



HAL
open science

A Non-Local Low-Rank Approach to Enforce Integrability

Hicham Badri, Hussein Yahia

► **To cite this version:**

Hicham Badri, Hussein Yahia. A Non-Local Low-Rank Approach to Enforce Integrability. IEEE Transactions on Image Processing, 2016. hal-01317151

HAL Id: hal-01317151

<https://inria.hal.science/hal-01317151>

Submitted on 18 May 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Non-Local Low-Rank Approach to Enforce Integrability

Hicham Badri and Hussein Yahia

Abstract—We propose a new approach to enforce integrability using recent advances in non-local methods. Our formulation consists in a sparse gradient data-fitting term to handle outliers together with a gradient-domain non-local low-rank prior. This regularization has two main advantages : 1) the low-rank prior ensures similarity between non-local gradient patches, which helps recovering high-quality clean patches from severe outliers corruption, 2) the low-rank prior efficiently reduces dense noise as it has been shown in recent image restoration works. We propose an efficient solver for the resulting optimization formulation using alternate minimization. Experiments show that the new method leads to an important improvement compared to previous optimization methods and is able to efficiently handle both outliers and dense noise mixed together.

Index Terms—Integrability, non-local methods, low-rank, surface reconstruction.

I. INTRODUCTION

Reconstruction from gradients is an important step in various computer vision and graphics applications such as shape-from-shading (SfS) [1] and photometric stereo (PS) [2]. PS and SfS methods first compute the surface normals. An estimated gradient is then calculated from the normals and used to reconstruct the depth map. However, due to perturbations in the input images such as noise, outliers, shadows and other sources, the estimated gradient field is subject to high-magnitude corruption (outliers) and dense corruption (noise). This downgrades the gradient field to a non-integrable vector field. As a result, a straightforward integration approach results in a deformed surface with various artifacts. The problem of integration is not limited only to surface reconstruction from a non-integrable field (SfG). For instance, in Adaptive Optics (AO) [3], the wavefront phase reconstruction consists in recovering the phase from low-resolution corrupted gradient measurements. Various low-level gradient-domain processing applications directly manipulate gradient fields via transformations or mixing various gradients [4], [5], [6]. In all these cases, the resulting transformed gradient field is no longer integrable and direct integration results in various disturbing artifacts.

This paper proposes a new optimization formulation for robust integration to handle dense and sparse high-magnitude gradient-domain corruptions. Motivated by recent advances in image restoration [7], [8], we propose to use a sparse gradient residual formulation coupled with a non-local regularization. The non-local regularization is substantially different from

previous integration works as it manipulates non-local patches of the vector field instead of independent gradient points. By iteratively gathering non-local gradient patches, stacking them in a matrix and forcing this matrix to be low-rank, our method is able to significantly correct imperfections in the vector field. We propose an efficient way to solve the corresponding problem via alternate minimization. To our knowledge, this is the first approach that uses a non-local prior in the context of integrability enforcement/surface-from-gradients. Experiments on synthetic and real data demonstrate the excellent performance of the proposed method even in extreme mixed sparse/dense corruptions.

II. RELATED WORK

A. Integrability Enforcement

The problem of integrability enforcement has received important attention in the vision community since early works on SfS and PS applications. The most popular and straightforward approach consists in formulating the problem in terms of a quadratic energy, which results in the Poisson equation as a solution [9]. Frankot and Chellappa propose to perform a projection of the non-integrable gradient field onto the set of Fourier basis. Kovessi [10] performed the projection on shapelets. The work in [11] uses loopy belief propagation scheme to deal with dense noise in the vector field. A generalization of the Poisson integration framework was proposed by Agrawal *et al.* [12], which uses various methods such as M-estimation, α -surface and diffusion. An algebraic approach to deal with sparse corruptions only in the vector field based on the zero-curl constraint was proposed in [13]. Spectral and Tikhonov regularization were proposed in [14] to improve robustness of the least-squares fitting. Sparsity-based methods have been proposed recently to deal with both sparse and dense corruptions. The work in [15] proposes to use the l_1 -norm instead of the l_2 -norm in the integrability formulation. Similarly, the method in [16] uses the l_1 -norm to model the sparsity of the residual gradient and a minimum curl constraint, which improves robustness to high-magnitude sparse corruptions. The method in [17] uses notions from Compressed Sensing to improve surface recovery. The method in [18] proposes to use a sparse gradient regularization together with a sparse gradient residual formulation, which has shown to significantly improve robustness to both outliers and noise corruptions. Another set of methods known as kernel methods [19] follow a different direction to deal with gradient-domain perturbations by using kernel basis functions for high-dimensional fitting. Finally, the approach proposed by Xie *et al.* in [20] formulates the surface reconstruction problem using discrete geometry and performs

