

Bayesian image restoration for mosaic active imaging

Nicolas Lermé*, François Malgouyres†, Dominique Hamoir‡
Emmanuelle Thouin§ ¶ ||

February 28, 2014

Abstract

In this paper, we focus on the restoration of images acquired with a new active imaging concept. This new instrument generates a mosaic of active imaging acquisitions. We first describe a simplified forward model of this so-called “mosaic active imaging”. We also assume a prior on the distribution of images, using the Total Variation (TV), and deduce a restoration algorithm. This algorithm is a two-stage iterative process which alternates between: i) the estimation of the restored image; ii) the estimation of the acquisition parameters. We then provide the details useful to the implementation of these two steps. In particular, we show that the image estimation can be performed with graph cuts. This allows a fast resolution of this image estimation step. Finally, we detail numerical experiments showing that acquisitions made with a mosaic active imaging device can be restored even under severe noise levels, with few acquisitions.

Keywords: active imaging, laser imaging, image reconstruction, image estimation, graph cuts.

1 Introduction

Flash laser imaging (also called flash active imaging, gated active viewing, or more commonly active imaging), illuminates the object to be observed with a very short laser flash (of typically 5-20 ns). It captures the image with a high-speed camera, sharply synchronized with the emission. The photons coming back to the sensor are selected according to their round-trip travel time. This allows us to reject the photons back-scattered by the foreground (e.g. by fog, dust or vegetation) and those back-scattered by the background. The controlled addition of photons and their temporal selection allow a better signal-to-noise ratio and a better contrast of the object over the background. It is of

*LTCI, CNRS UMR5141, Institut Mines-Télécom, Télécom ParisTech, Paris, France, nicolas.lerme@isep.fr

†IMT, CNRS UMR5219, Université de Toulouse, Toulouse, France, francois.malgouyres@math.univ-toulouse.fr

‡ONERA - The French Aerospace Lab, F-31055 Toulouse, France, dominique.hamoir@onera.fr

§ONERA - The French Aerospace Lab, F-31055 Toulouse, France, emmanuelle.thouin@onera.fr

¶Université de Toulouse, Institut Supérieur de l’Aéronautique et de l’Espace (ISAE), F-31055 Toulouse

||IMT, CNRS UMR5219, Université de Toulouse, Toulouse, France.

interest for surveillance and for target identification under bad weather conditions or at long ranges (several kilometers).

A discrimination in sub-meter distance can be obtained in some cases. The observed objects typically have metric dimensions (e.g. buildings, vehicles, personnel, animals, fences). Depending on the application, they are located at distances from the imaging system ranging from 10m to 20km. In the most demanding applications, including those requiring distances in kilometers, several physical limitations degrade the images [HVB⁺09, RHV⁺10].

First, atmospheric turbulence produces two types of degradation. On the one hand, the laser illumination is not uniform over the object and is not stationary due to the forward propagation of the laser beam through the turbulent atmosphere. We talk of turbulence-induced illumination speckle (also speckle). On the other hand, the image of the object is distorted by the backward propagation.

Second, the interaction of the laser spot with the object is accompanied by artifacts, in particular if the light may be scattered off several directions (e.g. when the laser hits the inner side of a dihedral).

Third, the maximum distance of observation is limited by the size, weight and power compatible with integration on a land or air vehicle, in particular that of the laser and that of the reception optical system. A first way to overcome this difficulty is to restore the information despite a low signal-to-noise ratio (currently of a few units). A second way is to improve the light sensor, for instance by switching to Avalanche Photodiodes (APD). A third approach is to restore the image from a mosaic of typically 100 to 1000 elementary thumbnails [Ham10]. In the latter case, that we will call mosaic laser imaging or mosaic active imaging, each thumbnail has strong gradients of illumination, and geometric readjustments may have to be considered. This is the option studied in this paper.

As such, the problem we are addressing is an inverse problem in the field of imaging and the literature on the subject is huge. The main element which is common to all the existing methods is the trade-off between the fidelity of the result to the data; and the enforcement of an expected regularity (or prior) to the result. We refer to [Dem89] for an historical perspective. Most of the current work is on the prior and usually, authors distinguish synthesis priors (see [Ela10], for an overview) and analysis priors. We have chosen a well established analysis prior, for which rapid minimization algorithms are available: the Total Variation (TV).

The TV prior has first been proposed for image denoising and deblurring (see [ROF92]) and has, since then, been applied in many other contexts of image restoration such as inpainting (see [CS05]), image zooming (see [MG01]), restoration of compression artifacts (see [CS00]), etc. Its success lies in its ability to properly restore sharp edges. Moreover, its minimization has been intensively studied, and fast and simple iterative algorithms have been developed (see, for instance [Cha04, BBFAC04, CP11]). Recently, algorithms using graph cuts have been developed and provide, for some models involving this prior, fast and exact (modulo a quantization) minimization methods (see [Cha05, DS06, CD09]). This is the numerical strategy we have chosen. Of course, the TV prior has known drawbacks. For instance, we are aware that images regularized with this prior tend to contain staircasing (see [Nik00]). It is also known that this prior removes textures and lowers the contrast (see [Mey01]).

In this paper, we investigate algorithmic ways to restore mosaic active images. The rest of this paper is organized as follows. In Section 2, we describe a simplified physical and mathematical model of the imaging process and describe the sketch of the restoration algorithm. In particular, this section exhibits that the image acquisition depends on imperfectly known acquisition parameters. The algorithm consists in alternating the estimation of these acquisition parameters and the estimation of the image. Next, we show in Section 3 how the estimation of the image can be formulated using level-sets and solved with graph cuts. Then, in Section 4, we give the details concerning the implementation of the algorithm used to estimate the acquisition parameters. Afterwards, we provide in Section 5 numerical experiments assessing the quality of the image estimate, the influence of the acquisition parameters, the convergence of the algorithm, an estimate of the expected image resolution and the results of the algorithm. Finally, we remind in Section 6 the contributions of this paper, propose some ideas to improve this work and discuss directions for future work.

2 Modeling of flash laser imaging

2.1 Overview

In flash laser imaging, a “light ball” is repeatedly sent towards the object to be observed. A time-gated camera synchronized with the laser is used to detect and select the light that is received within a brief time-interval or time-gate. This time-gate is typically of a few nano to micro seconds and starts after a chosen time delay of typically 10^{-7} to 10^{-4} second has elapsed. This allows the camera to record the photons coming back from the object (in the time-gate) and to reject those coming back from the foreground or from the background (before or after the time-gate). The wavelength of operation can be chosen according to the application but is usually in the so-called eye-safe region, between 1.5 and 1.6 micrometers.

Generally, the field of view of the camera is fully illuminated by the laser and is acquired at standard video rates, say 10 Hz. In mosaic laser imaging, we replace the low-repetition-rate 10Hz laser with optical parametric oscillator by a high-repetition-rate 10kHz fiber laser. The latter is expected to offer higher average power and plug-efficiencies within a few years. This concept presents additional advantages. As the repetition rate is larger by three orders of magnitude, the energy per pulse is lowered by the same ratio. In order to maintain the signal-to-noise ratio, only a reduced part of the field of view is illuminated at each laser flash. The corresponding region of interest of the sensor is read. The laser beam is then deflected in order to illuminate another region of interest. By repeating the process, we scan the field of view of the camera. This results in the successive acquisition of elementary images taken at a repetition-rate of 10 kHz that will tile as a mosaic in order to build the full-frame image at 10 Hz. The formation of each elementary image can be modeled as follows.

The object is illuminated with a Gaussian laser spot with position c_k and beam radius w_k , in the image (these quantities are expressed in terms of inter-pixel distance). This laser spot is affected by three perturbations due to pointing discrepancies and to the forward propagation of the laser beam through the inhomogeneous turbulent atmosphere: beam spreading, beam wandering, and turbulence-induced speckle. This illumination pattern is multiplied by the reflectance of the object




Figure 1: Acquisition process in mosaic laser imaging on an aerial image. The left image is the ideal image that we want to estimate. The middle image contains the reduced part of the field of view illuminated and acquired using a single laser flash. We call it laser shot. The right image contains the view composed of all laser shots. In this image, each pixel is assigned with its maximal intensity over all laser shots. Notice that the information is concentrated on illumination domes and is missing between them.

to form a luminance distribution. This travels through the atmosphere and is captured by the optical system of the camera to form an image near its focal plane. Shot noise and thermal noise are then added to the image. This is repeated for each elementary image (the images are indexed by k).

More precisely, we denote, for an integer $N > 0$, the set of all pixels by $\mathcal{P} = \{1, \dots, N\}^2$. We denote by $K \geq 1$, the number of elementary images. For every index $k \in \{1, \dots, K\}$, we denote by $\theta_k = (c_k, w_k) \in (\mathbb{R}^2 \times \mathbb{R}_+^*)$ the parameters of the Gaussian profile (here \mathbb{R}_+^* denotes the positive reals). We also consider the beam intensity profile G_{θ_k} defined, for every $p \in \mathcal{P}$, by

$$G_{\theta_k}(p) = \exp\left(-\frac{\|p - c_k\|^2}{2w_k^2}\right), \quad (1)$$

where, here and all along the paper, $\|\cdot\|$ denotes the Euclidean norm, whatever the dimension of the considered Euclidean space.

Notice that, when necessary, we denote the coordinates of the elements of \mathbb{R}^2 with subscript i and j (e.g. $c_k = (c_{k,i}, c_{k,j})$ and $p = (p_i, p_j)$).

For every k , the image $(v_p^k)_{p \in \mathcal{P}} \in \mathbb{R}^{\mathcal{P}}$ is obtained from the ideal image $(u_p)_{p \in \mathcal{P}} \in \mathbb{R}^{\mathcal{P}}$ (i.e. the one that would have been obtained with an ideal captor and an ideal illumination), using

$$v_p^k = u_p G_{\theta_k}(p) S_p^k + n_p^k, \quad \forall p \in \mathcal{P},$$

where $(S_p^k)_{p \in \mathcal{P}} \in \mathbb{R}^{\mathcal{P}}$ models the speckle pattern and $(n_p^k)_{p \in \mathcal{P}} \in \mathbb{R}^{\mathcal{P}}$ represents the noise. In the remaining of the paper, we will refer to the image containing the beam intensity profiles as the illumination domes and the image v as the laser shots. Examples of laser shots are provided in Figure 1.

The mathematical developments will be conducted on elementary images synthesized with this simplified model. We are aware that this model does not take into account neither the size of the reception pupil of the instrument nor the transverse sampling by the focal plane array. Leaving these degradations aside allows us indeed to use faster and more efficient restoration algorithms.

More precisely, improving the image creation model would force us to use more complex algorithms, requiring more computational resources. We therefore leave the study of more accurate degradation models, the development of adapted restoration algorithms, as well as the comparison of degradation models/algorithms couples for a future work.

