

Robust image reconstruction from multi-view measurements

Gilles Puy and Pierre Vandergheynst

Abstract. We propose a novel method to accurately reconstruct a set of images representing a single scene from few linear multi-view measurements. Each observed image is modeled as the sum of a background image and a foreground one. The background image is common to all observed images but undergoes geometric transformations, as the scene is observed from different viewpoints. In this paper, we assume that these geometric transformations are represented by a few parameters, e.g., translations, rotations, affine transformations, etc.. The foreground images differ from one observed image to another, and are used to model possible occlusions of the scene. The proposed reconstruction algorithm estimates jointly the images and the transformation parameters from the available multi-view measurements. The ideal solution of this multi-view imaging problem minimizes a non-convex functional, and the reconstruction technique is an alternating descent method built to minimize this functional. The convergence of the proposed algorithm is studied, and conditions under which the sequence of estimated images and parameters converges to a critical point of the non-convex functional are provided. Finally, the efficiency of the algorithm is demonstrated using numerical simulations for applications such as compressed sensing or super-resolution.

Keywords. compressed sensing, inverse problem, non-convex optimization, super-resolution, robust image alignment.

1. Introduction

Multi-view imaging has become more and more popular in the signal processing community during the past years with, for example, the design of camera arrays [1] and the rise of new applications such as 3D reconstruction of a scene [20].

In most multi-view imaging applications, the pose of the cameras, or of the sensing systems, is estimated by pre-calibration or by using side-information. In this paper, we study the problem of reconstructing a set of images, representing a single scene, from few measurements made at different viewpoints in absence of (or with inaccurate) prior pose estimation. This situation may occur in different scenarios.

For example, let us imagine that we have a camera that can be sent up to take pictures of the ground and that we are interested in obtaining higher resolution images from these acquisitions. Without more knowledge on the acquisition strategy, we can treat each image independently using a single-frame super-resolution technique. However, if we know that the duration between the capture of two images is not too long, it is highly probable that parts of the scene visible in one frame remain visible in the subsequent ones. Consequently, we can surely benefit from the high inter-correlation between observations to obtain better super-resolved images by jointly reconstructing several frames together, as in, e.g., [17, 43]. The inter-correlation may be modeled using geometric transformations, which depend on the camera pose, that register the images with respect to each other. In the absence of any side information, these transformations have to be estimated along with the high resolution images during the reconstruction process.

As another example, let us consider the problem of designing a similar camera as above but with strong power consumption constraints. To design such an energy-efficient system, compressed sensing (CS) is a powerful tool [12, 22]. Indeed, this theory states that sparse signals can be sampled with just a few linear and non-adaptive measurements. In scenarios where signals related to common phenomena are acquired, it is actually possible to reduce even further the number of measurements by jointly reconstructing the signals with a method exploiting the inter-correlation between them [6, 16]. A technique, as in [23, 44], able to estimate both the geometric transformations between the observed

This work was partly funded by the Hasler Foundation (project number 12080).

G. Puy and P. Vandergheynst are with the Institute of Electrical Engineering, Ecole Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland. E-mail: gilles.puy@epfl.ch and pierre.vandergheynst@epfl.ch.

images and the images themselves from the compressed measurements can thus help us to meet our power consumption constraints.

In the present work, we propose a novel method that estimates jointly a set of correlated images and the geometric transformations that align these images on each other. This technique also robustly handles the appearance of new objects in the scene and can be applied in different settings, such as the ones described above.

1.1. Problem formulation

Let $\mathbf{u} = (u_1, u_2) \in \mathbb{R}^2$ represent the Cartesian coordinates of a point on the Euclidean plane, and $G = \{\mathbf{u}_k\}_{1 \leq k \leq n}$ be a square uniform grid of n points. We model continuous images as functions $x: \mathbb{R}^2 \rightarrow \mathbb{R}$ living in a space $\{x(\mathbf{u}) = \sum_{i=1}^n x_i \varphi(\mathbf{u} - \mathbf{u}_i), x_i \in \mathbb{R}, \mathbf{u}_i \in G, i = 1, \dots, n\}$, where $\varphi: \mathbb{R}^2 \rightarrow \mathbb{R}$ is a generating function [42]. For simplicity, we suppose that $\varphi(\mathbf{u}_k - \mathbf{u}_i) = \delta(\mathbf{u}_k - \mathbf{u}_i)$ for all $\mathbf{u}_k, \mathbf{u}_i \in G$, so that the discrete image $\mathbf{x} \in \mathbb{R}^n$ obtained by sampling x on the grid G has samples x_1, \dots, x_n .

In our setting, several observers provide different observations $\mathbf{y}_1, \dots, \mathbf{y}_l \in \mathbb{R}^m$, with $m \leq n$, of a scene from different viewpoints. As a first approximation, one can consider that these observations describe a single image x_0 , referred as the *background* image hereafter. This image is “viewed” from different perspectives and thus undergoes geometric transformations. We consider here that these transformations are unknown and need to be estimated. However, we restrict ourselves to global transformations represented by few parameters, such as translations, rotations, or affine transformations. For each observer j , with $j = 1, \dots, l$, these transformations are modeled using a function $\tau_{\theta_j}: \mathbb{R}^2 \rightarrow \mathbb{R}^2$, depending on q parameters $\theta_j \in \mathbb{R}^q$, that maps the coordinates \mathbf{u} into new coordinates $\tau_{\theta_j}(\mathbf{u})$. The background image transformed by τ_{θ_j} is thus $x_0 \circ \tau_{\theta_j}$.

To complete the model, we also take into account possible occlusions of the scene. These occlusions may obviously be different from one observer to another. We model them here using l *foreground* images x_1, \dots, x_l and write the j^{th} observed image as $x_0 \circ \tau_{\theta_j} + x_j$.

Finally, we assume that the observations \mathbf{y}_j are obtained by linear projection of $x_0 \circ \tau_{\theta_j} + x_j$ onto m known functions $a_1^j, \dots, a_m^j: \mathbb{R}^2 \rightarrow \mathbb{R}$. These functions model the acquisition system. They can represent a blurring operator for a deconvolution problem, or random waveforms in a compressed sensing setting. The i^{th} entry of \mathbf{y}_j satisfies

$$\begin{aligned} y_{ij} &= \int_{\mathbb{R}^2} (x_0 \circ \tau_{\theta_j}(\mathbf{u}) + x_j(\mathbf{u})) a_i^j(\mathbf{u}) \, d\mathbf{u} \\ &= \int_{\mathbb{R}^2} \sum_{k=1}^n (x_k^0 \varphi(\tau_{\theta_j}(\mathbf{u}) - \mathbf{u}_k) + x_k^j \varphi(\mathbf{u} - \mathbf{u}_k)) a_i^j(\mathbf{u}) \, d\mathbf{u}, \end{aligned}$$

where x_k^0 and x_k^j , with $k = 1, \dots, n$, are the samples describing x_0 and x_j respectively. To facilitate the implementation of this observation model, we assume that the grid G has a sufficiently high resolution so that the integral above is “well” approximated by its Riemann sum on this grid. In the following, we consider that

$$y_{ij} = \sum_{k'=1}^n \sum_{k=1}^n x_k^0 \varphi(\tau_{\theta_j}(\mathbf{u}_{k'}) - \mathbf{u}_k) a_i^j(\mathbf{u}_{k'}) + \sum_{k=1}^n x_k^j a_i^j(\mathbf{u}_k),$$

where we supposed that the pixels of the grid have a unit square area. Therefore, the measurements vector \mathbf{y}_j is equal to $\mathbf{A}_j (\mathbf{S}(\theta_j) \mathbf{x}_0 + \mathbf{x}_j)$, where $\mathbf{A}_j = (a_i^j(\mathbf{u}_k))_{ik} \in \mathbb{R}^{m \times n}$, and $\mathbf{S}: \mathbb{R}^q \rightarrow \mathbb{R}^{n \times n}$ satisfies

$$(1.1) \quad \mathbf{S}(\theta_j) = \begin{bmatrix} \varphi(\tau_{\theta_j}(\mathbf{u}_1) - \mathbf{u}_1) & \dots & \varphi(\tau_{\theta_j}(\mathbf{u}_1) - \mathbf{u}_n) \\ \vdots & & \vdots \\ \varphi(\tau_{\theta_j}(\mathbf{u}_n) - \mathbf{u}_1) & \dots & \varphi(\tau_{\theta_j}(\mathbf{u}_n) - \mathbf{u}_n) \end{bmatrix}.$$

Concatenating all the observations in a single vector, we have

$$(1.2) \quad \begin{bmatrix} \mathbf{y}_1 \\ \vdots \\ \mathbf{y}_l \end{bmatrix} = \begin{bmatrix} \mathbf{A}_1 \mathcal{S}(\boldsymbol{\theta}_1) & \mathbf{A}_1 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{A}_l \mathcal{S}(\boldsymbol{\theta}_l) & 0 & \dots & \mathbf{A}_l \end{bmatrix} \begin{bmatrix} \mathbf{x}_0 \\ \vdots \\ \mathbf{x}_l \end{bmatrix} + \begin{bmatrix} \mathbf{n}_1 \\ \vdots \\ \mathbf{n}_l \end{bmatrix},$$

where we introduced l vectors $\mathbf{n}_1, \dots, \mathbf{n}_l \in \mathbb{R}^m$ to model additive measurement noise. Our goal is now to design a method that estimates the discrete images $\mathbf{x} = (\mathbf{x}_0^\top, \dots, \mathbf{x}_l^\top)^\top \in \mathbb{R}^{(l+1)n}$, and the transformation parameters $\boldsymbol{\theta}^\top = (\boldsymbol{\theta}_1^\top, \dots, \boldsymbol{\theta}_l^\top)^\top \in \mathbb{R}^{lq}$, using $\mathbf{y} = (\mathbf{y}_1^\top, \dots, \mathbf{y}_l^\top)^\top \in \mathbb{R}^{lm}$ as sole information.

The inter-correlation between the observed images appears explicitly in the measurement model (1.2). We consider that the observed images share a common signal - the background image - which lives on a low-dimensional manifold. In the literature, other models that exploit the inter-correlation between the measurements have been proposed.

1.2. Related works

The problem of reconstructing an image from multi-view measurements has been studied by several authors in contexts such as compressed sensing, super-resolution, or robust image alignment. We try here to give a brief overview of the different types of techniques used.

When the measurements are obtained by compressed sensing, some techniques, like ours, rely on the estimation of geometric transformations between images. For example, the authors in [13] propose to reconstruct the observed images by assuming that they are sparse in an over-complete dictionary created by scaling, translating and rotating several times a mother waveform. Then, they assume that the decompositions of these images in this dictionary are related by local geometrical transformations. In this method, one image is chosen as a reference image, and the final reconstruction quality depends on this choice. In [23, 44], the authors suppose that the images live on a low-dimensional manifold. Then they use a manifold lifting algorithm to obtain a first estimation of the transformation parameters and a standard compressed sensing algorithm for the estimation of the images. An alternate minimization method is then proposed to refine the estimation of the transformation parameters and of the images. However, no proof of convergence of the method is provided. Furthermore, to work efficiently, the algorithm requires the use of a specific measurement matrix, implemented with noiselets, to estimate the transformation parameters in a multiscale fashion. We also remark that the presence of occlusions in the images is not considered in this model.

Continuing with compressed sensing, other techniques rely on more general interpolation model between images. In [39, 40], the images are first estimated with a bidirectional interpolation using two known neighbor images and then corrected using the acquired measurements. An iterative process, alternating between interpolation and reconstruction, is then proposed to refine the estimation of the images. Note that a similar idea is used in [2] for dynamic MRI. In [21], a slightly simpler model is used. The authors assume that the difference between two images is sparse in a wavelet basis and that the background image is sparse in the gradient domain. Then, a convex minimization problem is solved to reconstruct jointly all the images. In [15, 18], the interpolation relies on ‘‘disparity maps’’ that indicate the correspondences between the pixels of two images. These maps are used to estimate the images from the measurements, and the newly estimated pictures can be used to refine the disparity maps. The process is then iterated several times. We remark that the authors do not provide any proof of convergence. We note that the authors deal with the problem of occlusions by supposing that they are sparse in image space. Let us also mention that a similar idea is proposed in [38]. Finally, in [33], the authors use a linear dynamical system to model the intercorrelation between frames, and reconstruct a video sequence.

When the matrices $\mathbf{A}_1, \dots, \mathbf{A}_m$ implement a blurring and downsampling operator, one can try to obtain a high-resolution background image from the images observed at low-resolution by solving (1.2). Super-resolution from multiple frames has also been widely studied and several image priors and correlation models between images have been proposed to solve this problem. As before, some

techniques model the correlation between images using geometric transformations [17, 43] while other ones use more general correlation models [27, 35].

Let us now highlight a similarity between our model and the one used in [24, 25] for robust image alignment, which consists in aligning a set of images despite the presence of large occlusions. With the measurement model (1.2), one can attempt to solve this problem by initializing the vectors $\mathbf{y}_1, \dots, \mathbf{y}_l$ with the observed images and the observations matrices $\mathbf{A}_1, \dots, \mathbf{A}_l$ with the identity. Then, the algorithm should estimate the background image, the transformation parameters that align this background image on the observed ones, and the foreground images. In [24, 25], a different approach, called RASL, is proposed. The authors concatenate the initial images $\mathbf{y}_1, \dots, \mathbf{y}_l$ in a matrix and explain that, after alignment of these images, this matrix can be decomposed as the sum of a low-rank matrix (modeling correlated components) and a sparse matrix (modeling occlusions). This model serves as the criterion to align the images. Note that in the limit where the rank of the final matrix modeling the correlated components is 1, their method, like ours, also decomposes the initial images into a single background image and l foreground images. Finally, let us mention that even though TILT [45] resembles RASL at first sight, their purpose is different. TILT is a method to correct and align low-rank textures and takes as input only one image. This is a different problem than the one studied here.

