

Image Enhancement by Gradient Distribution Specification

Yuanhao Gong^{1,2} and Ivo F. Sbalzarini^{1,2}

¹MOSAIC Group, Faculty of Computer Science, TU Dresden, Dresden, Germany;

²MOSAIC Group, Center for Systems Biology Dresden (CSBD), Max Planck Institute of Molecular Cell Biology and Genetics, Dresden, Germany

Abstract. We propose to use gradient distribution specification for image enhancement. The specified gradient distribution is learned from natural-scene image datasets. This enhances image quality based on two facts: First, the specified distribution is independent of image content. Second, the distance between the learned distribution and the empirical distribution of a given image correlates with subjectively perceived image quality. Based on those two facts, remapping an image such that the distribution of its gradients (and therefore also Laplacians) matches the specified distribution is expected to improve the quality of that image. We call this process “image naturalization”. Our experiments confirm that naturalized images are more appealing to visual perception. Moreover, “naturalness” can be used as a measure of image quality when ground-truth is unknown.

1 Introduction

Image enhancement plays a fundamental role in image processing. The usual means for image enhancement is histogram equalization or one of its variants. However, the intensity histogram greatly varies with image contents, so that there is no simple mathematical model for it. In the absence of such a model, histogram equalization assumes a uniform prior distribution. This assumption is not usually valid, causing histogram equalization to fail in many cases.

This problem is circumvented when replacing the intensity histogram with the gradient histogram, which is remarkably invariant across natural-scene images [1–3]. This fact has previously been exploited for image denoising [1], motion deblurring [2, 3], and image restoration [4]. In this paper, we propose to use gradient distribution specification for image enhancement in a novel process we call “image naturalization”.

The concept of image naturalization is illustrated in Fig. 1. Given a learned prior distribution, the quality of images can be improved by gradient remapping such that the new gradient and Laplace fields satisfy the prior. The result image can be reconstructed by solving a single Poisson equation.

1.1 First-Order Prior

It is well known that the gradient distributions of natural-scenes image have a heavy tail in log-scale [2–4]. Traditionally, such gradient distributions are modeled as generalized Laplace distributions:

$$\log(p(x)) = -k\|x\|^\alpha + \beta, \quad (1)$$

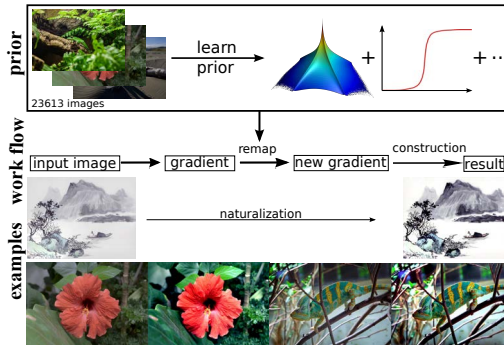


Fig. 1: Illustration of image naturalization.

where k , α , and β are parameters, x is the gradient, and p is the gradient probability distribution. This model includes several frequently used priors, including total variation (TV) ($\alpha = 1$) [5] and Hyper-Laplacian ($\alpha = 0.6$) [3].

This model, however, has several drawbacks. First, as shown in this paper, this model does not fit the data well. Second, all previous works treat k , α , and β as independent variables, violating the normalization of the distribution ($\int_{-\infty}^{+\infty} p dx \neq 1$). Third, the generalized Laplace model is computationally expensive to evaluate. Here, we present a new model that overcomes all of these issues.

1.2 Second-Order Prior

In addition to first-order (i.e., gradient) statistics, second-order statistics such as mean curvature (MC) [6] and Gaussian curvature (GC) [7, 8] can also be imposed as priors. Different orders of priors can be combined, e.g., in *Fields of Experts* models [9].

However, higher-order statistics may not be necessary in all cases. In Fig. 2, the correlation between $I(\mathbf{x})$ and $I(\mathbf{x} + \mathbf{r})$ is shown for different orders d of derivatives:

$$\text{Corr}(d, \mathbf{r}) := \text{correlation}(\nabla^d I(\mathbf{x}), \nabla^d I(\mathbf{x} + \mathbf{r})), \quad (2)$$

where $\mathbf{r} = (r, 0)$. The correlation reduces significantly for larger d . This explains why first- and second-order priors are so powerful for image processing, but higher-order derivatives do not necessarily improve the result. One reason is that the discrete image may not be higher-order differentiable. For image-processing tasks, second order has repeatedly been shown to be enough [8]. Therefore, we only consider derivatives up to the order of two.

Among all second order operators, the Laplace operator is the most attractive one since it leads to a Poisson equation, as does the gradient prior. Therefore, both priors can be imposed simultaneously by solving a single Poisson equation, for which efficient solvers are readily available.

