images. The images are converted to linear space, by applying the inverse gamma curve, and sub-sampled to reduce the residual quantization and noise. Then, each image is synthetically blurred using one of the kernels of Levin *et al.* [16], and saturated by clipping the pixels at the 98th percentile. Images are converted back to the non-linear color space, where additive white Gaussian noise of $\sigma^2 = 5$ is added. Finally a quantization with $q = \frac{1}{256}$ is applied.

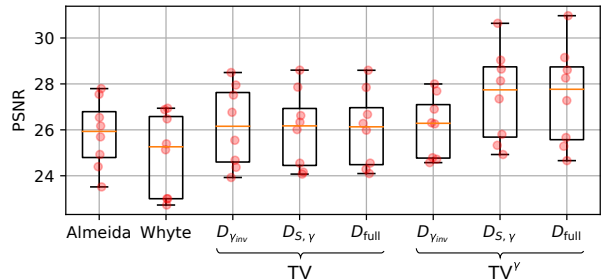


Fig. 5: PSNR statistics for the different models on a realistic dataset. $D_{S,\gamma}$ is a combination of Eq. (6) and Eq. (12). The orange bar indicates the median PSNR and the highest and lowest bar indicates the maximum and minimum PSNR obtained over the eight images.

Fig. 5 shows the PSNR results obtained with three models. From these results, we observe that considering all the degradations improves the results. Modeling the quantization does not always improve the PSNR but makes the minimization less sensitive to the regularization weight. We also report the results obtained by using the gamma corrected total variation. When combined with our model it yields a large gain in PSNR as well as in image quality (see Fig. 1). Due to space constraints, the dataset and the full resolution results are available on the project webpage: <https://goo.gl/oids7H>.

5. CONCLUSION

We proposed a non-blind image deconvolution method that handles non-linear degradations including, saturation, noise, gamma correction, and quantization. The optimization is possible thanks to a relaxed formulation of the quantization data term. The minimization of the resulting energy is performed by stochastic deconvolution [10]. Our experiments highlight the importance of modeling these realistic degradations present in the image processing pipeline. For the gamma correction, we show that the usual gamma inversion might introduce errors when the image is quantized in a non-linear color space, such as sRGB.

As future work, we would like to extend the method to blind image deconvolution, and study if kernel estimation can benefit from a more accurate image formation model. We also plan to explore the incorporation of other regularization terms, which could further improve the results.

6. REFERENCES

- [1] R. Wang and D. Tao, "Recent Progress in Image Deblurring," *ArXiv 1409.6838*, pp. 1–53, 2014.
- [2] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithm," *Physica D*, vol. 60, no. 1-4, pp. 259–268, 1992.
- [3] O. Whyte, J. Sivic, and A. Zisserman, "Deblurring Shaken and Partially Saturated Images," *International Journal of Computer Vision*, vol. 110, no. 2, pp. 185–201, 2014.
- [4] F. J. Anscombe, "The transformation of poisson, binomial and negative binomial data," *Biometrika*, vol. 35, pp. 246–254, 1948.
- [5] D. Kundur and D. Hatzinakos, "Blind image deconvolution," *IEEE signal processing magazine*, vol. 13, no. 3, pp. 43–64, 1996.
- [6] S. Cho, J. Wang, and S. Lee, "Handling outliers in non-blind image deconvolution," in *Proceedings of the IEEE International Conference on Computer Vision*, 2011, pp. 495–502.
- [7] J. Gregson, F. Heide, M. B. Hullin, M. Rouf, and W. Heidrich, "Stochastic Deconvolution," in *2013 IEEE Conference on Computer Vision and Pattern Recognition*. jun 2013, pp. 1043–1050, IEEE.
- [8] J. Gregson, M. Krimerman, M. B. Hullin, and W. Heidrich, "Stochastic tomography and its applications in 3d imaging of mixing fluids.," *ACM Trans. Graph.*, vol. 31, no. 4, pp. 52–1, 2012.
- [9] L. Xiao, J. Gregson, F. Heide, and W. Heidrich, "Stochastic Blind Motion Deblurring," *Image Processing, IEEE Transactions on*, vol. 24, no. 10, pp. 3071–3085, 2015.
- [10] J. Gregson, F. Heide, M. B. Hullin, M. Rouf, and W. Heidrich, "Stochastic deconvolution," in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2013, pp. 1043–1050.
- [11] Y.-W. Tai, X. Chen, S. Kim, S. J. Kim, F. Li, J. Yang, J. Yu, Y. Matsushita, and M. S. Brown, "Nonlinear camera response functions and image deblurring: Theoretical analysis and practice," *IEEE transactions on pattern analysis and machine intelligence*, vol. 35, no. 10, pp. 2498–2512, 2013.
- [12] M. S. C. Almeida and M. Figueiredo, "Deconvolving images with unknown boundaries using the alternating direction method of multipliers," *IEEE Transactions on Image Processing*, vol. 22, no. 8, pp. 3074–3086, 2013.
- [13] R. L. White, "Image Restoration Using the Damped Richardson-Lucy Method," *The Restoration of HST Images and Spectra II*, pp. 104–110, 1994.
- [14] W. H. Richardson, "Bayesian-Based Iterative Method of Image Restoration," *Journal of the Optical Society of America*, vol. 62, no. 1, pp. 55, 1972.
- [15] D. Martin, C. Fowlkes, D. Tal, and J. Malik, "A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics," in *Proc. 8th Int'l Conf. Computer Vision*, July 2001, vol. 2, pp. 416–423.
- [16] A. Levin, Y. Weiss, F. Durand, and W. T. Freeman, "Understanding and evaluating blind deconvolution algorithms," *2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2009*, pp. 1964–1971, 2009.