1.3. Main contribution and organization of the paper

In the present work, we propose a novel algorithm for *joint registration and reconstruction* of a set of misaligned images from linear measurements and, in contrast to the methods described above, provide a *convergence proof* of our algorithm. Furthermore, the model used permits us to easily deal with the problem of occlusions. It is also general enough to allow us to handle many image reconstruction problems. In order to study the efficiency of our algorithm, we use extensive numerical simulations in several applications. These numerical experiments confirm that our algorithm is competitive with well-known methods for robust image alignment, compressed sensing, and super-resolution.

In Section 2, we identify the solution of the multi-view imaging problem with the global minimizer of a non-convex functional, and propose an algorithm for the minimization of this functional. In Section 3, we study the convergence of the reconstruction method, and prove that the sequence of *estimates* (images and transformation parameters) converges to a critical point of the functional. The efficiency of the algorithm is then demonstrated experimentally in Section 4 for several problems. In the spirit of reproducible research, our code¹ and data needed to reproduce the results presented in this paper can be downloaded openly at lts2www.epfl.ch/people/gilles/software/. We finally conclude in Section 5. Some results required to prove the convergence of the algorithm are gathered in the appendix.

Note that we already briefly studied the problem of reconstructing a set of images from multi-view measurements in [28], where the measurement model (1.2) was also used. However, the technique proposed to reconstruct the images was based on the Bregman iterative regularization procedure, which is different from the one used here. Furthermore, the convergence results presented in [28] essentially concern the decrease of the objective value and, unlike here, do not ensure the convergence of the sequence of estimated images and parameters.

1.4. Notations and definitions

The Euclidean scalar product of \mathbb{R}^n is denoted by $\langle \cdot, \cdot \rangle$ and $\|\cdot\|_2$ is the corresponding ℓ_2 -norm. The ℓ_1 -norm of a vector $\mathbf{x} = (x_i)_{1 \leq i \leq n} \in \mathbb{R}^n$ is defined as $\|\mathbf{x}\|_1 = \sum_{i=1}^n |x_i|$.

Let C be a non-empty closed convex subset of \mathbb{R}^n . The indicator function of C is denoted by $i_C: \mathbb{R}^n \rightarrow \mathbb{R} \cup \{+\infty\}$. It is a proper lower semicontinuous convex function that satisfies $i_C(\mathbf{x}) = 0$ if $\mathbf{x} \in C$, and $i_C(\mathbf{x}) = +\infty$ otherwise.

Let $f: \mathbb{R}^n \rightarrow \mathbb{R} \cup \{+\infty\}$ be a proper lower semicontinuous function. The domain of f is denoted and defined by $\text{dom} f = \{\mathbf{x} \in \mathbb{R}^n : f(\mathbf{x}) < +\infty\}$. The subdifferential of f at $\mathbf{x} \in \text{dom} f$ is denoted by

¹We used the SPARCO toolbox available at www.cs.ubc.ca/labs/sc1/sparco/ [9].

$\partial f(\mathbf{x})$ (see, e.g., Section 2 of [3] for the definition). Note that $\partial f(\mathbf{x}) = \emptyset$ if $\mathbf{x} \notin \text{dom} f$. A necessary (but not sufficient) condition for $\mathbf{x}^* \in \mathbb{R}^n$ to be a minimizer of f is $\partial f(\mathbf{x}^*) \ni \mathbf{0}$. A point that satisfies the last condition is called a critical point of f .

Finally, \mathcal{C}^k denotes the set of functions that are k times continuously differentiable.

2. A non-convex optimization technique

In this section, we first identify the solution of the compressed multi-view imaging problem with the minimizer of a non-convex functional and propose an algorithm to minimize this functional.

2.1. Solution as minimizer of a non-convex functional

Let $\epsilon \geq 0$ be an upper bound on the energy of the measurement noise, i.e., $\|\mathbf{n}\|_2 \leq \epsilon$. Our objective is to find a set of parameters $\boldsymbol{\theta}^*$ and images \mathbf{x}^* that satisfy the constraint $\|A(\boldsymbol{\theta}^*) \mathbf{x}^* - \mathbf{y}\|_2 \leq \epsilon$, where

$$A(\boldsymbol{\theta}) = \begin{bmatrix} A_1 S(\boldsymbol{\theta}_1) & A_1 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ A_l S(\boldsymbol{\theta}_l) & 0 & \dots & A_l \end{bmatrix} \in \mathbb{R}^{lm \times (l+1)n}.$$

As infinitely many solutions satisfy this constraint, prior information is needed to restrict the set of admissible solutions. To regularize this ill-posed inverse problem, we can, for example, search for images sparse in a wavelet basis by minimizing the ℓ_1 -norm of their wavelet coefficients, or seek a solution that minimizes the Total Variation (TV) norm if the images are piecewise constant [32]. We may also require the transformation parameters $\boldsymbol{\theta}_j$, with $j = 1, \dots, l$, to belong to some convex sets $\Theta_j = \{\boldsymbol{\theta}_j \in \mathbb{R}^q : \underline{\boldsymbol{\theta}}_j \leq \boldsymbol{\theta}_j \leq \bar{\boldsymbol{\theta}}_j\}$ where $\underline{\boldsymbol{\theta}}_j \in \mathbb{R}^q$ and $\bar{\boldsymbol{\theta}}_j \in \mathbb{R}^q$ are pre-defined upper and lower bounds². Therefore, an estimate of the images and transformation parameters can be obtained by solving a minimization problem of the form

$$(2.3) \quad \min_{(\mathbf{x}, \boldsymbol{\theta})} f(\mathbf{x}) + \kappa \|A(\boldsymbol{\theta}) \mathbf{x} - \mathbf{y}\|_2^2 \quad \text{subject to } \boldsymbol{\theta} \in \Theta,$$

where $f: \mathbb{R}^{(l+1)n} \rightarrow \mathbb{R} \cup \{+\infty\}$ is a proper lower-semicontinuous convex function, $\kappa^{-1} > 0$ is a regularizing parameter that should be adjusted with the noise level ϵ , and $\Theta = \{\boldsymbol{\theta} = (\boldsymbol{\theta}_1^\top, \dots, \boldsymbol{\theta}_l^\top)^\top \in \mathbb{R}^{lq} : \boldsymbol{\theta}_j \in \Theta_j, j = 1, \dots, l\}$.

Problem (2.3) is non-linear and finding a global minimizer is not trivial. Nevertheless, one can still try to solve this problem using an alternating minimization technique such as the block coordinate descent method [41]. There is however no guarantee to converge to a critical point of Problem (2.3) with this technique. Recently, Attouch *et al.* proposed in [3] a new type of algorithms for the minimization of non-convex functionals and, more importantly, managed to prove that the sequence generated by these algorithms converges to a critical point for a wide class of functionals. As stated by the authors, their algorithm can be interpreted as a proximal regularization of the Gauss-Seidel method where cost-to-move functions are added in the minimization procedure. These convergence results were then generalized to inexact descent methods that satisfy a sufficient decrease condition [4]. Based on this work, we develop a minimization method for problem (2.3) and prove that the generated sequence converges to a critical point $(\mathbf{x}^*, \boldsymbol{\theta}^*)$ of the functional $L: \mathbb{R}^{(l+1)n} \times \mathbb{R}^{lq} \rightarrow \mathbb{R} \cup \{+\infty\}$ defined as

$$(2.4) \quad L(\mathbf{x}, \boldsymbol{\theta}) = f(\mathbf{x}) + \kappa \|A(\boldsymbol{\theta}) \mathbf{x} - \mathbf{y}\|_2^2 + i_\Theta(\boldsymbol{\theta}).$$

Let us recall that a global minimizer of L is also a critical point but that a critical point of L is not necessarily a global minimizer. The point $(\mathbf{x}^*, \boldsymbol{\theta}^*)$ might thus only be a local minimizer or a saddle point of Problem (2.3).

The proposed algorithm is a descent method. It generates a sequence of estimates $(\mathbf{x}^k, \boldsymbol{\theta}^k)_{k \in \mathbb{N}}$ such that $L(\mathbf{x}^{k+1}, \boldsymbol{\theta}^{k+1}) \leq L(\mathbf{x}^k, \boldsymbol{\theta}^k)$ for all $k \in \mathbb{N}$, and consists of two steps. We start by describing each of these steps and study the convergence of the algorithm in Section 3.

²Let $\bar{\boldsymbol{\theta}} = (\bar{\theta}_i)_{1 \leq i \leq q} \in \mathbb{R}^q$, $\boldsymbol{\theta} = (\theta_i)_{1 \leq i \leq q} \in \mathbb{R}^q$, $\boldsymbol{\theta} \leq \bar{\boldsymbol{\theta}}$ means that $\theta_i \leq \bar{\theta}_i$ for all $i \in \{1, \dots, q\}$.

2.2. First step of the algorithm

Let $(\mathbf{x}^k, \boldsymbol{\theta}^k) \in \mathbb{R}^{(l+1)n} \times \Theta$ be the estimates obtained after k iterations of the algorithm. The first step consists in finding a new estimate $\mathbf{x}^{k+1} \in \mathbb{R}^{(l+1)n}$ that decreases the value of the objective function L while keeping $\boldsymbol{\theta}^k$ fixed. Let us choose \mathbf{x}^{k+1} as a solution of

$$(2.5) \quad \min_{\mathbf{x} \in \mathbb{R}^{(l+1)n}} f(\mathbf{x}) + \kappa \|A(\boldsymbol{\theta}^k) \mathbf{x} - \mathbf{y}\|_2^2 + \frac{\lambda_{\mathbf{x}}}{2} g(\mathbf{x} - \mathbf{x}^k),$$

where $\lambda_{\mathbf{x}} > 0$, and $g: \mathbb{R}^{(l+1)n} \rightarrow \mathbb{R}$ is proper lower-semicontinuous convex function such that $g(\mathbf{x}) \geq 0$ for all $\mathbf{x} \in \mathbb{R}^{(l+1)n}$ and $g(\mathbf{0}) = 0$. It is clear that (2.5) implies

$$(2.6) \quad L(\mathbf{x}^{k+1}, \boldsymbol{\theta}^k) + \frac{\lambda_{\mathbf{x}}}{2} g(\mathbf{x}^{k+1} - \mathbf{x}^k) \leq L(\mathbf{x}^k, \boldsymbol{\theta}^k).$$

Hence $L(\mathbf{x}^{k+1}, \boldsymbol{\theta}^k) \leq L(\mathbf{x}^k, \boldsymbol{\theta}^k)$ with a decrease of $\lambda_{\mathbf{x}} g(\mathbf{x}^{k+1} - \mathbf{x}^k)/2$, at least. Note that this first minimization problem is convex in \mathbf{x} and can be solved using, e.g., the algorithms presented in [8,14,31].

In Problem (2.5), the function $\lambda_{\mathbf{x}} g(\cdot)$ acts as a proximal term. It ensures a sufficient decrease of the functional L at each iteration and is essential to prove the convergence of the sequence $(\mathbf{x}^k, \boldsymbol{\theta}^k)_{k \in \mathbb{N}}$ to a critical point. This function also provides, to some extent, a control on the evolution of the generated sequence, and we realized that it must be chosen carefully to achieve good and stable results in all the scenarios tested. Note that in the work [3] of Attouch *et al.*, the function g is a squared ℓ_2 -norm.

From the multi-view measurements $\mathbf{y}_1, \dots, \mathbf{y}_l$, the algorithm should reconstruct the background image, separate the occlusions (foreground), and provide the transformation parameters that register the initial images between them. This problem is solved here by alternatively refining the estimation of the images \mathbf{x} (background and foreground) and of the transformation parameters $\boldsymbol{\theta}$ (image registration). We noticed that to improve the quality of the registration, it is better to have a reconstruction of the images in a mutiscale fashion, as in [23,44]. In these papers, the authors use a dedicated measurement matrix that allows this type of reconstruction. In the current paper, this goal is achieved by choosing a function g that favors the reconstruction of the coarse scales first and then of the fine scales. There is thus no constraint on the choice of the measurement matrices anymore.

In order to have a reconstruction of the images from coarse to fine scales, we use a wavelet tight-frame $\mathbf{W} \in \mathbb{R}^{n \times p}$, with $p \geq n$, i.e., $\mathbf{W}\mathbf{W}^T$ is equal to $\mathbf{I}_n \in \mathbb{R}^{n \times n}$, the identity matrix. One should remark that most of the largest wavelet coefficients of a natural image live at the coarsest scales and that the amplitude of these coefficients decreases when going to finer scales (except at the few singularities). To achieve our goal, we need a function g that favors a reconstruction of the coarse scales first, by selecting the biggest wavelet coefficients, and then of the fine scales, by selecting the smallest wavelet coefficients. Let $\boldsymbol{\Psi} \in \mathbb{R}^{(l+1)n \times (l+1)p}$ be the block-diagonal matrix built by repeating $l+1$ times the matrix \mathbf{W} on the diagonal. We have $\boldsymbol{\Psi}\boldsymbol{\Psi}^T = \mathbf{I}_{(l+1)n}$. Take \mathbf{x}^{k+1} as a minimizer of $f(\mathbf{x}) + \kappa \|A(\boldsymbol{\theta}^k) \mathbf{x} - \mathbf{y}\|_2^2 + \lambda_{\mathbf{x}} \|\boldsymbol{\Psi}^T(\mathbf{x} - \mathbf{x}^k)\|_1/2$, and remember that the ℓ_1 -norm favors solutions with few large coefficients. Hence, the function f forces \mathbf{x}^{k+1} to fit our prior, the quadratic term forces \mathbf{x}^{k+1} to be consistent with the observations, and the cost-to-move function $\mathbf{x} \mapsto \|\boldsymbol{\Psi}^T(\mathbf{x} - \mathbf{x}^k)\|_1$ imposes that \mathbf{x}^{k+1} differs from \mathbf{x}^k by only a few large wavelet coefficients. Starting from $\mathbf{x}^0 = \mathbf{0} \in \mathbb{R}^{(l+1)n}$ at $k=0$, the first estimated images \mathbf{x}^1 have thus a sparse decomposition in the wavelet tight-frame \mathbf{W} . The sparsity increases with the parameter $\lambda_{\mathbf{x}}$. As we are reconstructing natural images, these few large coefficients mainly live at large scales and \mathbf{x}^1 is a coarse approximation of the solution. The estimation of the images is then refined at the next iteration, and the cost-to-move function favors the selection of the next biggest wavelet coefficients living at finer scales. Thanks to this procedure, we have a reconstruction of the images from coarse to fine scales. This behavior will be illustrated in Section 4.

