Sampling high-dimensional Gaussian distributions for general linear inverse problems

François F. Orieux, Olivier Féron, Jean-François Giovannelli

► To cite this version:

François F. Orieux, Olivier Féron, Jean-François Giovannelli. Sampling high-dimensional Gaussian distributions for general linear inverse problems. IEEE Signal Processing Letters, 2012, 19 (5), pp.251. 10.1109/LSP.2012.2189104 . hal-00779449

HAL Id: hal-00779449

<https://hal.science/hal-00779449>

Submitted on 22 Jan 2013

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.

Sampling high-dimensional Gaussian distributions for general linear inverse problems

F. Orieux*, O. Féron and J.-F. Giovannelli

Abstract—This paper is devoted to the problem of sampling Gaussian distributions in high dimension. Solutions exist for two specific structures of inverse covariance: sparse and circulant. The proposed algorithm is valid in a more general case especially as it emerges in linear inverse problems as well as in some hierarchical or latent Gaussian models. It relies on a perturbation-optimization principle: adequate stochastic perturbation of a criterion and optimization of the perturbed criterion. It is proved that the criterion optimizer is a sample of the target distribution. The main motivation is in inverse problems related to general (non-convolutive) linear observation models and their solution in a Bayesian framework implemented through sampling algorithms when existing samplers are infeasible. It finds a direct application in myopic/unsupervised inversion methods as well as in some non-Gaussian inversion methods. An illustration focused on hyperparameter estimation for super-resolution method shows the interest and the feasibility of the proposed algorithm.

Index Terms—Stochastic sampling, high-dimensional sampling, inverse problem, Bayesian strategy, unsupervised, myopic

I. INTRODUCTION

This work deals with simulation of high-dimensional Gaussian and conditional Gaussian distributions. The difficulty of the problem is directly related to handling high-dimensional covariances $\mathbf{R}$ and precision matrices $\mathbf{Q} = \mathbf{R}^{-1}$. The problem has already been investigated and solutions exist in two cases.

- When $\mathbf{Q}$ is sparse, two strategies are available. The first one [1, chap. 8], relies on a parallel Gibbs sampler based on a chessboard-like decomposition. It takes advantage of the sparsity of $\mathbf{Q}$ to update simultaneously large blocks of variables. The second strategy [2, 3] relies on a Cholesky decomposition $\mathbf{Q} = \mathbf{L}^t \mathbf{L}$: a sample $\mathbf{x}$ is obtained by solving the linear system $\mathbf{L}\mathbf{x} = \boldsymbol{\varepsilon}$, where $\boldsymbol{\varepsilon}$ is a zero-mean white Gaussian vector. The sparsity of $\mathbf{Q}$ ensures feasible numerical factorization and the sparsity of $\mathbf{L}$ ensures feasible numerical resolution of the linear system.
- [4, 5] propose a solution for circulant matrix $\mathbf{Q}$, even non-sparse. In this case, the covariance is diagonal in the Fourier domain: the sampling is based on independent sampling of the Fourier coefficients. Finally, the sample is computed by FFT and it has been used in [6–10].

To our knowledge there is no solution for more general structure in high dimension because factorization (Cholesky, QR, square root,...), diagonalization and inversion of $\mathbf{Q}$ and $\mathbf{R}$ are numerically infeasible. The obstacle is due to both computational cost and memory footprint. The proposed

algorithm overcomes this obstacle when $\mathbf{Q}$ is of the form

$$\mathbf{Q} = \sum_{k=1}^K \mathbf{M}_k^t \mathbf{R}_k^{-1} \mathbf{M}_k \quad (1)$$

as it appears in inverse problems [11]. Indeed, let us consider the general linear forward model $\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{n}$, where $\mathbf{y}$, $\mathbf{n}$ and $\mathbf{x}$ are the observation, the noise and the unknown image and $\mathbf{A}$ is a linear operator. Consider, again, two prior distributions for $\mathbf{n}$ and $\mathbf{x}$ that are Gaussian conditionally on a parameter $\boldsymbol{\theta}$. This framework is very general: it includes linear inverse problems [11] as well as some hierarchical or latent Gaussian models [12] and it can be used in many applications. In image reconstruction, it covers a majority of current problems, e.g. unsupervised [8] or myopic (semi-blind) [9] inverse problems, by including acquisition parameters and hyperparameters in $\boldsymbol{\theta}$. Moreover, the framework also includes non-linear models, based on conditional linear models such as bilinear or multilinear ones (see Section III-B). The framework also covers some non-stationary or inhomogeneous Gaussian priors and non-Gaussian priors involving auxiliary/latent variables [6, 8, 13–15] (e.g., location or scale mixtures of Gaussian), by including these variables in $\boldsymbol{\theta}$.

