



HAL
open science

Stochastic Super-Resolution For Gaussian Textures

Émile Pierret, Bruno Galerne

► **To cite this version:**

Émile Pierret, Bruno Galerne. Stochastic Super-Resolution For Gaussian Textures. ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), IEEE; IEEE, pp.1-5, 2024, 10.1109/ICASSP49357.2023.10096585 . hal-03863009v2

HAL Id: hal-03863009

<https://hal.science/hal-03863009v2>

Submitted on 2 Mar 2023

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.

STOCHASTIC SUPER-RESOLUTION FOR GAUSSIAN TEXTURES

Émile Pierret^a, Bruno Galerne^{a,b}

^a Institut Denis Poisson – Université d’Orléans, Université de Tours, CNRS

^b Institut universitaire de France (IUF)

ABSTRACT

Super-resolution (SR) is an ill-posed inverse problem which consists in proposing high-resolution images consistent with a given low-resolution one. While most SR algorithms are deterministic, stochastic SR deals with designing a stochastic sampler generating any realistic SR solution. The goal of this paper is to show that stochastic SR is a well-posed and solvable problem when restricting to Gaussian stationary textures. Using Gaussian conditional sampling and exploiting the stationarity assumption, we propose an efficient algorithm based on fast Fourier transform. We also demonstrate the practical relevance of the approach for SR with a reference image. Although limited to stationary microtextures, our approach compares favorably in terms of speed and visual quality to some state of the art methods designed for a larger class of images.

Index Terms— stochastic super-resolution, Gaussian textures, conditional simulation, kriging, super-resolution with a reference image

1. INTRODUCTION

The super resolution (SR) problem consists in generating a high-resolution (HR) image corresponding to a given low-resolution (LR) input image. This is a very ill-posed inverse problem that necessitates an image prior model or additional information to create images having sharp edges and rich texture details. This is a very important problem for the entertainment industry due to the increase in resolution of display screens as well as other imaging industries where resolution is critical, notably satellite imaging and microscopy.

Different frameworks are considered in the literature: the Single Image SR (SISR) consists in using only an LR input image and a generic database of HR images, while SR with a reference image considers an LR input image accompanied with an HR image that presents similarity with the unknown HR input. Most contributions in SISR are based on conditional generative neural networks [1, 2, 3, 4], adversarial [5] or not, and trained using the “perceptual loss” [6], that is, the distance between pretrained VGG features [7]. These references tackle zoom factor x4 or even x8. Contrary to a classical x2 problem that can be seen as an image sharpener problem, for such high zoom factors new image content must be generated in accordance with the LR input, and the space of possible images becomes very large. All references then insist on the importance of generating local textures and avoiding the “regression to the mean problem”: when favoring PSNR the optimal result consists in a blurry image close to the mean of plausible sharp images. While SISR can be seen as local conditional texture synthesis, a close inspection of the state-of-the-art techniques shows however that texture modeling is absent from GANs. Indeed, the perception loss does not favor the statistics of the textures one wishes to reconstruct. In addition, the proposed models are generally deterministic,

which is not desirable for texture generation since it does not allow checking the statistical consistency of the proposed solution. More recent models propose generative networks with stochastic response for SISR: SRFlow [8] use invertible generative flow [9] and [10] use Denoising Diffusion Probabilistic Models (DDPM) [11], also called score-based models [12]. They propose practical solutions for stochastic SR, that is providing a stochastic sampler that outputs any realistic solution to the SISR problem. In addition, let us note that several recent contributions solely focus on texture SR, using patch statistics from a reference image [13, 14] or a prior as spectral content [15].

In this paper, we solve the problem of stochastic SR when restricting to stationary Gaussian textures [16]. Following a similar approach as for Gaussian texture inpainting [17, 18], assuming that the LR input is stationary and Gaussian with a known covariance in HR space, we propose an exact SR sampler relying on conditional Gaussian simulation. We exploit the stationarity of both the texture model and the zoom-out operator to obtain an efficient algorithm based on Fast Fourier Transform (FFT). We further demonstrate experimentally the interest of the approach in the context of SR with a reference image.

The plan of the paper is as follows. We present our framework of SR and remind results about conditional Gaussian simulation. Then, we focus on Gaussian stationary textures and detail how to implement efficiently the stochastic simulation for these models. Finally, we extend the approach for SR with a reference image and compare our results with the state of the art before concluding.

