



HAL
open science

Stochastic super-resolution for Gaussian microtextures

Émile Pierret, Bruno Galerne

► **To cite this version:**

Émile Pierret, Bruno Galerne. Stochastic super-resolution for Gaussian microtextures. 2024. hal-04560456v1

HAL Id: hal-04560456

<https://hal.science/hal-04560456v1>

Preprint submitted on 26 Apr 2024 (v1), last revised 29 May 2024 (v2)

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 [†]

Bruno Galerne [‡]

April 26, 2024

Abstract

Super-Resolution (SR) is the problem that consists in reconstructing images that have been degraded by a zoom-out operator. This is an ill-posed problem that does not have a unique solution, and numerical approaches rely on a prior on high-resolution images. While optimization-based methods are generally deterministic, with the rise of image generative models more and more interest has been given to stochastic SR, that is, sampling among all possible SR images associated with a given low-resolution input. In this paper, we construct an efficient, stable and provably exact sampler for the stochastic SR of Gaussian microtextures. Even though our approach is limited regarding the scope of images it encompasses, our algorithm is competitive with deep learning state-of-the-art methods both in terms of perceptual metric and execution time when applied to microtextures. The framework of Gaussian microtextures also allows us to rigorously discuss the limitations of various reconstruction metrics to evaluate the efficiency of SR routines.

Keywords: stochastic super-resolution, Gaussian textures, conditional simulation, kriging, super-resolution with a reference image

1 Introduction

Super-Resolution (SR) algorithms aim at producing a High Resolution (HR) image corresponding to a Low Resolution (LR) one, the main challenge being to restore sharp edges as well as high frequency texture content that are lost when applying the zoom-out operator. Since the set of HR images compatible with a given LR input is an affine subspace of high dimension, a strong prior on realistic images is necessary to recover HR images of high visual quality within this subspace. In recent works, this prior is generally conveyed either by exploiting a large datasets of HR images via deep learning models or using a reference HR image to specify an adapted model for the unknown HR image. The former solution is rather generic while the later is particularly adapted for the SR of texture images that benefits from a particular prior on local statistics of the HR unknown image. Besides, rather than being deterministic, several contributions propose to tackle the one to many dilemma via stochastic SR which consists in sampling among all acceptable HR images associated with an LR input. Motivated by such recent contributions, in this paper we solve the stochastic SR problem when the HR texture image is assumed to follow a stationary Gaussian distribution.

1.1 State of the art and related work

First deep learning approaches for SR propose to optimize the L_2 or L_1 reconstruction loss using Convolutional Neural Network (CNN) trained in a end-to-end fashion [11, 40, 35, 37]. However, the CNN outputs tend to be blurry as they correspond to the mean of plausible SR solutions [56]. To avoid this regression to the mean issue, other approaches rely on a perceptual loss [4, 59, 52], that is, an L_2 loss between pre-trained VGG features [53] introduced in [31], or Generative Adversarial Network (GAN) [23] conditioned on the LR input image [39, 60]. More generally SR is a linear inverse problem that can be tackled by most general deep learning techniques for inverse problems in imaging [47] that go beyond end-to-end learning, notably variational frameworks such as deep image prior [57, 58] and Plug-and-Play

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

[†]Institut Denis Poisson – Université d’Orléans, Université de Tours, CNRS (emile.pierret@univ-orleans.fr, <https://www.idpoisson.fr/pierret/>, bruno.galerie@univ-orleans.fr, <https://www.idpoisson.fr/galerie/>)

[‡]Institut Universitaire de France (IUF) .

methods [32]. We do not explore this important body of literature in this work since we are more interested in stochastic methods.

In contrast with the deterministic approaches mentioned above, Stochastic SR algorithms sample within the set of plausible HR images consistent with the LR input. Lugmayr *et al.* [42] use a conditional normalizing flows [36] (a network with invertible layers) to learn this distribution via the change-of-variable formula. Diffusion or score-based models are generative networks which gradually learn to transform noise into samples of a distribution data [55, 26]. The diffusion framework can be used to train conditional diffusion models specific to SR [51]. Alternatively generic diffusion models can be used to provide approximate conditional sampling procedure for inverse problems such as SR [55, 34, 7, 51, 54], an approach that can be extended to latent diffusion models [50]. Alternatively, Jiang *et al.* [30] propose to wrap a module to a classical network to generate stochastic solutions of the SR samples. In general diffusion models have impressive performance to generate natural images but can have a lack of realism to generate high-frequency textures (as confirmed below by our comparative experiments). Other networks can provide results with some artefacts, even though this may be overcome using parametric priors [6].

Another approach is to apply reference-based SR [14, 13, 65, 5, 41]. In this setting the LR input is provided with a companion HR image that conveys information on the content of the unknown HR solution, e.g. by providing a HR image of the content from another viewpoint. Particularly relevant for this paper, in the context of super-resolution of stationary texture images, the patch distribution of the reference image can be exploited by minimizing the Wasserstein distance between HR output patches and the reference patches [25]. This variational approach, coined Wasserstein Patch Prior (WPP), exploits important statistics proved to be relevant for texture synthesis [24, 18, 28, 29]. Let us mention that WPP can be used as a training loss for conditional generative models [2] and made more robust using semi-unbalanced optimal transport [44, 45]. The WPP framework provides adaptive statistical guarantees to the proven efficiency of patch-based approaches [66, 63], which is critical for the SR of texture images.

In this work, we will solely focus on a precise class of textures, the so-called Gaussian microtextures [19]. This is a simple texture model that assumes that the texture is stationary (statistically invariant by translation) and follows a Gaussian distribution. Even though the class of texture images well-reproduced by this model is restricted, the Gaussian modeling allows for many probabilistic and statistical tools to be used for processing these textures. The textures can be summarized by a local texton [9, 16], can be generated on arbitrary large continuous domains [15, 17], extended to video and visually mixed using Wasserstein barycenters between Gaussian distributions [62]. Closely related to our work, the Gaussian modeling allows for solving texture inpainting via Gaussian conditional sampling based on kriging formulas [22, 21]. This work provides an iterative algorithm based on conjugate gradient descent last work provides an iterative algorithm based on conjugate gradient descent (CGD) that we extend to SR and use as a reference for evaluating our faster “Gaussian SR” solution. A preliminary version of the present work has been presented at ICASSP 2023 [49].

1.2 Contributions

Our approach builds on the kriging framework used for Gaussian inpainting [22, 21]. We further exploit the stationarity of the zoom-out operator to obtain an exact and fast convolution-based algorithm for stochastic SR of grayscale Gaussian textures. Due to correlation between color channels, this algorithm does not translate straightforwardly to RGB textures, but we propose an approximate algorithm for RGB textures that proves to be experimentally close to the exact but time-consuming algorithm. Finally, on a practical viewpoint, we extend our approach to the context of SR given a reference HR image, and show that, when applied to Gaussian textures, our fast Gaussian SR algorithm is both perceptually better and order of magnitudes faster than recent deep learning or optimal transport-based methods. While our approach has several inherent limitations that we discuss and illustrate, our strong experimental results show that it is of practical interest when dealing with simple microtextures.

1.3 Plan of the paper

The plan of the paper is as follows. In 2 we present our framework and remind results about Gaussian conditional simulation. Then, we detail an exact, stable and efficient simulation procedure of conditional SR for grayscale Gaussian microtextures. In 3, we propose an approximate algorithm for stochastic SR of RGB micro textures and compare it with the CGD-based reference approach. The extension of our sampling algorithm in the practical settings of SR given a reference image is discussed in 4 where we also compare our results with state-of-the-art methods. This

comparison leads to a discussion on adapted metrics to evaluate textures SR. Finally, we document the shortcomings of our approach and show that it extends to more general linear operators in 5 and conclude our paper in 6.

1.4 Notation

Let $M, N > 2$ be the size of the HR images and $r > 1$ a zoom-out factor such that the integers $M/r, N/r$ are the size of the LR images (it is assumed that r divides both M and N). We write $[k] = \{0, \dots, k-1\}$ for any integer $k \geq 1$, $\Omega_{M,N} = [M] \times [N]$ denotes the pixel grid, and all the images are extended on \mathbb{Z}^2 by periodization. For $\mathbf{u}, \mathbf{v} \in \mathbb{R}^{\Omega_{M,N}}$, $\mathbf{u} \star \mathbf{v}$ designates the discrete and periodic convolution defined by $(\mathbf{u} \star \mathbf{v})(x) = \sum_{y \in \Omega_{M,N}} \mathbf{u}(x-y)\mathbf{v}(y)$ for $x \in \Omega_{M,N}$. Given a kernel $\mathbf{t} \in \mathbb{R}^{\Omega_{M,N}}$, let us denote $\mathbf{C}_t \in \mathbb{R}^{\Omega_{M,N} \times \Omega_{M,N}}$ the matrix associated with the convolution by \mathbf{t} such that for $\mathbf{u} \in \mathbb{R}^{\Omega_{M,N}}$, $\mathbf{C}_t \mathbf{u} = \mathbf{t} \star \mathbf{u} \in \mathbb{R}^{\Omega_{M,N}}$. $\mathbf{I}_{\Omega_{M,N}}$ denotes the identity matrix on $\Omega_{M,N}$ while $\mathbf{1}_{\Omega_{M,N}}$ denotes the constant image of $\mathbb{R}^{\Omega_{M,N}}$ with value 1. $\check{\mathbf{u}}$ stands for the symmetric image of \mathbf{u} defined by for $\check{\mathbf{u}}(x) = \mathbf{u}(-x)$ for $x \in \Omega_{M,N}$. We express by $\mathcal{F}(\mathbf{u})$ or $\hat{\mathbf{u}}$ the Discrete Fourier Transform (DFT) of \mathbf{u} defined by $\hat{\mathbf{u}}(x) = \sum_{y \in \Omega_{M,N}} \mathbf{u}(y) \exp(-\frac{2i\pi x_1 y_1}{M}) \exp(-\frac{2i\pi x_2 y_2}{N})$ for $x \in \Omega_{M,N}$. Let us recall that \mathcal{F} is invertible with $\mathcal{F}^{-1} = \frac{1}{MN} \overline{\mathcal{F}}$, $\mathcal{F}(\check{\mathbf{u}}) = \overline{\mathcal{F}(\mathbf{u})}$ and $\mathcal{F}(\mathbf{u} \star \mathbf{v}) = \mathcal{F}(\mathbf{u}) \odot \mathcal{F}(\mathbf{v})$ where \odot denotes the componentwise product. We denote by \mathcal{F}_r the Fourier transform of images of $\mathbb{R}^{\Omega_{M/r, N/r}}$.

2 The conditional super-resolution of grayscale Gaussian microtextures

In this section we solve the problem of Stochastic super-resolution of grayscale Gaussian microtextures using a known zoom-out operator given by a convolution. Exploiting both the stationarity of the LR measurements and the HR texture model, our main result states that one can exactly sample new HR microtextures by applying a specific convolution operator. We first recall the specifics of the zoom-out operator and the ADSN microtexture model. Then we recall how solving for kriging coefficients allows for conditional Gaussian simulation and establish our main result.