1.3 Equivalence between Gradient Field and Original Image

The gradient field of an image is equivalent to the original image plus a single point constraint. The original image can hence be exactly reconstructed from its gradient




Fig. 2: $Corr(d, r)$ rapidly reduces for higher-order derivatives.

field [10–12]. An excellent review about signal processing in the gradient domain can be found in Ref. [11]. Some recent works in this area are described in Refs. [13, 14]. More details are given here in Section 4.2.

1.4 Our Contributions

- ★ We provide here new parametric models for the gradient and Laplace distributions of images. Instead of modeling the probability density functions (PDF) in log-scale, we model the cumulative distribution functions (CDF). We show that the resulting models are more accurate and computationally more efficient than previous models.
- ★ We show that the distance between the gradient/Laplace distributions of a given image and the learned prior distributions is correlated with image quality. Therefore, imposing these priors is expected to improve image quality.
- ★ We provide an algorithm for image enhancement by gradient/Laplace distribution specification through a remapping function. A nonlinear remapping function can be approximated by a linear function, which has only one scalar variable.
- ★ The proposed image enhancement process is parameter-free, which avoids manual parameter tuning.
- ★ First- and second-order priors are imposed simultaneously by solving a single Poisson equation.
- ★ We provide a simple scalar number that measures how close the gradient distribution of an image is to the prior. This number can be used to evaluate image quality in cases where ground-truth is unknown.

2 The Naturalization Prior

The naturalization prior proposed here is a linear combination of a gradient distribution prior and the consistent Laplace operator prior, i.e., the one corresponding to the divergence of the gradient. We learn these priors from a large dataset of natural-scene images. We then study the variability of the data around the priors. We provide novel parametric models for both priors and an efficient algorithm for imposing them on any given image. Finally, we define an “image naturalness factor” based on our models.

We use seven datasets of natural-scene images as shown in Table 1. Each image $I(x, y)$ was converted to 8-bit gray-scale. The gradient field \mathbf{G} is defined as

$$\mathbf{G}(x, y) = (\nabla_x I(x, y), \nabla_y I(x, y)), \quad (3)$$

where we use the first-order finite-difference approximations: $\nabla_x I = I(x + 1, y) - I(x, y)$ and $\nabla_y I = I(x, y + 1) - I(x, y)$. We use homogeneous Dirichlet boundary

Table 1: Natural-scene image datasets: sources and sizes.

Dataset	1 ^a	2 ^b	3 ^c	4 ^d	5 ^e	6 ^f	7 ^g	Total
#images	1005	1000	5063	832	1491	6033	8189	23613

conditions at the image borders. Due to the use of 8-bit gray-scale images, possible gradients are in the integer domain $[-255, 255] \times [-255, 255]$, where we can easily construct the two-dimensional histogram of \mathbf{G} . We use G^x and G^y to denote the respective components of \mathbf{G} .

The Laplace field L is defined as

$$L(x, y) = \Delta I(x, y), \quad (4)$$

where Δ is the Laplace operator, which is discretized using the second-order 5-point finite-difference stencil. Possible values are in the integer domain $[-1020, 1020]$.

In order to turn the histograms into probability distributions, we divide all bins by the total number of pixels mn in the image where m and n are the numbers of pixels along the x and y axes of the image. After aggregating data from all images in the dataset, we further normalize the histogram by the total number of images in the dataset. This yields the average distributions p_1^{pr} and p_2^{pr} of the gradient and Laplace operators. For color images, the priors are learned and applied separately for each color channel, performing all operations channel-wise.

2.1 Variability

We analyze how closely the natural-scene images in the training dataset match the average gradient and Laplace priors learned from them. The first row of Fig. 3 show histograms of the Root Mean Square (RMS) distances (left), Hellinger distance (middle), and Kullback Leibler divergence (right) between each image’s individual distribution and the prior. The second row of Fig. 3 shows the RMS distribution for the Laplace operator for grayscale (left) and color (right) images, respectively.

2.2 Gradient Distribution Models

Considering G^x and G^y individually, both satisfy the heavy-tail characteristic in log-scale, which can be modeled as a hyper-Laplacian distribution [3]. Those traditional one-dimensional models fit $p(G^x)$ in log-scale, which is why $\int_{-\infty}^{+\infty} p(G^x) dG^x = 1$

