



HAL
open science

A statistically constrained internal method for single image super-resolution

Pierrick Chatillon, Yann Gousseau, Sidonie Lefebvre

► **To cite this version:**

Pierrick Chatillon, Yann Gousseau, Sidonie Lefebvre. A statistically constrained internal method for single image super-resolution. International Conference on Pattern Recognition, IAPR, Aug 2022, Montréal, Canada. pp.1322-1328, 10.1109/ICPR56361.2022.9956498 . hal-03766523

HAL Id: hal-03766523

<https://hal.science/hal-03766523>

Submitted on 1 Sep 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A statistically constrained internal method for single image super-resolution

Pierrick Chatillon

pierrick.chatillon@onera.fr

DOTA, ONERA, Université Paris Saclay

F-91123 Palaiseau-France,

LTCI, Télécom Paris, IP Paris

Yann Gousseau

yann.gousseau@telecom-paris.fr

LTCI, Télécom Paris, IP Paris

19 place Marguerite Perey

91120 Palaiseau, France

Sidonie Lefebvre

sidonie.lefebvre@onera.fr

DOTA, ONERA, Université Paris Saclay

F-91123 Palaiseau-France

Abstract—Deep learning based methods for single-image super-resolution (SR) have drawn a lot of attention lately. In particular, various papers have shown that the learning stage can be performed on a single image, resulting in the so-called internal approaches. The SinGAN method is one of these contributions, where the distribution of image patches is learnt on the image at hand and propagated at finer scales. Now, there are situations where some statistical a priori can be assumed for the final image. In particular, many natural phenomena yield images having power law Fourier spectrum, such as clouds and other texture images. In this work, we show how such a priori information can be integrated into an internal super-resolution approach, by constraining the learned up-sampling procedure of SinGAN. We consider various types of constraints, related to the Fourier power spectrum, the color histograms and the consistency of the upsampling scheme. We demonstrate on various experiments that these constraints are indeed satisfied, but also that some perceptual quality measures can be improved by the proposed approach.

I. INTRODUCTION

Single image super-resolution has recently benefited from tremendous progresses thanks to the use of convolutional neural networks, from the early SRCNN architecture [1], through many methods improving the PSNR of SR, including EDSR [2], RDN [3] or the recent state-of-the-art work SwinIR [4]. Such methods usually rely on an extensive training on generic image databases and are sometimes referred to as *external methods*. More recently, it has been shown that the mere structure of convolutional neural networks provides a good enough prior for several image restoration tasks [5]. Further, and taking advantage of the learning power of Generative Adversarial Networks (GAN), methods have been proposed to train a super-resolution architecture on a single image, most notably through the SinGAN approach [6]. Such methods are often coined as *internal methods*, and can be seen as an extension to neural networks of the general principles underlying patch-based approaches exploiting the statistical property of patches from the input image, such as [7] or [8]. While being slightly less efficient than external methods, internal methods have the advantage of not necessitating an auxiliary image database that should have the same statistical property as the images to be restored.

In this work, we take interest in the application of internal super-resolution methods to a specific type of images,

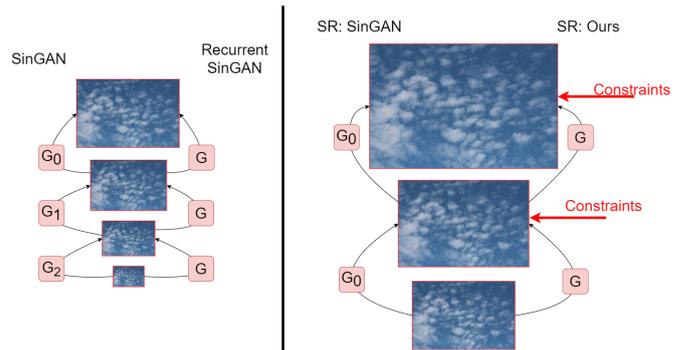


Fig. 1. Overview of the training (left) and SR procedure (right) for SinGAN and recurrent SinGAN [14]. This illustrates how we stabilize SinGAN’s SR method by adding statistical constraints.

namely “colored noise” images that is, images for which the power spectrum decays as a power function $1/f^\beta$. A common such example is pink noise, also known as $1/f$ noise. While such a decay is a general property of natural images [9], we here consider images for which this behavior is particularly accurate and isotropic. Colored noise processes are particularly abundant in nature [10], and can be encountered, as far as dimension two is concerned, in images of fracture patterns [11], glacier surface [12] and of clouds, which are our main application cases. Such processes have also been used as a simple yet very efficient source of textures in computer graphic applications [13].

Generating large, realistic and high resolution cloud fields is important either in the context of evaluating the performance of optics sensors or to benchmark radiative transfer models [15]. Turbulent flows simulations, such as the large-eddy ones, are expensive and stochastic models, based on fractal or Fourier framework, have been proposed in order to simulate more quickly realistic cloud fields. Examples include the bounded cascade model [16], the iterative amplitude adapted Fourier transform algorithm [15], and the tree-driven mass accumulation process (tdMAP) model [17]. They exploit the scale invariant properties observed in real clouds: the power spectra of the logarithm of their optical properties is characterized by a spectral slope of around $\frac{5}{3}$ [18] [17] [19]

from small scale (a few meters) to the outer scale (about one-hundred kilometers), where the spectrum becomes flat due to decorrelation. The goal of the present work is to develop an SR method in the spirit of these generative approaches, allowing the increase of resolution of colored noise images while preserving physically pertinent statistical property such as the power spectrum. To the best of our knowledge, such a model has not been proposed in the literature, except for [20], where a multi-fractal model is leveraged for the super-resolution of images from the quiet sun.