2.1 The zoom-out operator

The most common zoom-out operator used in the SR literature is the Matlab bicubic zoom-out operator provided by the function `imresize` [39, 60, 3]. Let us denote \mathbf{A} this bicubic zoom-out operator by a factor r . It can be described as a convolution followed by a subsampling \mathcal{S} with stride r such that for $\mathbf{u} \in \mathbb{R}^{\Omega_{M,N}}$, $\mathbf{S}\mathbf{u} \in \mathbb{R}^{\Omega_{M/r, N/r}}$ and for $x \in \Omega_{M/r, N/r}$, $(\mathbf{S}\mathbf{u})(x) = \mathbf{u}(rx)$ and for $x \in \Omega_{M/r, N/r}$, $\mathbf{v} \in \mathbb{R}^{\Omega_{M/r, N/r}}$, $(\mathcal{S}^T \mathbf{v})(rx) = \mathbf{v}(x)$, see e.g. [3]. We denote by \mathbf{c} the kernel associated with the bicubic kernel $\mathbf{A} = \mathbf{S}\mathbf{C}_c$. In the remaining of the paper, we study the following inverse problem

$$\mathbf{u}_{\text{LR}} = \mathbf{A}\mathbf{u}_{\text{HR}} \quad (1)$$

where \mathbf{u}_{LR} is the observed low resolution version of an unknown HR image \mathbf{u}_{HR} . This problem is ill-posed, its space of solutions is a high-dimensional affine subspace, and most such solutions are not desirable such that blurry HR versions of the LR image \mathbf{u}_{LR} . Our main result provides an efficient sampler of solutions of Equation 1 under the assumption that \mathbf{u}_{HR} follows an ADSN distribution, that is, a Gaussian microtexture model we recall below.

2.2 The ADSN model associated with a grayscale image

Given a grayscale image $\mathbf{u} \in \mathbb{R}^{\Omega_{M,N}}$ with mean grayscale $m \in \mathbb{R}$, one defines the Asymptotic Discrete Spot Noise (ADSN) associated with \mathbf{u} , $\text{ADSN}(\mathbf{u})$ as the distribution of $m\mathbf{1}_{\Omega_{M,N}} + \mathbf{t} \star \mathbf{w}$ where $\mathbf{t} = \frac{1}{\sqrt{MN}}(\mathbf{u} - m\mathbf{1}_{\Omega_{M,N}})$ is called the texton associated with \mathbf{u} and $\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{\Omega_{M,N}})$ [19]. Let us recall that in practice it is preferable to replace \mathbf{u} by its periodic component [46] to avoid boundary issues [19]. $\text{ADSN}(\mathbf{u})$ is a Gaussian distribution with mean $m\mathbf{1}_{\Omega_{M,N}}$ and covariance matrix $\mathbf{\Gamma} = \mathbf{C}_t \mathbf{C}_t^T = \mathbf{C}_{t \star \check{t}} \in \mathbb{R}^{\Omega_{M,N} \times \Omega_{M,N}}$. Samples can be generated fastly executing the convolution with the Fast Fourier Transform (FFT). $\text{ADSN}(\mathbf{u})$ is stationary, that is, for $\mathbf{X} \sim \text{ADSN}(\mathbf{u})$ and any $y \in \mathbb{Z}^2$, $\mathbf{X}(\cdot - y)$ has the same distribution as \mathbf{X} . This implies that the pixelwise variance of this distribution is constant. This restricts the application of ADSN model to microtextures that do not have structures. Practically, subtracting and adding the common mean gray-level to LR and HR images is considered as a preprocessing and a post-processing. In what follows, we will often consider the centered distribution $\mathbf{t} \star \mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{\Gamma})$ that will be called $\text{ADSN}(\mathbf{u})$ with a slight abuse of notation.




Figure 1: Super resolution of a Gaussian texture using kriging. From left to right: LR image, original HR image, conditional Gaussian sample and its corresponding kriging and innovation components. HR images size is 416×640 and the zoom-out factor is $r = 8$. The sample is the sum of the kriging component and the innovation component. Grayscale mean is added to all components for visualization. The deterministic kriging component is a kind of pseudo-inversion of the zoom-out operator applied to \mathbf{u}_{LR} that benefits from the knowledge of the covariance of the texture model. The innovation component provides an independent texture grain that is sampled for each realization of \mathbf{u}_{SR} .

2.3 The kriging reasoning

Let $\mathbf{X} \in \mathbb{R}^{\Omega_{M,N}}$ be a random vector following a Gaussian multivariate law $\mathcal{N}(\mathbf{0}, \mathbf{\Gamma})$ with $\mathbf{\Gamma} \in \mathbb{R}^{\Omega_{M,N} \times \Omega_{M,N}}$ and let $\mathbf{A} \in \mathbb{R}^{\Omega_{M/r, N/r} \times \Omega_{M,N}}$ be a linear operator. The simulation of \mathbf{X} knowing $\mathbf{A}\mathbf{X}$ can be computed using the following theorem, recalled in [21], which expresses the link between orthogonality and independence for Gaussian vectors.

Theorem 1 (Conditional Gaussian simulation). *Let $\mathbf{X} \sim \mathcal{N}(\mathbf{0}, \mathbf{\Gamma})$ and \mathbf{A} be a linear operator. The two Gaussian vectors $\mathbb{E}(\mathbf{X}|\mathbf{A}\mathbf{X})$ and $\mathbf{X} - \mathbb{E}(\mathbf{X}|\mathbf{A}\mathbf{X})$ are independent. Consequently, if $\tilde{\mathbf{X}}$ is independent of \mathbf{X} with the same distribution then $\mathbb{E}(\mathbf{X}|\mathbf{A}\mathbf{X}) + [\tilde{\mathbf{X}} - \mathbb{E}(\tilde{\mathbf{X}}|\mathbf{A}\tilde{\mathbf{X}})]$ has the same distribution as \mathbf{X} knowing $\mathbf{A}\mathbf{X}$.*

Furthermore, in the Gaussian context, if the distribution is zero-mean, the relation between the conditional expectation $\mathbb{E}(\mathbf{X}|\mathbf{A}\mathbf{X})$ and $\mathbf{A}\mathbf{X}$ is linear, as detailed in the following.

Theorem 2 (Gaussian kriging). *Let $\mathbf{X} \sim \mathcal{N}(\mathbf{0}, \mathbf{\Gamma})$ be a Gaussian vector with zero mean and $\mathbf{A} \in \mathbb{R}^{\Omega_{M/r, N/r} \times \Omega_{M,N}}$ be a linear operator. There exists $\mathbf{\Lambda} \in \mathbb{R}^{\Omega_{M/r, N/r} \times \Omega_{M,N}}$ such that $\mathbb{E}(\mathbf{X}|\mathbf{A}\mathbf{X}) = \mathbf{\Lambda}^T \mathbf{A}\mathbf{X}$ and $\mathbb{E}(\mathbf{X}|\mathbf{A}\mathbf{X}) = \mathbf{\Lambda}^T \mathbf{A}\mathbf{X}$ if and only if $\mathbf{\Lambda}$ verifies the matrix equation*

$$\mathbf{A}\mathbf{\Gamma}\mathbf{A}^T \mathbf{\Lambda} = \mathbf{A}\mathbf{\Gamma}. \quad (2)$$

Theorem 2 is a standard [38]. Yet since solving Equation 2 is central in what follows, we give a detailed proof in A. The matrix $\mathbf{\Lambda}$ is called the kriging matrix and Equation 2 is called the kriging equation. Note that if \mathbf{A} is invertible, a trivial solution for $\mathbf{\Lambda}^T$ is \mathbf{A}^{-1} . In our context of super-resolution, \mathbf{A} is not invertible and we are given an LR image $\mathbf{u}_{\text{LR}} = \mathbf{A}\mathbf{u}_{\text{HR}}$ with \mathbf{u}_{HR} unknown. Supposing that \mathbf{u}_{HR} is a realization of a law $\text{ADSN}(\mathbf{u}) = \mathcal{N}(\mathbf{0}, \mathbf{\Gamma})$, we would like to sample HR images \mathbf{u}_{SR} such that \mathbf{u}_{SR} follows $\mathcal{N}(\mathbf{0}, \mathbf{\Gamma})$ conditioned on $\mathbf{A}\mathbf{u}_{\text{SR}} = \mathbf{u}_{\text{LR}}$. Using Theorem 1 and Theorem 2, we aim at computing a HR sample image

$$\mathbf{u}_{\text{SR}} = \mathbf{\Lambda}^T \mathbf{u}_{\text{LR}} + \tilde{\mathbf{u}} - \mathbf{\Lambda}^T \mathbf{A} \tilde{\mathbf{u}} \quad (3)$$

where $\tilde{\mathbf{u}} \sim \mathcal{N}(\mathbf{0}, \mathbf{\Gamma})$ and $\mathbf{\Lambda}$ verifies the kriging Equation 2. Note that \mathbf{u}_{SR} is the sum of a deterministic term $\mathbf{\Lambda}^T \mathbf{u}_{\text{LR}}$, called the *kriging component*, and a stochastic term $\tilde{\mathbf{u}} - \mathbf{\Lambda}^T \mathbf{A} \tilde{\mathbf{u}}$ not depending on \mathbf{u}_{LR} , called the *innovation component*. As presented in 1, the kriging component is a HR blurred version of the LR image with added covariance information while the innovation component provides independently the granular aspect of the microtexture. In general, the kriging equation needs to be solved by an iterative method, even with the stationarity assumption of the Gaussian law [21]. Due to the stationarity of the Gaussian law $\mathcal{N}(\mathbf{0}, \mathbf{\Gamma})$ and the convolution form of \mathbf{A} , we show below that it is possible to solve Equation 2 exactly and fastly in the Fourier domain.

2.4 Solving the kriging equation in the super-resolution context

We show in this section that in our SR context, the kriging equation can be solved efficiently with a non-iterative algorithm.

Equation 2 may have several solutions and we only consider the specific solution $\mathbf{\Lambda} = (\mathbf{A}\mathbf{\Gamma}\mathbf{A}^T)^\dagger \mathbf{A}\mathbf{\Gamma}$ where $(\mathbf{A}\mathbf{\Gamma}\mathbf{A}^T)^\dagger$ is the pseudo-inverse of $\mathbf{A}\mathbf{\Gamma}\mathbf{A}^T$. The main practical issue is now to compute efficiently $\mathbf{\Lambda}^T \mathbf{v}$ for a given LR image $\mathbf{v} \in \mathbb{R}^{\Omega_{M/r, N/r}}$. As shown by the following proposition, this simply corresponds to applying a specific convolution.




Figure 2: Observation of the modulus logscale of DFT of the different components from 1. λ is the kernel associated with the kriging matrix Λ^T . Note that the kriging component retrieves low frequencies from LR image that are completed with the high frequencies of the innovation component.