H. Badri and H. Yahia are with INRIA Bordeaux Sud-Ouest.
E-mail: hicham.badri@inria.fr, hussein.yahia@inria.fr.

processing on the normals not on the vector field to reconstruct a 3D mesh. Contrary to this approach, we study in this paper the problem of reconstruction from corrupted gradient fields in the general case. Our solution is thus not limited only to surface-from-gradients problems.

B. Non-Local Methods

Non-local methods have emerged as a powerful approach to various image restoration tasks. The idea consists in gathering similar patches via block-matching (called *non-local patches*), stack them in a matrix such as each column is a vectorized version of a patch, and then process this matrix using a transformation. The standard block-matching approach widely used in low-level vision consists in clustering patches that minimize the l_2 -distance in a neighborhood. More sophisticated and faster block-matching approaches were proposed as well [21]. Contrary to local methods that manipulate each pixel separately, non-local methods take into account the self-similarities within an image, which leads to high-quality recovery even in the case of high-level corruptions. The first works on non-local methods that were proposed tackle the problem of image denoising. Various non-local transformations have been proposed : weighted averaging [22], 3D collaborative filtering [23], sparse coding/dictionary learning [24] and low-rank estimation [7], [8]. Some of the methods have been extended to image deconvolution [25] and more recently to optical flow estimation [26] and image editing [27]. In this paper, we propose to use this powerful prior for the first time in the context of integrability enforcement.

C. Low-Rank Estimation

Low-rank estimation has led to many important improvements in various applications as it can efficiently recover missing data and deal with outliers [28], [29], [30]. The approach consists mainly in reducing the magnitude of the singular values of a matrix via a shrinkage operator, which forces the columns of the matrix to be more similar, and thus get rid of corruptions. Particularly, it has shown to produce interesting results in the context of image restoration via non-local denoising [7], [8]. Another work worth mentioning uses a low-rank formulation for Photometric Stereo normals estimation [31]. Note however that this work addresses a completely different problem than robust integration. Moreover, the work that we present in this paper uses low-rank estimation as a non-convex *regularization* that results in a new mathematical formulation, while the method in [31] uses a robust PCA formulation for matrix separation.

In a nutshell, we propose to use a new optimization approach to directly recover the surface/image from a severely corrupted vector field, motivated by recent advances in non-local methods/low-rank estimation. Contrary to previous work, we use a *non-local* regularization that manipulates patches of the vector field and imposes low-rankness on the corresponding non-local patch matrices. This approach has many advantages : 1) it can efficiently reduce dense corruptions such as noise, 2) it can efficiently deal with high-magnitude

corruptions such as outliers, 3) it can recover missing data. The corresponding optimization problem can be efficiently solved via alternate minimization. We evaluate the new technique on synthetic and real data and compare with previous works. Our method shows excellent performance even in challenging situations when the vector field contains a high amount of sparse and dense corruptions.

