Non-convex optimization for robust multi-view imaging

Gilles Puy* and Pierre Vandergheynst*
* Institute of Electrical Engineering, Ecole Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland.

***Abstract*—We study the multi-view imaging problem where one has to reconstruct a set of l images, representing a single scene, from a few measurements made at different viewpoints. We first express the solution of the problem as the minimizer of a non-convex objective function where one needs to estimate one reference image, l foreground images modeling possible occlusions, and a set of l transformation parameters modeling the inter-correlation between the observations. Then, we propose an alternating descent method that attempts to minimize this objective function and produces a sequence converging to one of its critical points. Finally, experiments show that the method accurately recovers the original images and is robust to occlusions.**

I. Problem formulation

In multi-view imaging, we have in hand l observations $\boldsymbol{y}_1, \ldots, \boldsymbol{y}_l \in \mathbb{R}^m$ of a reference image $\boldsymbol{x}_0 \in \mathbb{R}^n$. As these observations are done from different viewpoints, the image $\boldsymbol{x}_0$ undergoes geometric transformations. We consider here transformations represented by few parameters (e.g., homography) and denote $\boldsymbol{\theta}_j \in \mathbb{R}^q$ the parameters associated to the j^{th} observations. The reference image transformed according to $\boldsymbol{\theta}_j$ is estimated using, e.g., a cubic spline interpolation and is equal to $\mathsf{S}(\boldsymbol{\theta}_j)\boldsymbol{x}_0$, with $\mathsf{S}(\boldsymbol{\theta}_j) \in \mathbb{R}^{n \times n}$.

To handle realistic applications, we also assume that parts of the reference image might sometimes be occluded. We model these occlusions using l foreground images $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_l \in \mathbb{R}^n$, and assume that the image "viewed" by the j^{th} observer is $\mathsf{S}(\boldsymbol{\theta}_j)\boldsymbol{x}_0 + \boldsymbol{x}_j$.

Finally, we model the acquisition device using a linear operator $\mathsf{A} \in \mathbb{R}^{m \times n}$, and the observation model satisfy

$$\begin{bmatrix} \boldsymbol{y}_1 \\ \vdots \\ \boldsymbol{y}_l \end{bmatrix} = \begin{bmatrix} \mathsf{AS}(\boldsymbol{\theta}_1) & \mathsf{A} & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ \mathsf{AS}(\boldsymbol{\theta}_l) & 0 & \ldots & \mathsf{A} \end{bmatrix} \begin{bmatrix} \boldsymbol{x}_0 \\ \vdots \\ \boldsymbol{x}_l \end{bmatrix} + \begin{bmatrix} \boldsymbol{n}_1 \\ \vdots \\ \boldsymbol{n}_l \end{bmatrix}, \quad (1)$$

where $\boldsymbol{n}_1, \ldots, \boldsymbol{n}_l \in \mathbb{R}^m$ represent additive measurement noise.

To reconstruct the images $\boldsymbol{x}^\intercal = (\boldsymbol{x}_0^\intercal, \ldots, \boldsymbol{x}_l^\intercal)$ and the transformation parameters $\boldsymbol{\theta}^\intercal = (\boldsymbol{\theta}_1^\intercal, \ldots, \boldsymbol{\theta}_l^\intercal)$ from the observations $\boldsymbol{y}^\intercal = (\boldsymbol{y}_1^\intercal, \ldots, \boldsymbol{y}_l^\intercal)$, we wish to solve the following problem

$$\min_{(\boldsymbol{x}, \boldsymbol{\theta})} \{ L(\boldsymbol{x}, \boldsymbol{\theta}) = \|\Psi^\intercal \boldsymbol{x}\|_1 + \kappa \|\mathsf{A}(\boldsymbol{\theta})\,\boldsymbol{x} - \boldsymbol{y}\|_2^2 + \sum_{1 \leqslant j \leqslant l} i_{\Theta_j}(\boldsymbol{\theta}_j) \}, \quad (2)$$

which is non-convex. The matrix $\Psi \in \mathbb{R}^{(l+1)n \times (l+1)n}$ is block-diagonal and built by repeating $l+1$ times, e.g., the Haar wavelet basis on the diagonal, $\kappa > 0$ is a regularizing parameter, $\mathsf{A}(\boldsymbol{\theta}) \in \mathbb{R}^{lm \times (l+1)n}$ is the matrix appearing in (1), $(\Theta_j)_{1 \leqslant j \leqslant l}$ are closed convex subsets of $\mathbb{R}^q$, and i_{Θ_j} is the indicator function of Θ_j.

II. Non-convex optimization

To solve problem (2), we propose an alternating descent method producing a sequence of estimates $(\boldsymbol{x}^k, \boldsymbol{\theta}^k)_{k \in \mathbb{N}}$, which converges to a critical point of L. The algorithm is inspired by recent results in non-convex optimization [1], [2], and consists of two main steps that sequentially decrease the value of the objective function.