Proposition 1 (Kriging as a convolution). *Let $A = SC_c$ and let $\Gamma = C_t C_t^T = C_{t\check{t}}$ be the covariance of a Gaussian distribution $\text{ADSN}(\mathbf{u})$ with associated texton $\mathbf{t} = \frac{1}{\sqrt{MN}}(\mathbf{u} - m\mathbf{1}_{\Omega_{M,N}})$. $\Lambda = (A\Gamma A^T)^\dagger A\Gamma$ is an exact solution of Equation 2 and for all $\mathbf{v} \in \mathbb{R}^{\Omega_{M/r,N/r}}$,*

$$\Lambda^T \mathbf{v} = \lambda \star (\mathbf{S}^T \mathbf{v}) \quad (4)$$

where $\lambda = \mathbf{t} \star \check{t} \star \check{c} \star (\mathbf{S}^T \kappa^\dagger)$ with $\kappa = \mathbf{S}(\mathbf{t} \star \check{t} \star \mathbf{c} \star \check{c}) \in \mathbb{R}^{\Omega_{M/r,N/r}}$ and κ^\dagger the convolution kernel defined in Fourier domain by

$$\hat{\kappa}^\dagger(\omega) = \begin{cases} \frac{1}{\hat{\kappa}(\omega)} & \text{if } \hat{\kappa}(\omega) \neq 0, \\ 0 & \text{otherwise,} \end{cases} \quad \omega \in \mathbb{R}^{\Omega_{M/r,N/r}}.$$

Proof. Equation 2 is a normal equation associated with the least squares problem $\text{argmin}_{\Lambda \in \mathbb{R}^{\Omega_{M/r,N/r} \times \Omega_{M,N}}} \|C_t^T A^T \Lambda - C_t^T\|_2^2$. Consequently, $\Lambda = (A\Gamma A^T)^\dagger A\Gamma$ is one of the solutions of Equation 2. Then, we use the following elementary lemma (proved in B for completeness).

Lemma 1 (Convolution and subsampling). *1. If C_α is the convolution on $\Omega_{M,N}$ by the kernel $\alpha \in \Omega_{M,N}$, $SC_\alpha S^T$ is the convolution on $\Omega_{M/r,N/r}$ by the kernel $S\alpha$, that is, $SC_\alpha S^T = C_{S\alpha}$.*
2. If C_β is a convolution on $\Omega_{M/r,N/r}$ by the kernel $\beta \in \Omega_{M/r,N/r}$, $S^T C_\beta = C_{S^T \beta} S^T$.

By Lemma 1, $A\Gamma A^T = \mathbf{S}(C_c \Gamma C_c^T) \mathbf{S}^T$ is a convolution matrix associated with the kernel $\kappa = \mathbf{S}(\mathbf{t} \star \check{t} \star \mathbf{c} \star \check{c}) \in \mathbb{R}^{\Omega_{M/r,N/r}}$. Moreover, the pseudo-inverse of a convolution is also a convolution. $(A\Gamma A^T)^\dagger$ is a convolution associated with the kernel κ^\dagger such that for $\omega \in \mathbb{R}^{\Omega_{M/r,N/r}}$, $\hat{\kappa}^\dagger(\omega) = \begin{cases} \frac{1}{\hat{\kappa}(\omega)} & \text{if } \hat{\kappa}(\omega) \neq 0 \\ 0 & \text{otherwise.} \end{cases}$. Also by Lemma 1, $\mathbf{S}^T (A\Gamma A^T)^\dagger = \mathbf{S}^T C_{\kappa^\dagger} = C_{S^T \kappa^\dagger} \mathbf{S}^T$. Consequently,

$$\Lambda^T = \Gamma^T A^T (A\Gamma A^T)^\dagger = \Gamma^T C_c^T S^T C_{\kappa^\dagger} = \Gamma^T C_c^T C_{S^T \kappa^\dagger} S^T = C_\lambda S^T$$

where $\lambda = \mathbf{t} \star \check{t} \star \check{c} \star \mathbf{S}^T \kappa^\dagger$. □

In what follows we refer to λ as the *kriging kernel*. 2 shows the DFT of the kernel λ associated with the images of 1. In contrast with a non-adaptative pseudo-inversion of the bicubic kernel that would be isotropic, we can observe that the kernel λ is adapted to the covariance structure of the texture and that it amplifies its characteristic frequencies.

To compute SR samples \mathbf{u}_{SR} given by Equation 3, we need to compute $\lambda \in \mathbb{R}^{\Omega_{M,N}}$ the kernel associated with Λ^T as described in Proposition Proposition 1 and then apply Λ^T as described in Equation 3. These two steps can be done fastly in the Fourier domain. The exact and fast corresponding procedure is described in Algorithm 1. To generate several samples, one only needs to rerun the second part of the algorithm.

The application of Λ^T implies the division by $\hat{\kappa}$ in the Fourier domain. However, in practice, this division can be unstable since a non-zero low value of $\hat{\kappa}$ amplifies its corresponding frequency. Actually, in our context, Λ^T is applied to LR versions of realizations of an ADSN model that are $\mathbf{v} = \mathbf{A}(\mathbf{t} \star \mathbf{w})$ with $\mathbf{w} \in \mathbb{R}^{\Omega_{M,N}}$ an unknown white Gaussian noise. Proposition 2 ensures that the application of Λ^T is stable when applied to images that comply with this assumption.

Proposition 2 (Stability of the kriging operator on the subspace of the LR ADSN samples). *Let $\Lambda^T = \Gamma A^T (A\Gamma A^T)^\dagger \in \mathbb{R}^{\Omega_{M,N} \times \Omega_{M/r,N/r}}$. Then,*

$$\forall \mathbf{w} \in \mathbb{R}^{\Omega_{M,N}}, \|\Lambda^T \mathbf{A}(\mathbf{t} \star \mathbf{w})\|_2 \leq \|C_t\|_2 \|\mathbf{w}\|_2 \leq \|\mathbf{t}\|_1 \|\mathbf{w}\|_2. \quad (5)$$

Algorithm 1 Super-resolution sampling for grayscale images with known Gaussian model

Input: An image $\mathbf{u}_{\text{LR}} \in \mathbb{R}^{\Omega_{M/N/r}}$, 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 $\mathbf{A} = \mathbf{S}\mathbf{C}_c$.

Preprocessing:

Compute the grayscale mean m from \mathbf{u}_{LR} and set $\mathbf{u}_{\text{LR}} := \mathbf{u}_{\text{LR}} - m\mathbf{1}_{\Omega_{M,N}}$

Step 1: Computation of the kriging kernel

Store the DFT transform of the kernel $\lambda = \mathbf{t} \star \check{\mathbf{t}} \star \check{\mathbf{c}} \star \mathbf{S}^T(\kappa^\dagger)$

Step 2: Simulation of \mathbf{u}_{SR}

Sample $\tilde{\mathbf{u}} = \mathbf{t} \star \mathbf{w}$ where $\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{\Omega_{M,N}})$

Compute $\mathbf{u}_{\text{SR}} = \lambda \star \mathbf{S}^T(\mathbf{u}_{\text{LR}} - \mathbf{A}\tilde{\mathbf{u}}) + \tilde{\mathbf{u}}$

Postprocessing:

Output: $m\mathbf{1}_{\Omega_{M,N}} + \mathbf{u}_{\text{SR}}$

Proof. As a reminder [48], let \mathbf{M} being a real-valued matrix,

$$\mathbf{M}^\dagger = (\mathbf{M}^T \mathbf{M})^\dagger \mathbf{M}^T$$

With $\mathbf{M} = \mathbf{A}\mathbf{C}_t^T$,

$$\mathbf{A}\mathbf{C}_t^{T\dagger} = (\mathbf{A}\mathbf{C}_t(\mathbf{A}\mathbf{C}_t)^T)^\dagger \mathbf{A}\mathbf{C}_t = (\mathbf{A}\mathbf{\Gamma}\mathbf{A}^T)^\dagger \mathbf{A}\mathbf{C}_t$$

Consequently, for $\mathbf{w} \in \mathbb{R}^{\Omega_{M,N}}$,

$$\mathbf{\Lambda}^T \mathbf{A}(\mathbf{t} \star \mathbf{w}) = \mathbf{\Gamma}\mathbf{A}^T (\mathbf{A}\mathbf{\Gamma}\mathbf{A}^T)^\dagger \mathbf{A}\mathbf{C}_t \mathbf{w} = \mathbf{C}_t (\mathbf{A}\mathbf{C}_t)^T ((\mathbf{A}\mathbf{C}_t)^T)^\dagger \mathbf{w}$$

and

$$\|\mathbf{\Lambda}^T \mathbf{A}(\mathbf{t} \star \mathbf{w})\|_2 \leq \|\mathbf{C}_t\|_2 \underbrace{\|(\mathbf{A}\mathbf{C}_t)^T ((\mathbf{A}\mathbf{C}_t)^T)^\dagger\|_2}_{\leq 1} \|\mathbf{w}\|_2 \leq \|\mathbf{C}_t\|_2 \|\mathbf{w}\|_2 \leq \|\mathbf{t}\|_1 \|\mathbf{w}\|_2,$$

using that for any matrix \mathbf{M} one has $\|\mathbf{M}\mathbf{M}^\dagger\|_2 \leq 1$. □

3 The conditional super-resolution of RGB Gaussian microtextures

3.1 Description of the framework and notation

We denote by $3\Omega_{M,N}$ the space of the RGB images and by $\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3$ the three chanel of a given RGB image $\mathbf{u} \in \mathbb{R}^{3\Omega_{M,N}}$. $\hat{\mathbf{u}}$ designates the Discrete Fourier transform of \mathbf{u} defined as the 2D DFT of each channel. The associated texton of an RGB image \mathbf{u} is $\mathbf{t} = \frac{1}{\sqrt{NM}}(\mathbf{u} - \mathbf{m}) \in \mathbb{R}^{3\Omega_{M,N}}$ where $\mathbf{m} \in \mathbb{R}^{3\Omega_{M,N}}$ is the mean RGB color. The ADSN model

ADSN(\mathbf{u}) is defined as the distribution of $\mathbf{m} + \begin{pmatrix} \mathbf{t}_1 \star \mathbf{w} \\ \mathbf{t}_2 \star \mathbf{w} \\ \mathbf{t}_3 \star \mathbf{w} \end{pmatrix}$ where $\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{\Omega_{M,N}})$. Note that the white Gaussian noise is the same in each channel [19]. ADSN(\mathbf{u}) is a multivariate Gaussian law associated with the covariance matrix

$\mathbf{\Gamma} = \begin{pmatrix} \mathbf{C}_{t_1} \\ \mathbf{C}_{t_2} \\ \mathbf{C}_{t_3} \end{pmatrix} \begin{pmatrix} \mathbf{C}_{t_1} \\ \mathbf{C}_{t_2} \\ \mathbf{C}_{t_3} \end{pmatrix}^T \in \mathbb{R}^{3\Omega_{M,N} \times 3\Omega_{M,N}}$ which is the matrix of a multi-channel convolution, that is, each output channel is

a linear combination of 2D convolutions of the input RGB channels. More precisely, for all $\mathbf{u} \in \mathbb{R}^{3\Omega_{M,N}}$,

$$(\mathbf{\Gamma}\mathbf{u})_i = \mathbf{t}_i \star \check{\mathbf{t}}_1 \star \mathbf{u}_1 + \mathbf{t}_i \star \check{\mathbf{t}}_2 \star \mathbf{u}_2 + \mathbf{t}_i \star \check{\mathbf{t}}_3 \star \mathbf{u}_3, \quad 1 \leq i \leq 3. \quad (6)$$

Unfortunately, this matrix is not diagonal in the Fourier basis which makes a major difference with the grayscale framework. We still denote by $\mathbf{A} = \mathbf{S}\mathbf{C}_c$ the zoom-out operator that acts on each channel for an RGB image.

