Synthetic images as a regularity prior for image restoration neural networks

Authors

Raphaël Achddou

Yann Gousseau

Saïd Ladjal

Introduction 1

Deep Neural Networks have recently surpassed other image restoration methods which rely on hand-crafted priors.However, they usually require large databases and need to be retrained for each new acquisition modality or perturbation.

Our contributions :

- showing that we can reach near optimal performances by training them on a synthetic dataset made of realizations of a dead leaves model [1], both for image denoising and super-resolution.
- showing that training a network with a mix of natural and synthetic images does not affect results on natural images while improving the results on synthetic images, which are classically used to evaluate the preservation of textures[2].

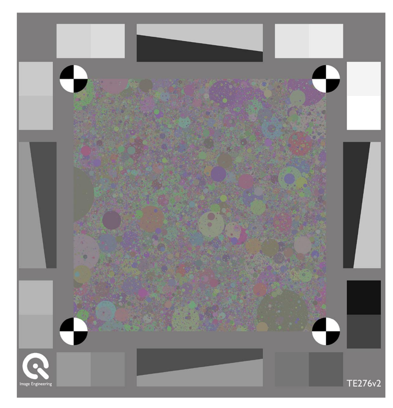

a dead leave target for camera evaluation

Dead leaves Image Generation 2

Image generation algorithm :

Parameters : a natural color image,r_min, r_max, alpha, *s*

Output : an image X

While the image plane is not covered :

- draw a random position (x,y) for the center of the next disk, draw a random color c from the natural image histogram, draw a radius r from the power law :

$$f(r_{min}, r_{max} r,) = \frac{r_{min}^{1-\alpha} - r_{max}^{1-\alpha}}{1-\alpha} * \mathbb{1}(r_{min} \le r \le r_{max}) r^{-\alpha}$$

- superimpose the associated color disk on the existing image X

Downscale X with a factor *s*

Example of the formation of a dead leave image

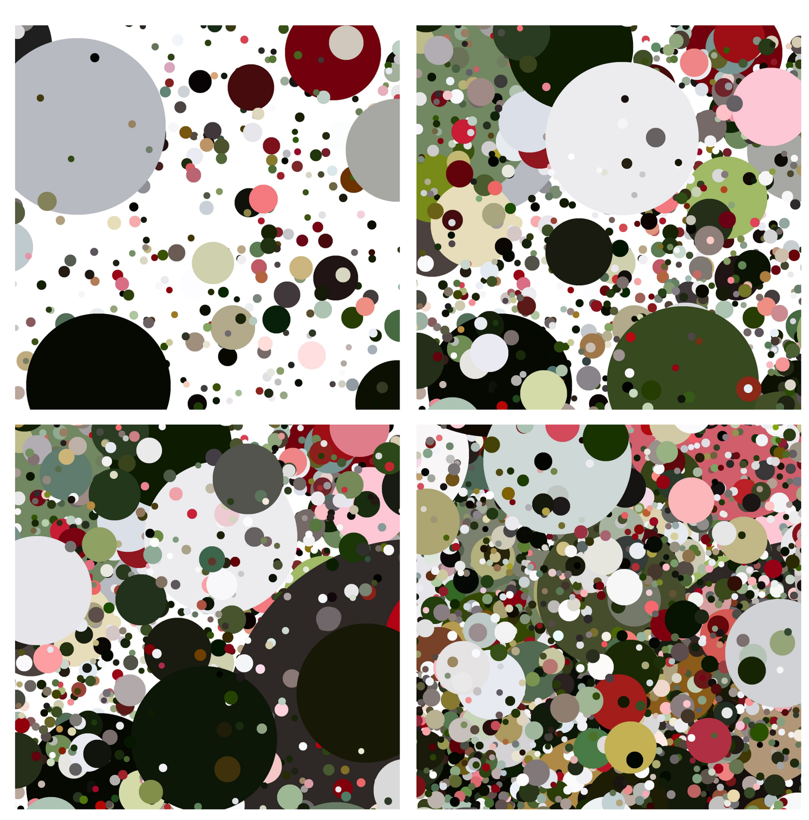

Role of the disk size parameters

- large r_min or small alpha : homogeneous zones, regular edges
- small r_min or large alpha : micro-textures and fine details