Moreover, the illumination speckle factor (S_p^k) is a colored noise that can be viewed as a textured illumination. It is a strong limitation in terrestrial applications but is negligible in airborne applications. We neglect this possible contribution in this first study and leave it for the near future.

By neglecting the speckle, we obtain the following simplified forward model:

$$v = M_{(\theta_k)_{1 \leq k \leq K}} u + n,$$

where $n = (n^k)_{1 \leq k \leq K}$ with $n^k \in \mathbb{R}^{\mathcal{P}}$ and $v = (v^k)_{1 \leq k \leq K}$ with $v^k \in \mathbb{R}^{\mathcal{P}}$ and

$$\begin{aligned} M_{(\theta_k)_{1 \leq k \leq K}} : \mathbb{R}^{\mathcal{P}} &\longrightarrow \mathbb{R}^{K\mathcal{P}}, \\ u &\longmapsto \left((G_{\theta_k}(p)u_p)_{p \in \mathcal{P}} \right)_{1 \leq k \leq K}. \end{aligned}$$

This is a linear model, once the acquisition parameters are fixed. We will see, however, that these parameters need to be estimated.

2.2 Beam spreading

The Gaussian beam does naturally spread along the propagation. We denote the distance between the laser source and the object by $d > 0$. The standard instrument-to-target range is $d = 1000$ meters. The minimum beam radius, at the laser source in our case, is called the beam waist radius and is denoted w_0 . In the absence of atmospheric turbulence, the spreading is only due to the diffraction of the beam:

$$\bar{w}^2 = \left(\frac{\lambda_{lw}d}{\pi w_0} \right)^2 + w_0^2,$$

where $\lambda_{lw} > 0$ is the laser wavelength. For every $k \in \{1, \dots, K\}$, we propose to model the variation, induced by the atmospheric turbulence, of the true beam spreading w_k around its expected value using the Gaussian law

$$\mathbb{P}(w_k) \propto \exp\left(-\frac{|w_k - \bar{w}|^2}{2\sigma_w^2}\right), \quad \text{for } 1 \leq k \leq K,$$

where $\sigma_w > 0$ is a known parameter.

In our experimental setup the expected beam radius is $\bar{w} = 16.2$ pixels, as seen by our 256×256 -pixels camera. The calculated standard deviations of the beam radius are given in Table 1 for standard turbulence levels, defined by their refractive-index-structure constant C_n^2 [Kol49]. Under these conditions, they are very small compared to the mean beam radius and to the pixel size.

2.3 Beam wandering

Beam wandering results from an angular deviation of the beam propagation axis, due to possible pointing discrepancies of the instrument and to the propagation through the turbulent atmosphere.

C_n^2	σ_w
$10^{-15}m^{-\frac{2}{3}}$	6×10^{-4} pixels
$10^{-14}m^{-\frac{2}{3}}$	8×10^{-3} pixels
$10^{-13}m^{-\frac{2}{3}}$	7×10^{-2} pixels

Table 1: Standard deviations of the beam radius w_k , in pixels.

C_n^2	σ_c
$10^{-15}m^{-\frac{2}{3}}$	0.09 pixels
$10^{-14}m^{-\frac{2}{3}}$	0.26 pixels
$10^{-13}m^{-\frac{2}{3}}$	0.81 pixels

Table 2: Standard deviations of the beam position c_k , in pixels.

When considering a perfect pointing, the statistics of this angular deviation follows a Gaussian law with zero mean. Its variance is given by (see [Fan75, Cha92])

$$\sigma_c^2 = 0.16(\lambda_{lw}d)^2(2w_0)^{-\frac{1}{3}}r_0^{-\frac{5}{3}},$$

where $2w_0$ is the beam waist diameter and r_0 is the Fried's coherence length of the turbulent atmosphere [Fri66], which is related to the refractive-index-structure constant C_n^2 [Rod81]. We remind that λ_{lw} is the laser wavelength and that d is the distance between the laser source and the object. Using simple geometry, and since the laser is almost orthogonal to the acquisition device, we approximate the probability density function of c_k with

$$\mathbb{P}(c_k) \propto \exp\left(-\frac{\|c_k - \bar{c}_k\|^2}{2\sigma_c^2}\right), \quad \text{for } 1 \leq k \leq K,$$

where $\bar{c}_k \in \mathbb{R}^2$ and σ_c^2 are known parameters. We assume moreover that the random variable c_k is independent of w_k . In our examples, the calculated standard deviations of the beam position c_k are given in Table 2 for standard turbulence levels.

2.4 Noise

The shot noise (or photon noise) is due to the statistics of emission of photons by the source. The number of photons received by the pixel follows a Poisson law. The thermal noise (or detection noise) is a white Gaussian noise with zero mean and standard deviation σ . The number of photoelectrons generated in the pixel writes as a weighted sum of these two noises.

In the applicative context we are interested in, the laser intensity needs to be lowered and hence we have a low signal level. In this challenging case, the thermal noise dominates the photon noise. Hence, in the following, we consider additive thermal noise only of normalized standard deviation σ .

More precisely, we assume that the data $v \in \mathbb{R}^{K\mathcal{P}}$ is obtained by corrupting the observation of an ideal image u through the operator $M_{(\theta_k)_{1 \leq k \leq K}}$ with an additive white Gaussian noise. We have

$$\mathbb{P}(v|u, (\theta_k)_{1 \leq k \leq K}) \propto \exp\left(-\frac{\|M_{(\theta_k)_{1 \leq k \leq K}}u - v\|^2}{2\sigma^2}\right),$$




Figure 2: Typical neighborhoods of $p = (p_i, p_j)$. These examples involve 4 pixels (dotted arrows), 8 pixels (dotted and dashed arrows) or 16 pixels (all arrows).

where $\sigma > 0$ is the known standard deviation of the noise. In the experimental section and for the purpose of comparison, we express σ with respect to the maximum intensity of the noise-free image.

2.5 Image prior and restoration principle

First, we denote by $\mathcal{N} \subset (\mathcal{P} \times \mathcal{P})$ a neighborhood system connecting pixels. The Figure 2 shows typical neighborhoods used on the lattice \mathcal{P} . We assume that the observed data $u \in \mathbb{N}^{\mathcal{P}}$ is a random variable following a law

$$\mathbb{P}(u) \propto \exp(-\beta TV(u)),$$

where $\beta > 0$ is an unknown parameter (which will later on be tuned by the user) and the TV is defined by

$$TV(u) = \sum_{(p,q) \in \mathcal{N}} d_q^p (u_p - u_q)^+, \quad (2)$$

where $(y)^+ = \max\{y, 0\}$ and, for all $(p, q) \in \mathcal{N}$, $d_q^p \geq 0$ is a known coefficient. Notice that the above definition allows to have $d_q^p \neq d_p^q$.

We also consider that the parameters c_k and w_k are all independent random variables. Therefore, their joint distribution satisfies

$$\mathbb{P}((\theta_k)_{1 \leq k \leq K}) \propto \prod_{k=1}^K \mathbb{P}(c_k) \mathbb{P}(w_k).$$

Applying Bayes' law and assuming that u is independent of the parameters $(\theta_k)_{1 \leq k \leq K}$, we obtain,

for any $v \in \mathbb{R}^{K\mathcal{P}}$, the posterior

$$\begin{aligned} \mathbb{P}(u, (\theta_k)_{1 \leq k \leq K} | v) &= \frac{\mathbb{P}(v | u, (\theta_k)_{1 \leq k \leq K}) \mathbb{P}(u, (\theta_k)_{1 \leq k \leq K})}{\mathbb{P}(v)}, \\ &\propto \exp\left(-\frac{\|M_{(\theta_k)_{1 \leq k \leq K}} u - v\|^2}{2\sigma^2}\right) \exp(-\beta TV(u)) \\ &\quad \times \prod_{k=1}^K \exp\left(-\frac{\|c_k - \bar{c}_k\|^2}{2\sigma_c^2}\right) \exp\left(-\frac{|w_k - \bar{w}|^2}{2\sigma_w^2}\right). \end{aligned}$$

We consider in the following a Maximum A Posteriori estimator (MAP) of u and $(\theta_k)_{1 \leq k \leq K}$. As is usual, we compute the MAP estimate by minimizing in u and $(\theta_k)_{1 \leq k \leq K}$

$$\begin{aligned} -\log(\mathbb{P}(u, (\theta_k)_{1 \leq k \leq K} | v)) &= C + \frac{\|M_{(\theta_k)_{1 \leq k \leq K}} u - v\|^2}{2\sigma^2} + \beta TV(u) \\ &\quad + \sum_{k=1}^K \frac{\|c_k - \bar{c}_k\|^2}{2\sigma_c^2} + \sum_{k=1}^K \frac{|w_k - \bar{w}|^2}{2\sigma_w^2}, \end{aligned}$$

where C is real and does not have any influence on the minimizer.

For simplicity, given a fixed $v \in \mathbb{R}^{K\mathcal{P}}$ and for any $u \in \mathbb{R}^{\mathcal{P}}$ and $(\theta_k)_{1 \leq k \leq K} \in (\mathbb{R}^2 \times \mathbb{R}_+^*)^K$, we denote the minimized function by

$$F(u, (\theta_k)_{1 \leq k \leq K}) = \frac{\|M_{(\theta_k)_{1 \leq k \leq K}} u - v\|^2}{2\sigma^2} + \beta TV(u) + \sum_{k=1}^K \frac{\|c_k - \bar{c}_k\|^2}{2\sigma_c^2} + \sum_{k=1}^K \frac{|w_k - \bar{w}|^2}{2\sigma_w^2}. \quad (3)$$

The definition of F involves four parameters: σ , β , σ_c , σ_w . Among these parameters, we remind that three are directly related to physical quantities that can be estimated by other means. The only parameter that needs to be tuned is β .

It is not difficult to see that, the function F is continuously differentiable over $\mathbb{R}^{\mathcal{P}} \times (\mathbb{R}^2 \times \mathbb{R}_+^*)^K$ and that both F and its partial derivatives ¹ can be continuously extended over $\mathbb{R}^{\mathcal{P}} \times (\mathbb{R}^2 \times \mathbb{R}_+)^K$. Throughout the paper, we abuse notations and assimilate F with its continuous and continuously differentiable extension over the closed domain $\mathbb{R}^{\mathcal{P}} \times (\mathbb{R}^2 \times \mathbb{R}_+)^K$. The minimization is also performed over this closed domain.

Finally, F is non-negative and coercive. Therefore, F reaches a global minimum over $\mathbb{R}^{\mathcal{P}} \times (\mathbb{R}^2 \times \mathbb{R}_+)^K$. However, F is non-convex and standard descent algorithms might get stuck in local minimas.