3.2 Approximating the kriging operator in the RGB case

To simulate conditional sample of the RGB ADSN model, the kriging equation 2 and Lemma 1 are still valid, except that now $\mathbf{A}\mathbf{\Gamma}\mathbf{A}^T$ and $(\mathbf{A}\mathbf{\Gamma}\mathbf{A}^T)^\dagger$ are multi-channel convolutions. However, while the pseudo-inverse of a 2D convolution is diagonal in the Fourier basis, this is not the case for a multi-channel convolution.

Computing the pseudo-inverse $(\mathbf{A}\mathbf{\Gamma}\mathbf{A}^T)^\dagger$ turns out to be a critical task. Even for small images, a direct computation of the pseudo-inverse of $\mathbf{A}\mathbf{\Gamma}\mathbf{A}^T$ with standard routines leads to instabilities. A more principled approach is to remark that for each $\mathbf{v} \in \mathbb{R}^{3\Omega_{M/r,N/r}}$ and each $\omega \in \Omega_{M/r,N/r}$, there exists a 3×3 matrix $\hat{K}(\omega)$ such that $\mathcal{F}_r(\mathbf{A}\mathbf{\Gamma}\mathbf{A}^T\mathbf{v})(\omega) = \hat{K}(\omega)\hat{\mathbf{v}}(\omega)$, reducing the problem to computing one 3×3 pseudo-inverse for each frequency.

Still, each matrix is close to be singular which leads to a high instability in practice.

In order to keep an algorithm which is as fast and as stable as in the grayscale case (Algorithm 1), we propose to make the following approximation: let $\mathbf{X} \in \mathbb{R}^{3\Omega_{M,N}}$, following the Gaussian law $\mathcal{N}(\mathbf{0}, \mathbf{\Gamma})$,

$$\mathbb{E}[\mathbf{X}_i \mid \mathbf{A}\mathbf{X}_1, \mathbf{A}\mathbf{X}_2, \mathbf{A}\mathbf{X}_3] \approx \mathbb{E}[\mathbf{X}_i \mid \mathbf{A}\mathbf{X}_i], \quad 1 \leq i \leq 3. \quad (7)$$

The rationale behind this approximation is that to reconstruct the channel \mathbf{X}_i from the three channels $\mathbf{A}\mathbf{X} = (\mathbf{A}\mathbf{X}_1 \mathbf{A}\mathbf{X}_2 \mathbf{A}\mathbf{X}_3)$, the more relevant information is in $\mathbf{A}\mathbf{X}_i$. This is a reasonable assumption with respect to the behavior of natural images and the form of the bicubic convolution kernel. This approximation amounts to use

$$\mathbf{\Lambda}_{\text{approx}} = \begin{pmatrix} \mathbf{\Lambda}_1 & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{\Lambda}_2 & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{\Lambda}_3 \end{pmatrix} \in \mathbb{R}^{3\Omega_{M/r,N/r} \times 3\Omega_{M,N}} \quad (8)$$

as an approximate solution of Equation 2 in the RGB setting, where for each $1 \leq i \leq 3$, $\mathbf{\Lambda}_i \in \mathbb{R}^{\Omega_{M/r,N/r} \times \Omega_{M,N}}$ is a solution of

$$\mathbf{A}\mathbf{\Gamma}_i\mathbf{A}^T\mathbf{\Lambda}_i = \mathbf{A}\mathbf{\Gamma}_i, \quad 1 \leq i \leq 3, \quad (9)$$

$\mathbf{\Gamma}_i$ being the covariance matrix of the i th channel. While $\mathbf{\Lambda}_{\text{approx}}$ is not a solution of the kriging equation, it is the kriging matrix associated with an ADSN law for which the RGB channels are uncorrelated (which is not an interesting model to generate textures). However, we apply this kriging operator $\mathbf{\Lambda}_{\text{approx}}^T$ to textures having properly correlated channels according to the ADSN model $\mathcal{N}(\mathbf{0}, \mathbf{\Gamma})$. More precisely, a RGB Gaussian SR sample \mathbf{u}_{SR} is defined as

$$\mathbf{u}_{\text{SR},i} = \mathbf{\Lambda}_i^T(\mathbf{u}_{\text{LR},i} - \mathbf{A}\tilde{\mathbf{u}}_i) + \tilde{\mathbf{u}}_i, \quad 1 \leq i \leq 3, \quad (10)$$

where, for $1 \leq i \leq 3$, $\tilde{\mathbf{u}}_i = \mathbf{t}_i \star \mathbf{w}$ with $\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{\Omega_{M,N}})$ the common noise used for each channel and the multiplication by $\mathbf{\Lambda}_i^T$ is done following Algorithm 1.

In the next section, we assess that our proposed approximation is harmless by evaluating it with an exact reference iterative alternative.

3.3 Comparison with the reference CGD algorithm

In [21], the authors propose to solve the kriging equation for Gaussian microtextures to resolve the inpainting by conditional simulation. To solve Equation 2, they apply a CGD to compute $(\mathbf{A}\mathbf{\Gamma}\mathbf{A}^T)^\dagger \varphi$ for a given $\varphi \in \mathbb{R}^{\Omega_{M/r,N/r}}$. For completeness, the algorithm is recalled in C. This is an iterative algorithm with a stopping criteria ε and a given number of iterations. The output ψ of CGD tends to minimize

$$\|(\mathbf{A}\mathbf{\Gamma}\mathbf{A}^T)^2\varphi - (\mathbf{A}\mathbf{\Gamma}\mathbf{A}^T)\psi\|_2 \quad (11)$$

which is the residual for the normal equation associated with Equation 2. We consider CGD with a high number of 10^6 steps as a reference. The recommended number of steps for the inpainting problem is 10^2 for fast results and 10^3 for high-quality results [20] and our experiments show that these recommendations remain valid for super-resolution. Results from our fast convolution-based algorithm and this reference CGD-based algorithm are presented in 3 for both the grayscale and RGB cases. The same realization of noise is used for the simulations so that both algorithm should produce the same image. Samples are visually similar for the two examples and when visualized as a sequence the texture details added by the CGD algorithm do not evolve after 10^3 steps.

In 1, we compare the samples generated with our algorithm and the CGD routine for different number of steps using the PSNR with respect to the CGD output after 10^6 iterations and the CGD residual as metrics. For the grayscale

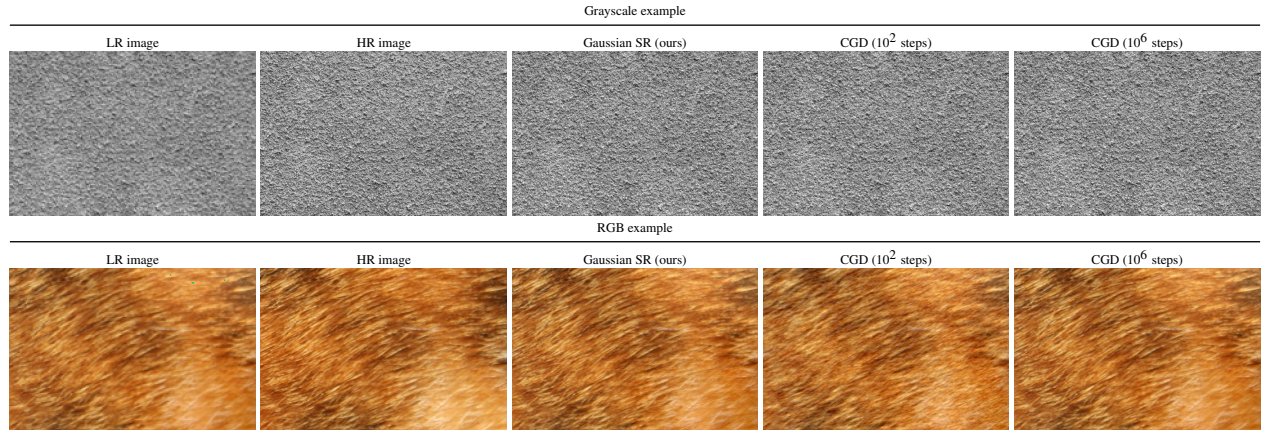


Figure 3: Comparison of our direct algorithm (Gaussian SR) and the iterative CGD-based algorithm for grayscale (first row) and RGB (second row) images. The HR image size is 512×768 , the SR factor is $r = 4$, and the same noise realization is used to compute the innovation components to allow visual comparison. The two methods provide similar results even for the RGB case for which our algorithm is not exact.

	Comparison with the reference CGD algorithm					
	Grayscale image			RGB image		
	Residual	Time(s) (CPU)	PSNR w.r.t CGD (10 ⁶ st.)	Residual	Time(s) (CPU)	PSNR w.r.t CGD (10 ⁶ st.)
Gaussian SR	1.32E-16	0.01	151.17	2.54E-1	0.01	37.94
CGD (10 ² steps)	2.74E-2	0.13	29.16	2.73E0	0.38	25.08
CGD (10 ³ steps)	1.03E-3	0.76	47.49	3.78E-1	2.49	30.19
CGD (10 ⁴ steps)	4.72E-5	7.44	67.60	1.36E-1	23.1	35.07
CGD (10 ⁵ steps)	1.30E-8	81.3	145.31	2.37E-2	258	39.71
CGD (10 ⁶ steps)	2.69E-42	749	-	4.97E-3	2588	-

Table 1: Quantitative evaluation of the results of 3. The residual is the norm from Equation 11. The CGD is executed with different numbers of steps where 10^2 is the pre-set parameter in the online demo [20] for inpainting. The grayscale example confirms that our method is exact in this context. The RGB example shows that our approximation is harmless for RGB images. The different execution times show that producing samples with similar quality to our Gaussian SR samples using the CGD algorithm is four to five orders of magnitude slower.

image, we observe that our method is indeed exact with very low residual value. The high PSNR indicates that our direct Gaussian SR algorithm is more precise than CGD 10^5 steps.

In the color case, as expected, Gaussian SR is not exact and provides higher residual and PSNR values. Still the Gaussian SR samples are closer to CGD outputs with a number of steps between 10^4 and 10^5 , which has been observed for several textures and for different zoom-out factors r . This validates our approach since the minor approximation we introduced is in practice negligible. Let us recall that the main objective of our approximation is to reduce the simulation time. As reported in 1, producing such an accurate RGB sample with the CGD algorithm is four to five orders of magnitude longer than with our direct convolution-based algorithm.