^a <http://www.vision.ee.ethz.ch/showroom/zubud/>

^b http://see.xidian.edu.cn/faculty/wsdong/Data/Flickr_Images.rar

^c <http://www.robots.ox.ac.uk/vgg/data/oxbuildings/>

^d <http://www.comp.leeds.ac.uk/scs6jwks/dataset/leedsbutterfly/>

^e <http://lear.inrialpes.fr/jegou/data.php>

^f <http://www.vision.caltech.edu/visipedia/CUB-200.html>

^g <http://www.robots.ox.ac.uk/vgg/data/flowers/102/index.html>

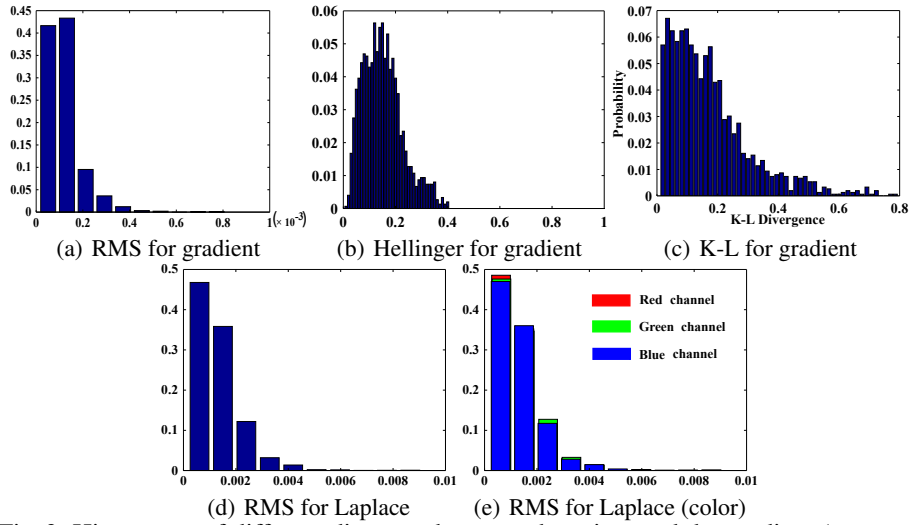


Fig. 3: Histograms of different distances between the priors and the gradient (top row) and Laplace (bottom row) distributions of individual training images.

cannot be guaranteed. Unlike previous works, we hence propose to do the modeling on the CDF instead.

The CDF of the gradient is defined as:

$$C(\mathbf{G}) = \int_{-255}^{G^x} \int_{-255}^{G^y} P((u, v)) \, du \, dv. \quad (5)$$

Observing the step property of $C(\mathbf{G})$, we propose to approximate the CDF with the parametric model:

$$\tilde{C}(\mathbf{G}) = \left(\frac{\text{atan}(T_1 G^x)}{\pi} + \frac{1}{2} \right) \left(\frac{\text{atan}(T_1 G^y)}{\pi} + \frac{1}{2} \right). \quad (6)$$

The choice of the atan function is motivated by the Student-T distribution or Cauchy distribution. This model has only one parameter to be fitted: T_1 . The fitting results are shown in Table 2.

Table 2: Fits of the parametric 2D CDF model to the image datasets.

Image set	1	2	3	4	5	6	7	all
T_1	0.37	0.26	0.38	0.35	0.56	0.37	0.7	0.46
SSE	20.71	23.11	19.08	23.7	22.94	19.64	22.97	18.75
R-square	0.9995	0.9995	0.9996	0.9995	0.9995	0.9996	0.9995	0.9996

The corresponding marginal model for the distribution of G^x (analogously G^y) is:

$$\log(P(G^x)) = \log\left(\frac{T_1}{\pi}\right) - \log\left(1 + (T_1 G^x)^2\right). \quad (7)$$




Fig. 4: Comparison of models in log scale, linear scale, and CDF.

Table 3: Comparison of our model (Eq. 7, top row) with generalized Laplace models (Eq. 1) for $\alpha = 0.6, 1, 2$ (rows two to four).

Image set	1	2	3	4	5	6	7
SSE	40.5	43.0	67.7	23.8	34.1	37.9	25.4
R^2	0.99	0.99	0.99	0.99	0.99	0.99	0.99
SSE	576	301	537	45.4	389	70.5	250
R^2	0.92	0.93	0.91	0.98	0.96	0.98	0.97
$SSE \times 10^{-3}$	1.86	3.01	3.02	3.95	2.34	3.90	3.95
R^2	0.74	0.30	0.52	0.13	0.81	0.10	0.57
$SSE \times 10^{-4}$	0.83	1.02	1.10	1.24	1.32	1.23	1.64
R^2	-0.12	-1.3	-0.72	-2.6	-0.046	-2.5	-0.75

We compare our model with other models in Fig. 4 and Table 3 using the optimal parameters for each model. In all cases, the present model describes the data best.

In Fig. 5, different images and their corresponding CDFs are shown. For the blurred image (Gaussian blur, $\sigma = 3$), the frequency of small gradients is increased. For the noisy image (10% Gaussian noise), the frequency of large gradients is increased. For the super-resolution (SR) image (upsampling factor 9), the frequency of small gradients is increased. For the bilateral filter ($w = 5$, $\sigma_s = 3$, $\sigma_c = 0.1$) and the guided filter ($r = 10$, $\epsilon = 0.01$), the frequency of small gradients is increased.

2.3 Laplace Distribution Models

Also for the distribution of the Laplace operator response, we use the CDF to model the statistics:

$$L(t) = \int_{-\infty}^t p(\Delta I(\mathbf{x})) d\Delta I(\mathbf{x}). \quad (8)$$

For the Laplace CDF, we propose the parametric model:

$$\tilde{L}(t) = \frac{\text{atan}(T_2 t)}{\pi} + \frac{1}{2}, \quad (9)$$

where T_2 is the only free parameter. The Laplace CDFs for different images are shown in Fig. 5. The figure also shows that the steps in the CDF indicate edge preservation in the spatial domain.