2.6 An alternate minimization algorithm

Notice that, considering the above properties of F , we cannot a priori provide guarantees that we compute a true minimizer of F . We propose an alternate minimization algorithm whose sketch is described in Table 3.

Notice that, for any $(c_k, w_k)_{1 \leq k \leq K} = (\theta_k)_{1 \leq k \leq K} \in (\mathbb{R}^2 \times \mathbb{R}_+)^K$, the function $u \mapsto F(u, (\theta_k)_{1 \leq k \leq K})$ is convex and coercive. It therefore achieves its minimum and one of its minimizers can be computed by

¹The partial derivatives of F are provided in Section 4.

<ul style="list-style-type: none"> • Initialize $(c_k^0, w_k^0)_{1 \leq k \leq K} = (\bar{c}_k, \bar{w})_{1 \leq k \leq K}$ • Repeat while $\ u^n - u^{n-1}\ \leq \varepsilon_a$ <ol style="list-style-type: none"> 1. Use a graph cuts algorithm to compute $u^n \in \underset{u \in \mathbb{R}^{\mathcal{P}}}{\operatorname{argmin}} F(u, (c_k^n, w_k^n)_{1 \leq k \leq K}).$ 2. Use a gradient-based algorithm to compute $(c_k^{n+1}, w_k^{n+1})_{1 \leq k \leq K} \in \underset{(c_k, w_k)_{1 \leq k \leq K} \in (\mathbb{R}^2 \times \mathbb{R}_+)^K}{\operatorname{argmin}} F(u^n, (c_k, w_k)_{1 \leq k \leq K}).$

Table 3: Structure of the algorithm used for approximating a minimizer of F .

a well thought optimization technique. In this paper, we propose a graph cuts-based minimization. Its main interest is to be fast and to provide an exact solution to the step 1 of the algorithm of Table 3, modulo a quantization step chosen by the user. All the details concerning this graph cuts algorithm are described in Section 3.

Similarly, when $u \in \mathbb{R}^{\mathcal{P}}$ is fixed, the function $(\theta_k)_{1 \leq k \leq K} \mapsto F(u, (\theta_k)_{1 \leq k \leq K})$ is continuous and coercive. It therefore reaches a global minimum. This function might however be non-convex. Therefore a gradient-based algorithm (such as the one we are using) is only guaranteed to converge to a stationary point (see [Ber03], p. 52). However, when σ_c and σ_w are small enough, we expect the second argument of the global minimizer of F to be close to $(\bar{c}_k, \bar{w})_{1 \leq k \leq K}$ and expect good convergence properties when using this initialization. In contrast to the step 1 of the algorithm, it is important to notice that the algorithm used for solving the step 2 relies on a stopping criterion. All the details concerning this gradient-based algorithm are given in Section 4.

It is not difficult to see that the sequence $(F(u^n, (c_k^n, w_k^n)_{1 \leq k \leq K}))_{n \in \mathcal{N}}$ generated by the algorithm of Table 3 decays. Moreover, since F is non-negative, this sequence converges to a minimum value F' . Moreover, since F is coercive, $(u^n, (c_k^n, w_k^n)_{1 \leq k \leq K})$ has limit points. Moreover, it is not difficult to see that any such limit point $(u', (c'_k, w'_k)_{1 \leq k \leq K})$ is a stationary point of F . We will see experimentally in Section 5 that the proposed algorithm has good convergence properties in the practical situations we are interested in.

Finally, the stopping criterion $\|u^n - u^{n-1}\| \leq \varepsilon_a$ controls that the variation of u , between two successive iterations, is smaller than a parameter ε_a .

3 Image estimation using graph cuts

In this section, we describe how the restored image u can be efficiently estimated when the acquisition parameters $(\theta_k)_{1 \leq k \leq K}$ are known (see the first step of Table 3). Using the approach of [CD09], we first show that minimizing the energy function F is equivalent to minimizing a sequence of problems which only involve binary variables. Then, we present some reminders about graph cuts optimization and show that each of these problems can be efficiently solved using this algorithmic strategy.

3.1 Leveled-energies decomposition

Since the acquisition parameters $(\theta_k)_{1 \leq k \leq K}$ are supposed to be known, the third and fourth term in (3) are constant, since they do not depend on u . We are therefore interested in minimizing among $u \in \{0, \dots, L-1\}^{\mathcal{P}}$ and for $L > 0$ ²

$$E(u) = \sum_{p \in \mathcal{P}} \underbrace{\left(\frac{1}{2\sigma^2} \sum_{k=1}^K (G_{\theta_k}(p)u_p - v_p^k)^2 \right)}_{E_p(u_p)} + \beta TV(u). \quad (4)$$

Recently, the authors of [CD09] proposed to decompose an energy similar to E , in (4), as a sum of energies on the level sets of u (see [CD09]). First, it is not difficult to see that for any pixel $p \in \mathcal{P}$, if we denote $u^\lambda = \mathbf{1}_{\{u \geq \lambda\}}$ the λ -level set of u , the term E_p in (4) can be decomposed as

$$E_p(u_p) = \sum_{\lambda=1}^{L-1} u_p^\lambda (E_p(\lambda) - E_p(\lambda-1)) + E_p(0). \quad (5)$$

Notice that the latter equation is consistent whatever $u_p \in \{0, \dots, L-1\}$. Similarly, the TV term in (4) can be decomposed as

$$TV(u) = \sum_{\lambda=1}^{L-1} \underbrace{\sum_{(p,q) \in \mathcal{N}} d_q^p(u_p^\lambda - u_q^\lambda)^+}_{TV(u^\lambda)}. \quad (6)$$

Notice that the last summation in the above term starts at $\lambda = 1$, since $u_p^0 = u_q^0 = 1, \forall (p, q) \in \mathcal{N}$.

Using (4), (5) and (6), the energy E becomes

$$E(u) = \sum_{\lambda=1}^{L-1} E^\lambda(u^\lambda) + C, \quad (7)$$

where C is a constant that does not depend on u and the energy E^λ is defined, for any $\lambda \in \{1, \dots, L-1\}$ and any $w \in \{0, 1\}^{\mathcal{P}}$, by

$$E^\lambda(w) = \sum_{p \in \mathcal{P}} w_p (E_p(\lambda) - E_p(\lambda-1)) + \beta TV(w). \quad (8)$$

For any level $\lambda \in \{1, \dots, L-1\}$, let us denote $\hat{u}^\lambda \in \{0, 1\}^{\mathcal{P}}$ a minimizer of E^λ . Notice that, if these minimizers satisfy

$$\hat{u}_p^\lambda \geq \hat{u}_p^{\lambda'}, \quad \forall 0 \leq \lambda \leq \lambda' \leq L-1, \quad \forall p \in \mathcal{P}, \quad (9)$$

then, provided (7), it is not difficult to check that the elements $\hat{u} \in \{0, \dots, L-1\}^{\mathcal{P}}$ defined for all $p \in \mathcal{P}$, by

$$\hat{u}_p = \max\{\lambda \in \{0, \dots, L-1\} \mid \hat{u}_p^\lambda = 1\} \quad \text{minimizes } E.$$

In words, if (9) holds, we can deduce a minimizer of E from all the minimizers, for all $\lambda \in \{1, \dots, L-1\}$, of E^λ . Let us prove that the monotone condition (9) holds. To do so, we adapt the proof given

² L denotes the number of grayscale levels of u . Typical values for L are 2^8 , 2^{16} or 2^{32} .

in the Appendix C of [CD09] by replacing the L^2 norm by the terms E_p given in (4). First, since \hat{u}^λ is a minimizer of E^λ , we have, for any level $\lambda \in \{1, \dots, L-2\}$,

$$\begin{aligned} E^{\lambda+1}(\hat{u}^{\lambda+1}) &\leq E^{\lambda+1}(\hat{u}^{\lambda+1} \wedge \hat{u}^\lambda), \\ E^\lambda(\hat{u}^\lambda) &\leq E^\lambda(\hat{u}^{\lambda+1} \vee \hat{u}^\lambda), \end{aligned}$$

where for any $x, x' \in \{0, 1\}^{\mathcal{P}}$ we define $(x \vee x')_p = \max\{x_p, x'_p\}$ and $(x \wedge x')_p = \min\{x_p, x'_p\}$, $\forall p \in \mathcal{P}$.

Summing the two above inequalities guarantees that, for any level $\lambda \in \{1, \dots, L-2\}$,

$$E^{\lambda+1}(\hat{u}^{\lambda+1}) + E^\lambda(\hat{u}^\lambda) \leq E^{\lambda+1}(\hat{u}^{\lambda+1} \wedge \hat{u}^\lambda) + E^\lambda(\hat{u}^{\lambda+1} \vee \hat{u}^\lambda). \quad (10)$$

Additionally, the authors of [CD09] proved that the TV is submodular, i.e. for any $x, x' \in \{0, 1\}^{\mathcal{P}}$

$$TV(x \vee x') + TV(x \wedge x') \leq TV(x) + TV(x'). \quad (11)$$

Using the definition of E^λ given by (8) in the inequality (10); and using (11) with $x = \hat{u}^\lambda$ and $x' = \hat{u}^{\lambda+1}$, we find that for any $\lambda \in \{1, \dots, L-2\}$

$$\begin{aligned} \sum_{p \in \mathcal{P}} \hat{u}_p^{\lambda+1} (E_p(\lambda+1) - E_p(\lambda)) + \hat{u}_p^\lambda (E_p(\lambda) - E_p(\lambda-1)) \\ \leq \sum_{p \in \mathcal{P}} (\hat{u}^{\lambda+1} \wedge \hat{u}^\lambda)_p (E_p(\lambda+1) - E_p(\lambda)) + (\hat{u}^{\lambda+1} \vee \hat{u}^\lambda)_p (E_p(\lambda) - E_p(\lambda-1)). \end{aligned}$$

As a consequence, for any $\lambda \in \{1, \dots, L-2\}$,

$$\begin{aligned} \sum_{p \in \mathcal{P}} (E_p(\lambda+1) - E_p(\lambda)) (\hat{u}_p^{\lambda+1} - (\hat{u}^{\lambda+1} \wedge \hat{u}^\lambda)_p) \\ \leq \sum_{p \in \mathcal{P}} (E_p(\lambda) - E_p(\lambda-1)) ((\hat{u}^{\lambda+1} \vee \hat{u}^\lambda)_p - \hat{u}_p^\lambda). \quad (12) \end{aligned}$$

Then, evaluating the four possible couples of values for $(\hat{u}_p^\lambda, \hat{u}_p^{\lambda+1})$, it is not difficult to check that

$$\hat{u}_p^{\lambda+1} - (\hat{u}^{\lambda+1} \wedge \hat{u}^\lambda)_p = (\hat{u}^{\lambda+1} \vee \hat{u}^\lambda)_p - \hat{u}_p^\lambda = (\hat{u}_p^{\lambda+1} - \hat{u}_p^\lambda)^+.$$

Combining the latter equalities with (12), we obtain, for any $\lambda \in \{1, \dots, L-2\}$,

$$\sum_{p \in \mathcal{P}} (E_p(\lambda+1) - 2E_p(\lambda) + E_p(\lambda-1)) (\hat{u}_p^{\lambda+1} - \hat{u}_p^\lambda)^+ \leq 0.$$

Beside, one can easily check, using the definition of E_p in (4), that for any pixel $p \in \mathcal{P}$, E_p is strictly convex and therefore whatever λ

$$E_p(\lambda+1) - 2E_p(\lambda) + E_p(\lambda-1) > 0.$$

This implies that, for any level $\lambda \in \{1, \dots, L-2\}$ and any pixel $p \in \mathcal{P}$, $(\hat{u}_p^{\lambda+1} - \hat{u}_p^\lambda)^+ = 0$ and so $\hat{u}_p^\lambda \geq \hat{u}_p^{\lambda+1}$. This permits to conclude that (9) holds.