4 Gaussian SR in practice: Reference image and comparison with state-of-the-art

4.1 Gaussian SR with a reference image

To apply Gaussian SR algorithm, the ADSN distribution followed by \mathbf{u}_{HR} has to be known (via the computation of the associated texton \mathbf{t}). In the experiments of the previous section, the texton is computed from the ground truth HR image \mathbf{u}_{HR} , making the algorithm impractical.

From now on, following recent contributions on texture super-resolution [25, 2], we consider that the input image

\mathbf{u}_{LR} is given with a companion reference HR image \mathbf{u}_{ref} such that both HR images \mathbf{u}_{ref} and the unknown \mathbf{u}_{HR} are supposed to be realizations of the same texture model. More precisely, in our context based on the ADSN model, Algorithm 1 will be called with input an LR image \mathbf{u}_{LR} to be upscaled and the texton $\mathbf{t}_{ref} = \frac{1}{\sqrt{MN}}(\mathbf{u}_{ref} - m\mathbf{1}_{\Omega_{M,N}})$ associated with the reference image \mathbf{u}_{ref} , meaning that \mathbf{u}_{HR} is supposed to be a realization of the ADSN(\mathbf{u}_{ref}) law. This makes our routine a stochastic super-resolution algorithm with reference image.

4 presents illustrative results that show the practicality of this reference-based approach when the reference images are HR images of another part of the same material. 4 also displays the kriging and innovation components and one can observe that the ADSN model from the reference image is good enough to compute an adapted kriging component.




Figure 4: Illustration of Gaussian SR with a reference image (HR size is 208×208 and $r = 8$). Our method performs well on fabric textures and we can observe the grain provided by the innovation component for different types of texture. Some details are not recovered by our algorithm due to the stationarity assumption.

	5					6				
	PSNR \uparrow	LR-PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	Time	PSNR \uparrow	LR-PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	Time
Gaussian SR (ours)	17.05 \pm 0.04	159.24 \pm 0.04	0.08 \pm 0.01	0.22\pm0.01	0.01¹	19.14 \pm 0.09	154.52\pm0.36	0.20 \pm 0.01	0.23\pm0.01	0.01¹
Kriging comp.	18.76	159.30	0.11	0.75	-	21.42	154.47	0.30	0.52	-
Bicubic	20.74	37.61	0.18	0.76	-	20.74	38.35	0.28	0.66	-
WPP	20.04	29.29	0.14	0.36	44.0 ²	17.68	18.84	0.19	0.30	64.0 ²
SRFlow ($\tau = 0$)	21.44	54.61	0.20	0.70	0.22 ²	21.64	51.63	0.29	0.54	0.23 ²
SRFlow ($\tau = 0.9$)	18.29 \pm 0.36	55.13 \pm 0.15	0.12 \pm 0.01	0.30 \pm 0.03	0.19 ²	18.21 \pm 0.53	54.02 \pm 0.23	0.16 \pm 0.02	0.39 \pm 0.06	0.20 ²
DDRM	21.57\pm0.05	56.59 \pm 0.17	0.20\pm0.00	0.70 \pm 0.02	1.66 ²	22.44\pm0.04	56.02 \pm 0.20	0.30\pm0.00	0.55 \pm 0.01	1.68 ²
DPS	20.59 \pm 0.22	7.61 \pm 0.02	0.14 \pm 0.01	0.65 \pm 0.08	132 ²	21.36 \pm 0.06	9.39 \pm 0.02	0.25 \pm 0.00	0.66 \pm 0.03	145 ²

¹CPU, ²GPU Nvidia A100

Table 2: Quantitative comparison with the state-of-the-art methods results from ?? realized on 100 samples for stochastic methods. Blurry results outperforms on SSIM and PSNR metrics while our method provides the best perceptual LPIPS metric. Our method provides also a strong attachment to data, illustrated by the LR-PSNR metric. Our method is also faster, working on CPU.

4.2 Comparison with other methods

Contrary to neural networks trained on a very large dataset of generic images, our method only works on the restricted area of Gaussian microtextures. This limited scope is a main drawback of Gaussian SR that cannot work on more generic images. Moreover, our method necessitates a HR reference image. Still it is interesting to compare our method with existing super-resolution algorithms, although one should keep in mind that these methods are designed to be used on a broader range of images.

We compare our method with other super-resolution methods on two microtextures in ?. Let us describe the concurrent algorithms and the metrics used for comparison.

Super-Resolution using Normalizing Flow (SRFlow¹) [42], Denoising Diffusion Restoration Models (DDRM²) [34] and Diffusion Posterior Sampling (DPS³) [7] are diverse super-resolution routines where several samples are proposed to solve the inverse problem 1. SRFlow is a normalizing flow network which has been trained during 5 days on a single NVIDIA V100 GPU with a general dataset of general images DIV2K [1] to solve the SR problem. It depends on a hyperparameter τ , the temperature, that modulates the variance of the latent space. DPS and DDRM work with a diffusion model that has been trained on the dataset ImageNet [8] for image generation during several days on GPU [10]. A step of attachment to degraded data is added in the generation to solve inverse problems such that SR. DPS proposes a gradient term and DDRM an adaptative reverse process of the score-based model depending on the singular values of the operator A . Similar to our method, Wasserstein Patch Prior (WPP⁴) [25] proposes to solve SR problem in the SISR setting using a reference image. It is a multiscale iterative algorithm that optimizes at several resolutions a optimal transport distance between patches of the proposed image and the ones of the reference image. In this comparative study, only WPP and Gaussian SR use the reference image while SRFlow, DDRM and DPS are trained on a larger dataset of generic images (and are not restricted to texture restoration).

To compare the methods, we use the four metrics Peak Signal to Noise Ratio (PSNR), LR-PSNR, Structural Similarity (SSIM) and Learned Perceptual Image Patch Similarity (LPIPS). The PSNR is the logarithmic scale of the mean-square error (MSE) between the HR image u_{HR} and the solution proposed by the method. LR-PSNR is the PSNR between the LR image and the LR version of the solution and quantifies fidelity to the input data u_{LR} . The SSIM [61] quantifies the similarity between images relying on a luminance, a contrast and a structure terms. LPIPS [64] is a perceptual metric that compute the weighted squared distance between features of the two images of a classification network. The weights have been trained with respect to perceptual criteria.

Quantitatively, as observed in ?, DDRM tends to produce blurred samples and DPS smooth images and consequently do not really recover the texture. SRFlow with a positive temperature gives good textured solutions but with artefacts generally observed in networks' outputs. WPP generates textured but too smooth image. We can also observe that the different luminance of the reference image in 6 influences the WPP result. Gaussian SR gives the desired perceptual grain of the textures. However, it does not retrieve correctly the details, as it can be observed with the white spot in 5.

2 shows that our method is faster and have the best LPIPS on each image, followed by WPP. Note that Gaussian SR works on CPU. Our method proposes also a strong data attachment, illustrated by the high LR-PSNR metric. It is

¹Code and weights from <https://github.com/andreas128/SRFlow>

²Code and weights from <https://github.com/bahjat-kawar/ddrm>

³Code and weights from <https://github.com/DPS2022/diffusion-posterior-sampling>

⁴https://github.com/johrertrich/Wasserstein_Patch_Prior

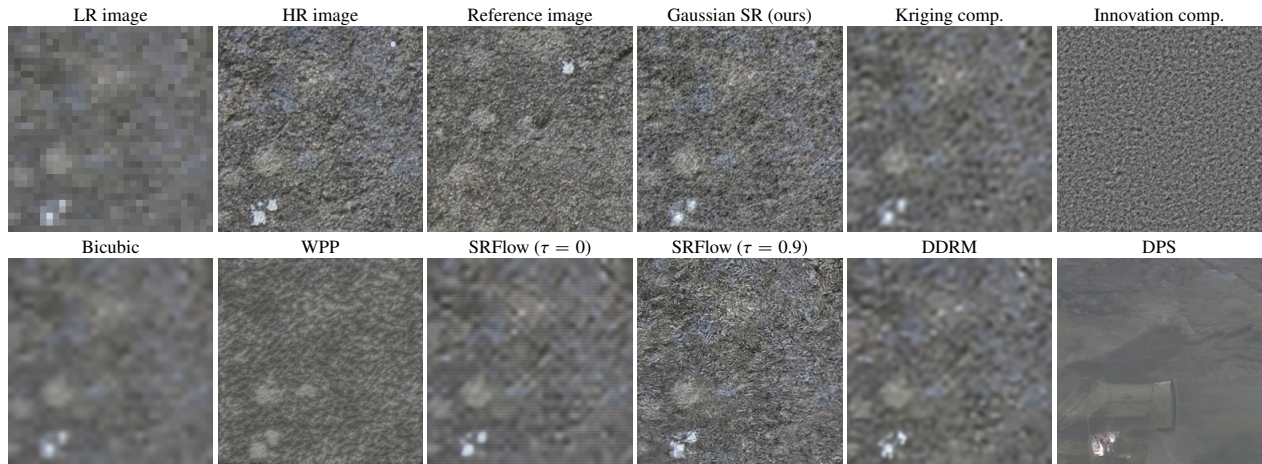


Figure 5: Comparison of our method with the stochastic SRflow, DPS and DDRM and with the SISR WPP that uses the reference image (The HR size is 256×256 and $r = 8$). Note that our method is unable to provide the white spot in the wall that does not respect the stationarity assumption. However it faithfully conveys the granular aspect of the texture given by the reference image.




Figure 6: Comparison of our methods with stochastic SRFlow, DPS and DDRM and with the SISR WPP that uses the reference image (The HR size is 256×256 and $r = 8$). Contrary to WPP, Gaussian SR is not influenced by the difference of luminance between the HR and the reference images.

well-known that PSNR promotes blurry images, a limitation called the “regression to the mean” problem [56] in the SR literature. Our image with the best PSNR comparison is the kriging component. We can explain it theoretically by the next Proposition 3 which shows that in expectation the best PSNR is provided by the mean of our samples, the kriging component.

Proposition 3 (Kriging component and MSE). *Let $\mathbf{u}_{\text{HR}} \in \mathbb{R}^{\Omega_{M,N}}$ be a HR image, $\mathbf{u}_{\text{LR}} = \mathbf{A}\mathbf{u}_{\text{HR}}$ its LR version, $\mathbf{\Lambda} \in \mathbb{R}^{\Omega_{M_r,N_r} \times \Omega_{M,N}}$ be the kriging operator and \mathbf{X}_{SR} the random image following the distribution of the SR samples generated with Equation 3 then*

$$\begin{aligned} \mathbb{E}_{\mathbf{X}_{\text{SR}}} (\|\mathbf{u}_{\text{HR}} - \mathbf{X}_{\text{SR}}\|_2^2) &= \|\mathbf{u}_{\text{HR}} - \mathbf{\Lambda}^T \mathbf{u}_{\text{LR}}\|_2^2 + \text{Tr} [(\mathbf{I}_{\Omega_{M,N}} - \mathbf{\Lambda}^T \mathbf{A}) \mathbf{\Gamma} (\mathbf{I}_{\Omega_{M,N}} - \mathbf{\Lambda}^T \mathbf{A})^T] \\ &\geq \|\mathbf{u}_{\text{HR}} - \mathbf{\Lambda}^T \mathbf{u}_{\text{LR}}\|_2^2. \end{aligned} \quad (12)$$

Simply put, the expected mean square error between the optimal HR image and Gaussian SR samples is always higher than the MSE between \mathbf{u}_{HR} and the associated kriging component $\mathbf{\Lambda}^T \mathbf{u}_{\text{LR}}$.