Let us focus on the joint estimation of $\mathbf{x}$ and $\boldsymbol{\theta}$ from the posterior $p(\mathbf{x}, \boldsymbol{\theta} | \mathbf{y})$. It commonly requires the handling of the conditional posterior $p(\mathbf{x} | \boldsymbol{\theta}, \mathbf{y})$ that is Gaussian with precision matrix $\mathbf{Q}$ of the form (1), as will be shown in section II-B. In the general case, $\mathbf{Q}$ is neither sparse nor circulant so existing sampling algorithms fail when the dimension of $\mathbf{x}$ is very large while the proposed one handles this case. It relies on a perturbation-optimization principle: adequate stochastic perturbation of a quadratic criterion and optimization of the perturbed criterion. A recent paper [16] briefly describes a similar algorithm for compressed sensing in signal processing. Our paper deepens and generalizes this contribution.

Subsequently, Section II presents the proposed algorithm and its direct application to linear inverse problems. Section III gives an illustration through an academic problem in super-resolution. Section IV presents conclusions and perspectives.

II. PERTURBATION-OPTIMIZATION ALGORITHM

A. Description

We focus on the problem of sampling from a target Gaussian distribution whose precision matrix $\mathbf{Q}$ is in the form (1). When $\mathbf{Q}$ is neither sparse nor circulant, existing algorithms fail in high dimension because of an excessive memory footprint as illustrated in section III. We propose a solution based on the Perturbation-Optimization (PO) algorithm described hereafter, whose memory footprint is far smaller.

F. Orieux is with Pasteur Institute, 75015 Paris, France, orieux@pasteur.fr. O. Féron is with EDF Research & Developments, OSIRIS, 92140 Clamart, France, olivier-2.feron@edf.fr. J.-F. Giovannelli is with Univ. Bordeaux, IMS, 33400 Talence, France, Giova@IMS-Bordeaux.fr.

Proposition 1: The optimizer $\hat{x}$ of criterion (5) resulting from Algorithm 1 is Gaussian

$$\hat{x} \sim \mathcal{N}\left(Q^{-1}\left(\sum_{k=1}^K M_k^t R_k^{-1} m_k\right), Q^{-1}\right). \quad (2)$$

Proof: The optimizer $\hat{x}$ of criterion (5) is explicit:

$$\begin{aligned} \hat{x} &= \left[\sum_{k=1}^K M_k^t R_k^{-1} M_k\right]^{-1} \left(\sum_{k=1}^K M_k^t R_k^{-1} \eta_k\right) \\ &= Q^{-1} \left(\sum_{k=1}^K M_k^t R_k^{-1} \eta_k\right). \end{aligned} \quad (3)$$

It is clearly a Gaussian vector as a linear combination of K Gaussian vectors. Its expectation and covariance are calculated below using elementary algebra: from (4) and (3), we have

$$\begin{aligned} \mathbb{E}[\hat{x}] &= Q^{-1} \left(\sum_k M_k^t R_k^{-1} \mathbb{E}[\eta_k]\right) \\ &= Q^{-1} \left(\sum_k M_k^t R_k^{-1} m_k\right) \end{aligned}$$

$$\begin{aligned} \mathbb{E}[\hat{x}\hat{x}^t] &= Q^{-1} \left(\sum_{k,k'} M_k^t R_k^{-1} \mathbb{E}[\eta_k \eta_{k'}^t] R_{k'}^{-1} M_{k'}\right) Q^{-1} \\ &= Q^{-1} \left(\sum_k M_k^t R_k^{-1} (R_k + m_k m_k^t) R_k^{-1} M_k\right) Q^{-1} \\ &= Q^{-1} + \mathbb{E}[\hat{x}] \mathbb{E}[\hat{x}]^t \end{aligned}$$

that completes the proof. ■

The feasibility of Step P clearly depends on the capability to sample from Gaussian distributions $\mathcal{N}(m_k, R_k)$. It is usually the case in inverse problems and it will be actually the case in super-resolution applications shown in section III-A and in other contributions shortly described in section III-B.