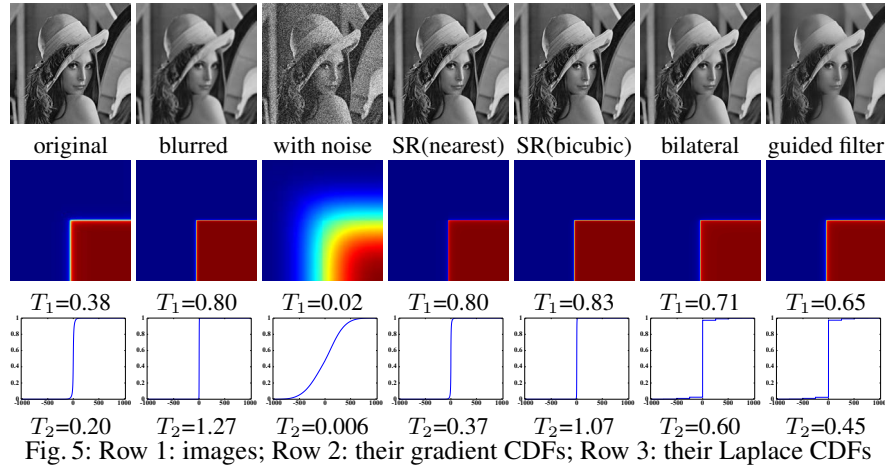


Fig. 5: Row 1: images; Row 2: their gradient CDFs; Row 3: their Laplace CDFs

We test the present parametric model by adding Gaussian noise with $\sigma = [0.02 : 0.02 : 0.8]$ to the Lena image and computing the difference $\tilde{L} - L$. The result is shown in Fig. 6. Even the maximum difference is small compared to the absolute value of L .




Fig. 6: Left: L , middle: \tilde{L} , right: $\tilde{L} - L$. The noise level σ is indicated by color.

2.4 Naturalness Factor and Image Naturalization

For a given image, it is easy to fit T_1 and T_2 using the above parametric distribution models. The corresponding values for the priors learned from the natural-scene dataset are: $T_1^{\text{pr}} = 0.38$ and $T_2^{\text{pr}} = 0.14$. Comparing the values of an image to these expected ones from natural-scene images tells how close the image is to a natural-scene one.

For any image I , the **naturalness factor** N_f is defined as

$$N_f = (1 - \theta) \frac{T_1}{T_1^{\text{pr}}} + \theta \frac{T_2}{T_2^{\text{pr}}}, \quad (10)$$

where $\theta \in [0, 1]$ is a weight parameter. The naturalness factor N_f^c of a color image is defined separately for each color channel c . The **naturalized image** I_n is generated from I such that $T_i \approx T_i^{\text{pr}}$ ($i \in \{1, 2\}$). This process is called **image naturalization**.

In our C++ implementation, we use the ternary search algorithm to find T_1 and T_2 . On a 2011 MacBook Pro, this code achieves 290 Mpixel/second for 8-bit three-channel color images. We provide a Matlab implementation in the supplementary material. Implementations in Matlab, C++, and Java will be publicly available from the MOSAIC Group web site mosaic.mpi-cbg.de after the conference.

3 Correlation with Image Quality

We show that the naturalization prior is correlated with subjectively perceived image quality. Therefore, image naturalization is expected to improve image quality. We test both the non-parametric priors and the above parametric models on the standard image quality assessment dataset LIVE [15], containing 779 images with five different types of distortions (degradations). The subjective image quality score (DMOS, difference mean opinion score) for each image is provided by LIVE. Our objective score is the average Hellinger distance (HD) between the ground truth image and distorted image, defined as:

$$score = \frac{1}{2}HD(p(\nabla I^{\text{true}}), p(\nabla I^{\text{distort}})) + \frac{1}{2}HD(p(\Delta I^{\text{true}}), p(\Delta I^{\text{distort}})). \quad (11)$$

We also define a score from the naturalization factor N_f . We compute N_f for both the ground-truth image and the distorted image and use the absolute difference as an objective quality score:

$$score_{N_f} = |N_f^{\text{true}} - N_f^{\text{distort}}|. \quad (12)$$

In the case where ground truth is unknown or unavailable, the score is defined with respect to the naturalization prior:

$$score_{pr} = \frac{1}{2}HD(p_1^{\text{pr}}, p(\nabla I^{\text{distort}})) + \frac{1}{2}HD(p_2^{\text{pr}}, p(\Delta I^{\text{distort}})). \quad (13)$$

Table 4: Pearson, Spearman, and Kendall correlation coefficients between DMOS and the different image quality measures on the LIVE benchmark dataset.