We also argue that imposing statistical constraint to an SR method is a way to overcome the limitation of PSNR training/evaluation. As pointed out in [21], PSNR alone is not enough, and two kind of metrics actually compete against each other: distortion metrics (e.g. PSNR) and perception metrics (e.g. perceptual quality [22], NIQE [23]). Optimizing for PSNR often leads to smooth images, when a perceptually good image needs some sharpness. Following this principle, many methods try to produce realistic SR outputs, using the GAN framework (introduced in [24]): SRGAN [25], ESRGAN [26] and EnhanceNet[27] combine the GAN training with distortion losses, perceptual quality losses and even texture aware losses [28].

In the experimental part of this paper, Section III, we will show that the chosen constraints indeed improves such perceptual (NIQE) and texture metrics.

II. METHOD

A. Statistical constraints

In this section, we describe the image statistics we impose to the SR framework. Some of these constraints are related to the specific SR task, while others are specific to colored noise images and related to desired physical property of the result.

Throughout the article, we denote as u the low resolution image (LR) and w the result of super-resolution by a factor $f > 1$. The notation ξ is for the frequency variable in the Fourier domain.

1) *Spectral constraints*: As explained in the introduction, we are interested in images whose spectrum decreases as a power function. Moreover, it is known that quasi-stationary signals such as cloud images yield uncorrelated discrete Fourier transform coefficients [29], and as a consequence, that the modulus follows a Rayleigh distribution. Therefore, the spectrum of the generated image is constrained so that :

- Past a radius r_0 in the Fourier domain, the mean modulus of \hat{w} taken on a circle \mathcal{C}_r of radius r decreases like $E_{\mathcal{C}_{r_0}}(\|\hat{w}\|)(\frac{r}{r_0})^{-p}$.
- Fourier coefficients with frequencies of equal modulus are Rayleigh-distributed.

To achieve both these objectives, we take the ordered moduli of all coefficients lying on a circle of radius r , then modify them so that they correspond to ordered samples from a Rayleigh distributed variable with the desired mean $\frac{E_{\mathcal{C}_{r_0}}}{r^p}$. We write ξ_i^r for the $|\mathcal{C}_r|$ frequencies lying on \mathcal{C}_r ordered by

magnitude of $\|\hat{w}_\xi\|$, and $Q_\beta^{\mathcal{R}}$ for the quantile function of the Rayleigh distribution of parameter β :

$$Q_\beta^{\mathcal{R}}(u) = \beta\sqrt{-2\log(1-u)}$$

Finally, we set β_r such that the modulus has the desired mean over \mathcal{C}_r : $\beta_r = \sqrt{\frac{2}{\pi} \frac{E_{\mathcal{C}_{r_0}}(\|\hat{w}\|)}{E_{\mathcal{C}_r}(\|\hat{w}\|)}} (\frac{r}{r_0})^{-p}$. Eventually, the spectrum is imposed using the projection $proj_{spectrum}$, defined as:

$$\begin{cases} proj_{spectrum}(\hat{w}) & = \hat{w}_\xi & \text{if } r \leq r_0 \\ proj_{spectrum}(\hat{w})_{\xi_i^r} & = \frac{\hat{w}_{\xi_i^r}}{\|\hat{w}_{\xi_i^r}\|} Q_{\beta_r}^{\mathcal{R}}(\frac{i}{|\mathcal{C}_r|}) & \text{if } r > r_0 \end{cases} \quad (1)$$

Additionally, to avoid artifacts and control only the relevant part of the spectrum, we rely on the 'periodic plus smooth decomposition' of an image [30]. This article proposes to get rid of spurious horizontal and vertical structures in the Fourier domain, induced by the implicit periodization of the image during the Fourier transform. The approach consists in decomposing the image into the sum of a smooth non-periodic image, and a periodized version of the image that contains most of the initial information.

Following the author's guideline, we perform our spectrum operations on the periodic image before adding back the untouched smooth image.

2) *Histogram constraint*: In order to counteract some possible limitations of the generator, we want the color histogram of w to match u 's. We adopt a sliced optimal transport approach [31].

This approach consists in iterating projection steps: for each step, we sample a random unit vector θ of the color space. We then modify the color histogram of w along the direction θ such that the histograms projected onto θ of u and w are closer.

Each step reads: $w_{\sigma_k^w} = w_{\sigma_k^w} + \epsilon\theta \cdot \langle u'_{\sigma_k^{u'}} - w_{\sigma_k^w}, \theta \rangle$,

where u' is a version of u that has been stretched to match the size of w , and σ^{image} is the permutation that rearranges $\langle image, \theta \rangle$ in increasing order.