Regarding Step O, J being quadratic, a large literature [17] is available about its numerical optimization, e.g. gradient procedure (standard, corrected, conjugate, optimal step size...). Such algorithms require the computation of criterion (5) and its gradient. The feasibility of Step O clearly depends on the capability to compute that without the storage of large matrices. It is usually the case in inverse problems and it will be actually the case in applications shown in section III-A and described in section III-B.

However, the desired sample is the exact optimizer, so, Step O could require N iterations of a conjugate gradient algorithm for a problem of dimension N . Therefore the complexity could be $O(N^3)$ that is equivalent to the one of a Cholesky decomposition. However, the optimization procedure can be stopped earlier without practical loss of precision and the complexity falls down to $O(PN^2)$ for P iterations. In addition, for a band matrix, the complexity of the proposed algorithm becomes $O(MPN)$ and the one of the Cholesky decomposition becomes $O(MN^2)$. Anyway, the main advantage of the proposed algorithm is its reduced memory footprint: it avoids the storage of neither Q nor its (Cholesky, QR, square root,...) factors.

Algorithm 1 : Perturbation-Optimization algorithm.

1: **Step P (Perturbation):** Generate independent vectors

$$\eta_k \sim \mathcal{N}(m_k, R_k), \quad \text{for } k = 1, \dots, K \quad (4)$$

2: **Step O (Optimization):** Compute $\hat{x}$ as the minimizer of

$$J(x) = \sum_{k=1}^K (\eta_k - M_k x)^t R_k^{-1} (\eta_k - M_k x) \quad (5)$$

Remark 1: Still regarding Step O, it would be awkward if Q was badly scaled, but it is not the case here for the following reason. In usual ill-conditioned inverse problems, A is badly-scaled but the aim of regularization is precisely to overcome this difficulty and to produce a well-scaled matrix Q .

B. Application to inverse problems

The purpose is to solve an inverse problem, stated by the forward model $y = Ax + n$, in a Bayesian framework based on the following models:

- A describes any observation system that can depend on unknown acquisition parameters,
- priors for the noise n and the object x are Gaussian $\mathcal{N}(m_n, R_n)$ and $\mathcal{N}(m_x, R_x)$, conditionally on a set of hyperparameters and auxiliary variables.

In a general statement, acquisition parameters, hyperparameters and auxiliary variables are collected in θ . The general inverse problem then consists in estimating x and θ through the posterior $p(x, \theta|y)$. Its exploration can be achieved by means of a Gibbs sampler which iteratively samples from $p(\theta|x, y)$ and $p(x|\theta, y)$. The conditional posterior $p(x|y, \theta)$ is a correlated Gaussian distribution: $\mathcal{N}(m_x^{\text{post}}, R_x^{\text{post}})$ with

$$\begin{aligned} R_x^{\text{post}} &= (A^t R_n^{-1} A + R_x^{-1})^{-1} \\ m_x^{\text{post}} &= R_x^{\text{post}} (A^t R_n^{-1} [y - m_n] + R_x^{-1} m_x) \end{aligned}$$

where θ is embedded in A, R_n and R_x for simpler notations.

If A has no particular properties, $Q = (R_x^{\text{post}})^{-1}$ is neither sparse nor circulant, and existing sampling algorithms are not applicable. The PO algorithm makes it possible to sample from $\mathcal{N}(m_x^{\text{post}}, R_x^{\text{post}})$ by applying Algorithm 1 with $K = 2$, $M_1 = A$, $M_2 = I$, $R_1 = R_n$, $R_2 = R_x$, $m_1 = m_n$ and $m_2 = m_x$. In this context, it can be said that the optimization procedure converts prior samples into a posterior one.

III. ILLUSTRATION

The proposed PO algorithm makes it possible to resort to stochastic sampling algorithms in inverse problems providing two main advances:

- capability to jointly estimate extra unknowns included in θ (acquisition parameters, hyperparameters, ...),
- access to the entire unknown distribution providing uncertainties (standard deviation, credibility interval,...).

These advances are illustrated in the present section.