2. GAUSSIAN CONDITIONAL SIMULATION FOR SUPER-RESOLUTION

2.1. Framework

Notations Let $M, N > 2$ be the size of the images. We extend all the images on \mathbb{Z}^2 by periodization. We write $[M] = \{0, \dots, M - 1\}$ and $\Omega_{M,N} = [M] \times [N]$. For $\mathbf{X}, \mathbf{Y} \in \mathbb{R}^{\Omega_{M,N}}$, $\mathbf{X} \star \mathbf{Y}$ designates the discrete and periodic convolution defined by $(\mathbf{X} \star \mathbf{Y})(\mathbf{x}) = \sum_{\mathbf{y} \in \Omega_{M,N}} \mathbf{X}(\mathbf{x} - \mathbf{y})\mathbf{Y}(\mathbf{y})$, $\mathbf{x} \in \Omega_{M,N}$. For $\lambda \in \mathbb{R}^{\Omega_{M,N}}$, we denote the symmetric $\check{\lambda}$ such that for $\mathbf{x} \in \Omega_{M,N}$, $\check{\lambda}(\mathbf{x}) = \lambda(-\mathbf{x})$. For $\mathbf{X} \in \mathbb{R}^{\Omega_{M,N}}$, we express by $\mathcal{F}_2(\mathbf{X})$ or $\hat{\mathbf{X}}$ its discrete Fourier transform defined by $\hat{\mathbf{X}}(\mathbf{x}) = \sum_{\mathbf{y} \in \Omega_{M,N}} \mathbf{X}(\mathbf{y}) \exp(-\frac{2i\pi\mathbf{x}_1\mathbf{y}_1}{M}) \exp(-\frac{2i\pi\mathbf{x}_2\mathbf{y}_2}{N})$, $\mathbf{x} \in \Omega_{M,N}$. Let us recall that \mathcal{F}_2 is invertible with $\mathcal{F}_2^{-1} = \frac{1}{MN} \check{\mathcal{F}_2}$ and that $\mathcal{F}_2(\mathbf{X} \star \mathbf{Y}) = \mathcal{F}_2(\mathbf{X}) \odot \mathcal{F}_2(\mathbf{Y})$ where \odot denotes the componentwise product.

Zoom-out operator We denote by r the factor of zoom-out ($= 4, 8, \dots$). From now on we assume that M and N are multiples of r . The loss of pixels in an image can be caused by multiple operators. We will assume in our experiments that the degradation is caused by the imresize function of Matlab which is used in many

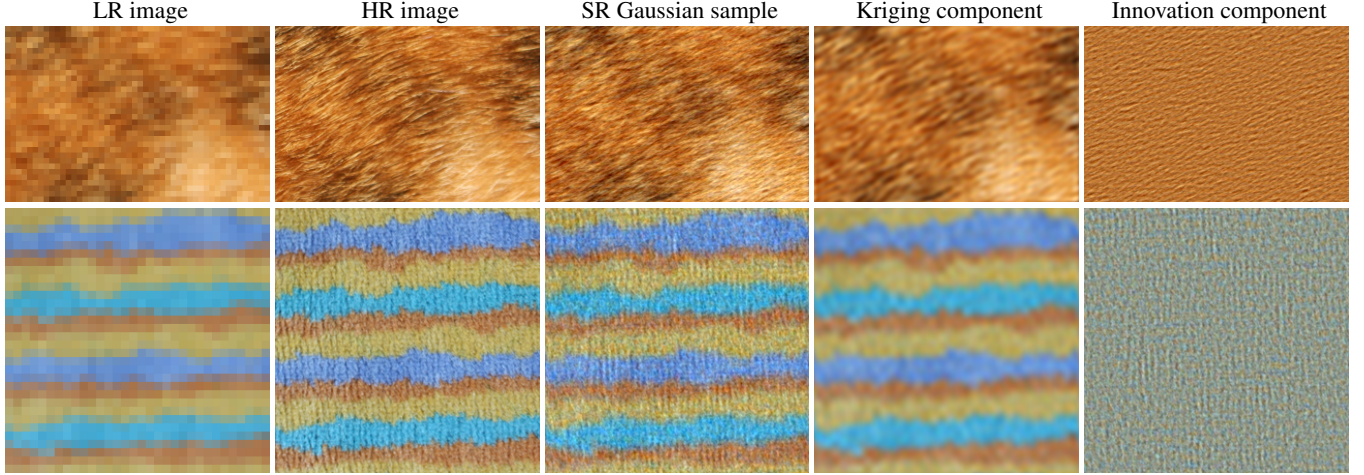


Fig. 1. Different components of the Gaussian SR model. Top row: SR with a factor $\times 16$, the images HR size is 512×768 ; Bottom row: SR with a factor $\times 8$, the images HR size is 256×256 . The last columns present respectively a Gaussian SR sample and its associated kriging and innovation components. Observe the complementarity of the the kriging component which deblurs the LR input and the innovation component which adds independent high-frequency details.