III. PROBLEM FORMULATION

Let $S(x, y)$ be the desired discrete surface to recover. We denote by s its vectorized form (which corresponds to concatenating each column in one vector). Let $\vec{v} = (v_x, v_y)$ be the given corrupted gradient field in the vectorized form. Typically, the given vector field \vec{v} is an estimate of the true gradient of the surface that we denote by $\nabla s = (s_x, s_y)$, which is subject to dense noise and high-magnitude outliers. Due to the present of high-magnitude outliers in \vec{v} , the residual gradient ($\nabla s - \vec{v}$) tends to be sparse. This observation was made in previous works [16], [15], [18]. Non-convex sparsity via the $l_{p<1}$ -norm has shown to produce better results than the l_1 -norm in various situations as it promotes more sparsity [32], [33], [18]. We use in this paper the non-convex $l_{p<1}$ -norm to model sparsity of the residual gradient as it reflects better the statistics of the residual. Figure 1 plots the empirical distribution of residual gradients coming from a Photometric Stereo experiment. As can be seen, the kurtotic hyper-Laplacian distribution with a low p value (in red) reflects better the empirical distribution (in black). This suggests that the $l_{p<1}$ -norm is a better choice than the l_1 -norm to model residual gradients.

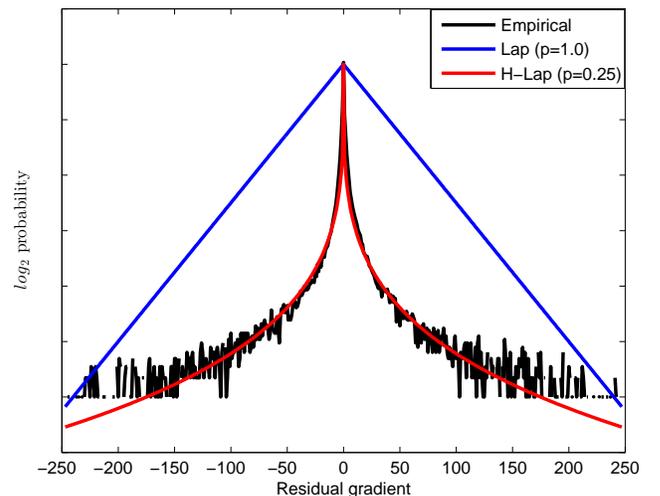


Fig. 1: log distribution of the residual gradient error in a Photometric Stereo experiment. A kurtotic distribution is needed to better model the empirical distribution of the residual.

Due to the high amount of corruptions that may affect the estimated gradient field, a regularization is necessary. We propose to use a new regularization via non-local low-rank estimation to recover a good quality surface even in severe corruption situations. This prior consists in gathering patches

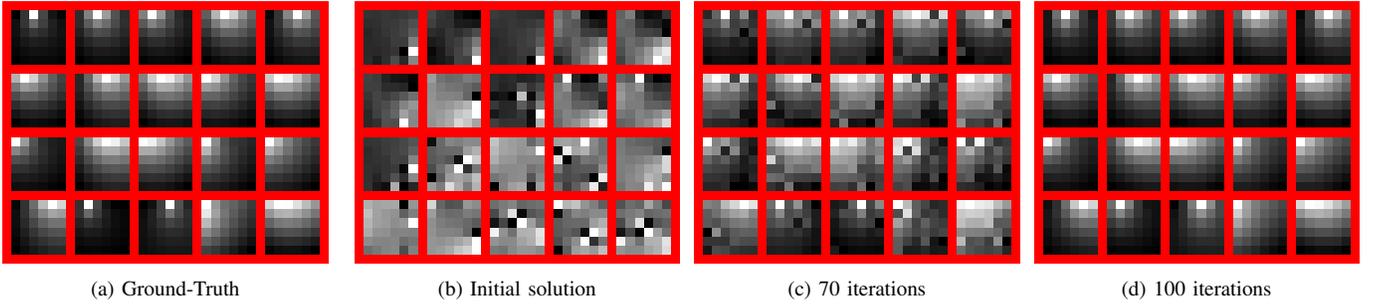


Fig. 2: Progressive non-local gradient patch recovery example. Our method starts with an initial solution calculated via the Poisson equation and iteratively improves the result by forcing low-rankness of block-matched patches. In this experiment, the corrupted gradient field contains 10% of outliers and the surface used is the popular Mozart dataset.