First, we update the images while keeping the parameters fixed. Let $(\boldsymbol{x}^k, \boldsymbol{\theta}^k)$ be the estimates obtained after k iterations, and $(\lambda_{\boldsymbol{x}}^k)_{k \in \mathbb{N}} > 0$ be a decreasing sequence. The next estimate satisfies

$$\boldsymbol{x}^{k+1} \in \operatorname*{argmin}_{\boldsymbol{x}} L(\boldsymbol{x},\ \boldsymbol{\theta}^k) + \frac{\lambda_{\boldsymbol{x}}^k}{2}\ h(\Psi^\intercal(\boldsymbol{x} - \boldsymbol{x}^k)),$$

Fig. 1. From left to right: 3rd initial image; 3rd reconstructed image; estimated reference image $\boldsymbol{x}_0$; SNR vs. m/n for the proposed method (red) and the BP problem (black). The curves represent the mean SNR over 30 simulations, and the vertical lines represent the error at 1 standard deviation.

where h is a smooth approximation of the ℓ_1-norm. We noticed that the addition of the cost term $\lambda_{\boldsymbol{x}}^k\, h(\Psi^\intercal(\boldsymbol{x} - \boldsymbol{x}^k))/2$ produces a reconstruction of the images in coarse-to-fine scales fashion and improves the accuracy of the estimated transformation parameters.

Then, we update the transformation parameters by minimizing a quadratic approximation of $\|\mathsf{A}(\boldsymbol{\theta})\,\boldsymbol{x} - \boldsymbol{y}\|_2^2$. To simplify notations, we introduce l new functions $Q_j(\boldsymbol{\theta}_j) = \|\mathsf{A}_j \mathsf{S}(\boldsymbol{\theta}_j)\boldsymbol{x}_0^{k+1} + \mathsf{A}_j \boldsymbol{x}_j^{k+1} - \boldsymbol{y}\|_2^2$, with $j = 1, \ldots, l$. Let $\mathsf{I} \in \mathbb{R}^{q \times q}$ be the identity matrix, and $\lambda_{\boldsymbol{\theta}} > 0$. Assuming that the entries of $\mathsf{S}(\boldsymbol{\theta}_j)$ are differentiable with respect to $\boldsymbol{\theta}_j$, the next estimates $\boldsymbol{\theta}_j^{k+1}$ is chosen as the minimizer of

$$P_j(\boldsymbol{\theta}_j) = \nabla Q_j(\boldsymbol{\theta}_j^k)^\intercal(\boldsymbol{\theta}_j - \boldsymbol{\theta}_j^k) + (\boldsymbol{\theta}_j - \boldsymbol{\theta}_j^k)^\intercal \frac{\mathsf{H}_j^k + 2^i \lambda_{\boldsymbol{\theta}} \mathsf{I}}{2} (\boldsymbol{\theta}_j - \boldsymbol{\theta}_j^k),$$

where i is the smallest positive integer such that $Q_j(\boldsymbol{\theta}_j^{k+1}) + \lambda_{\boldsymbol{\theta}} \|\boldsymbol{\theta}_j^{k+1} - \boldsymbol{\theta}_j^k\|_2^2/2 \leqslant Q_j(\boldsymbol{\theta}_j^k) + P_j(\boldsymbol{\theta}_j^{k+1})$. In the above equations, the matrix $\mathsf{H}_j^k = 2\left(\mathsf{AJ}_j^k\right)^\intercal\left(\mathsf{AJ}_j^k\right)$ with $\mathsf{J}_j^k = \left(\partial_{\theta_{1j}} \mathsf{S}(\boldsymbol{\theta}_j^k), \ldots, \partial_{\theta_{qj}} \mathsf{S}(\boldsymbol{\theta}_j^k)\right) \boldsymbol{x}_0^{k+1} \in \mathbb{R}^{n \times q}$.

III. Experiments and conclusion

We test the proposed method using 5 images of the same scene, taken from different viewpoints, and containing occlusions. We generate 5 measurement vectors using the compressed sensing technique of [3]. Fig. 1 shows the 3rd initial image next to the corresponding reconstructed image from $m = 0.3n$ measurements. The estimated reference image, free of occlusions, is also presented. We also show the curves of the reconstruction SNR as a function of m/n obtained with our method and by solving the Basis Pursuit (BP) problem, which does not benefit from the correlation between measurements. Our method exhibits better reconstruction qualities.

We have presented a method for the joint reconstruction of a set of misaligned images. Our algorithm is an alternating descent method that produces a sequence converging to a critical point of L. Experiments show that the method correctly estimates the underlying reference image $\boldsymbol{x}_0$, is robust to occlusions, and benefits from the inter-correlation between measurements.

References

[1] Attouch et al., "Proximal alternating minimization and projection methods for nonconvex problems: An approach based on the kurdyka-lojasiewicz inequality," *Math. Oper. Res.*, vol. 35(2), pp. 438–457, 2010.

[2] Attouch et al., "Convergence of descent methods for semi-algebraic and tame problems: proximal algorithms, forward backward splitting, and regularized gauss-seidel methods," *Math. Program.*, 2011.

[3] Puy et al., "Universal and efficient compressed sensing by spread spectrum and application to realistic fourier imaging techniques," *EURASIP J. Adv. Signal Process.*, vol. 2012(6), 2012.