models [3, 4, 8]. Concretely, this operator reduces the size of an image by weighting the pixels. Supposing that the image is periodic, we can split it off in a separable convolution which weights all the pixels followed by a subsampling with stride r [19]. We denote $\mathbf{A} = \mathbf{SC}$ this linear zoom-out operator, with \mathbf{C} the convolution and \mathbf{S} the subsampling operator with stride r such that for $\mathbf{U} \in \mathbb{R}^{\Omega_{M,N}}$, $\mathbf{SU} \in \mathbb{R}^{\Omega_{M/r, N/r}}$ and for $\mathbf{x} \in \Omega_{M/r, N/r}$, $(\mathbf{SU})(\mathbf{x}) = \mathbf{U}(r\mathbf{x})$.

2.2. Gaussian conditional simulation

Let $\mathbf{X} \in \mathbb{R}^{\Omega_{M,N}}$ be a stationary Gaussian process with distribution $\mathcal{N}(\mathbf{0}, \mathbf{\Gamma})$ and $\mathbf{X}_{\text{LR}} = \mathbf{AX} \in \mathbb{R}^{\Omega_{M/r, N/r}}$ be its LR version. \mathbf{A} reduces the size of the images and thus is not invertible. We would like to sample $\mathbf{X}|\mathbf{X}_{\text{LR}}$. In our further assumptions, \mathbf{X} is supposed to be Gaussian, consequently we will use the classical theorem recalled in [18]:

Theorem 1 (Gaussian simulation and Gaussian kriging). *Let \mathbf{X} be a Gaussian vector, and \mathbf{A} be a linear operator, $\mathbb{E}(\mathbf{X}|\mathbf{AX})$ and $\mathbf{X} - \mathbb{E}(\mathbf{X}|\mathbf{AX})$ are independent. Consequently, if $\tilde{\mathbf{X}}$ is independent of \mathbf{X} with the same distribution, then $\mathbb{E}(\mathbf{X}|\mathbf{AX}) + [\tilde{\mathbf{X}} - \mathbb{E}(\tilde{\mathbf{X}}|\mathbf{A}\tilde{\mathbf{X}})]$ has the same distribution as \mathbf{X} , knowing \mathbf{AX} . Furthermore, if \mathbf{X} is zero-mean, there exists $\mathbf{\Lambda} \in \mathbb{R}^{\Omega_{M/r, N/r} \times \Omega_{M,N}}$ such that $\mathbb{E}(\mathbf{X}|\mathbf{AX}) = \mathbf{\Lambda}^T \mathbf{AX}$ and $\mathbb{E}(\mathbf{X}|\mathbf{AX}) = \mathbf{\Lambda}^T \mathbf{AX}$ if and only if $\mathbf{\Lambda}$ verifies the matrix equation:*

$$\mathbf{\Lambda} \mathbf{\Gamma} \mathbf{\Lambda}^T \mathbf{\Lambda} = \mathbf{\Lambda} \mathbf{\Gamma}. \quad (1)$$

This theorem is a consequence of the fact that the space of square-integrable random variables is a Hilbert space. The assumption that \mathbf{X} follows a Gaussian law is crucial. As for now, the resolution of our problem requires the computation of $\mathbb{E}(\mathbf{X}|\mathbf{X}_{\text{LR}}) = \mathbb{E}(\mathbf{X}|\mathbf{AX}) = \mathbf{\Lambda}^T \mathbf{AX} = \mathbf{\Lambda}^T \mathbf{X}_{\text{LR}}$, where $\mathbf{\Lambda}$ is called the kriging matrix. $\mathbf{\Lambda}^T \mathbf{X}_{\text{LR}}$ is called the kriging component and $\tilde{\mathbf{X}} - \mathbf{\Lambda}^T \mathbf{AX}$ the innovation component. Therefore, stochastic SR for Gaussian textures corresponds to simulate $\mathbf{\Lambda}^T \mathbf{X}_{\text{LR}} + (\tilde{\mathbf{X}} - \mathbf{\Lambda}^T \mathbf{AX})$. Figure 1 presents the realization of $\mathbf{X}|\mathbf{X}_{\text{LR}}$, titled SR Gaussian sample. This is the sum of the kriging component which is an attachment to data and the innovation component that adds stochastic independent high-frequency variations. For the given LR image \mathbf{X}_{LR} and the HR covariance $\mathbf{\Gamma}$, it is just necessary to compute $\mathbf{\Lambda}$ to sample HR versions of \mathbf{X}_{LR} .