A. Unsupervised super-resolution

We detail an application of the proposed PO algorithm to the super-resolution (SR) academic problem: several blurred and down-sampled (low resolution) images of a scene are available in order to retrieve the original (high resolution) scene [18, 19]. It is shown that the crucial novelty, enabled by the proposed PO algorithm, is to allow the use of sampling algorithms in SR methods and thus to provide hyperparameter estimation.

We resort to the standard forward model in SR: $\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{n} = \mathbf{P}\mathbf{H}\mathbf{x} + \mathbf{n}$. In this equation, $\mathbf{y} \in \mathbb{R}^M$ collects the pixels of the low resolution images, here 5 images of 128×128 pixels ($M = 81920$) and $\mathbf{x} \in \mathbb{R}^N$ collects the pixels of the original image, here 256×256 pixels ($N = 65536$). The noise $\mathbf{n} \in \mathbb{R}^M$ accounts for measurement and modeling errors. $\mathbf{H}$ is a $N \times N$ circulant convolution matrix that accounts for the convolution part of the observation system. Practically, the impulse response is a Laplace shape with FWHM of 4 pixels. $\mathbf{P}$ is a $M \times N$ decimation matrix: it is a binary matrix indicating which pixel is observed. Finally, $\mathbf{A}$ is a $M \times N$ matrix (that is to say 81920×65536). The prior distribution for $\mathbf{n}$ is $\mathcal{N}(\mathbf{0}, \gamma_n^{-1}\mathbf{I})$ and the one for $\mathbf{x}$ is $\mathcal{N}(\mathbf{0}, \gamma_x^{-1}\mathbf{D}^t\mathbf{D})$ where $\mathbf{D}$ is the $N \times N$ circulant convolution matrix of the Laplacian filter. The hyperparameters γ_n and γ_x are *unknown* and their prior law are Jeffreys'. The posterior [9] is

$$p(\mathbf{x}, \gamma_n, \gamma_x | \mathbf{y}) \propto \gamma_n^{M/2-1} \gamma_x^{(N-1)/2-1} \exp \left[-\gamma_n \|\mathbf{y} - \mathbf{P}\mathbf{H}\mathbf{x}\|^2 / 2 - \gamma_x \|\mathbf{D}\mathbf{x}\|^2 / 2 \right]. \quad (6)$$

It is explored by a Gibbs sampler: iteratively sampling γ_n , γ_x and $\mathbf{x}$ under their respective posterior conditional distribution

$$p(\gamma_n^{(k)} | \mathbf{x}, \gamma_x, \mathbf{y}) = \mathcal{G} \left(1 + M/2, 2 / \left\| \mathbf{y} - \mathbf{P}\mathbf{H}\mathbf{x}^{(k-1)} \right\|^2 \right)$$

$$p(\gamma_x^{(k)} | \mathbf{x}, \gamma_n, \mathbf{y}) = \mathcal{G} \left(1 + (N-1)/2, 2 / \left\| \mathbf{D}\mathbf{x}^{(k-1)} \right\|^2 \right)$$

$$p(\mathbf{x}^{(k)} | \gamma_x, \gamma_n, \mathbf{y}) = \mathcal{N}(\mathbf{m}_x^{\text{post}}, \mathbf{R}_x^{\text{post}})$$

$$\text{with } \mathbf{R}_x^{\text{post}} = \left(\gamma_n^{(k)} \mathbf{H}^t \mathbf{P}^t \mathbf{P} \mathbf{H} + \gamma_x^{(k)} \mathbf{D}^t \mathbf{D} \right)^{-1}$$

$$\text{and } \mathbf{m}_x^{\text{post}} = \gamma_n^{(k)} \mathbf{R}_x^{\text{post}} \mathbf{P}^t \mathbf{H}^t \mathbf{y}.$$

The conditional posteriors for the hyperparameters are Gamma distributions so they are easy to sample.

The conditional posterior for $\mathbf{x}$ is Gaussian, but the use of existing algorithms is impossible due to the structure and the size of $\mathbf{R}_x^{\text{post}}$. Regarding the structure, according to Section II-B, with $\mathbf{A} = \mathbf{P}\mathbf{H}$: $\mathbf{A}$ is non-circulant due to the decimation and $\mathbf{A}$ is non-sparse due to large support of the impulse response. Regarding the size, $\mathbf{R}_x^{\text{post}}$ (and its Cholesky factor) is a huge $N \times N$ matrix, that is to say 65536×65536 and its footprint in memory would be 32 GB. As a consequence, neither the precision matrix nor its Cholesky factor can be stored on standard computers.