Fixed alpha = 3, r_max = 2000

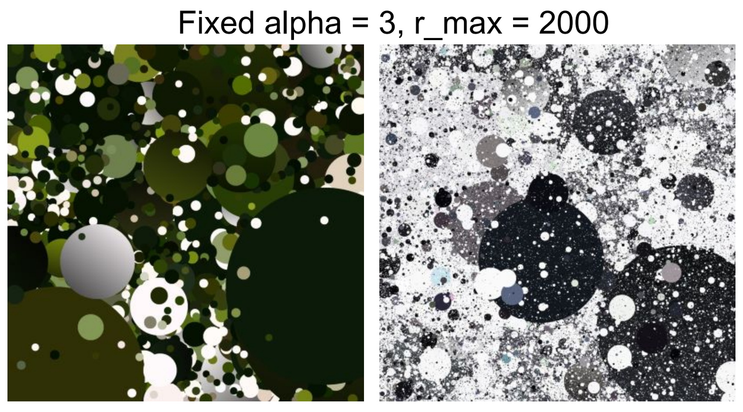

r_min = 20 r_min = 1

Fixed r_min = 1, r_max = 2000

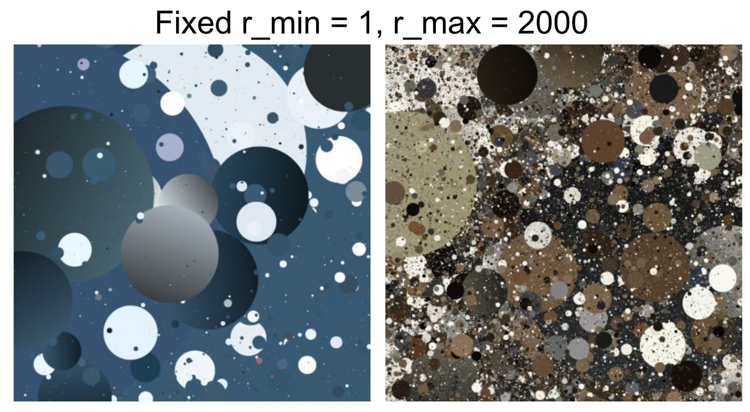

alpha = 2 alpha = 3

Importance of the color space

- sampling colors from a natural image : more natural looking images, and better performances
- sampling the RGB cube uniformly : to artificial colors

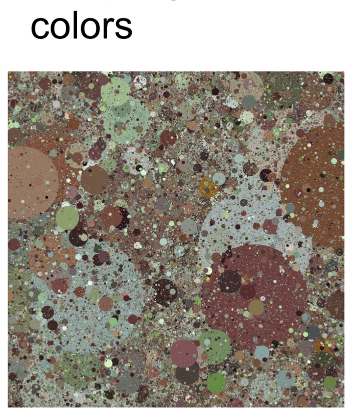

DL image with natural colors

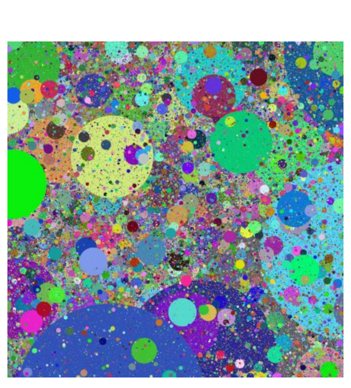

DL image with random colors

Justification of the downscaling step

- without downscaling : artificially sharp edges, no sub-pixel sized objects

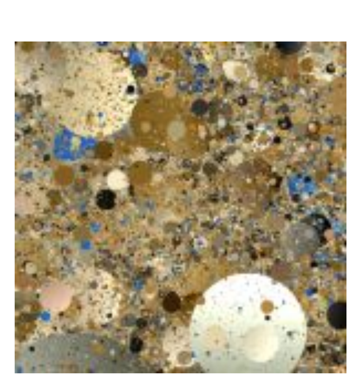

(150,150) crop on a downscaled dead leaves image

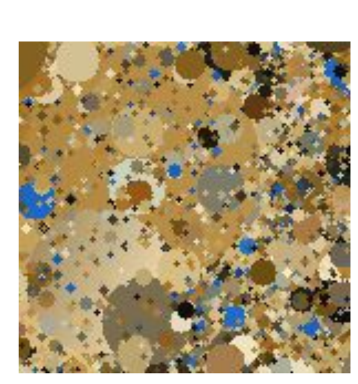

(150,150) crop on a dead leave image with no downscaling

Experiments 3

Experimental set-up:

- Architecture : FFDNet[3] for the denoising task and RDN[4] for the super-resolution, both state-of-the-art in their respective tasks
- Same training procedure
- Different training sets :
  - Synthetic sets : gaussian noise, gaussian random fields[5], single radius dead leaves, dead leaves dataset with : (⅓ r_min = 16,⅔ r_min = 1, alpha = 3)
  - Natural image sets : Waterloo database or ⅓ dead leaves ⅔ Waterloo
- Test on three datasets : 2 natural image sets (CBSD68,Kodak24) and a synthetic dead leaves test set

Denoising results

- The best synthetic image model is our dead leaves model both visually and numerically(0.6dB PSNR gap with SOTA).
- Training with a mix of dead leaves and natural images leads to the same performances on natural images but better performances on dead leaves images.