As a conclusion, we can compute a minimizer of E from the minimizers of E^λ . In the next section, we describe how each energy function E^λ can be (efficiently) minimized using graph cuts.

3.2 Leveled-energies minimization

To our best knowledge, it was first noticed by [PR75] that binary energies of the form

$$TV(x) = \sum_{p \in \mathcal{P}} \alpha_p x_p,$$

with $x \in \{0, 1\}^{\mathcal{P}}$, can be represented by a graph and minimized using maximum-flow (or by duality, minimum-cut) algorithms. Until the nineties, because of limited resources and limited algorithmic development, the application of this minimization procedure in image processing only focused on binary image denoising [GPS89]. A few years later, it has been shown that the submodularity of the pairwise terms³ of an energy is necessary for exactly minimizing this energy. The same condition is also sufficient to allow the minimization of the energy using a minimum-cut in an appropriate graph [KZ04]. Together with the arrival of a fast maximum-flow algorithm [BK04] designed for the typical graph problems encountered in image processing, this enabled to efficiently solve a wide range of problems such as image segmentation, denoising, reconstruction, optical flow or texture synthesis.

In what follows, we describe how this strategy can be applied for solving the energy E^λ (see (8)) for a fixed level $\lambda \in \{1, \dots, L-1\}$. The construction of the graph presented below is detailed in [KZ04].

In order to build the oriented and capacited graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, c)$, we first consider the set of nodes

$$\mathcal{V} = \mathcal{P} \cup \{s, t\},$$

where s and t are terminal nodes, respectively named the source and the sink. We also consider a set of directed edges

$$\mathcal{E} = (\{s\} \times \mathcal{P}) \cup (\mathcal{P} \times \{t\}) \cup \mathcal{N} \subset (\mathcal{V} \times \mathcal{V}).$$

As usual, the edges in $(\{s\} \times \mathcal{P}) \cup (\mathcal{P} \times \{t\})$ are called the t-links and the edges in \mathcal{N} are called the n-links. Moreover, we associate a non-negative capacity to any couple $(p, q) \in \mathcal{E}$, according to

$$\begin{cases} c(s, p) &= (E_p(\lambda) - E_p(\lambda - 1))^- , \forall p \in \mathcal{P}, \\ c(p, t) &= (E_p(\lambda) - E_p(\lambda - 1))^+ , \forall p \in \mathcal{P}, \\ c(p, q) &= \beta d_q^p , \forall (p, q) \in \mathcal{N}, \end{cases} \quad (13)$$

where $(y)^- = \max\{-y, 0\}$ and

$$\begin{aligned} E_p(\lambda) - E_p(\lambda - 1) &= \frac{1}{2\sigma^2} \sum_{k=1}^K (G_{\theta_k}(p)\lambda - v_p^k)^2 - (G_{\theta_k}(p)(\lambda - 1) - v_p^k)^2 \\ &= \frac{1}{2\sigma^2} \left[(2\lambda - 1) \sum_{k=1}^K G_{\theta_k}^2(p) - 2 \sum_{k=1}^K G_{\theta_k}(p)v_p^k \right]. \end{aligned}$$

Notice that, for all $(p, q) \in \mathcal{E}$,

$$c(p, q) \geq 0.$$

and, for all $p \in \mathcal{P}$,

$$c(s, p) \neq 0 \Rightarrow c(p, t) = 0.$$

³Here, this condition clearly holds for E^λ (see (8)) since TV is submodular (see (11)).

Notice also that $E_p(\lambda) - E_p(\lambda - 1)$ is computed at a limited cost once $\sum_{k=1}^K G_{\theta_k}^2(p)$ and $\sum_{k=1}^K G_{\theta_k}(p)v_p^k$ have been computed and stored. Moreover, for any p , the contribution to these sums of most of the indexes k can be neglected.

Next, we denote by $\mathcal{C} = (\mathcal{S}, \mathcal{T})$ an s-t cut in the graph \mathcal{G} ; i.e. a partition of the nodes \mathcal{V} such that $s \in \mathcal{S}$ and $t \in \mathcal{T}$. For any s-t cut \mathcal{C} , we remind that its value is given by

$$\text{val}_{\mathcal{G}}(\mathcal{C}) = \sum_{(p,q) \in (\mathcal{S} \times \mathcal{T})} c(p, q).$$

The s-t cut of minimum weight is called minimum s-t cut. For an s-t cut \mathcal{C} , we also define $x^{\mathcal{C}} \in \{0, 1\}^{\mathcal{P}}$ by

$$x_p^{\mathcal{C}} = \begin{cases} 0 & \text{if } p \in \mathcal{T} \\ 1 & \text{if } p \in \mathcal{S} \end{cases}, \quad \forall p \in \mathcal{P}.$$

From the latter equation, it is not difficult to see that the application $\mathcal{C} \mapsto x^{\mathcal{C}}$ establishes a one-to-one correspondence between the s-t cuts and the elements of $\{0, 1\}^{\mathcal{P}}$. Given the graph \mathcal{G} and the capacities (13), since $s \in \mathcal{S}$ and $t \in \mathcal{T}$, the contribution to $\text{val}_{\mathcal{G}}(\mathcal{C})$ of the t-links involving any $p \in \mathcal{P}$ is as follows:

$$\begin{cases} \text{if } p \in \mathcal{S} & \text{its contribution is } c(p, t) = (E_p(\lambda) - E_p(\lambda - 1))^+ \\ \text{if } p \in \mathcal{T} & \text{its contribution is } c(s, p) = (E_p(\lambda) - E_p(\lambda - 1))^- \end{cases}$$

Altogether, the contribution of the t-links involving $p \in \mathcal{P}$ to $\text{val}_{\mathcal{G}}(\mathcal{C})$ is

$$x_p^{\mathcal{C}} (E_p(\lambda) - E_p(\lambda - 1))^+ + (1 - x_p^{\mathcal{C}}) (E_p(\lambda) - E_p(\lambda - 1))^-.$$

Also, for any n-link $(p, q) \in \mathcal{N}$, we have

$$(p, q) \in (\mathcal{S} \times \mathcal{T}) \text{ if and only if } (x_p^{\mathcal{C}} - x_q^{\mathcal{C}})^+ = 1.$$

It is then straightforward to check that, for any s-t cut \mathcal{C} ,

$$\text{val}_{\mathcal{G}}(\mathcal{C}) = \sum_{p \in \mathcal{P}} x_p^{\mathcal{C}} (E_p(\lambda) - E_p(\lambda - 1))^+ + (1 - x_p^{\mathcal{C}}) (E_p(\lambda) - E_p(\lambda - 1))^- + \sum_{(p,q) \in \mathcal{N}} \beta d_q^p (x_p^{\mathcal{C}} - x_q^{\mathcal{C}})^+.$$

This can be rearranged under the form

$$\begin{aligned} \text{val}_{\mathcal{G}}(\mathcal{C}) = \sum_{p \in \mathcal{P}} x_p^{\mathcal{C}} [(E_p(\lambda) - E_p(\lambda - 1))^+ - (E_p(\lambda) - E_p(\lambda - 1))^-] + (E_p(\lambda) - E_p(\lambda - 1))^- \\ + \beta \sum_{(p,q) \in \mathcal{N}} d_q^p (x_p^{\mathcal{C}} - x_q^{\mathcal{C}})^+. \end{aligned}$$

Finally, using the definition of TV , in (2), and the definition of E^λ , in (8), we obtain

$$\begin{aligned} \text{val}_{\mathcal{G}}(\mathcal{C}) &= \sum_{p \in \mathcal{P}} x_p^{\mathcal{C}} (E_p(\lambda) - E_p(\lambda - 1)) + \beta TV(x^{\mathcal{C}}) + \sum_{p \in \mathcal{P}} (E_p(\lambda) - E_p(\lambda - 1))^-, \\ &= E^\lambda(x^{\mathcal{C}}) + C', \end{aligned}$$

where the constant $C' = \sum_{p \in \mathcal{P}} (E_p(\lambda) - E_p(\lambda - 1))^-$ has no impact on the minimizer. We immediately obtain that if \mathcal{C}^* is a minimum s-t cut in \mathcal{G} then $x^{\mathcal{C}^*}$ minimizes the energy E^λ defined by (8).

The interest of such an approach is that the minimum s-t cut of a graph \mathcal{G} can be computed in a finite time using a maximum-flow algorithm such as [BK04]. This time is moreover bounded from above by a polynomial function of the number of nodes $\#\mathcal{V}$ and edges $\#\mathcal{E}$. Empirically, this algorithm behaves like having a near-linear complexity on typical image processing problems (see [BK04]).

To minimize (4) using this graph cut strategy, we need to compute $(L - 1)$ minimum-cuts. Notice however that the nodes and the edges of the graph \mathcal{G} are independent of the level λ . This leads to an algorithmic scheme of worst-case complexity $O(T(\#\mathcal{V}, \#\mathcal{E}) \times L)$ where $T(\#\mathcal{V}, \#\mathcal{E})$ is the complexity of the maximum-flow algorithm. Such a scheme is particularly time consuming when L is large. As observed in [DS06], since (9) holds and given the knowledge of the minimum-cut for some levels, some pixels have a straightforward assignment at other levels. A dyadic scheme exploiting this observation is proposed by the same authors. Its complexity is $O(T(\#\mathcal{V}, \#\mathcal{E}) \times \log_2(L - 1))$. Finally, a slightly faster algorithm is obtained in [CD09] by reusing the maximum-flow at a level λ , to initialize the next level. It dynamically updates the edge weights in the same graph. This is possible since the set of nodes connected to the sink t in the graph \mathcal{G} grows when the level λ increases (see [CD09] for details). The implementation of the graph cuts used in the experiments (see Section 5) uses these algorithmic optimizations.