We call $proj_{hist}(w)$ the result of this process.

3) *Reversibility constraint*: The reversibility condition is a consistency test between the input low resolution image u and the super-resolution result w :

we want u to be the result of blurring then subsampling w . This condition can be seen as a fidelity term, making the upsampled result closer in the L_2 sense to the ground truth. In other words, we want our method to be balanced in the perception-distortion tradeoff ([21]). This condition can be computed as the following loss:

$$\mathcal{L}_{rev} = \frac{1}{2} \|D_f(w * g) - u\|^2 \quad (2)$$

where D_f is the decimation operator (from the dimension of w to the dimension of u), and the smoothing filter g is chosen as a gaussian filter of standard deviation $\sigma = .7 * f$, which is a reasonable filter size regarding Shannon's theorem. We call the resulting image $proj_{rev}(w)$.

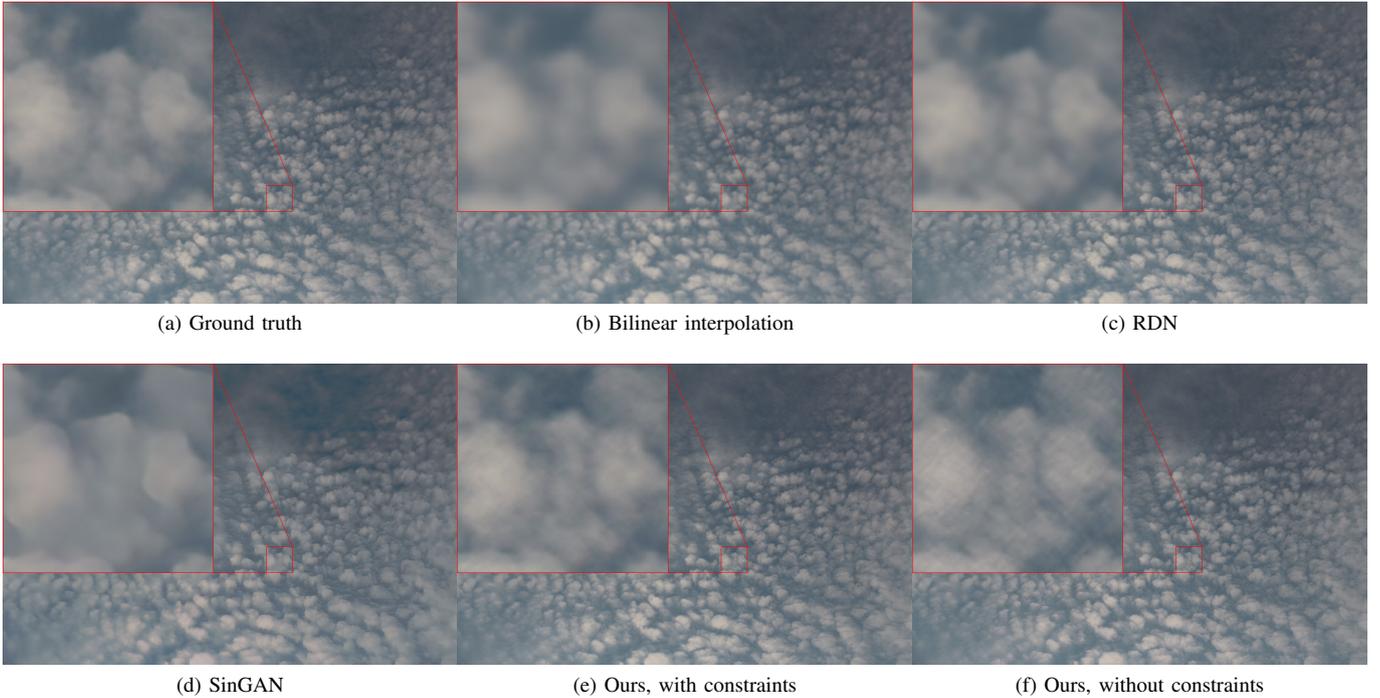


Fig. 2. Visual comparison of all studied methods for an image taken from our Cloud pictures dataset. The highlighted zone is 100*100 pixels.

B. SinGAN

In order to be able to adapt to the statistical specificity of a single image, we chose as a backbone the classical SinGAN architecture. We first quickly recall the architecture of this network and the way it can be applied to SR.

1) *Architecture*: First, a single training image is turned into a multi-resolution pyramid $\{x_N, \dots, x_0\}$ where x_0 is the original image, and x_n is x_0 downsampled n times by a factor $r > 1$.

SinGAN is based on a pyramidal architecture of generators $\{G_N, \dots, G_0\}$ and of discriminators $\{D_N, \dots, D_0\}$, each generator G_i being trained with the help of a patch-discriminator [32] D_i in the WGAN-GP framework [33].

The coarsest scale generator G_N is unconditional and maps a Gaussian noise z_N onto a plausible low-resolution image \tilde{x}_N . Then, at each step n of the pyramid, the same operation is performed by the residual generator G_n : \tilde{x}_{n+1} is upsampled by a factor r ($(\tilde{x}_{n+1}) \uparrow^r$) and some noise z_n is added. The result is fed through the generator, which outputs details that are added to $(\tilde{x}_{n+1}) \uparrow^r$, to get \tilde{x}_n .