Note that to be able to prove the convergence of the sequence $(\mathbf{x}^k, \boldsymbol{\theta}^k)_{k \in \mathbb{N}}$ to a critical point of L with the results presented in [3,4], we actually substitute the Huber function for the ℓ_1 -norm as cost-to-move function. The Huber function is a smooth approximation of the ℓ_1 -norm which depends on a smoothing parameter $\mu > 0$. This parameter can be chosen small so that both functions are

nearly identical. Let $\boldsymbol{\alpha} = (\alpha_i)_{1 \leq i \leq p} \in \mathbb{R}^{(l+1)p}$, the Huber function $h_\mu : \mathbb{R}^{(l+1)p} \rightarrow \mathbb{R}$ satisfies

$$h_\mu(\boldsymbol{\alpha}) = \sum_{i=1}^{(l+1)p} h_i, \quad \text{where} \quad h_i = \begin{cases} \alpha_i^2/(2\mu), & \text{if } |\alpha_i| < \mu, \\ |\alpha_i| + \mu/2, & \text{otherwise,} \end{cases} \quad \forall i \in \{1, \dots, (l+1)p\}.$$

From now on, the term $g(\mathbf{x} - \mathbf{x}^k)$ in Problem (2.5) is replaced by $h_\mu(\Psi^\top(\mathbf{x} - \mathbf{x}^k))$ (including in our implementation of the algorithm).

2.3. Second step of the algorithm

The second step consists in updating the transformation parameters to further decrease the value of the objective function L . As the functions $\boldsymbol{\theta} \mapsto \|\mathbf{A}(\boldsymbol{\theta}^k)\mathbf{x} - \mathbf{y}\|_2^2$ and i_Θ are separable in $\boldsymbol{\theta}_j$, with $j = 1, \dots, l$, we optimize the transformation parameters separately for each observation.

Let $(\mathbf{x}^k, \boldsymbol{\theta}^k) \in \mathbb{R}^{(l+1)n} \times \Theta$ be the estimates obtained after k iterations of the algorithm and $\mathbf{x}^{k+1} \in \mathbb{R}^{(l+1)n}$ be the solution of problem (2.5). To simplify the notations, we introduce l new functions $Q_j^{k+1} : \mathbb{R}^q \rightarrow \mathbb{R}$, with $j = 1, \dots, l$, defined as

$$(2.7) \quad Q_j^{k+1}(\boldsymbol{\theta}_j) = \|\mathbf{A}_j \mathbf{S}(\boldsymbol{\theta}_j) \mathbf{x}_0^{k+1} + \mathbf{A}_j \mathbf{x}_j^{k+1} - \mathbf{y}_j\|_2^2.$$

Note that $\|\mathbf{A}(\boldsymbol{\theta})\mathbf{x}^{k+1} - \mathbf{y}\|_2^2 = \sum_{j=1}^l Q_j^{k+1}(\boldsymbol{\theta}_j)$ and that $i_\Theta(\boldsymbol{\theta}) = \sum_{j=1}^l i_{\Theta_j}(\boldsymbol{\theta}_j)$. We choose to update the transformation parameters with the following projected Newton-like method [10, 34].

Let us assume that the entries of the matrix $\mathbf{S}(\boldsymbol{\theta}_j)$ are differentiable with respect to the transformation parameters. The first order Taylor expansion of $\mathbf{S}(\boldsymbol{\theta}_j)\mathbf{x}_0^{k+1}$ at $\boldsymbol{\theta}_j^k$ is $\mathbf{S}(\boldsymbol{\theta}_j^k)\mathbf{x}_0^{k+1} + \mathbf{J}(\boldsymbol{\theta}_j^k)(\boldsymbol{\theta}_j - \boldsymbol{\theta}_j^k)$ with

$$\mathbf{J}(\boldsymbol{\theta}_j^k) = (\partial_{\theta_{1j}} \mathbf{S}(\boldsymbol{\theta}_j^k) \mathbf{x}_0^{k+1}, \dots, \partial_{\theta_{qj}} \mathbf{S}(\boldsymbol{\theta}_j^k) \mathbf{x}_0^{k+1}) \in \mathbb{R}^{n \times q}.$$

Therefore,

$$Q_j^{k+1}(\boldsymbol{\theta}_j^k) + \langle \nabla Q_j^{k+1}(\boldsymbol{\theta}_j^k), \boldsymbol{\theta}_j - \boldsymbol{\theta}_j^k \rangle + \|\mathbf{A}_j \mathbf{J}(\boldsymbol{\theta}_j^k)(\boldsymbol{\theta}_j - \boldsymbol{\theta}_j^k)\|_2^2,$$

with

$$\nabla Q_j^{k+1}(\boldsymbol{\theta}_j^k) = 2 (\mathbf{A}_j \mathbf{J}(\boldsymbol{\theta}_j^k))^\top (\mathbf{A}_j \mathbf{S}(\boldsymbol{\theta}_j^k) \mathbf{x}_0^{k+1} + \mathbf{A}_j \mathbf{x}_j^{k+1} - \mathbf{y}_j),$$

is a second order approximation of Q_j^{k+1} at $\boldsymbol{\theta}_j^k$. The positive semi-definite matrix

$$(2.8) \quad \mathbf{H}(\boldsymbol{\theta}_j^k) = 2 (\mathbf{A}_j \mathbf{J}(\boldsymbol{\theta}_j^k))^\top (\mathbf{A}_j \mathbf{J}(\boldsymbol{\theta}_j^k)),$$

can be viewed as an approximation of the Hessian of the function Q_j^{k+1} . To update the transformation parameters, we choose to minimize this quadratic approximation to which we add another quadratic term that ensures a decrease of the objective function. We take as next estimate of the transformation parameters

$$(2.9) \quad \boldsymbol{\theta}_j^{k+1} = \underset{\boldsymbol{\theta}_j \in \Theta_j}{\operatorname{argmin}} \langle \nabla Q_j^{k+1}(\boldsymbol{\theta}_j^k), \boldsymbol{\theta}_j - \boldsymbol{\theta}_j^k \rangle + \frac{1}{2} \langle \boldsymbol{\theta}_j - \boldsymbol{\theta}_j^k, [\mathbf{H}(\boldsymbol{\theta}_j^k) + 2^i \lambda_\theta \mathbf{l}_q] (\boldsymbol{\theta}_j - \boldsymbol{\theta}_j^k) \rangle,$$

where $\mathbf{l}_q \in \mathbb{R}^{q \times q}$ is the identity matrix, $\lambda_\theta > 0$, and i is the smallest positive integer such that

$$(2.10) \quad \begin{aligned} Q_j^{k+1}(\boldsymbol{\theta}_j^{k+1}) &\leq Q_j^{k+1}(\boldsymbol{\theta}_j^k) + \langle \nabla Q_j^{k+1}(\boldsymbol{\theta}_j^k), \boldsymbol{\theta}_j^{k+1} - \boldsymbol{\theta}_j^k \rangle \\ &+ \frac{1}{2} \langle \boldsymbol{\theta}_j^{k+1} - \boldsymbol{\theta}_j^k, [\mathbf{H}(\boldsymbol{\theta}_j^k) + (2^i - 1) \lambda_\theta \mathbf{l}_q] (\boldsymbol{\theta}_j^{k+1} - \boldsymbol{\theta}_j^k) \rangle. \end{aligned}$$

The existence of such an i is discussed in the next section³. One can remark that if the minimization (2.9) was performed over all \mathbb{R}^q instead of Θ_j , the solution would be $\boldsymbol{\theta}_j^{k+1} = \boldsymbol{\theta}_j^k - [\mathbf{H}(\boldsymbol{\theta}_j^k) + 2^i \lambda_\theta \mathbf{l}_q]^{-1} \nabla Q_j^{k+1}(\boldsymbol{\theta}_j^k)$, as in a Newton method with Hessian $\mathbf{H}(\boldsymbol{\theta}_j^k) + 2^i \lambda_\theta \mathbf{l}_q$. One can also note that if we had chosen $\mathbf{H}(\boldsymbol{\theta}_j^k) = \mathbf{0}$ then $\boldsymbol{\theta}_j^{k+1}$ would be obtained with a simple projected gradient update.

³Note that if one is able to compute the Lipschitz constant Λ_j^{k+1} of ∇Q_j^{k+1} , then Λ_j^{k+1} can be substituted for $2^i \lambda_\theta$ and the algorithm also converges.

We noticed however that the Newton-like update (2.9), which can, for example, be solved using the algorithms presented in [8], yields more accurate results.

To check that our choice of new transformation parameters decreases the value of the objective function L , one just has to combine (2.9) and (2.10) to realize that

$$Q_j^{k+1}(\boldsymbol{\theta}_j^{k+1}) + \frac{\lambda_{\boldsymbol{\theta}}}{2} \|\boldsymbol{\theta}_j^{k+1} - \boldsymbol{\theta}_j^k\|_2^2 \leq Q_j^{k+1}(\boldsymbol{\theta}_j^k) + \langle \nabla Q_j^{k+1}(\boldsymbol{\theta}_j^k), \boldsymbol{\theta}_j - \boldsymbol{\theta}_j^k \rangle + \frac{1}{2} \langle \boldsymbol{\theta}_j - \boldsymbol{\theta}_j^k, [\mathbf{H}(\boldsymbol{\theta}_j^k) + 2^i \lambda_{\boldsymbol{\theta}} \mathbf{1}_q] (\boldsymbol{\theta}_j - \boldsymbol{\theta}_j^k) \rangle,$$

for any $\boldsymbol{\theta}_j \in \Theta_j$. Choosing $\boldsymbol{\theta}_j = \boldsymbol{\theta}_j^k$ in the last inequality, multiplying it by κ , and summing all the inequalities obtained for j from 1 to l yields

$$(2.11) \quad L(\mathbf{x}^{k+1}, \boldsymbol{\theta}^{k+1}) + \frac{\kappa \lambda_{\boldsymbol{\theta}}}{2} \|\boldsymbol{\theta}^{k+1} - \boldsymbol{\theta}^k\|_2^2 \leq L(\mathbf{x}^{k+1}, \boldsymbol{\theta}^k).$$

Therefore $L(\mathbf{x}^{k+1}, \boldsymbol{\theta}^{k+1}) \leq L(\mathbf{x}^{k+1}, \boldsymbol{\theta}^k)$ with a decrease of $\kappa \lambda_{\boldsymbol{\theta}} \|\boldsymbol{\theta}^{k+1} - \boldsymbol{\theta}^k\|_2^2 / 2$, at least.

3. Convergence analysis

In this section, we analyze the convergence of the optimization method described in Section 2. We first present sufficient conditions under which the generated sequence converges to a critical point of L . Then, we describe several cases where these conditions are satisfied. In particular, we study different types of geometric transformations which meet these requirements. Finally, we discuss the influence of the cost-to-move parameters on the final reconstruction quality.

Algorithm 1

Inputs: measurements $\mathbf{y} \in \mathbb{R}^{lm}$, wavelet tight-frame $\Psi \in \mathbb{R}^{(l+1)n \times (l+1)p}$, regularization parameter $\kappa > 0$, parameters $(\lambda_{\mathbf{x}}^k)_{k \in \mathbb{N}}$, $\lambda_{\boldsymbol{\theta}} > 0$, and bounds $\boldsymbol{\theta} \geq \underline{\boldsymbol{\theta}}$.

Initializations: set $k = 0$, $\mathbf{x}^0 = \mathbf{0} \in \mathbb{R}^{(l+1)n}$, and $\boldsymbol{\theta}^0 \in \Theta$.

repeat

1) Set

$$\mathbf{x}^{k+1} \leftarrow \operatorname{argmin}_{\mathbf{x} \in \mathbb{R}^{(l+1)n}} L(\mathbf{x}, \boldsymbol{\theta}^k) + \frac{\lambda_{\mathbf{x}}^k}{2} h_{\mu}(\Psi^{\top}(\mathbf{x} - \mathbf{x}^k)).$$

2) For all $j = 1, \dots, l$, set

$$\boldsymbol{\theta}_j^{k+1} \leftarrow \operatorname{argmin}_{\boldsymbol{\theta}_j \in \Theta_j} \langle \nabla Q_j^{k+1}(\boldsymbol{\theta}_j^k), \boldsymbol{\theta}_j - \boldsymbol{\theta}_j^k \rangle + \frac{1}{2} \langle \boldsymbol{\theta}_j - \boldsymbol{\theta}_j^k, [\mathbf{H}(\boldsymbol{\theta}_j^k) + 2^i \lambda_{\boldsymbol{\theta}} \mathbf{1}] (\boldsymbol{\theta}_j - \boldsymbol{\theta}_j^k) \rangle,$$

where Q_j^{k+1} is defined in (2.7), $\mathbf{H}(\boldsymbol{\theta}_j^k)$ is defined in (2.8), and i is the smallest positive integer such that (2.10) is satisfied.

3) $k \leftarrow k + 1$

until convergence or $k \geq k_{\max}$

Outputs: Estimated images \mathbf{x}^* and transformation parameters $\boldsymbol{\theta}^*$.

3.1. General convergence result

Let us consider Algorithm 1 which is a summary of the optimization method described in the previous section. We are now in position to state our main convergence result.

Theorem 1. *Let L be the objective function defined in (2.4) with $\kappa > 0$. Assume that L is bounded below, that the entries of the matrix \mathbf{S} defined in (1.1) are twice continuously differentiable with respect to the transformation parameters, that $\Psi \in \mathbb{R}^{(l+1)n \times (l+1)p}$ satisfies $\Psi \Psi^{\top} = \mathbf{I}_{(l+1)n}$, and that*

$0 < \underline{\lambda} \leq \lambda_{\mathbf{x}}^k, \lambda_{\boldsymbol{\theta}} \leq \bar{\lambda}$ for all $k \in \mathbb{N}$. Then, the sequence of estimates $(\mathbf{x}^k, \boldsymbol{\theta}^k)_{k \in \mathbb{N}}$ generated by Algorithm 1 is correctly defined and the following statements hold.