On the contrary, the proposed PO algorithm only requires the storage of four 256×256 matrices and its footprint in memory is only 2MB that is easy to manage on standard computers. Regarding the computational cost:

- Step P requires a sample under each prior distribution: $\mathbf{x}$ is computed by FFT (see item 2 of Section I) and $\mathbf{n}$ is trivially computed since it is a white noise.

- Step O is achieved by a conjugate gradient procedure with optimal step size. It only requires computations of convolutions (by FFT), decimation and zero-padding.

So, the proposed PO algorithm is feasible and it easily provides a desired sample. Practically, it takes¹ about one second (i.e. around $P = 50$ gradient iterations) to obtain one sample.

Fig. 1. Chains and histograms of hyperparameters γ_n and γ_x .

Fig. 1 shows the iterates and illustrates the operation and convergence. After a burn-in period of about 25 iterations, the algorithm is in its converged state and the total number of iterations is 59 to ensure a good exploration of the distribution. Histograms approximating marginal posteriors are also given and the posterior means are $\hat{\gamma}_n \approx 7.7$ and $\hat{\gamma}_x \approx 2.2 \times 10^{-3}$.

Concerning the images themselves, results are shown in Fig. 2: the estimated image in 2(c) clearly shows a better resolution than the data in 2(b) and it is visually close to the original image in 2(a). Nevertheless, it is important to keep in mind that, w.r.t. other SR methods, the proposed PO algorithm does not improve image quality itself but the crucial novelty is to allow for hyperparameter estimation. In this sense, it is clear that the approach produces correct hyperparameters i.e. correct balance between data and prior. Moreover, uncertainties are derived from the samples through the posterior standard deviation. It is illustrated in Fig. 2(d): the true image is inside the 99% credibility interval around the estimate. As a conclusion, the proposed PO algorithm makes it possible to resort to sampling algorithms in SR method whereas it was not possible before. It then enables hyperparameter estimation while other SR methods require hand-made hyperparameter tuning. In addition, it enables to compute uncertainties based on posterior standard deviation.

B. Three other examples

The PO algorithm has been used in three other contexts: electromagnetic inverse scattering [14], fluorescent microscopy through structured illumination [20] and super-resolution from data provided the Herschel observatory in astronomy [?].

¹The algorithm is implemented within the computing environment Matlab on a PC with a 3 GHz CPU and 3 GB of RAM.

Fig. 2. Image reconstruction: true image 2(a), one of the low resolution images 2(b) and the proposed estimate 2(c). The plot 2(d) is a true image slice inside the 99% credibility interval around the estimate.

The problems are tackled in a Bayesian framework and implemented by means of stochastic sampling. In these contexts, the distribution for the object given the other variables is Gaussian with large size precision matrix. Its structure is neither sparse nor circulant making the use of existing algorithms impossible. This is due to non-linearity and label variables in [14] and non-invariance of the observation model in [20, ?]. Nevertheless, the precision matrix is in the form (1), so, the proposed PO is applicable.

IV. CONCLUSION

The paper presents an algorithm for sampling high-dimensional Gaussian distributions when existing algorithms are infeasible. It relies on a perturbation-optimization principle: adequate stochastic perturbation of a criterion and optimization of the perturbed criterion. It is shown that the criterion optimizer is a sample of the target distribution. The algorithm is applicable for a particular decomposition of the precision matrix that emerges in general linear inverse problems.

There is a wide class of applications, in particular any processing problem based on a conditional linear forward model and conditional Gaussian priors for noise and object. The interest and the feasibility of the proposed algorithm have been illustrated in [14, 20, ?] and in this paper on a more academic super-resolution problem allowing automatic tuning of hyperparameters.

An interesting perspective deals with the case of stopped optimization procedure. It is a question under consideration to prove that, embedded in a Gibbs loop, a finite number (maybe one) of iteration of the optimization step is enough to guarantee convergence towards the target law.

V. ACKNOWLEDGMENT

The authors would like to thank J. IDIER (IRRCyN) for inspiration of this work [21], T. RODET and A. DJAFARI (L2S), for fruitful discussions, C. VACAR (IMS) and P. SZACHERSKI (IMS and CEA-LETI) for carefully reading the paper.