As the authors put it, if ψ_n is the convolutional part of the generator G_n , then G_n performs:

$$\tilde{x}_n = (\tilde{x}_{n+1}) \uparrow^r + \psi_n(z_n + (\tilde{x}_{n+1}) \uparrow^r) \quad (3)$$

The goal of the generators G_n is to learn the patch distribution of x_n . Moreover, a reconstruction loss is added, and weighted by a parameter α .

The total loss reads:

$$\min_{G_n} \max_{D_n} \mathcal{L}_{\text{adv}}(G_n, D_n) + \alpha \mathcal{L}_{\text{rec}}(G_n) \quad (4)$$

where \mathcal{L}_{adv} and \mathcal{L}_{rec} are described in SinGAN ([6] Section 2.2).

2) *Super Resolution*: Once the finest resolution generator G_0 is trained (using a reconstruction weight $\alpha = 100$ for SR), the authors propose to iterate this generator, yielding new image resolutions.

This step is arguably the weak point of SinGAN SR, and can yield strong color and directional patterns artifacts, as can be seen in the case of cloud images in Fig. 3.

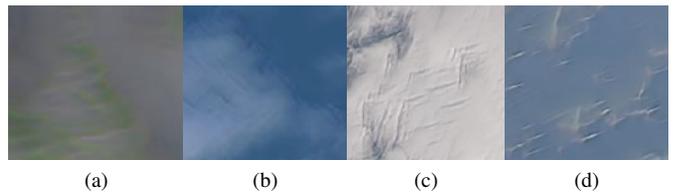


Fig. 3. Examples of SR artifacts using SinGAN on cloud images

We argue that these artifacts come from the iteration of a network G_0 which has never been trained to receive its own output as input. We tackle this issue in two ways:

- We use a scale-invariant generator, as initially proposed in [14]
- We use the previously described constraints to the generation procedure, keeping some statistical properties while iteratively upsampling the image in a similar way to SinGAN's.

C. Our approach

1) *Training procedure*: We train a unique residual generator G on the same pyramid as SinGAN, with 3 scales. Since our aim is to upscale texture and not structure, we do not care for a deep multi-scale pyramid (which allows structural variability). Instead of generating a low resolution image with a GAN (as in SinGAN [6] and Recurrent SinGAN [14]), we use a blurred and downsampled version of our training image as starting step. As a consequence, SinGAN uses as many generators as there are scales in the pyramid, Recurrent SinGAN uses 2, and we only use 1.

Next, following the principle of SinGAN, the image is iteratively passed through the generator, until original resolution is reached. In the original article SinGAN, the training takes place one generator at a time, from coarse resolution to high, because in order to train a given generator, we need to infer all the generators from coarser scales. This length of training is reduced by only training all stages of the pyramid at each training step. This is made possible by the fact that we use only one generator. As in SinGAN, we use a reconstruction loss, between the ground truth and the image obtained with our technique, using no noise.

2) *SR procedure*: In order to perform SinGAN-like SR under constraint, we alternate the application of our recurrent generator G and 3 steps of alternate projection onto our 3 constraints as described in Algorithm 1 and illustrated in Fig. 1. This provides both a multi-scale upsampling in the same fashion as SinGAN, while respecting all constraints along the way.

```

Input:  $x_0$ 
 $w = x_0$ 
while desired zoom factor not reached do
   $w = G(w)$ ; // upscales the image, and
  adds detail
  for step in range(3) do
     $w = proj_{spectrum}(w)$ 
     $w = proj_{hist}(w)$ 
     $w = proj_{rev}(w)$ 
  end
end

```

Algorithm 1: Alternate projection algorithm used to apply constraints to the image at each step of the iterative upsampling.

3) *Settings*: We choose the scaling factor used by [6] for super resolution $r = 0.793701$. We set the noise level $\sigma = 0.01$ and $\alpha = 100$. The Adam optimizer [34] is used for the generator and the discriminators with learning rate 10^{-4} , and parameters $(\beta_1, \beta_2) = (0.5, 0.999)$. We train our network for 2000 iterations, with a learning rate decay of factor 0.1 at the 1600th iteration (like SinGAN). At each GAN training step, the discriminator is trained for 3 iterations while the generator is only updated once, in the spirit of [33].

III. EXPERIMENTS

A. Datasets

1) *Synthetic dataset*: We create a synthetic dataset of 10 1200*1200 colored noise images. To forge them, we use the following method: first, we sample a Gaussian noise image, then take its discrete Fourier Transform, and divide each coefficient by its modulus. This yields an array of unit complex numbers which we multiply by scalars having the desired power decay behavior.

We set the the exponent of this power decay to be $\beta = 1.7$, approximating the theoretical exponent $\frac{5}{3}$ encountered in cloud images.

2) *Cloud dataset*: We create our own cloud pictures dataset from 12 1728*1152 photos of clouds taken from the ground.

3) *BSD100 dataset*: The BSD100 dataset is a classical natural images dataset, used by the authors of SinGAN to evaluate their SR method. We'll also use it to compare our results with theirs.