3. AN EFFICIENT IMPLEMENTATION IN THE STATIONARY CASE

3.1. ADSN model and super-resolution

Our objective is to study the SR of Gaussian stationary textures. We will focus on textures following the asymptotic discrete spot noise (ADSN) model [16]. Given a grayscale image $\mathbf{U} \in \mathbb{R}^{\Omega_{M,N}}$ with mean grayscale $m \in \mathbb{R}$, one defines ADSN(\mathbf{U}) as the distribution of $\mathbf{X} = \mathbf{t} \star \mathbf{W}$ where $\mathbf{t} = \frac{1}{\sqrt{MN}}(\mathbf{U} - m)$ is called the texton associated with \mathbf{U} . ADSN(\mathbf{U}) is a Gaussian distribution with mean zero and covariance matrix $\mathbf{\Gamma}$ that represents the convolution by the kernel $\boldsymbol{\gamma} = \mathbf{t} \star \tilde{\mathbf{t}}$. $\mathbf{X} \sim \text{ADSN}(\mathbf{U})$ is stationary since for all $\mathbf{x} \in \Omega_{M,N}$ and $\mathbf{y} \in \mathbb{Z}^2$, $\mathbf{X}(\mathbf{x} - \mathbf{y}) \stackrel{\mathcal{L}}{=} \mathbf{X}(\mathbf{x})$.

For a given input image $\mathbf{U} \in \mathbb{R}^{\Omega_{M,N}}$ and its LR version $\mathbf{U}_{\text{LR}} = \mathbf{AU}$, we would like to sample $\mathbf{X} \sim \text{ADSN}(\mathbf{U})$ conditioned on $\mathbf{X}_{\text{LR}} = \mathbf{U}_{\text{LR}}$, that is, by Theorem 1,

$$\mathbf{X}_{\text{SR}} = \mathbf{\Lambda}^T \mathbf{U}_{\text{LR}} + (\tilde{\mathbf{X}} - \mathbf{\Lambda}^T \mathbf{A}\tilde{\mathbf{X}})$$

where $\tilde{\mathbf{X}} \sim \text{ADSN}(\mathbf{U})$, and $\mathbf{\Lambda}$ is the kriging matrix associated with ADSN(\mathbf{U}) = $\mathcal{N}(\mathbf{0}, \mathbf{\Gamma})$. We will explain how to implement it efficiently in the following.

3.2. Structure of the problem induced by stationarity

For $k, \ell \in [r]$, let $\Omega_{M,N}^{k,\ell,r} = \{(k + ir, \ell + jr), i, j \in [M/r] \times [N/r]\} \subset \Omega_{M,N}$ be the subgrid of $\Omega_{M,N}$ having stride r and starting at (k, ℓ) . Note that each subgrid $\Omega_{M,N}^{k,\ell,r}$ has the same number of pixels as the LR image domain $\Omega_{M/r, N/r}$.

Proposition 1 (Structure of the kriging matrix). *There exists $\mathbf{\Lambda} \in \mathbb{R}^{\Omega_{M/r, N/r} \times \Omega_{M,N}}$ solution of Equation 1 such that $\mathbf{Y} \in \mathbb{R}^{\Omega_{M/r, N/r}} \mapsto \mathbf{\Lambda}^T \mathbf{Y} \in \mathbb{R}^{\Omega_{M,N}}$ corresponds to a convolution on each of the shifted subgrids $\Omega_{M,N}^{k,\ell,r}$, $k, \ell \in [r]$. More precisely, $\mathbf{\Lambda}$ is fully determined by its r^2 first columns $\lambda(k, \ell) = \mathbf{\Lambda}_{\Omega_{M/r, N/r} \times (k,\ell)}$, $k, \ell \in [r]$ and*

$$(\mathbf{\Lambda}^T \mathbf{Y}) (\Omega_{M,N}^{k,\ell,r}) = \check{\lambda}(k, \ell) \star \mathbf{Y}.$$