Proposition 3, proved in D, means that adding the high frequency content of the innovation component is penalized by the PSNR since it is not aligned with the original high frequency content. Yet this addition is obviously perceptually important. We can also observe that SSIM has a similar behaviour as PSNR, as already observed for other degradation problems [27]. We argue that our study of Gaussian super resolution shows that LPIPS is the best metric to study the performance of stochastic SR algorithms.

Surprisingly, the three deep learning-based approaches give relatively poor results when applied to simple textures with high frequency content. This suggests that Imagenet and/or DIV2K datasets should probably be enriched with texture images to adapt the models for texture SR. Given its simplicity, computational efficiency, and qualitative superiority in comparison with state of the art SR methods, we believe our Gaussian SR algorithm is of interest for practitioners interested in stationary texture SR.

5 Limitations and extensions

As shown in the previous section, when applied to Gaussian textures our Gaussian SR algorithm outperforms state-of-the-art methods regarding results quality and execution time. However, Gaussian SR has several inherent limitations that we now highlight.

5.1 Limited scope

Gaussian microtextures represent a limited scope of images. The images are supposed to be stationary with no geometric structures, excluding a lot of textures such as the brick wall in 7. When applying our Gaussian SR algorithm to such a structured texture, one can observe that the independent stationary stochastic innovation component is not aligned with the brick lines recovered in the kriging component. Also, textures with too big individual objects are unadapted to the Gaussian microtexture assumption, such as the bark texture in the second row of 7. The kriging component is attached to the LR data but the innovation component provides a stationary grainy texture that is not in accordance with the piecewise flat nature of the texture. This is a well-known limitation of the ADSN model [19].

5.2 Non-adapted reference image

The choice of the reference image has a strong impact on the final samples. The texture is essentially provided by the innovation component which is built with $\text{ADSN}(\mathbf{u}_{\text{ref}})$. 8 presents two different samples we can obtain depending on the choice of the reference image. The first reference image is the perfect HR image that retrieves the correct alignment of the texture lines. On the contrary, a reference image with opposite alignment modifies the texture of the Gaussian samples.

More generally, with a reference image, the stability result of Proposition 2 does not hold for the simulation of the kriging component. Algorithm 1 can lead to instabilities as illustrated in 9. Here we observe that a frequency has been intensified in the kriging component and provokes artificial parallel lines in the final sample. It is essentially due to a slight misalignment between the HR texture and the reference texture. Note that there is no artefact in the innovation component which fits the assumption of Proposition 2. This type of observations is marginal but can occur in some rare cases.

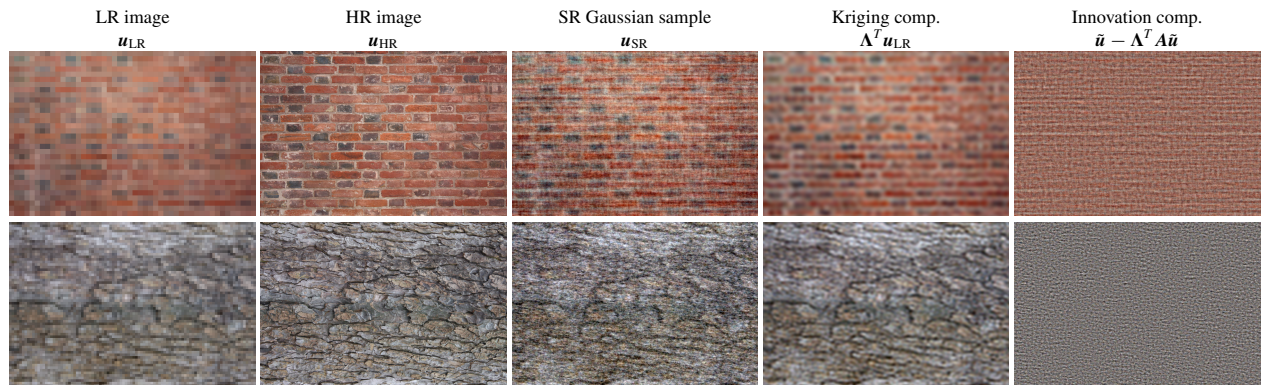


Figure 7: Examples of the application of Gaussian SR on images that are not Gaussian microtextures (The HR size is 512×768 and $r = 16$). The first image has a too structured texture to be treated as a Gaussian texture. The second image cannot be considered as a microtexture due to the presence of too big individual patterns.




Figure 8: Illustration of the importance of the choice of the reference image. Each row presents the samples of Gaussian SR obtained with the same LR image and a different reference image (The HR size is 208×320 and $r = 8$). The covariance information provided by the reference image guides the final texture sample.

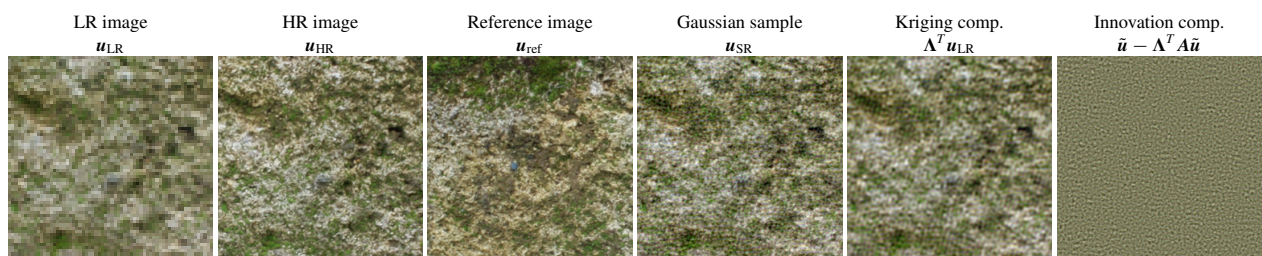


Figure 9: Example of Gaussian SR that illustrates the non-application of Proposition 2 in the setting with reference image (HR size is 256×256 and $r = 4$). A frequency has been amplified and provokes parallel lines in the kriging component generated by Gaussian SR. This artefact occurs in rare cases. This phenomenon is not observable in the innovation component which verifies the assumptions of Proposition 2.




Figure 10: Variance maps of the different algorithms for stochastic super-resolution applied to the HR image of 5, with $r = 8$. The variance of Gaussian SR has been computed theoretically while it has been estimated with 100 samples for the others, and the sum of the pixelwise variance over the three channels is displayed. While DDRM is close to being deterministic, the SRFlow and DPS algorithms have an adaptive variance to the HR image, showing the white spot of the image. In comparison, the Gaussian SR has a cyclostationary non-adaptive variance.

5.3 Non-adaptive variance

We can study the variance of the samples generated by our method with Proposition 4. As a reminder, in our approximation $\Lambda_{\text{approx}}^T$ is in the form $\Gamma A^T C_{\eta}$ where C_{η} is a diagonal multi-channel convolution but this is also the case with the perfect solution of Equation 2 $\Lambda = \Gamma A^T (A \Gamma A^T)^{\dagger}$. This proposition shows that the images sampled with Gaussian SR have a constant variance on sub-grids of the image domain. The law of the SR samples \mathbf{u}_{SR} inherits the stationarity of $\text{ADSN}(\mathbf{u}_{\text{ref}})$ being invariant by the translations by $(kr, \ell r)$ where $k, \ell \in \mathbb{Z}$. In other words, the law of our SR samples is a cyclostationary law [43].

Proposition 4 (Non-adaptive variance of Gaussian SR). *Let $\Lambda^T \in \mathbb{R}^{3\Omega_{M,N} \times 3\Omega_{M/r,N/r}}$ being in the form $C_{\eta} \mathbf{S}^T$ where the kernels $\eta \in \mathbb{R}^{3\Omega_{M,N}}$. Then, the law of the SR samples generated by Equation 3 is invariant by translations by $(kr, \ell r)$ where $k, \ell \in \mathbb{Z}$.*

Proposition 4 implies that the samples generated by Gaussian SR have a pixelwise variance constant by translations $(kr, \ell r)$ where $k, \ell \in \mathbb{Z}$. This variance can be computed theoretically and compared to the empirical pixelwise variance of the other stochastic algorithms, as done in 10. The white spot detail in the textures is taking into account in the variety of the samples proposed by the other stochastic super-resolution algorithms but not by our Gaussian SR. The variance of SRFlow is quite high and allow us to see the input image, it is clearly adapted to the LR image. DPS variance map show the white spot detail but one can guess that the texture is not really retrieved because the variance is smooth on the rest of the image. DDRM is quasi deterministic with a very low variance value. Gaussian SR presents a constant variance on sub-grids of the image. Indeed, one of the main limitation of the model is that the innovation component is independent of the kriging component.

5.4 Extension to other degradation operators

Our method can be applied to reconstruct images degraded by other degradation operators in the form $\mathbf{S}\mathbf{C}$ where \mathbf{C} is a convolution preserving the mean. All the proofs and the properties still hold. To illustrate it, we consider motion blurs randomly generated by the code⁵ with kernel size 61×61 and intensity value 0.5, following [7]. We propose to reconstruct images degraded by a motion blur followed by a subsampling operator \mathbf{S} with stride $r = 4$ and images degraded by a motion blur and the bicubic zoom-out operator \mathbf{A} with factor $r = 4$. These are not standard problems but they can be seen as multiple problems as studied in [54]. In 11, we present the results provided by our method. We can again observe the complementarity of the deterministic kriging component and the stochastic innovation component that retrieves the perceptual texture grain.

⁵<https://github.com/LeviBorodenko/motionblur>




Figure 11: Illustration of the application of Gaussian kriging to other linear operators. The images have square size which are respectively 208, 208, 176 and 256. The two first images are degraded by a motion blur, followed by a subsampling with stride $r = 4$ and the two others are degraded by a motion blur followed by the zoom-out bicubic operator with factor $r = 4$.

6 Conclusion

This work provides an efficient sampling algorithm for solutions of the stochastic SR problem for Gaussian microtextures. Our solution is exact for grayscale images and we propose a fast and reliable approximation for RGB images using a well-known mathematical kriging reasoning. Gaussian SR is a fast and qualitatively better than involved generic methods when applied to Gaussian microtextures accompanied with a reference HR image. Besides, of a more general interest, our comparative study shows that LPIPS is a pertinent metric to evaluate stochastic SR results that do not suffer from excessive blur. While our Gaussian SR has inherent limitations discussed in the paper, it could be of interest for fast and reliable restoration of HR stationary microtextures, even in the presence of motion blur.