B. Evaluation metrics

We compare SR results using different quality indices: PSNR, NIQE [23], and $D_{texture}$.

The NIQE index uses the distribution features extracted on patches to assess the quality of an image. Contrarily to the PSNR, it is a no-reference measure.

As done in [35], and inspired by the article [28], we define a texture distance D_{GATYS} :

given an image I , we compute VGG features $F_{c,i,j}^l \in \mathcal{R}^{C \times N_l \times M_l}$ from the layer $l = pool4$ (the deepest representation of the 5 used layers). Then we compute the Gram matrix of these features: $G_{c,d}^l = \frac{1}{N_l M_l} \sum_{i,j} F_{c,i,j}^l F_{d,i,j}^l$. Finally, D_{GATYS} is the L^2 distance between Gram matrices from both images.

$$D_{GATYS}(I, \hat{I}) = \frac{1}{4} \sum (G^l - \hat{G}^l)^2 \quad (5)$$

This distance varies a lot in magnitude upon taking different images. In order to compute meaningful statistics across multiple images, we normalize this distance for each image by the distance $d(I_{GT}, I_{bilinear})$ between the ground truth image I_{GT} and the 4x bilinear zoom of the downsampled image $I_{bilinear}$:

$$D_{texture}(I_{GT}, I) = \frac{D_{GATYS}(I_{GT}, I)}{D_{GATYS}(I_{GT}, I_{bilinear})} \quad (6)$$

As a consequence this distance is 1 when evaluating the bilinear interpolation against the ground truth. This distance can be seen as a ratio of performance when compared to the baseline bilinear interpolation.

All the experiments are run for a 4x zoom. We compare our results against the bilinear interpolation, and RDN [3], an external SR method aiming at maximizing the PSNR. Some results can be visually compared in Fig. 2, with a 100*100 section of the image magnified by a factor 8.

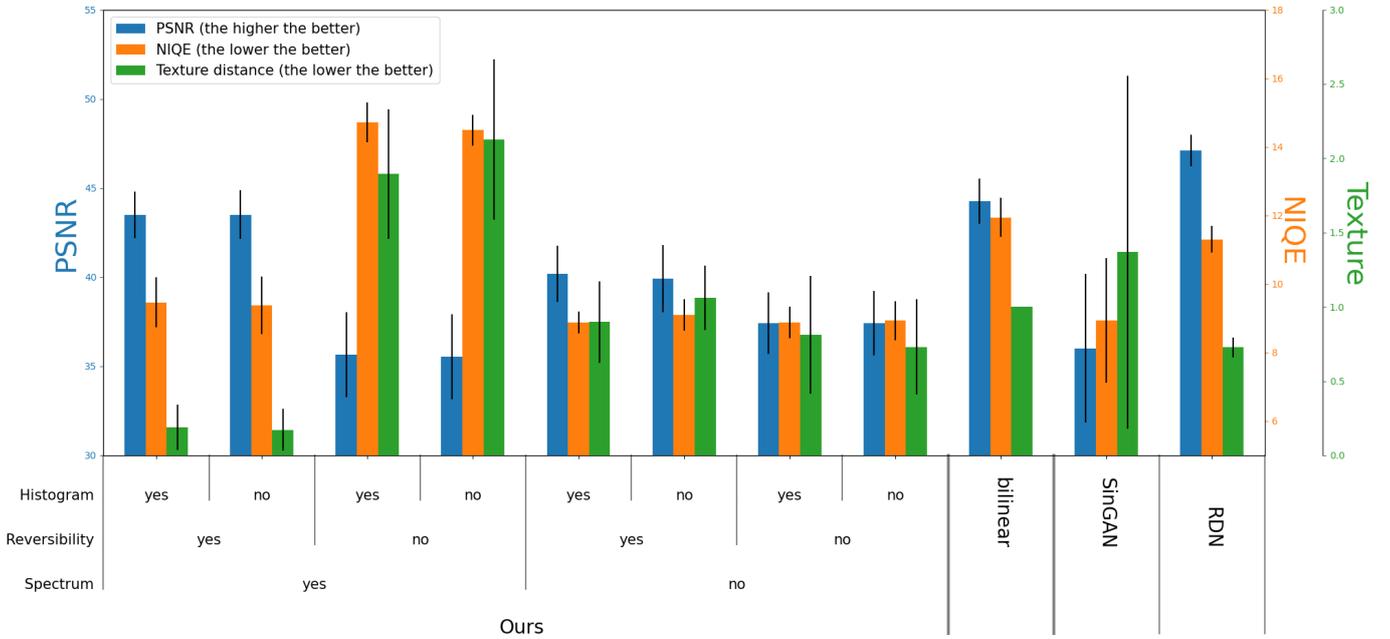


Fig. 4. Ablation study realized on the whole Cloud dataset. Performances of bilinear interpolation, SinGAN and RDN are plotted for comparison purposes.

C. Ablation study

First we verify that the choice of a recurrent generator helps to improve the results of SinGAN on auto-similar data, such as clouds. We notice indeed a slight improvement in PSNR and texture distance, when using our generator without any additional constraint (see Fig. 5).