Input: An image $\mathbf{U}_{\text{LR}} \in \mathbb{R}^{\Omega_{M/r, N/r}}$, r the zoom factor, \mathbf{t} the convolution kernel of the ADSN model, \mathbf{c} the kernel of the convolution of the zoom-out operator $\mathbf{A} = \mathbf{S}\mathbf{C}$

- 1 **Step 1: Computation of kriging matrix Λ**
- 2 Store $\mathbf{per}(\mathbf{t})$ the periodic component of \mathbf{t} ;
- 3 Store the convolution kernels $\boldsymbol{\gamma} = \mathbf{per}(\mathbf{t}) \star \check{\mathbf{per}}(\mathbf{t})$, $\mathbf{c} \star \boldsymbol{\gamma}$ and $\boldsymbol{\kappa} = \mathbf{c} \star \boldsymbol{\gamma} \star \check{\mathbf{c}}$ (computed in Fourier);
- 4 **for** $(k, \ell) \in [r]^2$ **do**
- 5 $\hat{\mathbf{b}} = \mathcal{F}_2\left(\mathbf{S}((\mathbf{c} \star \boldsymbol{\gamma})(\cdot - k, \cdot - \ell))\right)$;
- 6 $\hat{\lambda}(k, \ell) [\hat{\kappa} \neq 0] \leftarrow \frac{\hat{\mathbf{b}}[\hat{\kappa} \neq 0]}{\hat{\kappa}[\hat{\kappa} \neq 0]}$;
- 7 **end**
- 8 **Step 2: Sampling of one SR version of \mathbf{U}_{LR} :**
- 9 Generate $\mathbf{W} \in \mathbb{R}^{\Omega_{M, N}}$ following a Gaussian standard law;
- 10 $\check{\mathbf{X}} \leftarrow \mathbf{t} \star \mathbf{W}$;
- 11 $\check{\mathbf{X}}_{\text{LR}} \leftarrow \mathbf{A}\check{\mathbf{X}}$;
- 12 **for each shifted subgrid by** $(k, \ell) \in [r]^2$ **do**
- 13 $\mathbf{X}_{\text{SR}}(\Omega_{M, N}^{k, \ell, r}) \leftarrow$
 $\mathcal{F}_2^{-1}\left(\left(\hat{\mathbf{U}}_{\text{LR}} - \check{\mathbf{X}}_{\text{LR}}\right) \odot \overline{\hat{\lambda}(k, \ell)}\right) + \check{\mathbf{X}}(\Omega_{M, N}^{k, \ell, r})$;
- 14 **end**

Output: \mathbf{X}_{SR}

Algorithm 1: Pseudo-code of the Gaussian simulation for SR. To generate several samples, only the step 2 should be re-run.

The proof of Proposition 1 relies on the fact that both \mathbf{X} and $\mathbf{X}_{\text{LR}} = \mathbf{A}\mathbf{X}$ are stationary on their respective domains. Consequently, it is sufficient to solve the r^2 equations

$$\mathbf{A}\boldsymbol{\Gamma}\mathbf{A}^T \lambda(k, \ell) = \mathbf{A}\boldsymbol{\Gamma}_{\Omega_{M, N} \times (k, \ell)}, \quad k, \ell \in [r] \quad (2)$$

to fully determine the kriging component of $\mathbf{X}|\mathbf{X}_{\text{LR}}$. Remark that the matrix Λ is determined by MN values, the size of the HR image.

3.3. Resolution of the systems

To determine the kriging component, it is sufficient to solve the r^2 systems of Equation (2). We will use the Lemma 1.

Lemma 1 (Convolution and subsampling). *If \mathbf{B} is a 2D convolution on $\Omega_{M, N}$ by the kernel $\boldsymbol{\beta} \in \Omega_{M, N}$, \mathbf{SBS}^T is a convolution on $\Omega_{M/r, N/r}$ by the kernel $\mathbf{S}\boldsymbol{\beta} \in \Omega_{M/r, N/r}$.*

By Lemma 1, $\mathbf{A}\boldsymbol{\Gamma}\mathbf{A}^T$ is a convolution matrix with kernel $\boldsymbol{\kappa} = \mathbf{S}(\mathbf{c} \star \boldsymbol{\gamma} \star \check{\mathbf{c}})$ where \mathbf{c} is the kernel of \mathbf{C} . The Equations (2) become

$$\boldsymbol{\kappa} \star \lambda(k, \ell) = \mathbf{A}\boldsymbol{\Gamma}_{\Omega_{M, N} \times (k, \ell)}, \quad k, \ell \in [r] \quad (3)$$

and we can easily pseudo-invert the convolution by component-wise division in the Fourier domain.