1) For all $k \geq 0$,

$$(3.12) \quad L(\mathbf{x}^{k+1}, \boldsymbol{\theta}^{k+1}) + \frac{\lambda}{2} \left[\kappa \|\boldsymbol{\theta}^{k+1} - \boldsymbol{\theta}^k\|_2^2 + h_\mu(\Psi^\top(\mathbf{x}^{k+1} - \mathbf{x}^k)) \right] \leq L(\mathbf{x}^k, \boldsymbol{\theta}^k).$$

Hence $L(\mathbf{x}^k, \boldsymbol{\theta}^k)$ does not increase.

2) The sequences $(\mathbf{x}^{k+1} - \mathbf{x}^k)_{k \in \mathbb{N}}$ and $(\boldsymbol{\theta}^{k+1} - \boldsymbol{\theta}^k)_{k \in \mathbb{N}}$ converge. Indeed,

$$(3.13) \quad \lim_{k \rightarrow +\infty} \|\mathbf{x}^{k+1} - \mathbf{x}^k\|_2 + \|\boldsymbol{\theta}^{k+1} - \boldsymbol{\theta}^k\|_2 = 0.$$

3) Assume that L has the Kurdyca-Lojasiewicz property (see Definition 3.2 in [3]). Then, if the sequence $(\mathbf{x}^k)_{k \in \mathbb{N}}$ is bounded, the sequence $(\mathbf{x}^k, \boldsymbol{\theta}^k)_{k \in \mathbb{N}}$ converges to a critical point $(\mathbf{x}^*, \boldsymbol{\theta}^*)$ of L .

Proof. Let us start by showing that the sequence generated by Algorithm 1 is well-defined. For the first step of the algorithm, note that, for all $(\tilde{\mathbf{x}}, \tilde{\boldsymbol{\theta}}) \in \mathbb{R}^{(l+1)n} \times \mathbb{R}^{lq}$ and $\tilde{\lambda} > 0$, $\mathbf{x} \mapsto L(\mathbf{x}, \tilde{\boldsymbol{\theta}})$ is a proper lower semicontinuous convex function which is bounded below and that $\mathbf{x} \mapsto \tilde{\lambda} h_\mu(\Psi^\top(\mathbf{x} - \tilde{\mathbf{x}}))/2$ is a proper lower semicontinuous convex function which is coercive. Therefore, the function $\mathbf{x} \mapsto L(\mathbf{x}, \tilde{\boldsymbol{\theta}}) + \tilde{\lambda} h_\mu(\Psi^\top(\mathbf{x} - \tilde{\mathbf{x}}))/2$ has a minimizer (Corollary 11.15, [7]). For the second step of the algorithm, remark that, for all $(\tilde{\mathbf{x}}_0, \tilde{\mathbf{x}}_j) \in \mathbb{R}^n \times \mathbb{R}^n$, the function $Q_j: \boldsymbol{\theta}_j \mapsto \|\mathbf{A}_j \mathbf{S}(\boldsymbol{\theta}_j) \tilde{\mathbf{x}}_0 + \mathbf{A}_j \tilde{\mathbf{x}}_j - \mathbf{y}_j\|_2^2$ is \mathcal{C}^1 with Lipschitz continuous gradient on the closed and bounded convex set Θ_j . Let Λ_j be this Lipschitz constant, whose value *a priori* depends on $(\tilde{\mathbf{x}}_0, \tilde{\mathbf{x}}_j)$. Using the descent lemma (Lemma 3.1, [4]), we have

$$Q_j(\boldsymbol{\theta}_j^1) \leq Q_j(\boldsymbol{\theta}_j^2) + \langle \nabla Q_j(\boldsymbol{\theta}_j^2), \boldsymbol{\theta}_j^1 - \boldsymbol{\theta}_j^2 \rangle + \frac{\Lambda_j}{2} \|\boldsymbol{\theta}_j^1 - \boldsymbol{\theta}_j^2\|_2^2,$$

for any two points $(\boldsymbol{\theta}_j^1, \boldsymbol{\theta}_j^2) \in \Theta_j^2$. This proves the existence of an integer i such that (2.10) is satisfied at each second step of Algorithm 1. The existence of a minimizer in Problem (2.9) is then proved by using, e.g., Proposition 11.14 in [7]. An induction finally shows that sequence $(\mathbf{x}^k, \boldsymbol{\theta}^k)_{k \in \mathbb{N}}$ is well-defined.

Inequality (3.12) follows by combination of (2.6) and (2.11) and by using the fact that $\lambda_{\mathbf{x}}^k, \lambda_{\boldsymbol{\theta}} \geq \underline{\lambda}$ for all $k \in \mathbb{N}$. Then, summing these inequalities from $k = 0$ to K yields

$$\frac{\lambda}{2} \sum_{k=0}^K \kappa \|\boldsymbol{\theta}^{k+1} - \boldsymbol{\theta}^k\|_2^2 + h_\mu(\Psi^\top(\mathbf{x}^{k+1} - \mathbf{x}^k)) \leq L(\mathbf{x}^0, \boldsymbol{\theta}^0) - L(\mathbf{x}^{K+1}, \boldsymbol{\theta}^{K+1})$$

As L is bounded below and $L(\mathbf{x}^0, \boldsymbol{\theta}^0) = \kappa \|\mathbf{y}\|_2^2$, it is clear that the righthand side of the previous inequality is also bounded. Consequently,

$$\sum_{k=0}^{+\infty} h_\mu(\Psi^\top(\mathbf{x}^{k+1} - \mathbf{x}^k)) + \kappa \|\boldsymbol{\theta}^{k+1} - \boldsymbol{\theta}^k\|_2^2 < +\infty,$$

and, using the definition of h_μ and the fact that $\Psi \Psi^\top = \mathbf{I}_{(l+1)n}$, the second point of the theorem holds.

The proof of the third point make use of results established by Attouch *et al.* in [4] and can be found in Appendix A. \square

Note that to study the convergence of Algorithm 1, we considered, for simplicity, that the minimization problems (2.5) and (2.9) can be minimized exactly. However, the results presented by Attouch *et al.* in [4] also hold with inexact descent methods. This suggests that Algorithm 1 also converges with inexact minimizations of (2.5) and (2.9), upon the introduction of few additional requirements on the “quality” of these minimizations.

The third point of Theorem 1 applies if L has the Kurdyca-Lojasiewicz property. As explained in [3], this property is satisfied by several types of functions including semi-algebraic ones. A function $g: \mathbb{R}^n \rightarrow \mathbb{R} \cup \{+\infty\}$ is semi-algebraic if its graph $\{(\mathbf{x}, t) \in \mathbb{R}^n \times \mathbb{R} : g(\mathbf{x}) = t\}$ is a semi-algebraic

subset of \mathbb{R}^{n+1} , i.e., if it can be written as a finite union of sets of the form $\{\mathbf{x} \in \mathbb{R}^{n+1} : p_i(\mathbf{x}) = 0, q_i(\mathbf{x}) < 0, i = 1, \dots, r\}$, where p_i and q_i are polynomials. The ℓ_1 -norm is thus an example of such a function. Note that the indicator function of a semi-algebraic set is semi-algebraic, and that the set of semi-algebraic functions is stable under basics operations such as sums, products or compositions (Section 4.3, [3]).

In the definition of the objective function L , the set Θ is semi-algebraic. Consequently, L satisfies the Kurdyca-Lojasiewicz property if, e.g., f and the fidelity term $(\boldsymbol{\theta}, \mathbf{x}) \mapsto \|\mathbf{A}(\boldsymbol{\theta})\mathbf{x} - \mathbf{y}\|_2^2$ are semi-algebraic. For f , one can for example use $f(\mathbf{x}) = \|\Phi\mathbf{x}\|_1$ with any real matrix Φ . This choice also ensures that L is bounded below. Also note that if Φ has full column rank, then, using condition (3.12), one can show that the sequence $(\mathbf{x}_k)_{k \in \mathbb{N}}$ is bounded. For the fidelity term, its semi-algebraicity depends on the properties of the generating function φ and the type of geometric transformations. In the next section, we present examples of functions φ and geometric transformations such that this term is semi-algebraic and such that the entries of the matrix \mathbf{S} are twice continuously differentiable, as required by Theorem 1.

3.2. Examples of admissible geometric transformations

Let us start by studying the simple case of translations. This type of transformation can be represented by 2 parameters $\boldsymbol{\theta} = (\theta_1, \theta_2)^\top \in \mathbb{R}^2$ and a function $\tau_{\boldsymbol{\theta}}$ defined as

$$\tau_{\boldsymbol{\theta}}(\mathbf{u}) = (u_1 + \theta_1, u_2 + \theta_2)^\top.$$

Combined, for example, with the cubic interpolator defined in [19] as generating function $\varphi(\mathbf{u}) = \phi(u_1)\phi(u_2)$ where

$$(3.14) \quad \phi(u) = \begin{cases} \frac{3}{2}|u|^3 - \frac{5}{2}|u|^2 + 1, & \text{if } 0 \leq |u| < 1, \\ -\frac{1}{2}|u|^3 + \frac{5}{2}|u|^2 - 4|u| + 2, & \text{if } 1 \leq |u| < 2, \\ 0, & \text{otherwise,} \end{cases}$$

the entries of the matrix \mathbf{S} are twice continuously differentiable with respect to the transformation parameters, and the function $(\boldsymbol{\theta}, \mathbf{x}) \mapsto \|\mathbf{A}(\boldsymbol{\theta})\mathbf{x} - \mathbf{y}\|_2^2$ is semi-algebraic. Theorem 1 thus applies in this situation. Remark that with this choice of generating function the matrix \mathbf{S} is sparse, which permits us to transform rapidly the background image. Note that, instead of the cubic interpolator, we could also have chosen any B-spline interpolators of degree larger than 3 [42].

Then, with the same choice of generating function, we can also handle the case of affine transformations. Indeed, such transformations can be represented by 6 parameters $\boldsymbol{\theta} = (\theta_1, \dots, \theta_6)^\top \in \mathbb{R}^6$ with

$$(3.15) \quad \tau_{\boldsymbol{\theta}}(\mathbf{u}) = (\theta_1 u_1 + \theta_2 u_2 + \theta_3, \theta_4 u_1 + \theta_5 u_2 + \theta_6)^\top,$$

and the requirements to apply Theorem 1 are also met.

Next, let us show how to handle the case of homographies. An homography is usually represented using 8 parameters $\boldsymbol{\theta} = (\theta_1, \dots, \theta_8)^\top \in \mathbb{R}^8$ with $\tau_{\boldsymbol{\theta}}$ satisfying

$$\tau_{\boldsymbol{\theta}}(\mathbf{u}) = \left(\frac{\theta_1 u_1 + \theta_2 u_2 + \theta_3}{\theta_7 u_1 + \theta_8 u_2 + 1}, \frac{\theta_4 u_1 + \theta_5 u_2 + \theta_6}{\theta_7 u_1 + \theta_8 u_2 + 1} \right)^\top.$$

Unfortunately, the function $(\boldsymbol{\theta}, \mathbf{x}) \mapsto \|\mathbf{A}(\boldsymbol{\theta})\mathbf{x} - \mathbf{y}\|_2^2$ is not semi-algebraic in this case. However, if $|\theta_7 u_1 + \theta_8 u_2| \ll 1$, the transformation function can be approximated by

$$(3.16) \quad \tau_{\boldsymbol{\theta}}(\mathbf{u}) = \begin{pmatrix} \theta_1 u_1 + \theta_2 u_2 + \theta_3 \\ \theta_4 u_1 + \theta_5 u_2 + \theta_6 \end{pmatrix} (1 - \theta_7 u_1 - \theta_8 u_2),$$

using a first order Taylor expansion. With this approximation, all the conditions are now fulfilled to apply Theorem 1. Note that the condition $|\theta_7 u_1 + \theta_8 u_2| \ll 1$ can be enforced during the reconstruction by using appropriate bounds $\bar{\boldsymbol{\theta}}$ and $\underline{\boldsymbol{\theta}}$.

Beyond these examples, let us mention that if the generating function φ and the transformation function $\tau_{\boldsymbol{\theta}}$ are twice continuously differentiable, Algorithm 1 can also be applied and the first two points of Theorem 1 still hold.

3.3. Influence of the algorithm parameters

Let us highlight that, even though the sequence $(\mathbf{x}^k, \boldsymbol{\theta}^k)_{k \in \mathbb{N}}$ generated by Algorithm 1 converges to a critical point of L for all strictly positive parameter $\lambda_{\boldsymbol{\theta}}$ and strictly positive bounded sequence $(\lambda_{\mathbf{x}}^k)_{k \in \mathbb{N}}$, different values of these parameters yield different results. The choice of these parameters is thus important.

For the sequence $(\lambda_{\mathbf{x}}^k)_{k \in \mathbb{N}}$, we start from large values to constrain \mathbf{x}^{k+1} to stay “close” to \mathbf{x}^k when the estimated images are still inaccurate. Starting with large values is also essential to obtain images \mathbf{x}^k made of coarse wavelets only during the first iterations. As a rule of thumb, we have noticed that the best reconstructions were obtained when $\lambda_{\mathbf{x}}^0$ is chosen such that $\|\mathbf{A}(\boldsymbol{\theta}^0) \mathbf{x}^1 - \mathbf{y}\|_2^2 / \|\mathbf{y}\|_2^2 \in [0.1, 0.2]$, after the first step of the first iteration. Then, we slightly decrease the value of $\lambda_{\mathbf{x}}^k$ at each iteration with the recursion rule $\lambda_{\mathbf{x}}^{k+1} = \max(0.9\lambda_{\mathbf{x}}^k, 0.1)$. The parameter $\lambda_{\boldsymbol{\theta}}$ applying on the transformation parameters seems to have less influence on the final reconstruction quality, and we keep it fixed at 0.1 in all the experiments of the next section.

Finally, concerning the bounds $\underline{\boldsymbol{\theta}}$ and $\bar{\boldsymbol{\theta}}$ constraining the transformation parameters, we simply choose them large enough to ensure that the actual transformation parameters between images are in the allowed range, while making sure that uninteresting solutions, such as translations larger than the size of the images or scaling parameters tending to zero, are prevented.