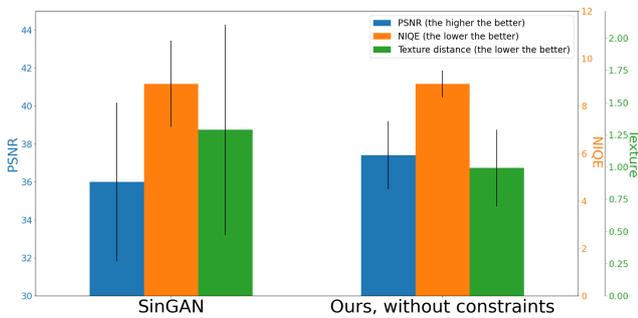


Fig. 5. Comparison between the original SinGAN and the use of a scale-invariant generator (without additional statistical constraints) on the Cloud dataset

We then perform an ablation study on the various statistical constraints considered in this paper. Fig. 4 reports the performances of the method on the Cloud dataset when each constraint is used ('yes') or not ('no').

We observe that, as expected, using the reversibility constraint improves the PSNR, and using both reversibility and spectrum yields the best texture score and PSNR of all choices.

Numerical results are reported in TABLE I.

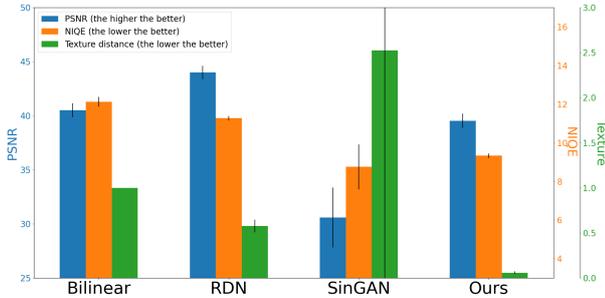
Spectrum	Reversibility	Histogram	Metrics		
			PSNR	NIQE	$D_{texture}$
yes	yes	yes	43.5 ± 1.3	9.5 ± 0.7	0.19 ± 0.15
		no	43.5 ± 1.4	9.4 ± 0.8	0.17 ± 0.14
	no	yes	35.6 ± 2.4	14.7 ± 0.6	1.89 ± 0.44
		no	35.5 ± 2.4	14.5 ± 0.4	2.13 ± 0.53
no	yes	yes	40.2 ± 1.6	8.9 ± 0.3	0.90 ± 0.27
		no	39.9 ± 1.9	9.1 ± 0.5	1.06 ± 0.22
	no	yes	37.4 ± 1.7	8.9 ± 0.5	0.81 ± 0.40
		no	37.4 ± 1.8	8.9 ± 0.6	0.73 ± 0.32
Bilinear interpolation			44.3 ± 1.3	11.9 ± 0.7	1
SinGAN			36.0 ± 4.1	8.9 ± 1.8	1.37 ± 1.18
RDN			47.1 ± 0.9	11.3 ± 0.4	0.73 ± 0.07

TABLE I
ABLATION RESULTS

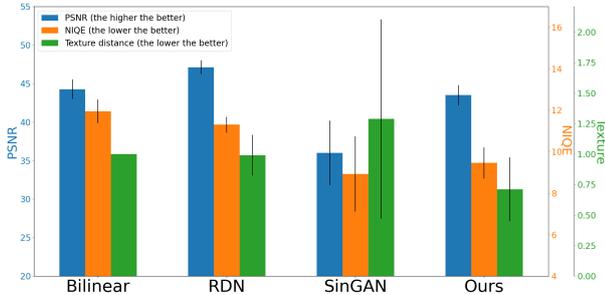
D. Quantitative analysis

We propose to compare bilinear interpolation, RDN, SinGAN, and our method using the metrics and datasets described above.

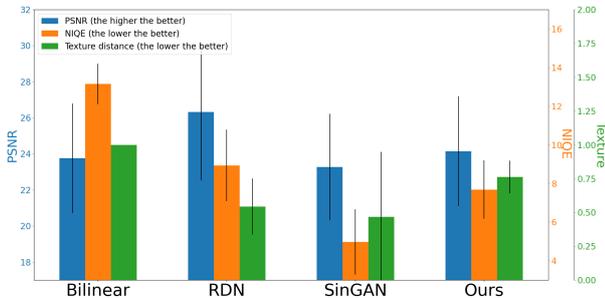
1) *Synthetic dataset*: As a sanity check, we verify that our method performs well on colored noise images. We use images from our synthetic dataset, downsample them by a factor 4, then evaluate several methods when zooming of a factor 4. We emphasize that our method assumes a spectral slope of 1.7 for its spectral constraint, and these images are created to have the same slope of 1.7, thus giving a lot of prior information to our method which indeed outperforms the others in texture quality while maintaining a good PSNR



(a) Comparison of different methods on our synthetic dataset



(b) Comparison of different methods on our Cloud dataset



(c) Comparison of different methods on the BSD100 dataset

Fig. 6. Evaluation on different datasets for our 3 metrics: PSNR, NIQE and texture distance

(see Fig.6a). Apart from NIQE scores comparable to ours, SinGAN performs particularly poorly on this dataset.