The full FFT-based procedure is given by Algorithm 1. Note that to avoid artefacts due to non-periodicity of the images, as classically done, we use the periodic plus smooth decomposition [20] to make the spectral computations. To work with non-zero mean images, the mean grayscale of the LR image should be subtracted at the beginning of the algorithm and added at the end. For RGB color images, the operations should be led on each channel. Note that the Gaussian noise \mathbf{W} to generate the ADSN textures should be the same for each channel [16], as done for the results presented in Figure 1.

4. SUPER-RESOLUTION WITH A REFERENCE IMAGE

In the previous theoretical sections we used the ground truth HR version \mathbf{U} to estimate the distribution of the Gaussian texture, making the process impractical. In this section we demonstrate that the approach extends to the context of SR with a reference image, the reference image being used in place of the ground truth for modeling the ADSN distribution. Given an LR texture $\mathbf{U}_{\text{LR}} = \mathbf{A}\mathbf{U} \in \mathbb{R}^{\Omega_{M/r, N/r}}$ and a reference HR texture $\mathbf{U}_{\text{ref}} \in \mathbb{R}^{\Omega_{M, N}}$, we simply replace the unknown ADSN kernel $\mathbf{t} = \frac{1}{\sqrt{MN}}(\mathbf{U} - m)$ by $\mathbf{t}_{\text{ref}} = \frac{1}{\sqrt{MN}}(\mathbf{U}_{\text{ref}} - m_{\text{ref}})$ in Algorithm 1, assuming that the two microtextures are similar. We sample $\mathbf{X}_{\text{SR}} = \mathbf{A}^T \mathbf{U}_{\text{LR}} + (\check{\mathbf{X}} - \mathbf{A}^T \mathbf{A}\check{\mathbf{X}})$ where $\check{\mathbf{X}}$ has the distribution $\text{ADSN}(\mathbf{U}_{\text{ref}})$ and Λ is the kriging matrix associated with $\text{ADSN}(\mathbf{U}_{\text{ref}})$.

Two SR experiments are given in Figure 2 where we evaluate and compare our results with two competitive methods: Wasserstein patch prior (WPP) [13] and SRFlow [8]. WPP¹ is deterministic and proposes to use optimal transport in patch space in a variational formulation for SR. The routine needs a reference image similar to ours. SRFlow² is a normalizing flow for stochastic SR trained on natural images. The network transforms a standard Gaussian latent variable into a SR sample given the LR image. It depends on a temperature hyperparameter $\tau \in [0, 1]$ that modulates the variance of the latent variable. Both methods require GPU and experiments were conducted using a NVIDIA V100 GPU. Visually, there are some issues in the borders of the image generated by WPP and the output texture tends to be blurry. The results of SRFlow with high temperature are too sharp for the Gaussian texture input, probably a bias induced by the training set of natural images that are often piecewise regular. The images obtained with SR Gaussian (ours) have a more granular appearance, due to the innovation component. However, the texture of details that breaks the stationary assumptions are not retrieved as in the white spot in the second image. Let us stress here that both WPP and SRFlow are designed for generic natural images while our method only works on Gaussian textures.

To evaluate the different methods we report in Figure 2 the metrics Peak Signal to Noise Ratio (PSNR), Structural SIMilarity (SSIM) and Learned Perceptual Image Patch Similarity (LPIPS). PSNR is the logarithm of the mean-squared error (MSE) between two images. SSIM is a metric which quantifies the similarity of two images studying their luminance, contrast and structure [22]. LPIPS is the norm between weighted features of a pre-trained classification network and quantifies the perceptual similarity between two images [23]. As illustrated by the results of Figure 2, PSNR and SSIM are not adapted to evaluate the quality of the texture samples in our context. Indeed the best solutions for this metric are always the most blurry solutions, namely the bicubic interpolation, SRFlow with temperature $\tau = 0$, and the kriging component. The LPIPS metric is more relevant for SR of textures. For this metric our Gaussian SR results are the best for both examples. Note also that our algorithm runs very quickly using only a CPU.