4 The estimate of the Gaussian parameters

In this section, we provide the details useful for the implementation of a gradient-based algorithm with an Armijo step size rule (see [Ber03]) solving the step 2 of the algorithm of Table 3.

Before giving the formula of the gradient of F , let us remind the notation $c_k = (c_{k,i}, c_{k,j}) \in \mathbb{R}^2$ and $p = (p_i, p_j) \in \mathcal{P}$.

For $v \in \mathbb{R}^{K\mathcal{P}}$, $u \in \mathbb{R}^{\mathcal{P}}$ and $(c_k, w_k)_{1 \leq k \leq K} \in (\mathbb{R}^2 \times \mathbb{R}_+)^K$, we obtain after some calculation

$$\begin{aligned} \frac{\partial F}{\partial c_{k,i}} &= \frac{c_{k,i} - \bar{c}_{k,i}}{\sigma_c^2} + \frac{1}{\sigma^2 w_k^2} \sum_{p \in \mathcal{P}} (p_i - c_{k,i}) e^{-\frac{\|p - c_k\|^2}{2w_k^2}} u_p \left[e^{-\frac{\|p - c_k\|^2}{2w_k^2}} u_p - v_p^k \right], \\ \frac{\partial F}{\partial c_{k,j}} &= \frac{c_{k,j} - \bar{c}_{k,j}}{\sigma_c^2} + \frac{1}{\sigma^2 w_k^2} \sum_{p \in \mathcal{P}} (p_j - c_{k,j}) e^{-\frac{\|p - c_k\|^2}{2w_k^2}} u_p \left[e^{-\frac{\|p - c_k\|^2}{2w_k^2}} u_p - v_p^k \right], \end{aligned}$$

and

$$\frac{\partial F}{\partial w_k} = \frac{w_k - \bar{w}}{\sigma_w^2} + \frac{1}{\sigma^2 w_k^3} \sum_{p \in \mathcal{P}} \|p - c_k\|^2 e^{-\frac{\|p - c_k\|^2}{2w_k^2}} u_p \left[e^{-\frac{\|p - c_k\|^2}{2w_k^2}} u_p - v_p^k \right].$$

The stopping criterion of the gradient-based algorithm controls that the variation of $(\theta_k)_{1 \leq k \leq K}$ between two successive iterations is smaller than a parameter ε_e . To avoid losing too much time during the first iterations of the algorithm of Table 3, we choose to express the level of accuracy ε_e as a decreasing function of the iteration number n and set:

$$\varepsilon_e = \left((\varepsilon_e^{max} - \varepsilon_e^{min}) \exp\left(-\frac{n}{\sigma_{\varepsilon_e}}\right) + \varepsilon_e^{min} \right) \in [\varepsilon_e^{min}, \varepsilon_e^{max}],$$

where the parameters ε_e^{min} , ε_e^{max} and σ_{ε_e} are empirically set (see Section 5). In this way, the estimation of the acquisition parameters is progressively more accurate as n increases. The strength of the decrease is controlled by the parameter $\sigma_{\varepsilon_e} > 0$. Empirically, we found that the form of ε_e has a limited influence on the convergence of the algorithm of Table 3.

5 Numerical experiments

5.1 Applicative framework

The camera is made of an optical system and a typical 256×256 pixels focal plane array (i.e. $N = 256$). As already mentioned, for low signal levels, we consider that the additive Gaussian thermal noise dominates. We consider several noise levels reflecting four possible illumination levels: $\sigma = 0.05, 0.1, 0.2$ and 0.4 ; while the intensities of the ideal images always range in $[0, 1]$.

As already said, for the imaging system considered in our application, the "radius" of the laser dome is $\bar{w} = 16.2$ pixels of the image. The standard deviation around this radius is typically of $\sigma_w = 0.07$ pixels (see Table 1). From laser shot to laser shot, the nominal beam axis is deviated over a regular grid of dimension $K = 9 \times 9$. After atmospheric perturbations, the expected location \bar{c}_k of the beam axis belongs to (in pixels) for any $k \in \{1, \dots, K\}$ is

$$\bar{c}_k \in \left\{ \frac{N}{18}, \frac{N}{18} + \frac{N}{9}, \dots, \frac{N}{18} + 8\frac{N}{9} \right\}^2.$$

The standard deviation around this expected value is typically of $\sigma_c = 0.81$ pixels (see Table 2).

5.2 Implementation details

Except in Figure 3 where the penalty parameter β varies, the parameter β is always set to a value that minimizes the Mean Square Error (MSE)⁴ between the image estimate and the ideal image. This minimization with respect to the parameter β is achieved by a golden section search algorithm [Kie53].

The estimation of the restored image is implemented with the maximum-flow algorithm v3.0 of [BK04]. Additionally, a neighborhood involving 16 neighbors is considered for each pixel p on the lattice \mathcal{P} (see Figure 2). The minimization is implemented with a dyadic parametric scheme and typically represents 10 percents of the overall computational times.

Nevertheless, we want to emphasize that our implementation is not optimized. We therefore do not provide detailed computing times since we believe they are not representative of the computing time for an optimized version of the algorithm of Table 3. In particular, a simple improvement with this regard would consist in extracting from each image v^k a small window containing the laser shot. Also, many computations could be parallelized. With the current implementation, the restoration of an image of size 256×256 from 81 laser shots requires between 1 and 6 minutes on a computer whose processor is clocked at 3.47GHz.

Concerning the stopping criterion for the estimation of the acquisition parameters, we empirically set $\varepsilon_e^{min} = 5 \times 10^{-3}$, $\varepsilon_e^{max} = 0.5$ and $\sigma_{\varepsilon_e} = 2.0$. We set the stopping criterion of the algorithm of Table 3 to $\varepsilon_a = 1$. This provides a good trade-off between time consumption and accuracy. In particular, such a value of ε_a corresponds to an error of one grayscale level for all pixels between two successive image estimates.

⁴The MSE and Peak Signal-to-Noise Ratio (PSNR) measures are both described at <http://megawave.cmla.ens-cachan.fr/stuff/guid3/node256.html#fmse>.

5.3 Measuring the influence of the parameters

This section focuses on the influence of the parameters β and \bar{w} on the quality of the image estimated by the algorithm of Table 3. We consider an intermediate noise level of $\sigma = 0.1$. An example of reconstruction is illustrated in Figure 3 with a varying $\beta = 1, 5, 10, 20$. To ease the visualization of available data, laser shots and illumination domes are each gathered into a single image where a pixel is assigned with its maximum intensity over all $k \in \{1, \dots, K\}$. In order to illustrate the influence of the parameter β , we also set the other parameters in such a way that the center and the width of all illumination domes do essentially not vary. We therefore set $\sigma_c = 10^{-4}$, $\sigma_w = 10^{-4}$ and $\bar{w} = 30$. Due to the particular values of these parameters and the level of accuracy ε_a , we have $w_k \simeq \bar{w}$ and $c_k \simeq \bar{c}_k, \forall k \in \{1, \dots, K\}$.

The strength of the regularization grows with the parameter β . This is consistent with the Equation (4) since this parameter is attached to the regularization term. The parameter β therefore needs to be adequately tuned to remove noise without losing too much details. Also, as expected when the parameter β increases, the image estimate progressively becomes a cartoon-like image with sharp boundaries separating large and flat regions. In particular, textures and thin details tend to disappear. The considered minimization strategy is indeed a TV+ L^2 model, we therefore observe that it suffers from its known drawbacks: staircasing and loss of contrast. While the staircasing effect can be reduced by using larger neighborhoods \mathcal{N} (at the expense of a larger computational cost), we believe that a Bregman iteration strategy could be adopted to improve the contrast of the image estimate. Notice that, whatever the value of β , we do not see on the restored image the dark stripes at poorly illuminated pixels.

Finally, the operator $M_{(\theta_k)_{1 \leq k \leq K}}$ is linear and goes from $\mathbb{R}^{\mathcal{P}}$ to $\mathbb{R}^{K\mathcal{P}}$. It could be represented by a $KN^2 \times N^2$ matrix. Its conditioning and therefore the difficulty of the considered inverse problem is characterized by the singular values of this matrix. Typically, singular values that decay rapidly correspond to difficult inverse problems. Notice that when \bar{w} is small the singular values of the matrix corresponding to $M_{(\theta_k)_{1 \leq k \leq K}}$ are approximately equal to K repetitions of the largest values of the Gaussian (1) when p varies. Therefore, they decay more rapidly when \bar{w} is small. As a consequence, the conditioning of the operator $M_{(\theta_k)_{1 \leq k \leq K}}$ is less favorable to the restoration of the image. The Figure 4 contains results when the parameter \bar{w} varies. The partial available data through laser shots is represented on the top row, in the same way as in Figure 3. The obtained results are depicted for $\bar{w} = 6, 9, 12, 20$. To account for the difficulty of the problem, we also provide

$$\min_{p \in \mathcal{P}} \max_{1 \leq k \leq K} G_{\theta_k}(p), \quad (14)$$

for each value of \bar{w} . The parameters σ_c , σ_w , and σ are set with the same values as in the Figure 3. As expected, when \bar{w} is small, the reconstruction is of poor quality and details of the image cannot be accurately recovered between illumination domes. Indeed, for such pixels, the data is lost and their intensities are assigned by regularization. Nevertheless, the results show that a better quality reconstruction can be obtained when \bar{w} is larger. At $\bar{w} = 12$, the reconstructed image still contains most details.