References

- [1] Eirikur Agustsson and Radu Timofte. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, July 2017.
- [2] Fabian Altekruiger and Johannes Hertrich. WPPNets and WPPFlows: The power of Wasserstein patch priors for superresolution. *SIAM Journal on Imaging Sciences*, 16(3):1033–1067, 2023.
- [3] Yuval Bahat and Tomer Michaeli. Explorable super resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2716–2725, 2020.
- [4] Joan Bruna, Pablo Sprechmann, and Yann LeCun. Super-resolution with deep convolutional sufficient statistics. In Yoshua Bengio and Yann LeCun, editors, *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016.
- [5] Hong Chang, Dit-Yan Yeung, and Yimin Xiong. Super-resolution through neighbor embedding. In *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004. CVPR 2004.*, volume 1, pages I–I, 2004.
- [6] 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*, pages 1322–1328. IEEE, 2022.
- [7] Hyungjin Chung, Jeongsol Kim, Michael Thompson Mccann, Marc Louis Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems. In *The Eleventh International Conference on Learning Representations*, 2023.
- [8] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [9] A. Desolneux, L. Moisan, and S. Ronsin. A compact representation of random phase and Gaussian textures. In *ICASSP’12*, pages 1381–1384, 2012.
- [10] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 8780–8794. Curran Associates, Inc., 2021.
- [11] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Learning a deep convolutional network for image super-resolution. In *European Conference on Computer Vision*, 2014.
- [12] J.L. Doob. *Stochastic Processes*. Probability and Statistics Series. Wiley, 1953.
- [13] Gilad Freedman and Raanan Fattal. Image and video upscaling from local self-examples. *ACM Trans. Graph.*, 28(3):1–10, 2010.
- [14] W.T. Freeman, T.R. Jones, and E.C. Pasztor. Example-based super-resolution. *IEEE Computer Graphics and Applications*, 22(2):56–65, 2002.

- [15] B. Galerne, A. Lagae, S. Lefebvre, and G. Drettakis. Gabor noise by example. *ACM Trans. Graph.*, 31(4):73:1–73:9, jul 2012.
- [16] B. Galerne, A. Leclaire, and L. Moisan. A texton for fast and flexible Gaussian texture synthesis. In *Proceedings of the 22nd European Signal Processing Conference (EUSIPCO)*, pages 1686–1690, 2014.
- [17] B. Galerne, A. Leclaire, and L. Moisan. Texton noise. *Computer Graphics Forum*, 2017.
- [18] B. Galerne, A. Leclaire, and J. Rabin. A texture synthesis model based on semi-discrete optimal transport in patch space. *SIAM Journal on Imaging Sciences*, 11(4):2456–2493, 2018.
- [19] Bruno Galerne, Yann Gousseau, and Jean-Michel Morel. Random Phase Textures: Theory and Synthesis. *IEEE Transactions on Image Processing*, 20(1):257–267, 2011.
- [20] Bruno Galerne and Arthur Leclaire. An Algorithm for Gaussian Texture Inpainting. *Image Processing On Line*, 7:262–277, 2017. <https://doi.org/10.5201/ipo1.2017.198>.
- [21] Bruno Galerne and Arthur Leclaire. Texture Inpainting Using Efficient Gaussian Conditional Simulation. *SIAM Journal on Imaging Sciences*, 10(3):1446–1474, 2017.
- [22] Bruno Galerne, Arthur Leclaire, and Lionel Moisan. Microtexture inpainting through Gaussian conditional simulation. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1204–1208, 2016.
- [23] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative Adversarial Nets. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014.
- [24] J. Gutierrez, B. Galerne, J. Rabin, and T. Hurtut. Optimal patch assignment for statistically constrained texture synthesis. In F. Lauze, Y. Dong, and A. B. Dahl, editors, *Scale Space and Variational Methods in Computer Vision: 6th International Conference, SSVM 2017, Kolding, Denmark, June 4-8, 2017, Proceedings*, pages 172–183, Cham, 2017. Springer International Publishing.
- [25] Johannes Hertrich, Antoine Houdard, and Claudia Redenbach. Wasserstein Patch Prior for Image Superresolution. *IEEE Transactions on Computational Imaging*, 8:693–704, 2022.
- [26] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- [27] Alain Horé and Djemel Ziou. Image quality metrics: Psnr vs. ssim. In *2010 20th International Conference on Pattern Recognition*, pages 2366–2369, 2010.
- [28] Antoine Houdard, Arthur Leclaire, Nicolas Papadakis, and Julien Rabin. Wasserstein generative models for patch-based texture synthesis. In Abderrahim Elmoataz, Jalal Fadili, Yvain Quéau, Julien Rabin, and Loïc Simon, editors, *Scale Space and Variational Methods in Computer Vision*, pages 269–280, Cham, 2021. Springer International Publishing.
- [29] Antoine Houdard, Arthur Leclaire, Nicolas Papadakis, and Julien Rabin. A generative model for texture synthesis based on optimal transport between feature distributions. *J. Math. Imaging Vis.*, 65(1):4–28, 2023.
- [30] H. Jiang, Y. Shen, J. Xie, J. Li, J. Qian, and J. Yang. Sampling network guided cross-entropy method for unsupervised point cloud registration. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 6108–6117, Los Alamitos, CA, USA, oct 2021. IEEE Computer Society.
- [31] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, editors, *Computer Vision – ECCV 2016*, pages 694–711, Cham, 2016. Springer International Publishing.

- [32] Ulugbek S. Kamilov, Charles A. Bouman, Gregory T. Buzzard, and Brendt Wohlberg. Plug-and-play methods for integrating physical and learned models in computational imaging: Theory, algorithms, and applications. *IEEE Signal Processing Magazine*, 40(1):85–97, 2023.
- [33] W. J. Kammerer and M. Z. Nashed. On the convergence of the conjugate gradient method for singular linear operator equations. *SIAM Journal on Numerical Analysis*, 9(1):165–181, 1972.
- [34] Bahjat Kawar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising diffusion restoration models. In *ICLR Workshop on Deep Generative Models for Highly Structured Data*, 2022.
- [35] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Accurate image super-resolution using very deep convolutional networks. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1646–1654, 2016.
- [36] Durk P Kingma and Prafulla Dhariwal. Glow: Generative Flow with Invertible 1x1 Convolutions. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.
- [37] Wei-Sheng Lai, Jia-Bin Huang, Narendra Ahuja, and Ming-Hsuan Yang. Deep laplacian pyramid networks for fast and accurate super-resolution. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5835–5843, 2017.
- [38] C. Lantuejoul. *Geostatistical Simulation: Models and Algorithms*. Collezione legale Pirola. Springer Berlin Heidelberg, 2001.
- [39] 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)*, pages 105–114, 2017.
- [40] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual networks for single image super-resolution. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, July 2017.
- [41] Z. Liu, W. Siu, and L. Wang. Variational autoencoder for reference based image super-resolution. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 516–525, Los Alamitos, CA, USA, jun 2021. IEEE Computer Society.
- [42] 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*, pages 715–732, Berlin, Heidelberg, 2020. Springer-Verlag. event-place: Glasgow, United Kingdom.
- [43] N. Lutz, B. Sauvage, and J-M. Dischler. Cyclostationary gaussian noise: theory and synthesis. *Computer Graphics Forum*, 40:xx–yy, 2021.
- [44] Simon Mignon, Bruno Galerne, Moncef Hidane, Cécile Louchet, and Julien Mille. Semi-unbalanced regularized optimal transport for image restoration. In *31st European Signal Processing Conference, EUSIPCO 2023, Helsinki, Finland, September 4-8, 2023*, pages 466–470. IEEE, 2023.
- [45] Simon Mignon, Bruno Galerne, Moncef Hidane, Cécile Louchet, and Julien Mille. Semi-unbalanced optimal transport for image restoration and synthesis. working paper or preprint, 2024.
- [46] Lionel Moisan. Periodic Plus Smooth Image Decomposition. *Journal of Mathematical Imaging and Vision*, 39(2):161–179, 2011.
- [47] Gregory Ongie, Ajil Jalal, Christopher A. Metzler, Richard G. Baraniuk, Alexandros G. Dimakis, and Rebecca Willett. Deep learning techniques for inverse problems in imaging. *IEEE Journal on Selected Areas in Information Theory*, 1(1):39–56, 2020.

- [48] R. Penrose. A generalized inverse for matrices. *Mathematical Proceedings of the Cambridge Philosophical Society*, 51(3):406–413, 1955.
- [49] Émile Pierret and Bruno Galerne. Stochastic super-resolution for gaussian textures. In *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5, 2023.
- [50] Litu Rout, Negin Raoof, Giannis Daras, Constantine Caramanis, Alex Dimakis, and Sanjay Shakkottai. Solving linear inverse problems provably via posterior sampling with latent diffusion models. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [51] 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*, pages 1–14, 2022.
- [52] Mehdi S. M. Sajjadi, Bernhard Schölkopf, and Michael Hirsch. EnhanceNet: Single Image Super-Resolution through Automated Texture Synthesis. In *Computer Vision (ICCV), 2017 IEEE International Conference on*, pages 4501–4510. IEEE, 2017.
- [53] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- [54] Jiaming Song, Arash Vahdat, Morteza Mardani, and Jan Kautz. Pseudoinverse-guided diffusion models for inverse problems. In *International Conference on Learning Representations (ICLR)*, May 2023.
- [55] Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-Based Generative Modeling through Stochastic Differential Equations. 2023.
- [56] 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.
- [57] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Deep image prior. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [58] Dmitry Ulyanov, Andrea Vedaldi, and Victor S. Lempitsky. Deep image prior. *Int. J. Comput. Vis.*, 128(7):1867–1888, 2020.
- [59] Xintao Wang, Ke Yu, Chao Dong, and Chen Change Loy. Recovering realistic texture in image super-resolution by deep spatial feature transform. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [60] 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 Laura Leal-Taixé and Stefan Roth, editors, *Computer Vision – ECCV 2018 Workshops*, pages 63–79, Cham, 2019. Springer International Publishing.
- [61] 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*, 13(4):600–612, 2004.
- [62] G.-S. Xia, S. Ferradans, G. Peyré, and J.-F. Aujol. Synthesizing and mixing stationary Gaussian texture models. *SIAM J. on Imaging Science*, 8(1):476–508, 2014.
- [63] G. Yu, G. Sapiro, and S. Mallat. Solving inverse problems with piecewise linear estimators: From Gaussian mixture models to structured sparsity. *Image Processing, IEEE Transactions on*, 21(5):2481–2499, May 2012.
- [64] 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)*, pages 586–595, Los Alamitos, CA, USA, June 2018. IEEE Computer Society.

- [65] Haitian Zheng, Mengqi Ji, Haoqian Wang, Yebin Liu, and Lu Fang. Crossnet: An end-to-end reference-based super resolution network using cross-scale warping. In Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss, editors, *Computer Vision – ECCV 2018*, pages 87–104, Cham, 2018. Springer International Publishing.
- [66] Daniel Zoran and Yair Weiss. From learning models of natural image patches to whole image restoration. In *2011 International Conference on Computer Vision*, pages 479–486, 2011.