	PSNR	SSIM	FSIM	SFF	<i>score</i>	$N_f(\theta=0)$	$N_f(\theta=1)$	<i>score_{pr}</i>
PCC	-0.8585	-0.8252	-0.8586	-0.8126	0.8687	0.6692	0.5199	0.761
SCC	-0.8756	-0.9104	-0.9634	-0.9649	0.8630	0.7118	0.5817	0.706
KCC	-0.6865	-0.7311	-0.8337	-0.8365	0.6745	0.5123	0.4134	0.522

The correlations between DMOS and all of these scores are shown in Fig. 7 and Table 4. Correlations are reported as Pearson linear correlation coefficients (PCC), Spearman rank-order correlation coefficients (SCC), and Kendall rank-order correlation coefficients (KCC). We use the HD in our score because it outperforms all other tested distances (L_2 , L_1 , \cos , χ^2) for these correlation coefficients. We compare the

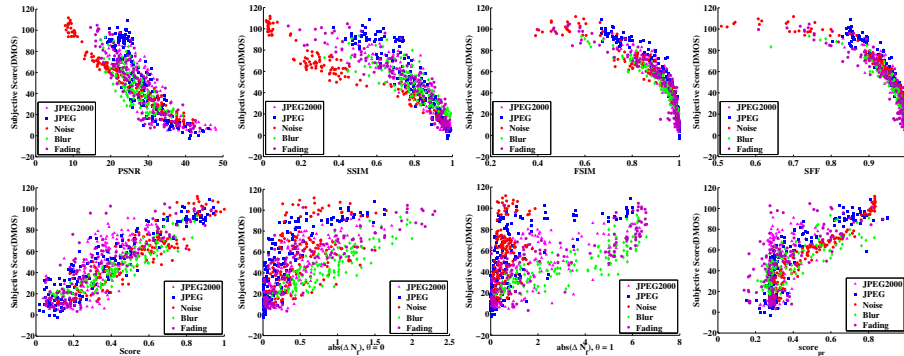


Fig. 7: Correlation of different objective image quality measures with subjectively perceived image quality (DMOS) on the LIVE benchmark dataset. First row: PSNR, SSIM, FSIM, SFF; second row: $score$, $score_{N_f}$ for $\theta = 0$ and 1, $score_{pr}$.

present score with the state-of-the-art image quality assessment methods SSIM [15], FSIM [16], and SFF [17]. Among all approaches, our nonparametric method shows the best linearity between $score$ and DMOS, which is preferable for image enhancement tasks. Our result is comparable with PSNR, but does not require knowing the noise level in the image, or any other geometric information about the ground truth.

Taken together, these results confirm that imposing the gradient and Laplace distribution priors is expected to improve image quality, therefore making them good candidates for image enhancement priors.

4 Image Naturalization Algorithm

The image naturalization process consists in remapping the gradient field of the image to satisfy the prior distributions, followed by reconstructing the naturalized output image from the remapped field.




Fig. 8: Two example images and their nonlinear Map functions (color indicates the different color channels).

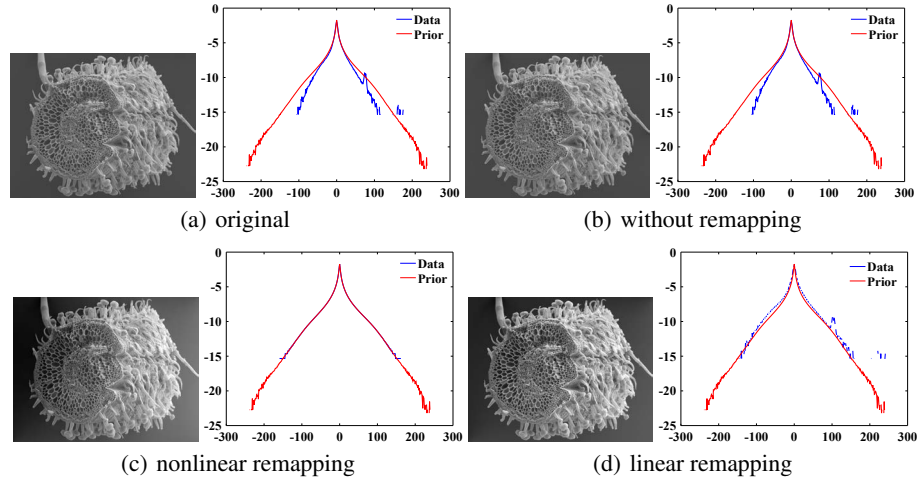


Fig. 9: Comparison of different types of remapping. The reconstruction errors (RMS) are 2.0, 33, and 23, respectively. The corresponding gradient distributions after remapping are shown in comparison with the naturalization prior next to the image.