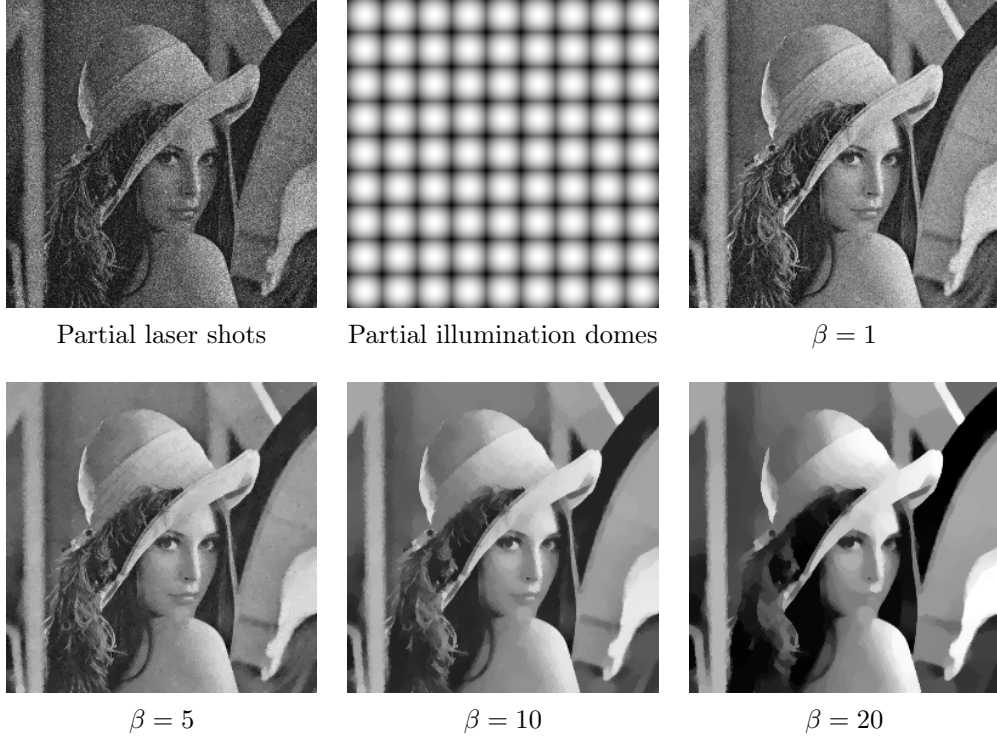


Figure 3: Reconstruction with a noise level $\sigma = 0.1$. The remaining parameters are set with $\sigma_c = 10^{-4}$, $\sigma_w = 10^{-4}$ and $\bar{w} = 30$. Each pixel of the top left and middle top images is assigned with its maximum intensity over all $k \in \{1, \dots, K\}$. Then we present the results for different β .




Figure 4: Influence of the parameter \bar{w} on the reconstruction with a noise level $\sigma = 0.1$. The remaining parameters are set to $\sigma_c = 10^{-4}$ and $\sigma_w = 10^{-4}$. On top row, each pixel of the images is assigned with its maximum intensity over all laser shots $k \in \{1, \dots, K\}$. On the bottom row, we provide the results of the algorithm of Table 3. As a measure of the difficulty of the problem, we also provide, next to \bar{w} , the quantity (14) within parenthesis.

5.4 Convergence of $(\theta_k)_{1 \leq k \leq K}$

In Figure 5, we empirically illustrate the behavior of the algorithm of Table 3 for two levels of noise and eight grayscale images. The remaining parameters are set as follows: $\bar{w} = 16.2$, $\sigma_c = 0.81$ and $\sigma_w = 0.07$. For each noise level, we measure the distance between the image estimate u^n obtained at iteration n (see Table 3) and the ideal image u^* with

$$\|u^n - u^*\|. \quad (15)$$

To illustrate the difficulty of the restoration problem as a function of the noise level, we also provide the PSNR between the noisy and the noise-free laser shots.

As expected, we see that the distance (15) strongly decreases for all images in the very first iterations and becomes relatively stable in the following ones. This means that the image estimate is mainly improved during the first iterations of the algorithm of Table 3. In particular, we see that the number of iterations of the latter could be reduced by slightly relaxing the parameter ε_a . As the noise level σ increases, we also observe that the average number of iterations becomes smaller and the decrease of (15) is less important. According to our experiments, when the standard deviation of the noise is larger than 0.1 the noise dominates the error induced by the inaccuracy of the parameters of the Gaussian. As a consequence, the estimation of these parameters is not very useful. Notice finally that the estimation of the acquisition parameters is fairly robust to the observed scene.

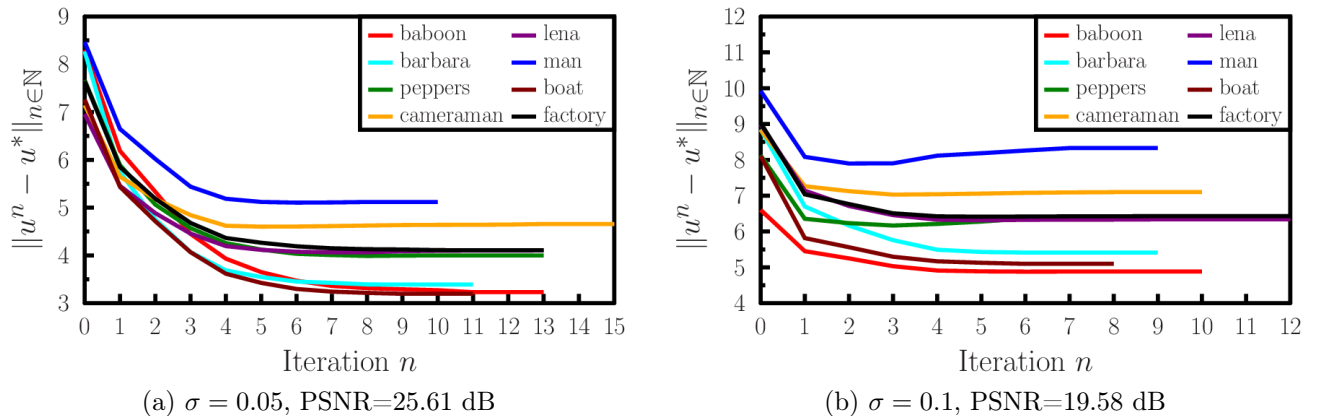


Figure 5: Convergence of the algorithm of Table 3 for two noise levels. On each figure, the distance between the image estimate u^n and the true image u^* is represented as a function of the iteration n . The PSNR between the noisy and the noise-free laser shots is also indicated as a measure of the difficulty of the problem. The remaining parameters are set as follows: $\bar{w} = 16.2$, $\sigma_c = 0.81$ and $\sigma_w = 0.07$.

5.5 Image resolution

Our instrument includes a reception optical system. We cannot expect to go beyond its resolution. However, depending on the noise level and the regularization induced by the restoration algorithm, we might obtain a final image of a much lower resolution. This section aims at evaluating empirically the resolution of our instrument. In particular, we would like to evaluate if all the frequencies

that can be captured by the reception optical system are properly captured by the mosaic active imaging system. The motivation for considering frequencies comes from the usual sampling theory and Shannon-Nyquist theorem. In order to do so, we simulate the acquisition of several simple targets made of pure cosine functions. Notice that this choice is also motivated by the fact that the mosaic active imaging system is not translation-invariant. For instance, the spoke target usually used to evaluate a blur level is not relevant for the non-linear, non translation-invariant acquisition system investigated in this paper.

Therefore, we propose to restore a sequence of $M = 25$ cosine patterns of increasing frequencies. We define these frequencies using a linear progression and set, for any $m \in \{2, \dots, M\}$,

$$(k_m, l_m) = m(k_1, l_1), \tag{16}$$

where $(k_1, l_1) = (5, 2)$.

Moreover, for any $m \in \{1, \dots, M\}$, the pixels $(p_i, p_j) \in \mathcal{P}$ of the related cosine pattern are assigned with

$$\frac{1}{2} \left(1 + \cos \left(2\pi \left(\frac{k_m \times p_i + l_m \times p_j}{N} \right) \right) \right) \in [0, 1],$$

where $(k_m, l_m) \in \{-\frac{N}{2} + 1, \dots, \frac{N}{2}\}^2$ is defined in (16).

Also, the illumination domes are generated according to the applicative context: We set $\sigma_c = 0.81$, $\sigma_w = 0.07$ and $\bar{w} = 16.2$. We consider an intermediate noise level: $\sigma = 0.1$. All the degraded cosines are restored using the same illumination domes and the same noise realization. For these illumination domes, we have $\min_{p \in \mathcal{P}} \max_{1 \leq k \leq K} G_{\theta_k}(p) = 0.43$. We distinguish the well and poorly illuminated areas with the criterion

$$\begin{cases} p \text{ is well illuminated} & \text{if } (\max_{1 \leq k \leq K} G_{\theta_k}(p)) \geq \gamma \\ p \text{ is poorly illuminated} & \text{otherwise,} \end{cases}$$

where the threshold $\gamma = 0.65$. Reconstruction results for two distinct frequencies are reported in Figure 6. On all images, we superimpose in green the boundary separating well and poorly illuminated pixels. Despite the important amount of noise, one can see that the cosine patterns are well preserved, even for high frequencies.

In Figure 7, we represent the MSE and PSNR between the ideal and the estimated cosine patterns, as a function of the frequency index $m \in \{1, \dots, M\}$. We also distinguish the performance of the algorithm of Table 3 in well and poorly illuminated areas. These results show that the quality of the image estimate diminishes as the frequency of the cosine increases but remains almost the same for $m \geq 8$. As expected, these results also show that the algorithm described in Table 3 behaves slightly better on well illuminated areas than on poorly ones. However, even for high frequencies, in poorly illuminated areas, the final error is comparable to the error due to the noise with a perfectly illuminated scene.

This suggest that, in the context studied in this paper, the resolution of the mosaic active imaging device is similar to the resolution of the reception optical system.




Figure 6: Reconstruction of pure cosine patterns (see (16)) for $m = 1$ (upper row) and $m = 25$ (lower row). The first column corresponds to the partial laser shots where each pixel is assigned with its maximum intensity over all $k \in \{1, \dots, K\}$. The second column contains a detail extracted from the images of the first column. The third and fourth columns respectively contain details extracted from the image estimate and the ideal image. Pixels in green are on the boundary between poorly and well illuminated pixels.




Figure 7: MSE and PSNR between the ideal and estimated cosine patterns for increasing frequencies. The MSE and PSNR are for the whole image and its restriction to the poorly and well illuminated pixels. The frequencies (k_m, l_m) are given in (16).

5.6 Accuracy

In this section, we study the quality of the image estimate as well as the performance of the algorithm of Table 3 with $\sigma = 0.05, 0.1, 0.2, 0.4$ and on the same images as in Section 5.4. We remind that the intensities range between 0 and 1 in all the ideal images and that a noise of standard deviation 0.4 is a very strong noise. Let us first briefly describe the experimental setting. First, for each image and level of noise σ , we independently generate 10 laser shots and illumination domes. Next, we restore each of them using the algorithm of Table 3 and measure the error between the restored image and the ideal image using two metrics: PSNR and MSE. We then compute the mean and standard deviation of these quantities.

We remind that the parameter β is automatically tuned by minimizing the MSE between the image estimate and the true image. As already indicated in Section 5.1, the remaining parameters are set according to the applicative context: $\sigma_c = 0.81$, $\sigma_w = 0.07$ and $\bar{w} = 16.2$. The results of these experiments are reported in Tables 4 and 5 and illustrated in Figures 8, 9, 10 and 11 for $\sigma = 0.05, 0.1, 0.2, 0.4$, respectively. For the sake of clarity, we have also split the results according to the amount of noise: moderate ($\sigma = 0.05$ and $\sigma = 0.1$, see Table 4) and severe ($\sigma = 0.2$ and $\sigma = 0.4$, see Table 5). For each image, we provide the ideal one, the image estimate as well as the partial laser shots to account for the difficulty of the problem. Let us now analyze the obtained results.