σ	Dataset	CBSD68	Kodak24	Dead leaves testset
$\sigma = 25$	White Noise	19.52/0.416/2.386	19.68/0.365/2.502	20.36/0.607/2.043
	Gaussian field	29.63/0.845/1.402	30.24/0.835/1.471	26.23/0.826/1.254
	DL $r = 100$	29.56/0.820/1.218	30.49/0.819/1.024	26.13/0.799/1.263
	Dead leaves	30.58/0.867/0.711	31.27/0.859/0.739	27.46/0.865/0.573
	Mix	31.07/0.881/0.639	31.98/0.876/0.603	27.33/0.860/0.567
	Natural Images	31.09/0.882/0.629	32.00/0.878/0.599	27.05/0.851/0.576
$\sigma = 50$	White Noise	15.58/0.247/4.682	15.71/0.209/4.785	16.24/0.387/2.932
	Gaussian field	26.68/0.738/2.203	27.41/0.737/2.353	23.31/0.694/2.158
	DL $r = 100$	26.85/0.720/1.563	27.91/0.739/1.314	23.24/0.654/2.005
	Dead leaves	27.40/0.762/1.088	28.21/0.765/1.154	24.21/0.737/1.020
	Mix	27.86/0.782/0.997	28.86/0.789/0.985	24.12/0.732/1.015
	Natural Images	27.87/0.786/0.991	28.89/0.792/0.978	23.90/0.722/1.053

Table 1 : Numerical results of our different trainings of FFDNet on 3 test sets. Best results in blue, second in red.3 metrics evaluated (PSNR/SSIM/Pieapp [6])

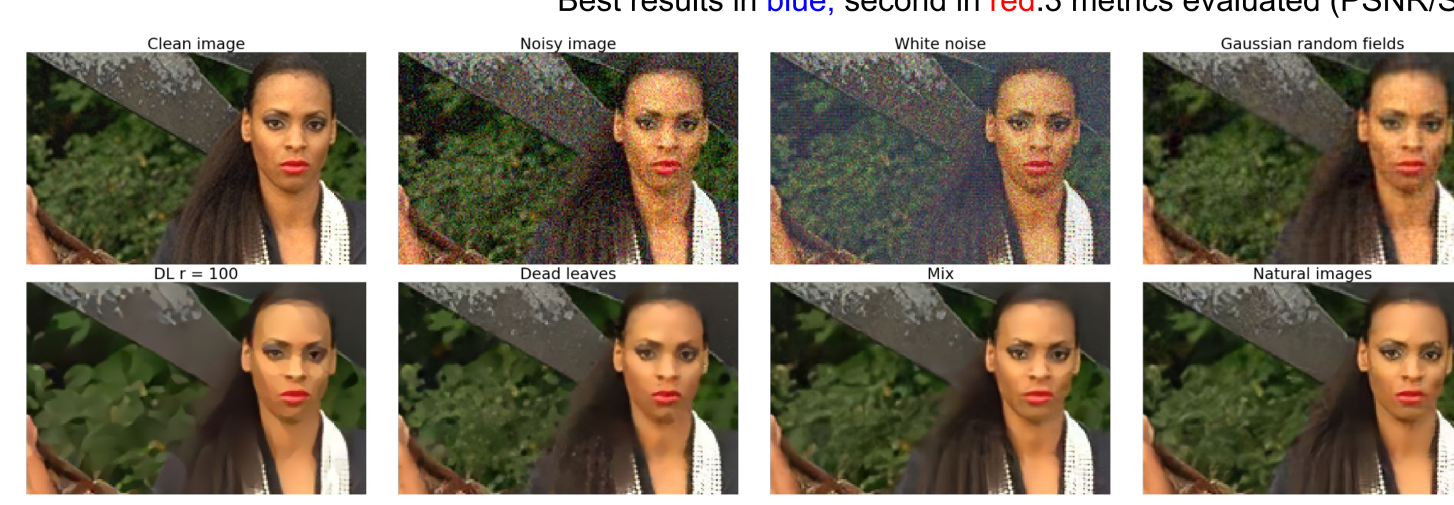

Visual comparison of the different denoising results

Super-resolution results

- Results behind but close to the state-of-the-art in PSNR
- Very close results visually, except some small artifacts on lines which are absent in the synthetic dataset

Dataset	Set 5		Set 14	
scale	×2	×3	×2	×3
Dead leaves	36.76	33.82	32.93	30.42
Natural Images	38.18	34.71	33.88	30.73

Table 2: Numerical results of our different trainings of RDN on 2 test sets. Best results in blue.