2) *Cloud and BSD100 datasets:* On the BSD100 dataset (Fig. 6c), SinGAN beats us in both perceptual quality metrics. This is logical since the natural images that form the BSD100 dataset are very different from colored noise like images. Our spectral prior is not valid on this dataset, which explains this degradation.

On the cloud dataset however (Fig. 6b), we report better texture and PSNR results than SinGAN, while beating RDN and bilinear interpolation on both perceptual metrics.

Let us note that our PSNR results (Fig. 6) are always worse than RDN and slightly worse or comparable with the bilinear interpolation. We do not claim to compete in terms of PSNR with an external method such as RDN optimized for PSNR performance. We aim at respecting our constraints, and by doing that we yield good perceptual results, while maintaining

acceptable PSNR scores.

E. Respect of the constraints

We verify on the Cloud dataset that the output of our method also respects the constraints we have set with three metrics, closely related to these constraints:

- **Color histogram distance** We compute a sliced transport distance between the color histograms of the ground truth and the output of the four evaluated methods.
- **Reversibility error** We compute the L_2 distance between the low resolution image, and the blurred and downsampled SR output w : $\|D_f(w * g) - u\|^2$
- **Straight slope error** We compute the standard error in the least square regression of the log-log graph (radial mean of Fourier modulus against frequency f). It measures how linearly the log-modulus decays in log-frequency. In other words, it measures how colored noise-like an image is.

The results are reported in TABLE II.

	Bilinear	RDN	SinGAN	Ours
Color histogram distance	39 ± 28	17 ± 3	938 ± 2202	47 ± 40
Reversibility error ($\cdot 10^{-4}$)	1.2 ± 0.4	0.5 ± 0.2	20.4 ± 49.8	0.6 ± 0.3
Straight slope error ($\cdot 10^{-3}$)	8.0 ± 1.2	8.6 ± 1.3	4.6 ± 1.2	3.3 ± 0.9

TABLE II
EVALUATION OF THE RESPECT OF THE 3 CONSTRAINTS FOR ALL 4 METHODS ON THE CLOUD DATASET

Keeping in mind that our Histogram constraint uses as reference the low resolution image (whose histogram differs slightly from the high resolution ground truth), we see that our method performs reasonably in this color evaluation. Color artifacts and lack of reversibility from SinGAN are particularly salient in this table with very poor measures of color histogram distance and reversibility error. Adding constraints allows to significantly reduce these problems. Our methods, as expected, yields the best performance in terms of spectral slope linearity, followed by SinGAN.

IV. CONCLUSION

In this paper, we have introduced an internal SR method which preserves statistics that are related to the specificity of SR tasks and to decreasing property of the Fourier spectrum. The interest of our approach is demonstrated on comparison results with SinGAN, which we improve in terms of PSNR and texture fidelity on colored noise images datasets.

ACKNOWLEDGMENTS

This research was partially supported by the Defence Innovation Agency (Pierrick Chatillon PhD grant) and the ANR-19-CE40-0005 MISTIC project.

REFERENCES

- [1] C. Dong, C. C. Loy, K. He, and X. Tang, "Image super-resolution using deep convolutional networks," *IEEE transactions on pattern analysis and machine intelligence*, vol. 38, no. 2, pp. 295–307, 2015.
- [2] B. Lim, S. Son, H. Kim, S. Nah, and K. Mu Lee, "Enhanced deep residual networks for single image super-resolution," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2017, pp. 136–144.