Under a moderate noise level, the algorithm of Table 3 behaves well: large flat areas are well denoised; thin structures and textures are well preserved even in poorly illuminated (see e.g. the “barbara” image and the “factory” image in Figure 9). The latter point is important and is due to the fact that the illumination domes are not too far from each other in the targeted application. Under severe noise levels, large flat areas are still well smoothed (see e.g. the “cameraman” image in Figure 10) but textures disappear in the residues (see e.g. the “man” image in Figure 11). These observations are also confirmed by the increase of the MSE and the decrease of the PSNR for all images in Tables 4 and 5. Notice that restoring images with such large levels of noise is very challenging since almost nothing is visible in the partial available data images. In such situations, the algorithm of Table 3 enables the distinction of the main contours of the images. Finally, the small standard deviations on MSE and PSNR measures depict a good robustness of the proposed approach, whatever the selected noise level σ .

6 Conclusion

In this paper, we address the problem of image restoration in mosaic active imaging using a simplified forward model. To solve this model, we propose a two-stage iterative process alternating between (i) the estimation of the restored image using a graph cuts-based algorithm and (ii) the estimation of the acquisition parameters using a gradient-based algorithm. The numerical results show that the proposed restoration algorithm quickly converges towards an image estimate of good quality, even under large noise levels. This validates the mosaic active imaging strategy.

Several ways are currently under investigation for improving the results. First, we should improve the forward model to better take into account the noise, the sampling and the speckle caused by the turbulence. Ideas from [MG12], could, for instance, be adapted to our framework. Taking into

Noise level	Image	MSE	PSNR (dB)
$\sigma = 0.05$	baboon	$1.89 \times 10^{-3} \pm 1.00 \times 10^{-5}$	$26.57 \pm 1.69 \times 10^{-2}$
	barbara	$1.44 \times 10^{-3} \pm 1.73 \times 10^{-5}$	$27.81 \pm 4.90 \times 10^{-2}$
	peppers	$8.57 \times 10^{-4} \pm 2.13 \times 10^{-11}$	$30.37 \pm 2.27 \times 10^{-2}$
	cameraman	$8.54 \times 10^{-4} \pm 1.00 \times 10^{-5}$	$30.53 \pm 4.26 \times 10^{-2}$
	lena	$8.76 \times 10^{-4} \pm 1.00 \times 10^{-5}$	$30.13 \pm 3.80 \times 10^{-2}$
	man	$1.39 \times 10^{-3} \pm 1.00 \times 10^{-5}$	$28.56 \pm 2.83 \times 10^{-2}$
	boat	$1.20 \times 10^{-3} \pm 1.00 \times 10^{-5}$	$29.18 \pm 3.40 \times 10^{-2}$
	factory	$1.18 \times 10^{-3} \pm 1.00 \times 10^{-5}$	$29.16 \pm 3.32 \times 10^{-2}$
$\sigma = 0.1$	baboon	$4.93 \times 10^{-3} \pm 2.83 \times 10^{-5}$	$22.39 \pm 2.41 \times 10^{-2}$
	barbara	$3.63 \times 10^{-3} \pm 4.00 \times 10^{-5}$	$23.80 \pm 4.71 \times 10^{-2}$
	peppers	$1.92 \times 10^{-3} \pm 1.41 \times 10^{-5}$	$26.89 \pm 3.37 \times 10^{-2}$
	cameraman	$1.98 \times 10^{-3} \pm 1.73 \times 10^{-5}$	$26.88 \pm 3.52 \times 10^{-2}$
	lena	$1.93 \times 10^{-3} \pm 1.73 \times 10^{-5}$	$26.65 \pm 3.85 \times 10^{-2}$
	man	$3.26 \times 10^{-3} \pm 2.24 \times 10^{-5}$	$24.85 \pm 3.02 \times 10^{-2}$
	boat	$2.28 \times 10^{-3} \pm 2.24 \times 10^{-5}$	$25.59 \pm 3.43 \times 10^{-2}$
	factory	$2.60 \times 10^{-3} \pm 2.83 \times 10^{-5}$	$25.71 \pm 4.59 \times 10^{-2}$

Table 4: Accuracy of the algorithm of Table 3 under a moderate level of noise $\sigma = 0.05$ and $\sigma = 0.1$. The remaining parameters are set as follows: $\sigma_c = 0.81$, $\sigma_w = 0.07$ and $\bar{w} = 16.2$. The MSE and PSNR measures are calculated over 10 runs and rounded to the nearest value.

Noise level	Image	MSE	PSNR (dB)
$\sigma = 0.2$	baboon	$9.25 \times 10^{-3} \pm 2.83 \times 10^{-5}$	$19.66 \pm 1.36 \times 10^{-2}$
	barbara	$6.74 \times 10^{-3} \pm 4.90 \times 10^{-5}$	$21.11 \pm 3.15 \times 10^{-2}$
	peppers	$3.95 \times 10^{-3} \pm 5.39 \times 10^{-5}$	$23.73 \pm 5.86 \times 10^{-2}$
	cameraman	$4.07 \times 10^{-3} \pm 7.07 \times 10^{-5}$	$23.75 \pm 7.46 \times 10^{-2}$
	lena	$3.81 \times 10^{-3} \pm 4.80 \times 10^{-5}$	$23.75 \pm 5.46 \times 10^{-2}$
	man	$6.36 \times 10^{-3} \pm 5.00 \times 10^{-5}$	$21.95 \pm 3.40 \times 10^{-2}$
	boat	$5.34 \times 10^{-3} \pm 6.78 \times 10^{-5}$	$22.71 \pm 5.48 \times 10^{-2}$
	factory	$4.91 \times 10^{-3} \pm 3.32 \times 10^{-5}$	$22.96 \pm 2.91 \times 10^{-2}$
$\sigma = 0.4$	baboon	$1.26 \times 10^{-2} \pm 1.06 \times 10^{-4}$	$18.30 \pm 3.65 \times 10^{-2}$
	barbara	$9.51 \times 10^{-3} \pm 1.07 \times 10^{-4}$	$19.62 \pm 4.88 \times 10^{-2}$
	peppers	$7.70 \times 10^{-3} \pm 1.41 \times 10^{-4}$	$20.84 \pm 7.97 \times 10^{-2}$
	cameraman	$7.30 \times 10^{-3} \pm 1.30 \times 10^{-4}$	$21.21 \pm 7.76 \times 10^{-2}$
	lena	$6.76 \times 10^{-3} \pm 9.43 \times 10^{-5}$	$21.25 \pm 6.08 \times 10^{-2}$
	man	$1.06 \times 10^{-2} \pm 1.22 \times 10^{-4}$	$19.74 \pm 5.00 \times 10^{-2}$
	boat	$8.67 \times 10^{-3} \pm 9.54 \times 10^{-5}$	$20.60 \pm 4.80 \times 10^{-2}$
	factory	$7.91 \times 10^{-3} \pm 1.08 \times 10^{-4}$	$20.88 \pm 5.92 \times 10^{-2}$

Table 5: Accuracy of the algorithm of Table 3 under a severe level of noise $\sigma = 0.2$ and $\sigma = 0.4$. The remaining parameters are set as follows: $\sigma_c = 0.81$, $\sigma_w = 0.07$ and $\bar{w} = 16.2$. The MSE and PSNR measures are calculated over 10 runs and rounded to the nearest value.