3.4. Limitations of our algorithm

We now conclude this section by discussing some limitations of our algorithm.

In this work, we considered geometric transformations globally described by *few* parameters. This restricts the number of applications where the algorithm can be applied. For example, when the observed scene is not planar, the geometric transformations between images are usually more complex and depend on the structure of the scene. In this situation, more general transformation models have to be used. One should note however that the following two properties of the proposed method may help to handle such a situation. First, the requirements of Theorem 1 hold for a large class of transformation models $\tau_{\boldsymbol{\theta}}$. We thus can use transformation models with more degrees of freedom to approximate more complex transformations. Second, the presence of foreground images in the measurement model can render our method robust to small transformation model mismatches by identifying the parts of the scene that cannot be aligned with the assumed model as occlusions.

Finally, another limitation concerns the number of transformation parameters that the algorithm can efficiently handle. Indeed, if the number q of transformation parameters is large, the minimization problem at the second step of Algorithm 1 might become intractable. First, the computation of the matrix $\mathbf{H} \in \mathbb{R}^{q \times q}$, defined in (2.8), will become time consuming (unless the transformation model used exhibits special structure). Second, solving the minimization problem (2.9) is computationally expensive for unstructured large matrix \mathbf{H} . Note however that one way to circumvent this issue is to approximate the matrix \mathbf{H} by computing only its most “significant” entries and setting the remaining ones to zero.

4. Experiments

In this section, we demonstrate numerically the efficiency of the algorithm in three different applications: robust image alignment, compressed sensing, and super-resolution. The code and data needed to reproduce the results presented in this section are available at lts2www.epfl.ch/people/gilles/software.

4.1. Robust image alignment

Robust image alignment consists in aligning several images of a given scene despite the presence of large occlusions, as, for example, in the images presented in the top panel⁴ of Fig. 1.

⁴Dataset available at perception.csl.illinois.edu/matrix-rank/rasl.html.




Figure 1. Top panel: first 5 input images. Bottom panels: estimated background images with the proposed method and RASL using $l = 6, 8, 10, 12$ or 16 input images (from left to right).




Figure 2. First 2 estimated foreground images with the proposed method and RASL using 8 (left) or 16 (right) input images.




Figure 3. *Left panel: superposition of the first 10 input images. Middle panel: superposition of the same 10 images after registration with our method. Right panel: superposition of the same 10 images after registration with RASL.*




Figure 4. *Sequence of estimated background images \mathbf{x}_0^{k+1} at $k = 0$ (first from the left), $k = 10$ (second), $k = 20$ (third), and at convergence (fourth).*

We propose here to solve this problem by using the measurement model defined in (1.2) and Algorithm 1 to minimize (2.4). In the measurement model, the matrices A_1, \dots, A_l , are set to the identity $I_n \in \mathbb{R}^{n \times n}$, and $\mathbf{y}_1, \dots, \mathbf{y}_l$, contain the occluded observed images. Ideally our algorithm should estimate the background image \mathbf{x}_0 , i.e., the wall and the windows in Fig. 1 (top panel), extract the occlusions $\mathbf{x}_1, \dots, \mathbf{x}_l$, i.e., the branches, and provide the transformation parameters $\theta_1, \dots, \theta_l$, which align the background image on the observed images.

To show the efficiency of our technique, we compare our results with the ones obtained using RASL [24, 25]. This method is based on low rank and sparse approximations. More precisely, the authors in [24, 25] concatenate the occluded input images in a matrix $\mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_l] \in \mathbb{R}^{m \times l}$ and propose to solve

$$(4.17) \quad (\tilde{\mathbf{B}}, \tilde{\mathbf{F}}, \tilde{\boldsymbol{\theta}}) \in \underset{\mathbf{B}, \mathbf{F}, \boldsymbol{\theta}}{\operatorname{argmin}} \|\mathbf{B}\|_* + \frac{1}{\sqrt{m}} \|\mathbf{F}\|_{1,1} \quad \text{subject to} \quad \mathbf{Y} \circ \tau_{\boldsymbol{\theta}} = \mathbf{B} + \mathbf{F},$$

where $\mathbf{B} \in \mathbb{R}^{m \times l}$, $\mathbf{F} \in \mathbb{R}^{m \times l}$, $\|\mathbf{B}\|_*$ denotes the nuclear norm of \mathbf{B} , $\|\mathbf{F}\|_{1,1}$ is the sum of the absolute value of the entries of \mathbf{F} , and $\mathbf{Y} \circ \tau_{\boldsymbol{\theta}}$ are the input images transformed according to the parameters $\boldsymbol{\theta}$ and the transformation model $\tau_{\boldsymbol{\theta}}$. This optimization problem decomposes the initial images \mathbf{Y} in a low-rank matrix \mathbf{B} which contains the background elements, and a sparse matrix \mathbf{F} which contains the occlusions. The transformation parameters $\boldsymbol{\theta}$ are adjusted so that the matrix \mathbf{F} is as sparse as possible and the rank of the matrix \mathbf{B} is as small as possible.

One can remark that the data constraint $\mathbf{Y} \circ \tau_{\boldsymbol{\theta}} = \mathbf{B} + \mathbf{F}$ is non-linear in $\boldsymbol{\theta}$ and minimizing (4.17) is thus not trivial. To get around this difficulty, the authors linearize the term $\mathbf{Y} \circ \tau_{\boldsymbol{\theta}}$ with respect to the transformation parameters. This approximation is only valid for small transformations but the optimization problem becomes convex and easier to solve. Then, to align images with large

transformations between them, the authors propose to repeatedly solve this convex problem and linearize the data constraint around the new estimated parameters to improve, little by little, the registration. Note that the authors do not provide a rigorous proof of convergence of this procedure.

Finally, it is interesting to note that, surprisingly, no constraints are imposed on the transformation parameters in (4.17). Therefore, if there exist a set of parameters $\theta_0 \in \mathbb{R}^{lq}$ such that $Y \circ \tau_{\theta_0} = 0$, then it is easy to prove that the global minimum of (4.17) is 0 and $\tilde{\mathbf{B}} = \tilde{\mathbf{F}} = 0$. In such a situation, finding a global minimizer of (4.17) obviously does not yield to the desired solution, and, if one correctly aligns the images, it is only because the procedure described in [24, 25] converges to a local minima of problem (4.17). Note that as soon as the class of transformations considered allows us to scale the images, we already have $Y \circ \tau_{\theta} \rightarrow 0$ when the scaling parameter tends to zero. Let us however highlight that this issue can trivially be corrected by adding a constraint on the transformation parameters in the original formulation (4.17) and modify accordingly the procedure described in (4.17).

4.1. Background/foreground separation quality

In this first set of experiments, we use $l = 6, 8, 10, 12$, and 16 occluded images of $n = 256 \times 256$ pixels as inputs. The first 5 images are presented in the top panels of Fig. 1. Then, for each l , we run Algorithm 1 and RASL to align the images and to estimate the background and foreground images.

In our algorithm, we set $\kappa = 100$, $\lambda_{\theta} = 0.1$, $\lambda_{\mathbf{x}}^k = \max(0.9^k (20\kappa), 0.1)$, and $\mu = 10^{-10}$. The matrix Ψ is built using the Haar wavelet basis, and $f(\mathbf{x})$ is set to $\|\Psi^T \mathbf{x}\|_1$. We also assume that the transformations between images are well modeled by homographies and thus use model (3.16) for τ_{θ} . The parameters are constrained as follows: $|1 - \theta_1|, |1 - \theta_5| \leq 0.2$, $|\theta_2|, |\theta_4| \leq 0.2$, $|\theta_3|, |\theta_6| \leq 20$, $|\theta_7|, |\theta_8| \leq 0.01$, with u_1 and u_2 belonging to $\{-127, -126, \dots, 128\}$.

In Fig. 1, we present for each l the background images⁵ estimated with our method and RASL. One can first remark that with a small number of input images, i.e., $l = 6, 8$ and 10, the background images estimated with our method contains much less artifacts due to the occlusions than the ones estimated with RASL. Finally, our method only needs 10 input images to obtain a background image free of occlusion artifacts while RASL needs more input images.

In Fig. 2, we present the first two foreground images estimated with both methods from $l = 8$ and 16 input images. One can notice that, for both methods, the foreground images essentially contain the occlusions, i.e., the branches, present in the initial input images. However, at $l = 8$, the foreground images estimated with RASL are contaminated by other objects not initially present in the corresponding input images. With a small number of input images, our method separates the background from the foreground better than RASL does. At $l = 16$, the foreground images estimated by both methods are quite similar. We note however that the windows are slightly less visible in the foreground images estimated by RASL.

To visualize the quality of the estimated transformation parameters, we present in Fig. 3 the superposition of $l = 10$ input images before and after registration with our method and RASL. One can easily remark that both methods are able to accurately register the set of input images.

Finally, we present in Fig. 4 the evolution of the estimated background image \mathbf{x}_0^{k+1} at iterations $k = 0, 10, 20, 30$, and at convergence with our method. This sequence of images illustrates the fact that, as explained in Section 2.2, the images are reconstructed in a multiscale fashion.

4.1. Registration quality

We now compare quantitatively the registration quality that one can obtain with RASL or our method.

In this second set of experiments, we use a “scaling+translational” model represented by 3 parameters $\theta = (s, t_x, t_y)^T \in \mathbb{R}^3$ with

$$(4.18) \quad \tau_{\theta}(\mathbf{u}) = (s u_1 + t_x, s u_2 + t_y)^T,$$

where s is a scaling factor, and (t_x, t_y) are translation parameters in the first and second dimensions of the images, respectively. To compare the registration performance, we run two different types

⁵For RASL, the background image is obtained by averaging the estimated background images in the matrix $\tilde{\mathbf{B}}$.




Figure 5. Registration error σ averaged over 30 simulations for the proposed method (left panels) and RASL (right panels) as a function of $(\Delta t_x, \Delta t_y)$ and of $(\Delta t_x, \Delta s)$. The colorbar goes from white for $\sigma = 0$ to black for $\sigma = 0.5$.

of simulations on images of $n = 128 \times 128$ pixels. The spatial coordinates u_1 and u_2 belong to $\{-63, -62, \dots, 64\}$. Let us mention that with this transformation model, all the requirements needed to apply Theorem 1 are met.

For the first type of simulations, we draw independently and uniformly at random l translation parameters, $\{(t_x)_1, \dots, (t_x)_l\}$ and $\{(t_y)_1, \dots, (t_y)_l\}$, in the intervals $[-\Delta t_x/2, \Delta t_x/2]$ and $[-\Delta t_y/2, \Delta t_y/2]$, respectively, while keeping s fixed at 1. Then, we transform a reference image⁶ according to these parameters to create l different input images. Finally, we register these images using RASL and our algorithm. This procedure is repeated for different sizes of the intervals, $\Delta t_x \in \{2, 8, 16, 24, 32\}$ and $\Delta t_y \in \{2, 8, 16, 24, 32, 40, 48\}$, and 30 independent simulations are performed for each pair of intervals.

For the second type of simulations, we draw independently and uniformly at random l translation parameters $\{(t_x)_1, \dots, (t_x)_l\}$ in the intervals $[-\Delta t_x/2, \Delta t_x/2]$ and l scaling parameters $\{s_1, \dots, s_l\}$ in the interval $[1 - \Delta s/2, 1 + \Delta s/2]$, while keeping t_y fixed at 0. This procedure is repeated with $\Delta t_x \in \{2, 8, 16, 24, 32\}$ and $\Delta s \in \{0, 0.13, 0.25, 0.38, 0.5\}$, and 30 independent simulations are also performed for each pair of intervals.

In our algorithm, we set $\kappa = 100$, $\lambda_{\theta} = 0.1$, $\lambda_{\mathbf{x}}^k = \max(0.9^k (50\kappa), 0.1)$, and $\mu = 10^{-10}$. The matrix Ψ is built using the Daubechies 8 wavelet basis, and $f(\mathbf{x})$ is set to $\|\Psi^T \mathbf{x}\|_1$. The parameters are constrained as follows: $|1 - s| \leq 0.5$, and $|t_x|, |t_y| \leq 64$.

To compare the registration quality of both methods, we compute the normalized root-mean-square-error $\sigma_{i \rightarrow j}$ between the ground truth parameters $\theta_{i \rightarrow j} \in \mathbb{R}^3$ that register the i^{th} image onto the j^{th} one and the corresponding parameters $\theta_{i \rightarrow j}^* \in \mathbb{R}^3$ estimated by either RASL or our method. These values are then averaged over all possible combinations of images to obtain the registration error

$$\sigma = \frac{1}{l(l-1)} \sum_{i=1}^l \sum_{j \neq i} \sigma_{i \rightarrow j} \quad \text{with} \quad \sigma_{i \rightarrow j} = \frac{\|\theta_{i \rightarrow j}^* - \theta_{i \rightarrow j}\|_2}{\|\theta_{i \rightarrow j}\|_2}.$$

Fig. 5 presents the registration error σ averaged over the 30 simulations obtained with RASL and our method for each type of simulations with $l = 10$. One can remark that for small transformations RASL is slightly more precise than our method. However, for large transformations, we obtain a smaller registration error. Our method is thus more robust to large displacements. Note that we also performed these simulations with $l = 5$ and the same conclusions hold.

4.2. Compressed sensing

In this section, we study the performance of our algorithm in a compressed sensing setting in comparison with other common algorithms.

We use $l = 5$ different images⁷ of the same scene to generate 5 different measurement vectors $\mathbf{y}_1, \dots, \mathbf{y}_l$. These images, presented in the top panels of Fig. 6, contain $n = 128 \times 128$ pixels. Note

⁶The image used is part of the castle-R20 dataset available at cvlab.epfl.ch/~strecha/multiview/rawMVS.html [37].

⁷All images are in the castle-R20 dataset available at cvlab.epfl.ch/~strecha/multiview/rawMVS.html [37].