To advocate for the irrelevance of the PSNR for stochastic texture SR, in our Gaussian framework it can be shown that the MSE is lower for the blurry kriging component than for the perfect SR samples.

Proposition 2 (Kriging component and MSE). *Let $\mathbf{U} \in \mathbb{R}^{\Omega_{M, N}}$ be a Gaussian texture, $\mathbf{U}_{\text{LR}} = \mathbf{A}\mathbf{U}$ its LR version, \mathbf{U}_{ref} a reference image,*

¹https://github.com/johertrich/Wasserstein_Patch_Prior

²Code and weights from <https://github.com/andreas128/SRFlow>




Fig. 2. Gaussian SR with a reference image for a factor $\times 8$ and comparison with other methods. Top: HR size 208×208 . Bottom: HR size 256×256 . For the stochastic methods SR Gaussian and SRFlow, the table has been realized on 200 samples. Note that our method outperforms in terms of the perceptual LPIPS metric and execution time, while PSNR and SSIM are optimal for blurry images, illustrating that these metrics are not relevant for texture SR. Images from [21].

$\mathbf{\Gamma}$ such that $\text{ADSN}(\mathbf{U}_{\text{ref}}) = \mathcal{N}(\mathbf{0}, \mathbf{\Gamma})$ and $\mathbf{\Lambda}$ associated to $\mathbf{\Gamma}$. Let \mathbf{X}_{SR} be a random image following the distribution of the SR samples obtained with Algorithm 1, then

$$\begin{aligned} \mathbb{E}_{\mathbf{X}_{\text{SR}}} (\|\mathbf{U} - \mathbf{X}_{\text{SR}}\|^2) &= \|\mathbf{U} - \mathbf{\Lambda}^T \mathbf{U}_{\text{LR}}\|^2 + \text{Tr}(\mathbf{\Gamma} - \mathbf{\Lambda}^T \mathbf{\Lambda} \mathbf{\Gamma} \mathbf{\Lambda}^T) \\ &\geq \|\mathbf{U} - \mathbf{\Lambda}^T \mathbf{U}_{\text{LR}}\|^2. \end{aligned}$$

In other words, the kriging component $\mathbf{\Lambda}^T \mathbf{U}_{\text{LR}}$ has a lower MSE in expectation than the perfect SR samples \mathbf{X}_{SR} .

The inequality is due to the positiveness of the covariance matrix $\mathbf{\Gamma} - \mathbf{\Lambda}^T \mathbf{\Lambda} \mathbf{\Gamma} \mathbf{\Lambda}^T$ of the innovation component. This property is consistent with the fact that computing $\mathbb{E}(\mathbf{X}|\mathbf{X}_{\text{LR}})$ aims to minimize the mean squared error. This is a specific illustration of the well-known regression to the mean problem in SR [24]. The same effect can be observed for the network SRFlow with temperature $\tau = 0$ which solves deterministically the MSE problem [8].

5. CONCLUSION

This study solves the problem of stochastic SR for stationary Gaussian textures. Such texture models constitute a base case for stochastic SR for which all the computations are accessible, leading to an efficient algorithm for stochastic SR and stochastic SR with a reference image. Our method outperforms some state of the art methods in terms of execution time and LPIPS metrics and we have demonstrated experimentally that the PSNR and the SSIM metrics are not compatible with the evaluation of the perceptual quality of SR samples for textures. This illustrates the necessity to study further the performance on microtextures of generic stochastic SR models. Furthermore, the presented method could be applied for other inverse problems involving non invertible operators of the form convolution followed by subsampling.

Acknowledgements: The authors acknowledge the support of the project MISTIC (ANR-19-CE40-005).