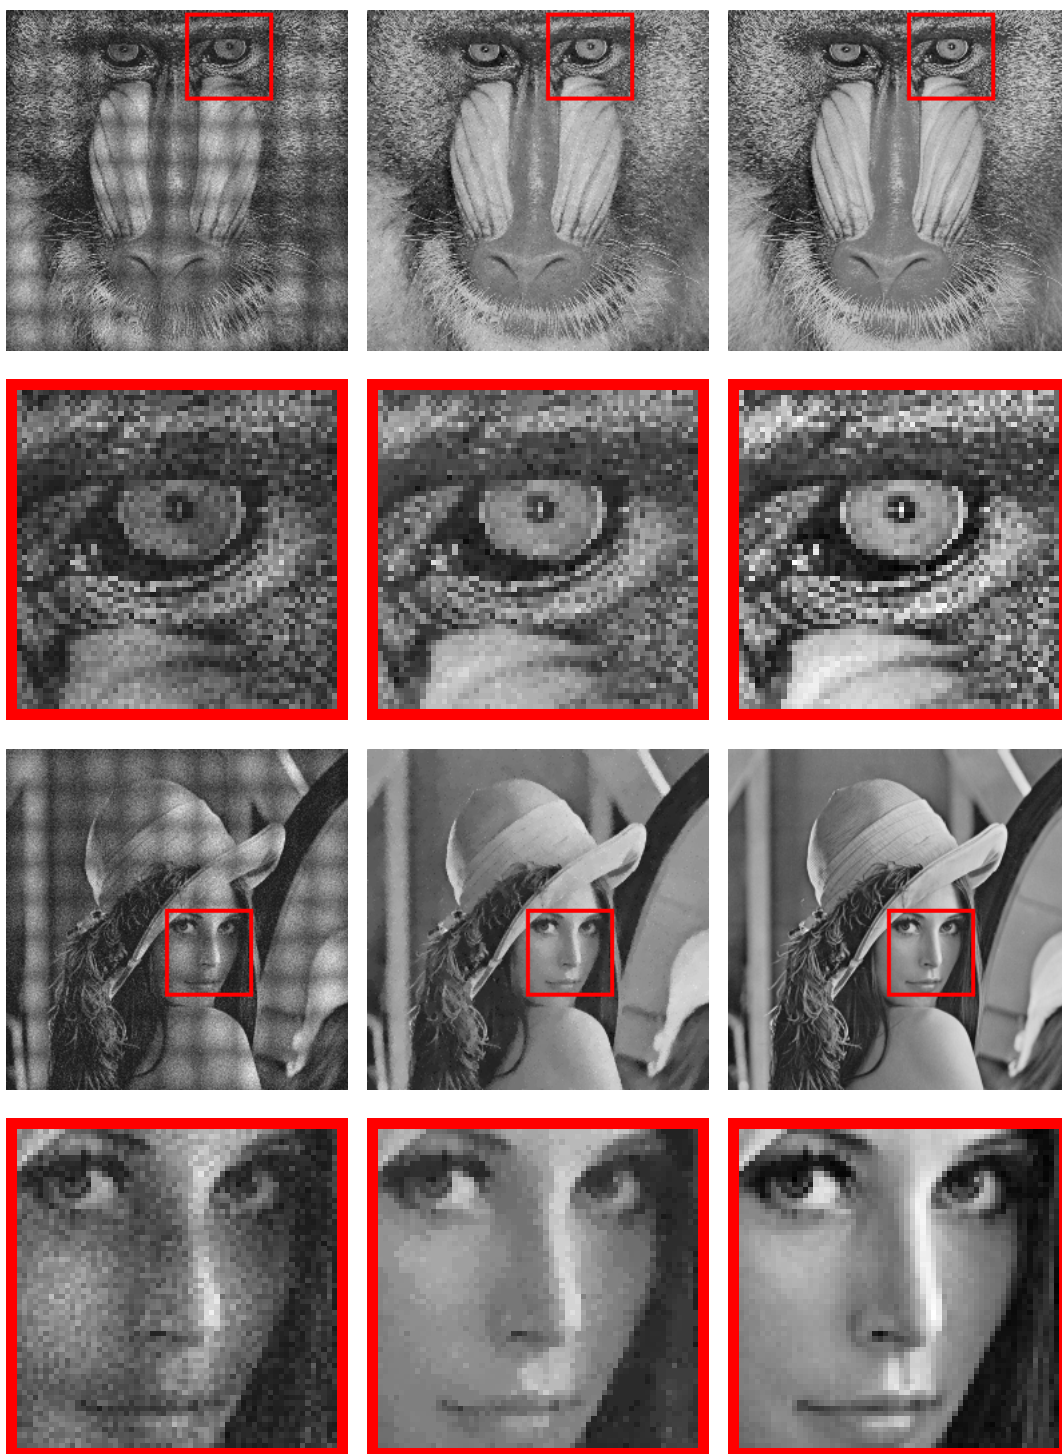


Figure 8: Reconstruction of the images “baboon” (first and second rows) and “lena” (third and fourth rows) with a noise level $\sigma = 0.05$. The remaining parameters are set as follows: $\bar{w} = 16.2$, $\sigma_c = 0.81$ and $\sigma_w = 0.07$. The left column correspond to partial available data where each pixel is assigned its maximum intensity over all $k \in \{1, \dots, K\}$. The middle and right columns correspond resp. to the image estimate and the ideal image. Detailed views of all these images are also provided on the second and fourth rows.



Figure 9: Reconstruction of the images “barbara” (first and second rows) and “factory” (third and fourth rows) with a noise level $\sigma = 0.1$. The remaining parameters are set as follows: $\bar{w} = 16.2$, $\sigma_c = 0.81$ and $\sigma_w = 0.07$. The left column correspond to partial available data where each pixel is assigned its maximum intensity over all $k \in \{1, \dots, K\}$. The middle and right columns correspond resp. to the image estimate and the ideal image. Detailed views of all these images are also provided on the second and fourth rows.



Figure 10: Reconstruction of the images “cameraman” (first and second rows) and “boat” (third and fourth rows) with a noise level $\sigma = 0.2$. The remaining parameters are set as follows: $\bar{w} = 16.2$, $\sigma_c = 0.81$ and $\sigma_w = 0.07$. The left column correspond to partial available data where each pixel is assigned its maximum intensity over all $k \in \{1, \dots, K\}$. The middle and right columns correspond resp. to the image estimate and the ideal image. Detailed views of all these images are also provided on the second and fourth rows.




Figure 11: Reconstruction of the images “peppers” (first and second rows) and “man” (third and fourth rows) with a noise level $\sigma = 0.4$. The remaining parameters are set as follows: $\bar{w} = 16.2$, $\sigma_c = 0.81$ and $\sigma_w = 0.07$. The left column correspond to partial available data where each pixel is assigned its maximum intensity over all $k \in \{1, \dots, K\}$. The middle and right columns correspond resp. to the image estimate and the ideal image. Detailed views of all these images are also provided on the second and fourth rows.

account the Poisson noise (together with the Gaussian noise) is another possible way to further improve the quality of the reconstructed images in the case of good lightning conditions. The removal of a Poisson-Gaussian noise increases the difficulty of the problem since the Poisson noise is signal-dependent. A common and efficient approach to tackle this problem is to proceed as follows [MF13]: (i) apply a nonlinear transformation to the image to make the noise approximately Gaussian with unitary standard deviation; (ii) compute the image estimate with an appropriate restoration algorithm under the Gaussian noise assumption; (iii) apply the inverse transform of the step (i). Clearly, this procedure is independent of the restoration algorithm and could be used without much efforts. Ideas from [LBU11] could also be adapted to our situation.

Other interesting ideas could be also investigated. For instance, one might want to improve the proposed model by taking into account the situations where a different contrast occur between laser shots. This could be easily managed by embedding multiplicative factors in Gaussian profiles and assuming an add-hoc distribution on them. Also, more efficient optimization methods could be used for estimating acquisition parameters. Indeed, it is not difficult to see that the Hessian of F is block diagonal. Also, we have empirically observed that the diagonal of this Hessian is dominant. Thus, second-order methods could potentially offer faster convergence.

Finally, we plan to evaluate the proposed approach on real acquisitions made with a mosaic active imaging device.

References

- [BBFAC04] J. Bect, L. Blanc-Féraud, G. Aubert, and A. Chambolle. A L1-unified variational framework for image restoration. *Lecture Notes in Computer Science*, 3024:1–13, 2004. Proceedings of European Conference on Computer Vision 2004.
- [Ber03] D.P. Bertsekas. *Nonlinear Programming*. Athena Scientific, second edition, 2003.
- [BK04] Y. Boykov and V. Kolmogorov. An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *Pattern Analysis And Machine Intelligence*, 26(9):1124–1137, 2004.
- [CD09] A. Chambolle and J. Darbon. On total variation minimization and surface evolution using parametric maximum flows. *International Journal of Computer Vision*, 84(3):288–307, 2009.
- [Cha92] F. Chassat. *Optical propagation through atmospheric turbulence: moral study and application of anisoplanatism in adaptive optics*. PhD thesis, University of Paris Sud, 1992.
- [Cha04] A. Chambolle. An algorithm for total variation minimization and applications. *Journal of Mathematical Imaging and Vision*, 20(1-2):89–97, January-March 2004.
- [Cha05] A. Chambolle. Total variation minimization and a class of binary MRF models. In *Lecture Notes in Computer Science, Energy Minimization Methods in Computer Vision and Pattern Recognition*, volume 3757, pages 136–152. Springer, 2005.

- [CP11] Antonin Chambolle and Thomas Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of Mathematical Imaging and Vision*, 40(1):120–145, 2011.
- [CS00] R. R. Coifman and A. Sowa. Combining the calculus of variations and wavelets for image enhancement. *Applied and computational harmonic analysis*, 9(1):1–18, 2000.
- [CS05] Tony Chan and Jackie Shen. Variational image inpainting. *Communications in Pure and applied Math.*, 58(3):579–619, 2005.
- [Dem89] G. Demoment. Image reconstruction and restoration: Overview of common estimation structures and problems. *IEEE, Transactions on Acoustics, Speech and Signal Processing*, pages 2024–2036, 1989.
- [DS06] J. Darbon and M. Sigelle. Image restoration with discrete constrained total variation part I: Fast and exact optimization. *Journal of Mathematical Imaging and Vision*, 26(3):261–276, 2006.
- [Ela10] M. Elad. *Sparse and redundant representations: From theory to applications in signal and image processing*. Springer, 2010.
- [Fan75] R.L. Fante. Electromagnetic beam propagation in turbulent media. *Proceedings of the IEEE*, 63(12):1669–1692, December 1975.
- [Fri66] D.L. Fried. Optical resolution through a randomly inhomogeneous medium for very long and very short exposures. *Journal of the Optical Society of America*, 56(10):1372–1379, October 1966.
- [GPS89] D.M. Greig, B.T. Porteous, and A.H. Seheult. Exact maximum a posteriori estimation for binary images. *Journal of the Royal Statistical Society*, 51(2):271–279, 1989.
- [Ham10] D. Hamoir. Procédé et système d’imagerie active à champ large : Method and system for active imaging with a large field. Patent WO 2010119225, October 2010.
- [HVB⁺09] L. Hespel, M.-T. Velluet, A. Bonnefois, N. Rivière, M. Fracès, D. Hamoir, B. Tanguy, B. Duchenne, and J. Isbert. Comparison of a physics-based BIL simulator with experiments. In *Society of Photo-Optical Instrumentation Engineers, International Symposium on Photoelectronic Detection and Imaging*, pages 73822T–73822T–10, 2009.
- [Kie53] J. Kiefer. Sequential minimax search for a maximum. *Proceedings of the American Mathematical Society*, 4(3):502–506, 1953.
- [Kol49] A.N. Kolmogorov. The local structure of turbulence in incompressible viscous fluid for very large Reynolds numbers. *Doklady Akademii Nauk SSSR*, 66(4):825, 1949. English translation in "Turbulence: Classic papers on statistical theory", edited by S. K. Friedlander and L. Topper (Interscience Publishers, Inc., New York, 1961), p. 151, 1961.
- [KZ04] V. Kolmogorov and R. Zabih. What energy functions can be minimized via graph cuts? *Pattern Analysis And Machine Intelligence*, 26(2):147–159, 2004.

- [LBU11] F. Luisier, T. Blu, and M. Unser. Image denoising in mixed Poisson-Gaussian noise. *IEEE Transactions on Image Processing*, 20(3):696–708, 2011.
- [Mey01] Y. Meyer. *Oscillating patterns in image processing and in some nonlinear evolution equation*. AMS, Boston, MA, USA, 2001. The Fifteenth Dean Jacqueline B. Lewis Memorial Lectures.
- [MF13] M. Mäkitalo and A. Foi. Optimal inversion of the generalized Anscombe transformation for Poisson-Gaussian noise. *IEEE Transactions on Image Processing*, 22(1):91–103, 2013.
- [MG01] F. Malgouyres and F. Guichard. Edge direction preserving image zooming: a mathematical and numerical analysis. *SIAM, Journal on Numerical Analysis*, 39(1):1–37, 2001.
- [MG12] Y. Mao and J. Gilles. Non rigid geometric distortions correction - application to atmospheric turbulence stabilization. *Inverse Problems and Imaging Journal*, 6(3):531–546, August 2012.
- [Nik00] M. Nikolova. Local strong homogeneity of a regularized estimator. *SIAM, Journal of Applied Mathematics*, 61(2):633–658, 2000.
- [PR75] J.C. Picard and H.D. Ratliff. Minimum cuts and related problems. *Networks Journal*, 5(4):357–370, 1975.
- [RHV⁺10] N. Rivière, L. Hespel, M.-T. Velluet, Y.-M. Frédéric, P. Barillot, and F. Hélias. Modeling of an active burst illumination imaging system: comparison between experimental and modelled 3d scene. In *Society of Photo-Optical Instrumentation Engineers, International Symposium on Photoelectronic Detection and Imaging*, volume 7382, pages 783509–783509–11, 2010.
- [Rod81] F. Roddier. The effects of atmospheric turbulence in optical astronomy. *Progress in Optics*, 19:281–376, 1981.
- [ROF92] L.I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60:259–268, 1992.