Figure 6. *Top panel: 5 initial images. Middle-top panel: reconstructed images with method (4.19). Middle-bottom panel: reconstructed images with method (4.20). Bottom panel: reconstructed images with the proposed method.*




Figure 7. *Left panel: Reconstruction SNR as a function of m/n . Right panel: Reconstruction SNR as a function of the noise level σ . All curves represent the mean SNR over the 30×5 reconstructed images. The vertical lines show the variation of the SNR at one standard deviation. The dashed black, continuous blue, and dot-dashed red curves represent the reconstruction SNR for method (4.19), (4.20), and our algorithm, respectively.*




Figure 8. *Left panel: superposition of the 5 initial images without registration. Middle panel: superposition of the 5 initial images after registration with the estimated parameters $\theta_1^*, \dots, \theta_5^*$. Right panel: estimated background image \mathbf{x}_0^* .*

that parts of the scene are sometimes occluded. The measurement vectors are obtained using the spread spectrum technique of [30], which consists of pre-modulating each image with a random ± 1 sequence before randomly selecting m Fourier coefficients. Let $\mathbf{c} = (c_1, \dots, c_n)^\top \in \mathbb{R}^n$ be a Rademacher sequence, $\mathbf{F} \in \mathbb{C}^{n \times n}$ be the matrix that implements the 2D discrete Fourier transform, and $\Omega_j \in \mathbb{N}^m$ be l sets of m indices selected independently and uniformly at random in $\{1, \dots, n\}$. As we are dealing with images living in \mathbb{R}^n , we can restrict the selection of the Fourier coefficients to one half of the Fourier plane. We denote $\tilde{\mathbf{F}} \in \mathbb{C}^{n/2 \times n}$ the restriction of \mathbf{F} to this half plane. The observation matrices satisfy

$$\mathbf{A}_j = \begin{pmatrix} \tilde{\mathbf{F}}^r \\ \tilde{\mathbf{F}}^i \end{pmatrix}_{\Omega_j} \mathbf{C} \in \mathbb{R}^{m \times n}, \quad j = 1, \dots, l,$$

where $\tilde{\mathbf{F}}^r, \tilde{\mathbf{F}}^i \in \mathbb{R}^{n/2 \times n}$ are respectively the real and imaginary parts of $\tilde{\mathbf{F}}$, $\mathbf{C} \in \mathbb{R}^{n \times n}$ is the diagonal matrix containing the sequence \mathbf{c} , and $(\cdot)_{\Omega_j}$ restricts an $n \times n$ matrix to its rows indexed by Ω_j .

We assume here that the transformations between images can be represented by homographies and use model (3.16) for τ_θ . For the prior term f , we assume that the images are sparse in the Haar wavelet basis $\mathbf{W} \in \mathbb{R}^{n \times n}$. We build the block-diagonal matrix Ψ by repeating $l+1$ times the matrix \mathbf{W} on the diagonal, and set $f(\mathbf{x}) = \|\Psi^\top \mathbf{x}\|_1$. The same matrix Ψ is used in the cost-to-move term h_μ . In all simulations, we use $\lambda_\theta = 0.1$, $\lambda_{\mathbf{x}}^k = \max(0.9^k (20\kappa), 0.1)$, $\mu = 10^{-10}$, and the following constraints apply on the transformation parameters: $|1 - \theta_1|, |1 - \theta_5| \leq 0.2$, $|\theta_2|, |\theta_4| \leq 0.2$, $|\theta_3|, |\theta_6| \leq 20$, $|\theta_7|, |\theta_8| \leq 0.01$, with u_1 and u_2 belonging to $\{-63, -62, \dots, 64\}$.

In a first set of simulations, we study the performance of Algorithm 1 in the absence of noise for different number of measurements: $m \in \{0.1n, 0.2n, 0.3n, 0.4n, 0.5n, 0.7n\}$. In a second set of simulations, we study the performance of Algorithm 1 at $m = 0.3n$ for different noise levels. The noise vector \mathbf{n} follows an i.i.d zero-mean Gaussian distribution of standard deviation $\sigma \in \{0.01, 0.0177, 0.0316, 0.0562, 0.1, 0.177, 0.316\}$. The squared ℓ_2 -norm of the noise vector \mathbf{n}/σ follows a chi-square distribution with lm degrees of freedom and we compute the bound on the noise level $\|\mathbf{n}\|_2^2 \leq \epsilon^2$ using the 99th percentile of this distribution. Then, κ is set to 100 in the noiseless case and is chosen so that $\|\mathbf{A}(\theta^*)\mathbf{x}^* - \mathbf{y}\|_2 \approx \epsilon$ in the presence of noise⁸.

For comparison, we also reconstruct the images by solving the following convex optimization problems:

$$(4.19) \quad \min_{\mathbf{X}} \|\mathbf{W}^\top \mathbf{X}\|_1 \quad \text{subject to} \quad \|\mathbf{Y} - \mathbf{A}\mathbf{X}\|_F \leq \epsilon,$$

and

$$(4.20) \quad \min_{\mathbf{X}} \|\mathbf{W}^\top \mathbf{X}\|_{2,1} \quad \text{subject to} \quad \|\mathbf{Y} - \mathbf{A}\mathbf{X}\|_F \leq \epsilon,$$

⁸For each noise level, we use one simulation to estimate a value of κ such that, at converge of the algorithm, $\|\mathbf{A}(\theta^*)\mathbf{x}^* - \mathbf{y}\|_2 \in [0.99\epsilon, 1.01\epsilon]$. We then use the same value of κ for all subsequent simulations.

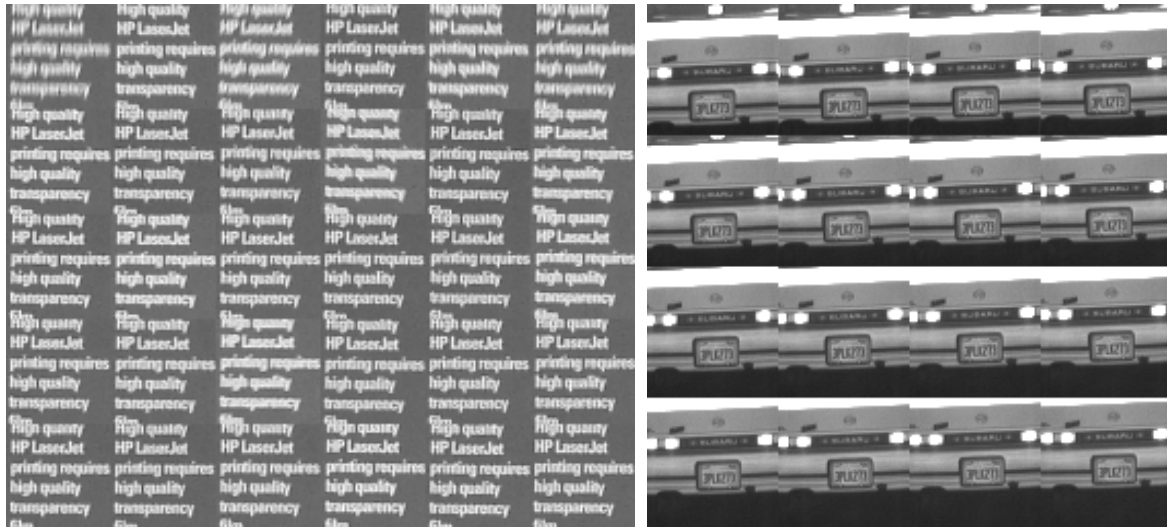


Figure 9. *Multi-frame datasets used to test the efficiency of our algorithm for super-resolution. The dataset on the right contains 30 low-resolution images of $m = 48 \times 48$ pixels. The dataset on the left contains 16 low-resolution images of $m = 64 \times 64$ pixels.*

where $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_l) \in \mathbb{R}^{n \times l}$, $\mathbf{Y} = (\mathbf{y}_1, \dots, \mathbf{y}_l) \in \mathbb{R}^{m \times l}$, $\|\cdot\|_F$ is the Frobenius norm, $\|\mathbf{X}\|_1 = \sum_{i=1}^n \sum_{j=1}^l |x_{ij}|$, and $\|\mathbf{X}\|_{2,1} = \sum_{i=1}^n (\sum_{j=1}^l x_{ij}^2)^{1/2}$, with x_{ij} the entry on the i^{th} line and j^{th} column of \mathbf{X} . In the first problem, the images $\mathbf{x}_1, \dots, \mathbf{x}_l$ are reconstructed with a sparse prior in the Haar wavelet basis. This is an extension of the Basis Pursuit problem, advocated by the compressed sensing theory [12], to multiple signals. In the second problem, the prior term favors reconstructions which have a similar sparsity pattern in the Haar wavelet basis [5], therefore imposing some correlation between images.

Fig. 7 shows the quality of the reconstructed images in term of SNR. Let $(\mathbf{x}^*, \boldsymbol{\theta}^*) \in \mathbb{R}^{(l+1)n} \times \mathbb{R}^{ln}$ be the images and transformation parameters recovered with our algorithm and $\mathbf{x} \in \mathbb{R}^{ln}$ be the l initial images. The reconstruction SNR satisfies $-20 \log_{10}(\|\mathbf{S}(\boldsymbol{\theta}_j^*)\mathbf{x}_0^* + \mathbf{x}_j^* - \mathbf{x}_j\|_2 / \|\mathbf{x}_j\|_2)$ in this case. For the two other methods (4.19) and (4.20), the reconstruction SNR satisfies $-20 \log_{10}(\|\mathbf{x}_j^* - \mathbf{x}_j\|_2 / \|\mathbf{x}_j\|_2)$ where $\mathbf{X}^* = (\mathbf{x}_1^*, \dots, \mathbf{x}_l^*) \in \mathbb{R}^{n \times l}$ are the reconstructed images with these methods. The first graph of Fig. 7 shows the reconstructions SNR as a function of m/n in the absence of noise, and the second graph shows the reconstructions SNR as a function of the noise level σ for $m/n = 0.3$. In each case, we perform 30 simulations with independent noise realization and choice of Ω_j . The curves show the mean SNR over the 30×5 reconstructed images. The vertical lines indicate the variation of the SNR at one standard deviation. One can deduce from these curves that the best reconstructions are obtained with our method.

In Fig. 6, we present the ground truth images as well as reconstructions obtained with the three different methods for $m/n = 0.3$ in the absence of noise. One can remark that the reconstructions obtained with our algorithm exhibits much finer details. Fig. 8 shows the estimated background image, next to the superposed initial images before and after registration with the estimated parameters. One can also note that the estimated background image is free of occlusions, and that the initial images are better aligned after registration with the estimated parameters.

4.3. Super-resolution

In this last experiment, we demonstrate the performance of our algorithm for super-resolution from multiple frames in comparison with the method described in [43], whose code is available at lcam.epl.ch/software/super-resolution. We test both algorithms on the datasets⁹ presented in Fig. 9.

⁹Datasets available at users.soe.ucsc.edu/~milanfar/software/sr-datasets.html (credit: Peyman Milanfar)

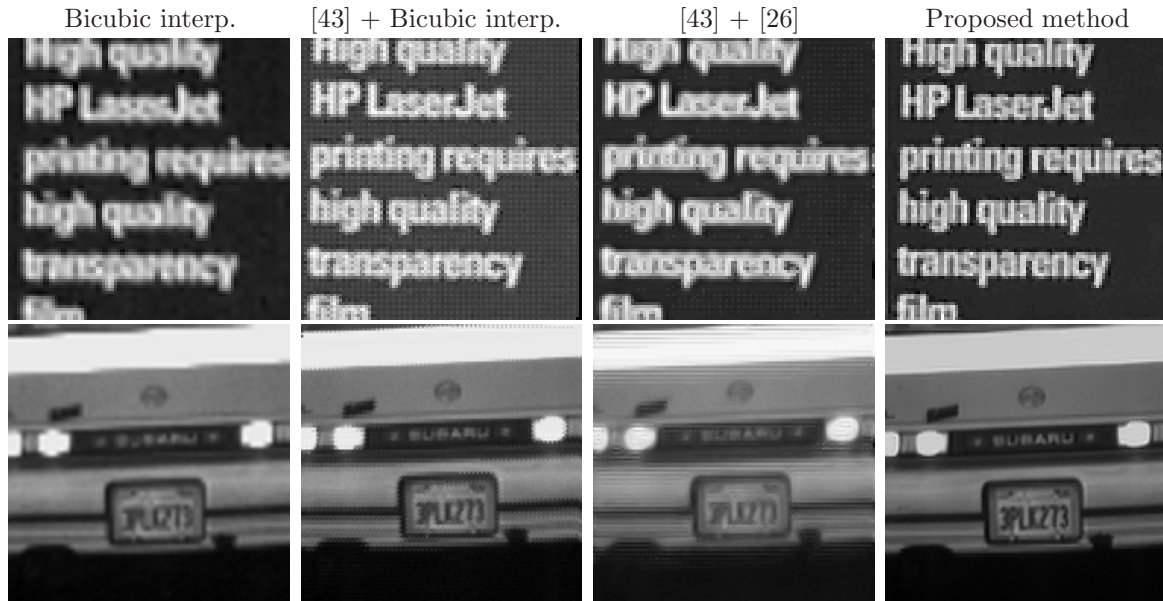


Figure 10. From left to right: one input image of the original dataset upsampled independently using a bicubic interpolation; super-resolved image using the method described in [43] in combination with a bicubic interpolation; super-resolved image using the method described in [43] in combination with [26]; super-resolved background image estimated by our algorithm.