REFERENCES

- [1] G. Winkler, *Image Analysis, Random Fields and Markov Chain Monte Carlo Methods*. Springer Verlag, BerlinGermany, 2003.
- [2] H. Rue, "Fast sampling of Gaussian Markov random fields," *J. R. Statist. Soc. B*, vol. 63, no. 2, 2001.
- [3] P. Lalanne, D. Prévost, and P. Chavel, "Stochastic artificial retinas: algorithm, optoelectronic circuits, and implementation," *Applied Optics*, vol. 40, 2001.
- [4] R. Chellappa and S. Chatterjee, "Classification of textures using Gaussian Markov random fields," *IEEE Trans. Acoust. Speech, Signal Processing*, vol. ASSP-33, pp. 959–963, Aug. 1985.
- [5] R. Chellappa and A. Jain, *Markov Random Fields: Theory and Application*. Academic Press Inc, 1992.
- [6] D. Geman and C. Yang, "Nonlinear image recovery with half-quadratic regularization," *IEEE Trans. Image Processing*, vol. 4, no. 7, pp. 932–946, July 1995.
- [7] P. Charbonnier, L. Blanc-Féraud, G. Aubert, and M. Barlaud, "Deterministic edge-preserving regularization in computed imaging," *IEEE Trans. Image Processing*, vol. 6, no. 2, pp. 298–311, Feb. 1997.
- [8] J.-F. Giovannelli, "Unsupervised Bayesian convex deconvolution based on a field with an explicit partition function," *IEEE Trans. Image Processing*, vol. 17, no. 1, pp. 16–26, Jan. 2008.
- [9] F. Orieux, J.-F. Giovannelli, and T. Rodet, "Bayesian estimation of regularization and point spread function parameters for Wiener–Hunt deconvolution," *J. Opt. Soc. Amer.*, vol. 27, no. 7, pp. 1593–1607, 2010.
- [10] H. Helgason, V. Pipiras, and P. Abry, "Fast and exact synthesis of stationary multivariate Gaussian time series using circulant embedding," *Signal Processing*, 2011.
- [11] J. Idier, Ed., *Bayesian Approach to Inverse Problems*. London: ISTE Ltd and John Wiley & Sons Inc., 2008.
- [12] H. Rue, S. Martino, and N. Chopin, "Approximate bayesian inference for latent gaussian models by using integrated nested laplace approximations," *J. R. Statist. Soc. B*, vol. 71, pp. 1–35, 2009.
- [13] S. Geman and D. Geman, "Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 6, no. 6, pp. 721–741, Nov. 1984.
- [14] O. Féron, B. Duchêne, and A. Mohammad-Djafari, "Microwave imaging of piecewise constant objects in a 2D-TE configuration," *International Journal of Applied Electromagnetics and Mechanics*, vol. 26, no. 6, pp. 167–174, IOS Press 2007.
- [15] H. Ayasso and A. Mohammad-Djafari, "Joint NDT image restoration and segmentation using Gauss-Markov-Potts prior models and variational Bayesian computation," *IEEE Trans. Image Processing*, vol. 19, no. 9, pp. 2265–2277, 2010.
- [16] X. Tan, J. Li, and P. Stoica, "Efficient sparse Bayesian learning via Gibbs sampling," in *Proc. IEEE ICASSP*, Mar. 2010, pp. 3634–3637.
- [17] J. Nocedal and S. J. Wright, *Numerical Optimization*, ser. Series in Operations Research. New York: Springer Verlag, 2000.
- [18] S. C. Park, M. K. Park, and M. G. Kang, "Super-resolution image reconstruction: a technical overview," *IEEE Trans. Signal Processing Mag.*, pp. 21–36, May 2003.
- [19] G. Rochefort, F. Champagnat, G. Le Besnerais, and J.-F. Giovannelli, "An improved observation model for SR under affine motion," *IEEE Trans. Image Processing*, vol. 15, no. 11, pp. 3325–3337, Nov. 2006.
- [20] F. Orieux, E. Sepulveda, V. Lorient, B. Dubertret, and J.-C. Olivo-Marin, "Bayesian estimation for optimized structured illumination microscopy," *IEEE Trans. Image Processing*, 2011, in press.
- [21] J. Idier, "Optimization and sampling in the Gaussian case," Personal communication, 2006.