4.1 Gradient Distribution Specification

Let Map be a function that maps the gradient field \mathbf{G} to a new gradient field \mathbf{G}_n , which satisfies the naturalization prior:

$$\mathbf{G}_n = Map(\mathbf{G}), \quad s.t. \quad p(\mathbf{G}_n) = p_1^{pr}, \quad p(\nabla \cdot \mathbf{G}_n) = p_2^{pr}. \quad (14)$$

In general, Map can be non-parametric and nonlinear. More specifically, we use here modified exact histogram specification [18] as the Map function. Two example images and their nonlinear non-parametric Map functions are shown in Fig. 8.

4.2 Image Reconstruction

We reconstruct the naturalized image I_n from the remapped gradient field by solving the variational model:

$$\min \left\{ \int_{\mathbf{x} \in \Omega} \|\nabla I_n - \mathbf{G}_n\|_2^2 d\mathbf{x} \right\}, \quad (15)$$

which leads to the Poisson equation

$$\Delta I_n = \nabla \cdot \mathbf{G}_n. \quad (16)$$

This equation can be solved efficiently by FFT-based algorithms or wavelet solvers. A short summary of commonly available Poisson solvers is shown in Table 5.

Reconstructing an image from its gradient field is accurate and computationally efficient. An example is shown in Fig. 9. The original image (Fig. 9a) is an 8-bit grayscale

^h dense Cholesky decomposition

ⁱ sparse Cholesky decomposition

Table 5: Summary of Poisson solvers (we implemented FFT and wavelet solvers).

Solver	Cholesky ^h	Jacobi	Gauss-Seidel	SOR
Type	direct	iterative	iterative	iterative
	$(mn)^3$	$(mn)^2$	$(mn)^2$	$(mn)^{3/2}$
Solver	Cholesky ⁱ	FFT	Multigrid	Wavelet
Type	direct	direct	iterative	direct
	$(mn)^{3/2}$	$(mn)\log(mn)$	(mn)	(mn)

image. The absolute pixel-wise RMS of the reconstruction without remapping (Fig. 9b) is 2.0 while the average intensity value is 105. The size of the image is 1881×2400 pixels and the reconstruction took 3.5 seconds using our MATLAB implementation of the wavelet Poisson solver on a 2 GHz Intel i7 processor. The naturalized image after nonlinear remapping is shown in Fig. 9c.

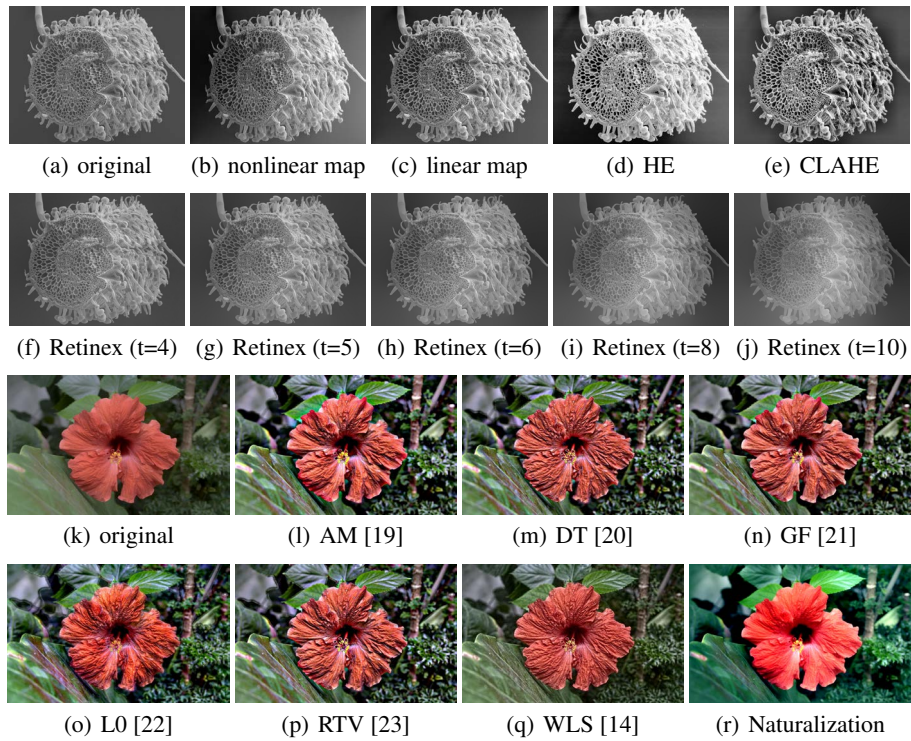


Fig. 10: Comparison of image enhancement methods: nonlinear and linear naturalization, HE (Histogram Equalization), CLAHE (Contrast Limited Adaptive Histogram Equalization), Retinex (for different parameters), AM (Adaptive Manifold), DT (Domain Transform), GF (Guided Filter), L0 norm, RTV (Relative Total Variation), WLS (Weighted Least Square), and Naturalization.