To increase the resolution of the available images by a factor of 2 in both dimensions with our method, we use the measurement model (1.2) where all the sensing matrices A_j , with $j = 1, \dots, l$, model the same blurring and downsampling operator that averages all the pixels in a rectangular window of 2×2 pixels and uniformly downsamples the resulting blurred image by a factor of 2 in both dimensions. The vectors $\mathbf{y}_1, \dots, \mathbf{y}_l$ are initialized with the low-resolution images. The transformations between images are assumed to be affine, and we use model (3.15) for τ_θ . The transformation parameters are constrained as follows: $|1 - \theta_1|, |1 - \theta_5| \leq 0.2$, $|\theta_2|, |\theta_4| \leq 0.2$, $|\theta_3|, |\theta_6| \leq 30$, with u_1 and u_2 belonging to $\{-n/2 + 1, \dots, n/2\}$. The matrix Ψ is built using the Haar wavelet basis. We also assume that the high-resolution images $\mathbf{x}_0, \dots, \mathbf{x}_l$ are sparse in the gradient domain and set $f(\mathbf{x}) = \sum_{j=0}^l g(\mathbf{x}_j)$ in (2.4), where $g: \mathbb{R}^n \rightarrow \mathbb{R}$ is the isotropic Total Variation (TV) norm. Let us recall that the vector \mathbf{x}_j contains the samples $x_j(\mathbf{u}_k)$, with $k = 1, \dots, n$ and $\mathbf{u}_k = (u_1^k, u_2^k)$ (see Section 1.1). The isotropic TV norm g satisfies

$$g(\mathbf{x}_j) = \sum_{k=1}^n \sqrt{|\nabla_1 x_j(\mathbf{u}_k)|^2 + |\nabla_2 x_j(\mathbf{u}_k)|^2}, \text{ where } \begin{cases} \nabla_1 x_j(\mathbf{u}_k) = x_j(u_1^{k+1}, u_2^k) - x_j(u_1^k, u_2^k), \\ \nabla_2 x_j(\mathbf{u}_k) = x_j(u_1^k, u_2^{k+1}) - x_j(u_1^k, u_2^k), \end{cases}$$

assuming that $x(u_1^{n+1}, \cdot) = x(\cdot, u_2^{n+1}) = 0$. Note that g and, consequently, f are semi-algebraic. Finally, we run Algorithm 1 to minimize (2.4) with $\kappa = 100$, $\lambda_\theta = 0.1$, $\lambda_{\mathbf{x}}^k = \max(0.9^k (20\kappa), 0.1)$, and $\mu = 10^{-10}$.

The first dataset contains $l = 30$ low-resolution images of $m = 48 \times 48$ pixels. The super-resolved images thus contain $n = 96 \times 96$ pixels. In the top panel of Fig. 10, we present one super-resolved image obtained by upsampling independently one of the original images using a bicubic spline interpolation, our super-resolved background image, and two super-resolved images obtained with the code of [43]. These last two images are obtained by first registering the low resolution images with the technique described in [43]. The first high-resolution image is then obtained by using a bicubic interpolation to combine the information of all registered low-resolution images. The second high-resolution image is obtained using the technique described in [26]. On this first dataset, one can easily remark that our



Figure 11. *Super-resolved background image estimated by our algorithm using the Haar wavelets (left) or the Daubechies wavelets (right) for Ψ in the cost-to-move function h_μ .*

super-resolved image exhibits much sharper details than the images obtained with the other methods. In particular, the contours of the letters are better defined and the text is easier to read.

The second dataset contains $l = 16$ low-resolution images of $m = 64 \times 64$ pixels. The super-resolved images thus contain $n = 128 \times 128$ pixels. In Fig. 10, we present one super-resolved image obtained by upsampling independently one of the original images using a bicubic spline interpolation, our super-resolved background image, and two super-resolved images obtained with the code of [43], as described before. On this second dataset, one can notice that the brand of the car and license plate number are much better reconstructed with our method. The details are sharper, confirming again the superiority of our technique.

Let us mention the fact that the registration quality obtained with the method described in [43] is very sensitive to the choice of the reference frame. The results presented here are the best we were able to obtain after a careful choice of the reference frame. For most of the other choices of reference frame, this method was not able to register properly the low-resolution images and the subsequent super-resolved images were suffering of strong artifacts. Note that our method does not suffer of this issue as the algorithm automatically selects a reference coordinate system within the allowed range Θ .

Finally, to show the effect of the choice of Ψ in the cost-to-move function h_μ , we run again Algorithm 1 by building Ψ with the Daubechies 8 wavelets instead of the Haar wavelets. We present in Fig. 11 the reconstructions obtained with both wavelet bases. We do not observe any significant differences between both reconstructions. The choice of Ψ in the cost-to-move term h_μ thus does not seem to affect significantly the quality of the final reconstruction.

4.4. Running time of the algorithm

We finally conclude this section by mentioning the computation times required to minimize the non-convex functional (2.3) using our Matlab implementation of Algorithm 1. The running times mentioned have been obtained using a desktop computer (Windows 7, Intel Core i7 3.50Ghz, 16GB of RAM) and the parallel toolbox of Matlab to estimate the transformation parameters simultaneously.

For the robust image alignment problem, our implementation of the algorithm has been running for about 14 minutes to align 6 images of 256×256 pixels. For the compressed sensing problem, around 7 minutes have been necessary to reconstruct 5 images of 128×128 pixels from 30 percent of measurements. Finally, for the super-resolution problem, it has taken about 21 minutes to estimate a high resolution image of the car with 128×128 pixels from the 16 low-resolution input images.

Compared to the other methods we tested, the running time of our algorithm is longer. However, we would like to highlight that our implementation is general to be able to handle all these examples. We have not used the specificities of each problem to optimize the code and reduce the computation time. Furthermore, we have noticed that only a few iterations of the algorithm are necessary to estimate

the transformation parameters. During the last iterations, only the estimation of the background and foreground images are refined. Therefore, an approximate solution can be obtained by stopping the algorithm after a few iterations and then solve one last time problem (2.5) with $\lambda_{\mathbf{x}} = 0$ to estimate the images. Note however that the sequence of estimates does not converge to a critical point of L with such an approach.

5. Conclusion

We presented a method to solve multi-view imaging problems where one has to reconstruct several images of a scene from only few linear observations made at different viewpoints. Each observed image is modeled as the sum of a geometrically transformed background image and of a foreground image, modeling possible occlusions. We considered here global geometric transformations represented by few parameters, such as translations, affine transformations, or homographies. The proposed reconstruction method jointly estimates the images (background and foreground) and the transformation parameters by minimizing a non-convex functional. We studied the convergence of the proposed algorithm and showed that the generated sequence of estimates converges to a critical point of the non-convex functional for commonly used priors and transformation models. Numerical experiments show that the proposed algorithm is competitive with well-known methods in robust image alignment, compressed sensing, and super-resolution. The method is able to robustly register a set of images in the presence of large displacements, and the separation background/foreground is also very accurate, even when only few random measurements of the images are accessible.

Finally, we would like to conclude by highlighting the potential interest of the proposed technique for free breathing coronary magnetic resonance imaging (MRI) [36]. One major challenge in this application is to handle respiratory motion properly. The low acquisition speed in MRI forces researchers to invent new techniques to compensate for the inevitable motions occurring during the acquisition. To avoid motions due to heart contractions, an ECG signal is acquired to ensure that the Fourier measurements are always taken at the same instant in the heart cycle. A few measurements are then taken at each cycle. As the number of measurements acquired during one cycle is too small to accurately reconstruct an image of the heart, one should combine the measurements of several cycles to gather enough information. It is thus mandatory to properly compensate for the respiratory motion occurring between heart beats in the reconstruction method. Note that simple geometric transformations, such as translations, are already sufficient to reach good image quality [11]. Consequently, as the technique developed in this paper automatically compensates for such geometric transformations, it is a promising method for such an application. Note that this hope is confirmed by the preliminary results presented in [29] on real MRI data.

Appendix A

To establish the third point of Theorem 1, we need the following theorem established by Attouch *et al.* in [4].

Theorem 2 (Theorem 6.2, [4]). *Let p be a positive integer larger than 2, $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_p)$ be a vector in $\mathcal{H} = \mathbb{R}^{n_1} \times \dots \times \mathbb{R}^{n_p}$, and $g: \mathcal{H} \rightarrow \mathbb{R} \cup \{+\infty\}$ a function of the form $g(\mathbf{x}) = Q(\mathbf{x}) + \sum_{i=1}^p g_i(\mathbf{x}_i)$, where $Q: \mathcal{H} \rightarrow \mathbb{R}$ is a \mathcal{C}^1 function with locally Lipschitz continuous gradient and $g_i: \mathbb{R}^{n_i} \rightarrow \mathbb{R} \cup \{+\infty\}$ is a proper lower semicontinuous function, $i = 1, \dots, p$. Assume that g is bounded from below and satisfies the Kurdyca-Lojasiewicz property. Let $(\mathbf{x}^k)_{k \in \mathbb{N}}$ and $(\boldsymbol{\nu}^k)_{k \in \mathbb{N}}$ be sequences such that*

$$\begin{aligned}
 & g_i(\mathbf{x}_i^{k+1}) + Q(\mathbf{x}_1^{k+1}, \dots, \mathbf{x}_{i-1}^{k+1}, \mathbf{x}_i^{k+1}, \dots, \mathbf{x}_p^k) + a \|\mathbf{x}_i^{k+1} - \mathbf{x}_i^k\|_2^2 \\
 & \leq g_i(\mathbf{x}_i^k) + Q(\mathbf{x}_1^{k+1}, \dots, \mathbf{x}_{i-1}^{k+1}, \mathbf{x}_i^k, \dots, \mathbf{x}_p^k), \\
 \text{(A.21)} \quad & \boldsymbol{\nu}_i^{k+1} \in \partial g_i(\mathbf{x}_i^{k+1}), \\
 & \|\boldsymbol{\nu}_i^{k+1} + \partial_{x_i} Q(\mathbf{x}_1^{k+1}, \dots, \mathbf{x}_i^{k+1}, \mathbf{x}_{i+1}^k, \dots, \mathbf{x}_p^k)\|_2 \leq b \|\mathbf{x}_i^{k+1} - \mathbf{x}_i^k\|_2,
 \end{aligned}$$

where i ranges over $\{1, \dots, p\}$ and $a, b > 0$. Then, if the sequence $(\mathbf{x}^k)_{k \in \mathbb{N}}$ is bounded, it converges to some critical point \mathbf{x}^* of g .

The objective function L defined in (2.4) is equal to

$$L(\mathbf{x}, \boldsymbol{\theta}) = \kappa Q(\mathbf{x}, \boldsymbol{\theta}) + f(\mathbf{x}) + i_{\Theta}(\boldsymbol{\theta}),$$

where $Q: \mathbb{R}^{(\ell+1)n} \times \mathbb{R}^{\ell q} \rightarrow \mathbb{R}$ satisfies $Q(\mathbf{x}, \boldsymbol{\theta}) = \|\mathbf{A}(\boldsymbol{\theta})\mathbf{x} - \mathbf{y}\|_2^2$. The objective function has the form required by Theorem 2. Furthermore, f and i_{Θ} are proper lower semicontinuous functions, and, using the assumptions of Theorem 1, Q is twice continuously differentiable and thus have locally Lipschitz continuous gradient. We also note that L is bounded below and satisfies the Kurdyca-Lojasiewicz property. To prove the third point of Theorem 1, it remains to show that the sequence generated by Algorithm 1 is bounded and satisfies conditions (A.21).

First, due to the constraints that apply on $\boldsymbol{\theta}$, it is obvious that the sequence $(\boldsymbol{\theta}^k)_{k \in \mathbb{N}}$ is bounded. Therefore, assuming that $(\mathbf{x}^k)_{k \in \mathbb{N}}$ is bounded, as in the third point of Theorem 1, is enough to ensure the boundedness condition of the entire sequence of estimates.

Second, $\|\mathbf{x}^{k+1} - \mathbf{x}^k\|_2$ tends to 0 as $k \rightarrow +\infty$, as established in the second point of Theorem 1. Therefore there exists $k_0 \in \mathbb{N}$ such that for all $k \geq k_0$, $h_{\mu}(\Psi^{\top}(\mathbf{x}^{k+1} - \mathbf{x}^k)) = \|\Psi^{\top}(\mathbf{x}^{k+1} - \mathbf{x}^k)\|_2^2 / (2\mu) = \|\mathbf{x}^{k+1} - \mathbf{x}^k\|_2^2 / (2\mu)$, as $\Psi\Psi^{\top} = \mathbf{I}_{(\ell+1)n}$. Let k be larger than k_0 in the following. Inequality (2.6) shows that

$$f(\mathbf{x}^{k+1}) + \kappa Q(\mathbf{x}^{k+1}, \boldsymbol{\theta}^k) + \frac{\underline{\lambda}}{4\mu} \|\mathbf{x}^{k+1} - \mathbf{x}^k\|_2^2 \leq f(\mathbf{x}^k) + \kappa Q(\mathbf{x}^k, \boldsymbol{\theta}^k).$$

Furthermore, the first order optimality condition of Problem (2.5) shows that there exists $\boldsymbol{\nu}_{\mathbf{x}}^{k+1} \in \partial f(\mathbf{x}^{k+1})$ such that

$$\boldsymbol{\nu}_{\mathbf{x}}^{k+1} + \kappa \partial_{\mathbf{x}} Q(\mathbf{x}^{k+1}, \boldsymbol{\theta}^k) + \frac{\underline{\lambda} \mathbf{x}}{2\mu} (\mathbf{x}^{k+1} - \mathbf{x}^k) = 0,$$

which implies,

$$\|\boldsymbol{\nu}_{\mathbf{x}}^{k+1} + \kappa \partial_{\mathbf{x}} Q(\mathbf{x}^{k+1}, \boldsymbol{\theta}^k)\|_2 \leq \frac{\bar{\lambda}}{2\mu} \|\mathbf{x}^{k+1} - \mathbf{x}^k\|_2.$$

Finally, we deduce from equation (2.11) that

$$i_{\Theta}(\boldsymbol{\theta}^{k+1}) + \kappa Q(\mathbf{x}^{k+1}, \boldsymbol{\theta}^{k+1}) + \frac{\kappa \lambda_{\boldsymbol{\theta}}}{2} \|\boldsymbol{\theta}^{k+1} - \boldsymbol{\theta}^k\|_2^2 \leq i_{\Theta}(\boldsymbol{\theta}^k) + \kappa Q(\mathbf{x}^{k+1}, \boldsymbol{\theta}^k).$$