- [3] Y. Zhang, Y. Tian, Y. Kong, B. Zhong, and Y. Fu, "Residual dense network for image super-resolution," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 2472–2481.
- [4] J. Liang, J. Cao, G. Sun, K. Zhang, L. Van Gool, and R. Timofte, "Swinir: Image restoration using swin transformer," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 1833–1844.
- [5] D. Ulyanov, A. Vedaldi, and V. Lempitsky, "Deep image prior," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 9446–9454.
- [6] T. R. Shaham, T. Dekel, and T. Michaeli, "Singan: Learning a generative model from a single natural image," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 4570–4580.
- [7] D. Glasner, S. Bagon, and M. Irani, "Super-resolution from a single image," in *Proceedings of the IEEE International Conference on Computer Vision*, 2009, pp. 349 – 356.
- [8] M. Zontak and M. Irani, "Internal statistics of a single natural image," in *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, vol. 2, 2011, pp. 977 – 984.
- [9] D. J. Field, "Relations between the statistics of natural images and the response properties of cortical cells," *JOSA A*, vol. 4, no. 12, pp. 2379–2394, 1987.
- [10] M. S. Keshner, "1/f noise," *Proceedings of the IEEE*, vol. 70, no. 3, pp. 212–218, 1982.
- [11] T. Hirata, T. Satoh, and K. Ito, "Fractal structure of spatial distribution of microfracturing in rock," *Geophysical Journal International*, vol. 90, no. 2, pp. 369–374, 1987.
- [12] N. S. Arnold and W. G. Rees, "Self-similarity in glacier surface characteristics," *Journal of Glaciology*, vol. 49, no. 167, pp. 547–554, 2003.
- [13] K. Perlin, "An image synthesizer," *ACM Siggraph Computer Graphics*, vol. 19, no. 3, pp. 287–296, 1985.
- [14] X. He and Z. Fu, "Recurrent singan: Towards scale-agnostic single image gans," in *Proceedings of the 2021 5th International Conference on Electronic Information Technology and Computer Engineering*, ser. EITCE 2021. New York, NY, USA: Association for Computing Machinery, 2021, p. 361–366. [Online]. Available: <https://doi.org/10.1145/3501409.3501476>
- [15] V. Venema, S. Meyer, S. G. García, A. Kniffka, C. Simmer, S. Crewell, U. Löhnert, T. Trautmann, and M. Andreas, "Surrogate cloud fields generated with the iterative amplitude adapted fourier transform algorithm," *Tellus A*, vol. 58A, pp. 104–120, 2006.
- [16] R. F. Cahalan, "Bounded cascade clouds: albedo and effective thickness," *Nonlin. Processes Geophys.*, vol. 1, pp. 156–167, 1994.
- [17] A. Benassi, F. Szczap, A. Davis, M. Masbou, C. Cornet, and P. Bleuyard, "Thermal radiative fluxes through inhomogeneous cloud fields: A sensitivity study using a new stochastic cloud generator," *Atmos. Res.*, vol. 72, pp. 291–315, 2004.
- [18] R. F. Cahalan and J. H. Joseph, "Fractal statistics of cloud fields," *Monthly Weather Review*, vol. 117, pp. 261–272, 1989.
- [19] F. Szczap, Y. Gour, T. Fauchez, C. Cornet, T. Faure, O. Jourdan, G. Penide, and P. Dubuisson, "A flexible three-dimensional stratocumulus, cumulus and cirrus cloud generator (3dcloud) based on drastically simplified atmospheric equations and the fourier transform framework," *Geosci. Model Dev.*, vol. 7, pp. 1779–1801, 2014.
- [20] P. Chainais, E. Koenig, V. Delouille, and J.-F. Hochedez, "Virtual Super Resolution of Scale Invariant Textured Images Using Multifractal Stochastic Processes," *Journal of Mathematical Imaging and Vision*, vol. 39, no. 1, pp. 28–44, 2011. [Online]. Available: <https://hal.archives-ouvertes.fr/hal-00707631>
- [21] Y. Blau and T. Michaeli, "The perception-distortion tradeoff," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 6228–6237.
- [22] J. Johnson, A. Alahi, and L. Fei-Fei, "Perceptual losses for real-time style transfer and super-resolution," in *Computer Vision – ECCV 2016*, B. Leibe, J. Matas, N. Sebe, and M. Welling, Eds. Cham: Springer International Publishing, 2016, pp. 694–711.
- [23] A. Mittal, R. Soundararajan, and A. Bovik, "Making a "completely blind" image quality analyzer," *Signal Processing Letters, IEEE*, vol. 20, pp. 209–212, 03 2013.
- [24] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in neural information processing systems*, 2014, pp. 2672–2680.
- [25] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, and W. Shi, "Photo-realistic single image super-resolution using a generative adversarial network," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 105–114.
- [26] X. Wang, K. Yu, S. Wu, J. Gu, Y. Liu, C. Dong, Y. Qiao, and C. C. Loy, "EsrGAN: Enhanced super-resolution generative adversarial networks," in *The European Conference on Computer Vision Workshops (ECCVW)*, September 2018.
- [27] M. S. M. Sajjadi, B. Schölkopf, and M. Hirsch, "Enhancenet: Single image super-resolution through automated texture synthesis," *2017 IEEE International Conference on Computer Vision (ICCV)*, pp. 4501–4510, 2017.
- [28] L. A. Gatys, A. S. Ecker, and M. Bethge, "Texture synthesis using convolutional neural networks," in *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1*, ser. NIPS'15. Cambridge, MA, USA: MIT Press, 2015, p. 262–270.
- [29] B. Rivet, L. Girin, and C. Jutten, "Log-rayleigh distribution: A simple and efficient statistical representation of log-spectral coefficients," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 15, no. 3, pp. 796–802, 2007.
- [30] L. Moisan, "Periodic plus smooth image decomposition," *Journal of Mathematical Imaging and Vision*, vol. 39, no. 2, pp. 161–179, Feb 2011. [Online]. Available: <https://doi.org/10.1007/s10851-010-0227-1>
- [31] J. Rabin, G. Peyré, J. Delon, and M. Berton, "Wasserstein barycenter and its application to texture mixing," in *International Conference on Scale Space and Variational Methods in Computer Vision*. Springer, 2011, pp. 435–446.
- [32] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 5967–5976.
- [33] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville, "Improved training of Wasserstein GANs," *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [34] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [35] A. Houdard, A. Leclaire, N. Papadakis, and J. Rabin, "Wasserstein Generative Models for Patch-based Texture Synthesis," in *Scale Space and Variational Methods in Computer Vision*, vol. LNCS 12679, Cabourg, France, May 2021, pp. 269–280. [Online]. Available: <https://hal.archives-ouvertes.fr/hal-02824076>