A Proof of Theorem 2

The fact that the conditional expectation is linear for zero-mean Gaussian multivariate law is a classical result (see e.g. [12]). Let us denote $(\lambda(x))_{x \in \Omega_{M,N}}$ the columns of $\mathbf{\Lambda}$. $\mathbf{\Lambda}$ minimizes

$$\begin{aligned} \mathbb{E} \left(\|\mathbf{\Lambda}^T \mathbf{A} \mathbf{X} - \mathbf{X}\|_2^2 \right) &= \mathbb{E} \left[\sum_{x \in \Omega_{M,N}} ((\mathbf{\Lambda}^T \mathbf{A} \mathbf{X})(x) - \mathbf{X}(x))^2 \right] \\ &= \sum_{x \in \Omega_{M,N}} \mathbb{E} \left[(\lambda(x)^T \mathbf{A} \mathbf{X} - \mathbf{X}(x))^2 \right]. \end{aligned}$$

This is a separable function of the columns $(\lambda(x))_{x \in \Omega_{M,N}}$. For each for $x \in \Omega_{M,N}$, let us minimize $\lambda(x) \in \mathbb{R}^{\Omega_{M/r, N/r}} \mapsto \mathbb{E} \left[(\lambda(x)^T \mathbf{A} \mathbf{X} - \mathbf{X}(x))^2 \right]$. One has

$$\begin{aligned} \mathbb{E} \left[(\lambda(x)^T \mathbf{A} \mathbf{X} - \mathbf{X}(x))^2 \right] &= \text{Var}(\lambda(x)^T \mathbf{A} \mathbf{X}) + \text{Var}(\mathbf{X}(x)) - 2 \text{Cov}(\lambda(x)^T \mathbf{A} \mathbf{X}, \mathbf{X}(x)) \\ &= \lambda(x)^T \mathbf{A} \mathbf{\Gamma} \mathbf{A}^T \lambda(x) + \text{Var}(\mathbf{X}(x)) - 2 \lambda(x)^T \mathbf{A} \text{Cov}(\mathbf{X}, \mathbf{X}(x)). \end{aligned}$$

This is a quadratic functional associated with the positive matrix $2\mathbf{A} \mathbf{\Gamma} \mathbf{A}^T$ and the vector $2\mathbf{A} \text{Cov}(\mathbf{X}, \mathbf{X}(x)) = 2\mathbf{A} \mathbf{\Gamma}_{\Omega_{M,N} \times \{x\}}$. It is minimal for $\lambda(x)$ any solution of the linear system $\mathbf{A} \mathbf{\Gamma} \mathbf{A}^T \lambda(x) = \mathbf{A} \mathbf{\Gamma}_{\Omega_{M,N} \times \{x\}}$. This is valid for any column $\lambda(x)$ of $\mathbf{\Lambda}$. As a consequence, $\mathbf{\Lambda}$ is a solution if and only if it verifies

$$\mathbf{A} \mathbf{\Gamma} \mathbf{A}^T \mathbf{\Lambda} = \mathbf{A} \mathbf{\Gamma}.$$

B Proof of Lemma 1

It is sufficient to identify the convolutions on the canonical basis. For $x \in \Omega_{M,N}$, we define $\delta_x \in \mathbb{R}^{\Omega_{M,N}}$ such that for $z \in \Omega_{M,N}$, $\delta_x(z) = \mathbb{1}_{x=z}$ and identically δ_x^r for $x \in \Omega_{M/r, N/r}$. Denoting T_x the translation by $x \in \Omega_{M,N}$, that is for all $z \in \mathbb{R}^{\Omega_{M,N}}$, $T_x \mathbf{u}(z) = \mathbf{u}(z - x)$. Remark that convolving with $\delta_x^{\Omega_{M,N}}$ corresponds to translating by x : for all $\mathbf{u} \in \mathbb{R}^{\Omega_{M,N}}$, $\mathbf{u} \star \delta_x = T_x \mathbf{u}$.

1. For $x \in \Omega_{M/r, N/r}$,

$$\mathbf{S} \mathbf{C}_\alpha \mathbf{S}^T \delta_x^r = \mathbf{S} \mathbf{C}_\alpha \delta_{rx} = \mathbf{S}(\alpha \star \delta_{rx}) = \mathbf{S} T_{rx} \alpha = T_x \mathbf{S} \alpha = (\mathbf{S} \alpha) \star \delta_x^r. \quad (13)$$

2. For $x \in \Omega_{M/r, N/r}$,

$$\mathbf{S}^T \mathbf{C}_\beta \delta_x^r = \mathbf{S}^T (\beta \star \delta_x^r) = \mathbf{S}^T T_x \beta = T_{rx} \mathbf{S}^T \beta = (\mathbf{S}^T \beta) \star \delta_{rx} = (\mathbf{S}^T \beta) \star (\mathbf{S}^T \delta_x^r). \quad (14)$$

Finally, the equality of the convolutions on the canonical basis is established.

C Conjugate Gradient Descent algorithm

The CGD algorithm is an iterative method to approximate the solution of a linear equation by solving the associated least-squares problem [33]. Algorithm 2 aims at solving the normal equation

$$\mathbf{B}^T \mathbf{B} \psi = \mathbf{B}^T \varphi \quad (15)$$

associated with the least-squares problem $\psi \mapsto \|\mathbf{B}\psi - \varphi\|_2$ and has good convergence properties when applied to the inpainting problem for Gaussian textures [21]. First we observed that this also applies in the case of super-resolution. In our SR context, our goal is to express $(\mathbf{A}\mathbf{\Gamma}\mathbf{A}^T)^\dagger \mathbf{v}$ for a given LR image $\mathbf{v} \in \mathbb{R}^{\Omega_{M,r,N/r}}$. It is equivalent to minimize $\psi \in \mathbb{R}^{\Omega_{M,r,N/r}} \mapsto \|(\mathbf{A}\mathbf{\Gamma}\mathbf{A}^T)^T \psi - \mathbf{v}\|_2$. Therefore, we apply CGD with $\mathbf{B} = \mathbf{A}\mathbf{\Gamma}\mathbf{A}^T$ and $\varphi = \mathbf{v}$.

Algorithm 2 CGD Algorithm CGD to compute $\mathbf{B}^\dagger \varphi$

Input: Initialize $k \leftarrow 0, \psi_0 \leftarrow 0, r_0 \leftarrow \mathbf{B}^T \varphi - \mathbf{B}^T \mathbf{B} \psi_0, d_0 \leftarrow r_0$

while $\|r_k\|_2 > \varepsilon$ **do**

$$\alpha_k \leftarrow \frac{\|r_k\|_2^2}{d_k^T \mathbf{B}^T \mathbf{B} d_k}$$

$$\psi_{k+1} \leftarrow \psi_k + \alpha_k d_k$$

$$r_{k+1} \leftarrow r_k - \alpha_k \mathbf{B}^T \mathbf{B} d_k$$

$$d_{k+1} \leftarrow r_{k+1} + \frac{\|r_{k+1}\|_2^2}{\|r_k\|_2^2} d_k$$

$$k \leftarrow k + 1$$

end while

Output: ψ_k

D Proof of Proposition 3

Let $\mathbf{u}_{\text{HR}} \in \mathbb{R}^{\Omega_{M,N}}$ be a HR image, $\mathbf{u}_{\text{LR}} = \mathbf{A}\mathbf{u}_{\text{HR}}$ its LR version, \mathbf{u}_{ref} a reference image, $\mathbf{\Gamma}$ such that $\text{ADSN}(\mathbf{u}_{\text{ref}}) = \mathcal{N}(\mathbf{0}, \mathbf{\Gamma})$, $\mathbf{\Lambda} \in \mathbb{R}^{\Omega_{M,r,N/r} \times \Omega_{M,N}}$ be the kriging operator and \mathbf{X}_{SR} the random image following the distribution of the SR samples generated with Equation 3. \mathbf{X}_{SR} has the same law as:

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

with $\tilde{\mathbf{X}} \sim \mathcal{N}(\mathbf{0}, \mathbf{\Gamma})$. Consequently, considering $\mathbb{E}_{\tilde{\mathbf{X}}}(\tilde{\mathbf{X}}) = \mathbf{0}$,

$$\begin{aligned} \mathbb{E}_{\mathbf{X}_{\text{SR}}}(\|\mathbf{u}_{\text{HR}} - \mathbf{X}_{\text{SR}}\|_2^2) &= \|\mathbf{u}_{\text{HR}} - \mathbf{\Lambda}^T \mathbf{u}_{\text{LR}}\|_2^2 + \mathbb{E}_{\tilde{\mathbf{X}}}(\|\tilde{\mathbf{X}} - \mathbf{\Lambda}^T \mathbf{A} \tilde{\mathbf{X}}\|_2^2) \\ &\quad + 2\langle \mathbf{u}_{\text{HR}} - \mathbf{\Lambda}^T \mathbf{u}_{\text{LR}}, \mathbb{E}_{\tilde{\mathbf{X}}}(\tilde{\mathbf{X}}) - \mathbf{\Lambda}^T \mathbf{A} \mathbb{E}_{\tilde{\mathbf{X}}}(\tilde{\mathbf{X}}) \rangle \\ &= \|\mathbf{u}_{\text{HR}} - \mathbf{\Lambda}^T \mathbf{u}_{\text{LR}}\|_2^2 + \text{Tr}[(\mathbf{I}_{\Omega_{M,N}} - \mathbf{\Lambda}^T \mathbf{A}) \mathbf{\Gamma} (\mathbf{I}_{\Omega_{M,N}} - \mathbf{\Lambda}^T \mathbf{A})^T] \end{aligned}$$

And, $(\mathbf{I}_{\Omega_{M,N}} - \mathbf{\Lambda}^T \mathbf{A}) \mathbf{\Gamma} (\mathbf{I}_{\Omega_{M,N}} - \mathbf{\Lambda}^T \mathbf{A})^T$ is a positive semi-definite matrix with non-negative trace which gives

$$\mathbb{E}_{\mathbf{X}_{\text{SR}}}(\|\mathbf{u}_{\text{HR}} - \mathbf{X}_{\text{SR}}\|_2^2) \geq \|\mathbf{u}_{\text{HR}} - \mathbf{\Lambda}^T \mathbf{u}_{\text{LR}}\|_2^2. \quad (16)$$

E Proof of Proposition 4

Let $\mathbf{\Lambda}^T \in \mathbb{R}^{3\Omega_{M,N} \times 3\Omega_{M,r,N/r}}$ being in the form $\mathbf{C}_\eta \mathbf{S}^T$ where the kernel $\eta \in \mathbb{R}^{3\Omega_{M,N}}$. The law of $\tilde{\mathbf{X}} - \mathbf{\Lambda}^T (\mathbf{A} \tilde{\mathbf{X}})$ is invariant by translation by $(rk, r\ell)$ for $k, \ell \in \mathbb{Z}$ since $\tilde{\mathbf{X}}$ follows a stationary law, $\mathbf{\Lambda}^T \mathbf{A} \tilde{\mathbf{X}} = \mathbf{C}_\eta \star \mathbf{S}^T (\mathbf{A} \tilde{\mathbf{X}})$ with $\mathbf{A} \tilde{\mathbf{X}} \in \mathbb{R}^{3\Omega_{M,r,N/r}}$ following a stationary law.