And, using the first order optimality condition of the problems (2.9), we conclude that there exists $\boldsymbol{\nu}_{\boldsymbol{\theta}_j}^{k+1} \in \partial i_{\Theta_j}(\boldsymbol{\theta}_j^{k+1})$ such that

$$\boldsymbol{\nu}_{\boldsymbol{\theta}_j}^{k+1} + \kappa \nabla Q_j^{k+1}(\boldsymbol{\theta}_j^k) + \kappa [\mathbf{H}(\boldsymbol{\theta}_j^k) + 2^i \lambda_{\boldsymbol{\theta}} \mathbf{I}_q](\boldsymbol{\theta}_j^{k+1} - \boldsymbol{\theta}_j^k) = 0,$$

for all $j = 1, \dots, l$. Let Γ denote the Lipschitz constant of ∇Q on a product of balls $\mathcal{B}_{\mathbf{x}} \times \mathcal{B}_{\boldsymbol{\theta}}$ containing the sequence $(\mathbf{x}_k, \boldsymbol{\theta}_k)_{k \in \mathbb{N}}$. By construction of i in Algorithm 1, we necessarily have $2^i \lambda_{\boldsymbol{\theta}} \leq \eta = \max(2\Gamma, \lambda_{\boldsymbol{\theta}})$. Remark also that, on the same products of balls, there exists Λ such that the singular values of the matrices $\mathbf{H}(\boldsymbol{\theta}_j^k)$, $j = 1, \dots, l$, are bounded above by Λ . Therefore,

$$\begin{aligned} \|\boldsymbol{\nu}_{\boldsymbol{\theta}_j}^{k+1} + \kappa \nabla Q_j^{k+1}(\boldsymbol{\theta}_j^{k+1})\|_2 &= \|\boldsymbol{\nu}_{\boldsymbol{\theta}_j}^{k+1} + \kappa \nabla Q_j^k(\boldsymbol{\theta}_j^k) + \kappa \nabla Q_j^{k+1}(\boldsymbol{\theta}_j^{k+1}) - \kappa \nabla Q_j^{k+1}(\boldsymbol{\theta}_j^k)\|_2 \\ &\leq \|\boldsymbol{\nu}_{\boldsymbol{\theta}_j}^{k+1} + \kappa \nabla Q_j^{k+1}(\boldsymbol{\theta}_j^k)\|_2 + \kappa \|\nabla Q_j^{k+1}(\boldsymbol{\theta}_j^{k+1}) - \nabla Q_j^{k+1}(\boldsymbol{\theta}_j^k)\|_2 \\ &\leq \kappa(\Lambda + \eta) \|\boldsymbol{\theta}_j^{k+1} - \boldsymbol{\theta}_j^k\|_2 + \kappa \Gamma \|\boldsymbol{\theta}_j^{k+1} - \boldsymbol{\theta}_j^k\|_2 \\ &= \kappa(\Lambda + \eta + \Gamma) \|\boldsymbol{\theta}_j^{k+1} - \boldsymbol{\theta}_j^k\|_2, \end{aligned}$$

which yields,

$$\|\boldsymbol{\nu}_{\boldsymbol{\theta}}^{k+1} + \kappa \partial_{\boldsymbol{\theta}} Q(\mathbf{x}^{k+1}, \boldsymbol{\theta}^{k+1})\|_2 \leq \kappa(\Lambda + \eta + \Gamma) \|\boldsymbol{\theta}^{k+1} - \boldsymbol{\theta}^k\|_2,$$

with $\boldsymbol{\nu}_{\boldsymbol{\theta}}^{k+1} = ((\boldsymbol{\nu}_{\boldsymbol{\theta}_1}^{k+1})^{\top}, \dots, (\boldsymbol{\nu}_{\boldsymbol{\theta}_l}^{k+1})^{\top})^{\top} \in \partial i_{\Theta}(\boldsymbol{\theta}^{k+1})$. The sequence $(\mathbf{x}^k, \boldsymbol{\theta}^k)_{k \geq k_0}$ thus satisfies conditions (A.21) with $a = \min(\underline{\lambda}/(4\mu), \kappa \lambda_{\boldsymbol{\theta}}/2)$ and $b = \max(\bar{\lambda}/(2\mu), \kappa(\Lambda + \eta + \Gamma))$. This terminates the proof of the third point of Theorem 1.

References

- [1] AFSHARI, H., JACQUES, L., BAGNATO, L., SCHMID, A., VANDERGHEYNST, P., AND LEBLEBICI, Y. The PANOPTIC camera: A plenoptic sensor with real-time omnidirectional capability. *J. Signal Process. Syst. Signal Image Video Technol.* *accepted* (2012).
- [2] ASIF, M. S., HAMILTON, L., BRUMMER, M., AND ROMBERG, J. Motion-Adaptive Spatio-Temporal Regularization (MASTeR) for accelerated dynamic MRI. *Magn. Reson. Med.*, *submitted* (2012).
- [3] ATTOUCH, H., BOLTE, J., REDONT, P., AND SOUBEYRAN, A. Proximal alternating minimization and projection methods for nonconvex problems: an approach based on the Kurdyka-Lojasiewicz inequality. *Math. Oper. Res.* *35*, 2 (2010), 438–457.
- [4] ATTOUCH, H., BOLTE, J., AND SVAITER, B. F. Convergence of descent methods for semi-algebraic and tame problems: proximal algorithms, forward-backward splitting, and regularized Gauss-seidel methods. *Math. Program.* (2011).
- [5] BACH, F., JENATTON, R., MAIRAL, J., AND OBOZINSKI, G. Optimization with sparsity-inducing penalties. *Foundations and Trends in Machine Learning* *4*, 1 (2012), 1–106.
- [6] BARON, D., WAKIN, M. B., DUARTE, M. F., SARVOTHAM, S., AND BARANIUK, R. G. Distributed compressed sensing. Tech. Rep. ECE-0612, Electrical and Computer Engineering Department, Rice University, 2006.
- [7] BAUSCHKE, H. H., AND COMBETTES, P. L. *Convex analysis and monotone operator theory in Hilbert spaces*. CMS Books in Mathematics. Springer New York, 2011.
- [8] BECK, A., AND TEOULLE, M. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J. Imaging Sci.* *2*, 1 (2009), 183–202.
- [9] BERG, E. V., FRIEDLANDER, M. P., HENNENFENT, G., HERRMANN, F., SAAB, R., AND YILMAZ, Ö. Sparco: A testing framework for sparse reconstruction. Tech. Rep. TR-2007-20, Dept. Computer Science, University of British Columbia, Vancouver, October 2007.
- [10] BERTSEKAS, D. P. Projected Newton methods for optimization problems with simple constraints. *SIAM J. Control and Optimization* *20*, 2 (1982), 221–246.
- [11] BONANNO, G., PUY, G., WIAUX, Y., VAN HEESWIJK, R., AND STUBER, M. About the performance of multi-dimensional radial self-navigation incorporating compressed sensing for free-breathing coronary MRI. In *International Society for Magnetic Resonance in Medicine (ISMRM) conference* (Melbourne, 2012).
- [12] CANDÈS, E. J. Compressive sampling. In *Proc. International Congr. of Mathematicians* (Madrid, Spain, 2006), vol. 3, pp. 1433–1452.
- [13] CHEN, X., AND FROSSARD, P. Joint reconstruction of compressed multi-view images. In *Proc. IEEE Conf. Acoustics, Speech and Signal Processing* (2009), pp. 1005 – 1008.
- [14] COMBETTES, P. L., AND J.-C-PESQUET. *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*. Springer Optimization and Its Applications. Springer New York, 2011, ch. Proximal Splitting Methods in Signal Processing, pp. 185–212.
- [15] DAI, Q., FU, C., JI, Y., AND ZHANG, Y. Robust joint reconstruction in compressed multi-view imaging. In *Proc. Picture Coding Symp.* (Poland, May 2012), pp. 13–16.
- [16] DUARTE, M. F., SARVOTHAM, S., WAKIN, M. B., BARON, D., AND BARANIUK, R. G. Joint sparsity models for distributed compressed sensing. In *Proc. Workshop on Signal Processing with Adaptive Sparse Structured Representations* (2005).
- [17] FARSIU, S., ROBINSON, D., ELAD, M., AND MILANFAR, P. Fast and robust multi-frame super-resolution. *IEEE Trans. Image Process.* *12*, 10 (2004), 1327–1344.
- [18] FU, C., JI, Y., AND DAI, Q. Compressed multi-view imaging with joint reconstruction. In *Proc. Data Compression Conf.* (2011), p. 448.
- [19] KEYS, R. Cubic convolution interpolation for digital image processing. *IEEE Trans. Acoustics, Speech, Signal Process.* *29*, 6 (1981).
- [20] KUBOTA, A., SMOLIC, A., MAGNOR, M., TANIMOTO, M., CHEN, T., AND ZHANG, C. Multiview imaging and 3DTV. *IEEE Signal Process. Mag.* *24*, 6 (2007), 10–21.
- [21] LI, X., WEI, Z., AND XIAO, L. Compressed sensing joint reconstruction for multi-view images. *Electronics Lett.* *46*, 23 (2010), 1548 – 1550.
- [22] MAMAGHANIAN, H., KHALED, N., ALONSO, D. A., AND VANDERGHEYNST, P. Design and exploration of low-power analog to information conversion based on compressed sensing. *IEEE J. Emerg. Sel. Topics in Circuits Syst.* (2012).
- [23] PARK, J. Y., AND WAKIN, M. B. A geometric approach to multi-view compressive imaging. *EURASIP J. Adv. Signal Process.* *2012*, 37 (2012).
- [24] PENG, Y., GANESH, A., WRIGHT, J., XU, W., AND MA, Y. RASL: Robust Alignment by Sparse and Low-rank decomposition for linearly correlated images. In *IEEE Conf. Computer Vision and Pattern Recognition* (2010), pp. 763 – 770.
- [25] PENG, Y., GANESH, A., WRIGHT, J., XU, W., AND MA, Y. RASL: Robust Alignment by Sparse and Low-rank decomposition for linearly correlated images. *IEEE Trans. Pattern Anal. Mach. Intell.* *34*, 11 (2012), 2233–2246.
- [26] PHAM, T. Q., VLIET, L. J. V., AND SCHUTTE, K. Robust fusion of irregularly sampled data using adaptive normalized convolution. *EURASIP J. Applied Signal Process.* *2006* (2006).

- [27] PROTTER, M., ELAD, M., TAKEDA, H., AND MILANFAR, P. Generalizing the nonlocal-means to super-resolution reconstruction. *IEEE Trans. Image Process.* 18, 1 (2009), 36–51.
- [28] PUY, G., AND VANDERGHEYNST, P. Robust joint reconstruction of misaligned images using semi-parametric dictionaries. In *ICML Workshop on Sparsity, Dictionaries and Projections in Machine Learning and Signal Processing* (Edinburgh, Scotland, 2012).
- [29] PUY, G., AND VANDERGHEYNST, P. Joint reconstruction of misaligned images from incomplete measurements for cardiac mri. In *International Conference on Sampling Theory and Applications (SAMPTA)* (Bremen, Germany, 2013).
- [30] PUY, G., VANDERGHEYNST, P., GRIBONVAL, R., AND WIAUX, Y. Universal and efficient compressed sensing by spread spectrum and application to realistic Fourier imaging techniques. *EURASIP J. Adv. Signal Process.* 2012, 6 (2012).
- [31] RAGUET, H., FADILI, J., AND PEYRÉ, G. A generalized forward-backward splitting. *SIAM J. Imaging Sci. to appear* (2013).
- [32] RUDIN, L. I., OSHER, S., AND FATEMI, E. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena* 60, 1-4 (1992), 259–268.
- [33] SANKARANARAYANAN, A. C., TURAGA, P. K., BARANIUK, R. G., AND CHELLAPPA, R. Compressive acquisition of dynamic scenes. In *Computer Vision - ECCV 2010* (2010), vol. 6311, pp. 129–142.
- [34] SCHMIDT, M., KIM, D., AND SRA, S. *Optimization for Machine Learning*. MIT press, 2011, ch. Projected Newton-type methods in machine learning.
- [35] SHEN, H., ZHANG, L., HUANG, B., AND LI, P. A MAP approach for joint motion estimation, segmentation, and super resolution. *IEEE Trans. Image Process.* 16, 2 (2007), 479–490.
- [36] STEHNING, C., BÖRNERT, P., NEHRKE, K., EGGERS, H., AND STUBER, M. Free-breathing whole-heart coronary MRA with 3D radial SSFP and self-navigated image reconstruction. *Magn. Reson. Med.* 54, 2 (2005), 476–480.
- [37] STRECHA, C., VON HANSEN, W., GOOL, L. V., FUA, P., AND THOENNESSEN, U. On benchmarking camera calibration and multi-view stereo for high resolution imagery. In *IEEE Conf. Computer Vision and Pattern Recognition* (2008), pp. 1 – 8.
- [38] THIRUMALAI, V., AND FROSSARD, P. Joint reconstruction of multi-view compressed images. *IEEE Trans. Image Process.*, submitted (2012).
- [39] TROCAN, M., MAUGEY, T., FOWLER, J. E., AND PESQUET-POPESCU, B. Disparity-compensated compressed-sensing reconstruction for multiview images. In *Proc. IEEE Conf. Multimedia and Expo* (2010), pp. 1225 – 1229.
- [40] TROCAN, M., MAUGEY, T., TRAMEL, E. W., FOWLER, J. E., AND PESQUET-POPESCU, B. Compressed sensing of multiview images using disparity compensation. In *Proc. IEEE Conf. Image Processing* (2010), pp. 3345 – 3348.
- [41] TSENG, P. Convergence of a block coordinate descent method for nondifferentiable minimization. *J. Optim. Theory Appl.* 109, 3 (2001), 475–494.
- [42] UNSER, M. Sampling—50 years after Shannon. *Proc. IEEE* 88, 4 (2000), 569– 587.
- [43] VANDEWALLE, P., SÜSSTRUNK, S., AND VERTELLI, M. A frequency domain approach to registration of aliased images with application to super-resolution. *EURASIP J. Applied Signal Process.* 2006 (2006), 71459.
- [44] WAKIN, M. B. A manifold lifting algorithm for multi-view compressive imaging. In *Proc. Picture Coding Symp.* (2009), pp. 1 – 4.
- [45] ZHANG, Z., GANESH, A., LIANG, X., AND MA, Y. TILT: Transform Invariant Low-Rank Textures. *Int. J. Comput. Vis.* 89, 1 (2012), 1–24.