6. REFERENCES

- [1] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei 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.
- [2] Joan Bruna, Pablo Sprechmann, and Yann LeCun, “Super-Resolution with Deep Convolutional Sufficient Statistics,” in *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, Yoshua Bengio and Yann LeCun, Eds., 2016.
- [3] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, and Wenzhe 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.
- [4] Xintao Wang, Ke Yu, Shixiang Wu, Jinjin Gu, Yihao Liu, Chao Dong, Yu Qiao, and Chen Change Loy, “ESRGAN: Enhanced Super-Resolution Generative Adversarial Networks,” in *Computer Vision – ECCV 2018 Workshops*, Laura Leal-Taixé and Stefan Roth, Eds., Cham, 2019, pp. 63–79, Springer International Publishing.
- [5] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio, “Generative Adversarial Nets,” in *Advances in Neural Information Processing Systems*, Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger, Eds. 2014, vol. 27, Curran Associates, Inc.
- [6] Justin Johnson, Alexandre Alahi, and Li Fei-Fei, “Perceptual losses for real-time style transfer and super-resolution,” in *European Conference on Computer Vision*, 2016.
- [7] Karen Simonyan and Andrew Zisserman, “Very Deep Convolutional Networks for Large-Scale Image Recognition,” in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, Yoshua Bengio and Yann LeCun, Eds., 2015.
- [8] Andreas Lugmayr, Martin Danelljan, Luc Van Gool, and Radu Timofte, “SRFlow: Learning the Super-Resolution Space with Normalizing Flow,” in *Computer Vision – ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part V*, Berlin, Heidelberg, 2020, pp. 715–732, Springer-Verlag, event-place: Glasgow, United Kingdom.
- [9] Durk P Kingma and Prafulla Dhariwal, “Glow: Generative Flow with Invertible 1x1 Convolutions,” in *Advances in Neural Information Processing Systems*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, Eds. 2018, vol. 31, Curran Associates, Inc.
- [10] Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J. Fleet, and Mohammad Norouzi, “Image Super-Resolution Via Iterative Refinement,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 1–14, 2022.
- [11] Jonathan Ho, Ajay Jain, and Pieter Abbeel, “Denoising Diffusion Probabilistic Models,” in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, Eds. 2020, vol. 33, pp. 6840–6851, Curran Associates, Inc.
- [12] Yang Song and Stefano Ermon, “Generative modeling by estimating gradients of the data distribution,” in *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, pp. 11918–11930. Curran Associates Inc., Red Hook, NY, USA, 2019.
- [13] Johannes Hertrich, Antoine Houdard, and Claudia Redenbach, “Wasserstein Patch Prior for Image Superresolution,” *IEEE Transactions on Computational Imaging*, vol. 8, pp. 693–704, 2022.
- [14] Johannes Hertrich, Lan Dang Phuong Nguyen, Jean-François Aujol, Dominique Bernard, Yannick Berthoumieu, Abdellatif Saadaldin, and Gabriele Steidl, “PCA Reduced Gaussian Mixture Models with Applications in Superresolution,” *Inverse Problems and Imaging*, vol. 16, no. 2, pp. 341–366, 2022.
- [15] Pierrick Chatillon, Yann Gousseau, and Sidonie Lefebvre, “A statistically constrained internal method for single image super-resolution,” in *26th International Conference on Pattern Recognition, ICPR 2022, Montreal, QC, Canada, August 21–25, 2022*, pp. 1322–1328, IEEE.
- [16] Bruno Galerne, Yann Gousseau, and Jean-Michel Morel, “Random Phase Textures: Theory and Synthesis,” *IEEE Transactions on Image Processing*, vol. 20, no. 1, pp. 257–267, 2011.
- [17] Bruno Galerne, Arthur Leclaire, and Lionel Moisan, “Micro-texture inpainting through Gaussian conditional simulation,” in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2016, pp. 1204–1208.
- [18] Bruno Galerne and Arthur Leclaire, “Texture Inpainting Using Efficient Gaussian Conditional Simulation,” *SIAM Journal on Imaging Sciences*, vol. 10, no. 3, pp. 1446–1474, 2017.
- [19] Yuval Bahat and Tomer Michaeli, “Explorable super resolution,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 2716–2725.
- [20] Lionel Moisan, “Periodic Plus Smooth Image Decomposition,” *Journal of Mathematical Imaging and Vision*, vol. 39, no. 2, pp. 161–179, 2011.
- [21] Bruno Galerne, Yann Gousseau, and Jean-Michel Morel, “Micro-texture synthesis by phase randomization,” *Image Processing On Line*, vol. 1, 2011.
- [22] Zhou Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli, “Image quality assessment: from error visibility to structural similarity,” *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, 2004.
- [23] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, “The Unreasonable Effectiveness of Deep Features as a Perceptual Metric,” in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Los Alamitos, CA, USA, June 2018, pp. 586–595, IEEE Computer Society.
- [24] C. Sønderby, J. Caballero, L. Theis, W. Shi, and F. Huszár, “Amortised map inference for image super-resolution,” in *International Conference on Learning Representations*, 2017.