4.3 Linear Approximation of Map

Our one-parameter model results in a linear approximation of the remapping function Map :

$$G_n = N_f G, \quad (17)$$

which is equivalent to scaling the original image:

$$I_n = N_f I. \quad (18)$$

Simply scaling the image with N_f significantly accelerates the naturalization process while still avoiding halo artifacts in the result. An example of a naturalized image after linear remapping is shown in Fig. 9d.

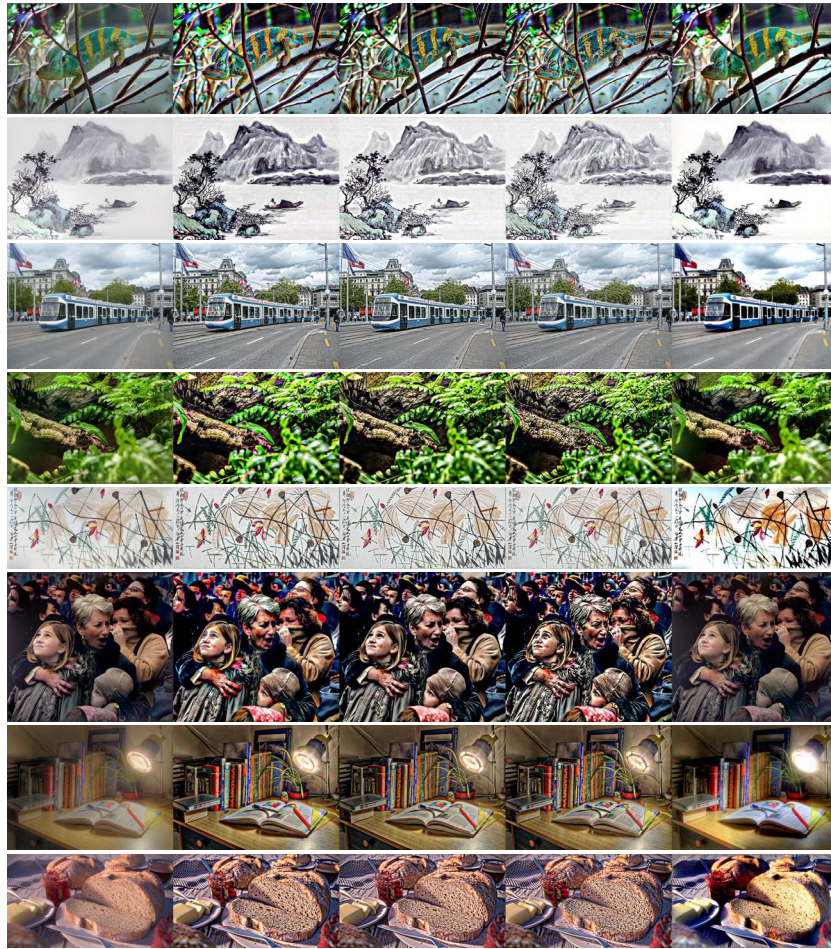


Fig. 11: Comparison of images (left column) enhanced with AM (column 2), GF (column 3), RTV (column 4), and naturalization ($\theta = 0.5$, column 5).



Fig. 12: Objective evaluation of image-processing results: N_f for each color channel (in RGB order). Better-quality results show values closer to 1.

5 Experiments

In the past decades, image smoothing and sharpening have received a lot of attention. However, there is to date no objective measure to decide if an image should be smoothed or sharpened. N_f can be used to automatically determine whether an image should be smoothed or sharpened. It is clear from Eq. 17 that an image gets smoothed for $N_f < 1$ and sharpened for $N_f > 1$. Two examples are shown in Fig. 10 and compared with other methods using their respective default parameters. More results and comparisons are shown in Fig. 11.

5.1 Objective Image Quality Evaluation

In addition to image naturalization, the prior can be used for objective evaluation of results from various algorithms, such as deblurring, image editing, image synthesis, etc. Significant progress has recently been made in image deblurring, High Dynamic Range (HDR) compression, Poisson image editing, etc. In all these applications, the ground truth is unknown. The naturalness factor can provide an objective evaluation of the results. To the best of our knowledge, this is the first proposal of an objective measure to evaluate such processing results.

In Fig. 12, the first two rows show deblurring results (original, [24], [25], [26]); the third row considers object adding and removing by Poisson editing (original, add, original, remove); the fourth row shows HDR results ([27]); the fifth row scene rendering results (realistic image synthesis); the sixth row gives an evaluation of a panoramastitched image. It is obvious that visually pleasing results satisfy the naturalization prior better.

6 Conclusion

We have presented novel parametric models for the gradient and Laplace distributions of natural-scene images. We have shown that the models are more accurate than previous models and are easy to compute. We have further shown that prior models learned from natural-scene images are correlated with image quality. Therefore, remapping an image to satisfy the learned gradient prior yields a quality-improved image. We call this novel image enhancement method “image naturalization”. The remapping function has a linear approximation, which further simplifies the algorithm. The corresponding linear factor is called the “naturalness factor” and can be used to evaluate the quality of image-processing result when ground truth is unavailable.