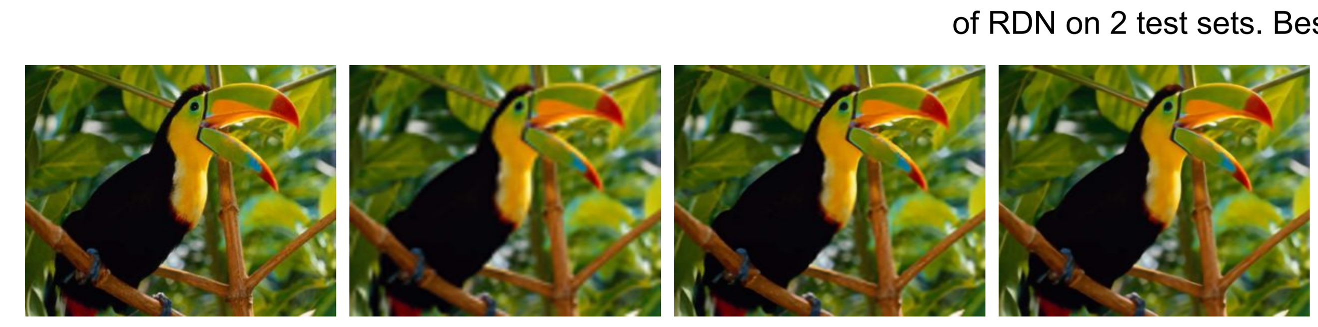

High resolution Image   Bicubic SR x3   RDN - DL x3   RDN Nat x3

Ablation study

To validate the choices we made for our generation's algorithm, we tried other parameters for the disks size and removed some steps to generated many ablated datasets. We then trained FFDNet on those sets to identify which were the key components of our method.

σ	DL-1	DL-2	DL-4	DL-8	DL-16	Rand. col	No sub	No blur	Final
25	31.03	31.03	31.09	31.07	30.98	29.99	30.79	31.25	31.27
50	27.98	27.96	28.04	28.06	28.05	27.16	27.74	28.20	28.21

Table 3: Numerical results of our different trainings of FFDnet on the ablated datasets. DL-i corresponds to r_min = i. We also tried random RGB colors, no down scaling, and no blur

References 5

[1] Cao, F., Guichard, F., Hornung, H.: Measuring texture sharpness of a digital camera. In: Digital Photography V. vol. 7250, p. 72500H. International Society for Optics and Photonics (2009)

[2] Bordenave, C., Gousseau, Y., Roueff, F.: The dead leaves model: a general tessellation modeling occlusion. Advances in applied probability 38(1), 31–46 (2006)

[3] Zhang, K., Zuo, W., Zhang, L.: Ffdnet: Toward a fast and flexible solution for cnnbased image denoising. IEEE Transactions on Image Processing 27(9), 4608–4622 (2018)

[4] Zhang, Y., Tian, Y., Kong, Y., Zhong, B., Fu, Y.: Residual dense network for image super-resolution. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2472–2481 (2018)

[5] Galerne, B., Gousseau, Y., Morel, J.M.: Micro-texture synthesis by phase randomization. Image Processing On Line 1, 213–237 (2011)

[6] Prashnani, E., Cai, H., Mostofi, Y., Sen, P.: Pieapp: Perceptual image-error assessment through pairwise preference. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1808–1817 (2018)

Conclusions and perspectives 4

- A first attempt to train image restoration neural networks on synthetic images
- Both for denoising and super-resolution, performances close to the original models, with an image generation algorithm with very few parameters
- Training on synthetic and natural data : same performances on natural images but much better performances on dead leaves images, opening the door to better imaging devices when tested on standard test patterns such as [1].
- **Perspectives :**
  - reduce the number of color parameters either with gaussian mixture models or by sampling the horseshoe color space
  - train a restoration model for a specific imaging device with dead leaves images

Contact raphael.achddou@telecom-paris.fr

Site web https://perso.telecom-paristech.fr/rachddou