References

1. Zhu, S., Mumford, D.: Prior learning and Gibbs reaction-diffusion. *IEEE Trans. Pattern Anal. & Machine Intell. (PAMI)* **19** (1997) 1236–1250
2. Shan, Q., Jia, J., Agarwala, A.: High-quality motion deblurring from a single image. *Acm Transactions On Graphics* **27** (2008) 73

3. Krishnan, D., Fergus, R.: Fast image deconvolution using hyper-Laplacian priors. *Adv. Neural Inform. Proc. Sys.* **22** (2009) 1–9
4. Cho, T.S., Zitnick, C.L., Joshi, N., Kang, S.B., Szeliski, R., Freeman, W.T.: Image restoration by matching gradient distributions. *IEEE Trans. Pattern Anal. & Machine Intell. (PAMI)* **34** (2012) 683–694
5. Rudin, L.I., Osher, S., Fatemi, E.: Nonlinear total variation based noise removal algorithms. *Physica D* **60** (1992) 259–268
6. El-Fallah, A.I., Ford, G.E.: Mean curvature evolution and surface area scaling in image filtering. *IEEE Trans. Image Proc.* **6** (1997) 750–753
7. Lee, S.H., Seo, J.K.: Noise removal with Gauss curvature-driven diffusion. *IEEE Trans. Image Proc.* **14** (2005) 904–909
8. Gong, Y., Sbalzarini, I.F.: Local weighted Gaussian curvature for image processing. *Intl. Conf. Image Proc. (ICIP)* (2013) 534–538
9. Roth, S., Black, M.J.: Fields of experts. *IJCV* **82** (2009) 205–229
10. Fattal, R., Lischinski, D., Werman, M.: Gradient domain high dynamic range compression. *ACM Trans. Graph.* **21** (2002) 249–256
11. Agrawal, A., Raskar, R.: Gradient domain manipulation techniques in vision and graphics. *ICCV short course* (2007)
12. Gong, Y., Paul, G., Sbalzarini, I.F.: Coupled signed-distance functions for implicit surface reconstruction. In: *IEEE Intl. Symp. Biomed. Imaging (ISBI)*. (2012) 1000–1003
13. Fattal, R.: Edge-avoiding wavelets and their applications. *ACM Trans. Graph.* **28** (2009) 1–10
14. Farbman, Z., Fattal, R., Lischinski, D., Szeliski, R.: Edge-preserving decompositions for multi-scale tone and detail manipulation. *ACM Trans. Graph.* **27** (2008) 67:1–67:10
15. Wang, Z., Bovik, A., Sheikh, H., Simoncelli, E.: Image quality assessment: from error visibility to structural similarity. *Image Processing, IEEE Transactions on* **13** (2004) 600–612
16. Zhang, L., Zhang, L., Mou, X., Zhang, D.: Fsim: A feature similarity index for image quality assessment. *Ieee Transactions On Image Processing* **20** (2011) 2378–2386
17. Chang, H.W., Yang, H., Gan, Y., Wang, M.H.: Sparse feature fidelity for perceptual image quality assessment. *Ieee Transactions On Image Processing* **22** (2013) 4007–4018
18. Coltuc, D., Bolon, P., Chassery, J.M.: Exact histogram specification. *IEEE Transactions on Image Processing* **15** (2006) 1143–1152
19. Gastal, E.S.L., Oliveira, M.M.: Adaptive manifolds for real-time high-dimensional filtering. *ACM TOG* **31** (2012) 33:1–33:13
20. Gastal, E.S.L., Oliveira, M.M.: Domain transform for edge-aware image and video processing. *ACM TOG* **30** (2011) 69:1–69:12 *Proceedings of SIGGRAPH 2011*.
21. He, K., Sun, J., Tang, X.: Guided image filtering. *ECCV 2010* (2010) 1–14
22. Xu, L., Lu, C., Xu, Y., Jia, J.: Image smoothing via L0 gradient minimization. *ACM Trans. Graph.* (2011)
23. Xu, L., Yan, Q., Xia, Y., Jia, J.: Structure extraction from texture via relative total variation. *ACM Trans. Graph.* **31** (2012) 139:1–139:10
24. Hirsch, M., Schuler, C.J., Harmeling, S., Schölkopf, B.: Fast removal of non-uniform camera shake. In: *ICCV*. (2011) 463–470
25. Gupta, A., Joshi, N., Zitnick, C.L., Cohen, M.F., Curless, B.: Single image deblurring using motion density functions. In: *ECCV* (1). (2010) 171–184
26. Xu, L., Zheng, S., Jia, J.: Unnatural l0 sparse representation for natural image deblurring. In: *CVPR. CVPR '13*, Washington, DC, USA, IEEE Computer Society (2013) 1107–1114
27. Farbman, Z., Fattal, R., Lischinski, D.: Convolution pyramids. *ACM Trans. Graph.* **30** (2011) 175:1–175:8