of the estimated vector field via block-matching, stack them as columns in a matrix and force this matrix to be low-rank. Reducing the rank of this matrix efficiently attenuates dense corruptions in the estimated vector field, and also helps to correct more outliers or missing data. We use a standard block-matching that clusters patches in a neighborhood based on their l_2 -distance similar to what is used in [7], [8]. We start with an initial solution given by the Poisson equation and iteratively improve recovery. Figure 2 shows an example of a non-local gradient patch cluster and how the proposed method iteratively improves the result. As can be seen, the clustered patches tend to be very similar. Low-rankness is thus very effective to get rid of corruptions.

Let X be a matrix. Its Singular Values Decomposition (SVD) is given by $X = U \text{diag}(\sigma) V^T$, where σ are the singular values of X and $\text{diag}(\cdot)$ is a diagonal operator. Low-rank estimation consists in forcing σ to be sparse typically by applying a shrinkage operator. The most popular way to promote sparsity on the singular values is known as the so-called nuclear norm [34], which comes to use the l_1 -norm on σ ($\|X\|_* = \|\sigma\|_1$). Considering the $l_{p < 1}$ -norm that promotes more sparsity than the l_1 case, we propose to use the “ l_p -nuclear norm” that we denote by $\|X\|_{*,p} = \|\sigma\|_p^p$ as a prior in our problem.

Putting it all together, the proposed optimization formulation is given as follows

$$\underset{s}{\operatorname{argmin}} \|\nabla s - \vec{v}\|_{p_1}^{p_1} + \frac{\lambda}{g} \sum_{j=1}^g \left\{ \|R_x^j D_x s\|_{*,p_2} + \|R_y^j D_y s\|_{*,p_2} \right\}, \quad (1)$$

where λ is a positive regularization term, g is the number of clusters of non-local patches, D_x, D_y and are discrete differential operators and R_x^j, R_y^j are binary matrices that select block-matched patches for cluster j and $p_{1,2} \leq 1$.

Optimization

Problem (1) is hard to evaluate directly. We use alternate minimization to estimate the latent surface/image s . This consists in introducing additional variables w_s, w_x^j, w_y^j and split the problem into sub-problems that are easier to solve. We use a Half-Quadratic (HQ) solver [35] that is widely used

in the computer vision community. Applying this approach to problem (1) results in the following

$$\underset{s, w_s, w_x^j, w_y^j}{\operatorname{argmin}} \left\{ \|w_s\|_{p_1}^{p_1} + \frac{\beta}{2} \|(\nabla s - \vec{v}) - w_s\|_2^2 + \frac{\lambda}{g} \sum_{j=1}^g \left\{ \|w_x^j\|_{*,p_2} + \frac{\beta}{2} \|R_x^j D_x s - w_x^j\|_2^2 \right\} + \frac{\lambda}{g} \sum_{j=1}^g \left\{ \|w_y^j\|_{*,p_2} + \frac{\beta}{2} \|R_y^j D_y s - w_y^j\|_2^2 \right\} \right\}, \quad (2)$$

where β is a new positive regularization parameters. The corresponding subproblems for alternate minimization are given by considering one variable a time as follows

$$\begin{aligned} (p_1) : w_s^{(k+1)} &\leftarrow \underset{w_s}{\operatorname{argmin}} \|w_s\|_{p_1}^{p_1} + \frac{\beta}{2} \|(\nabla s^{(k)} - \vec{v}) - w_s\|_2^2 \\ (p_2) : w_x^{j,(k+1)} &\leftarrow \underset{w_x^j}{\operatorname{argmin}} \|w_x^j\|_{*,p_2} + \frac{\beta}{2} \|R_x^j D_x s^{(k)} - w_x^j\|_2^2 \\ (p_3) : w_y^{j,(k+1)} &\leftarrow \underset{w_y^j}{\operatorname{argmin}} \|w_y^j\|_{*,p_2} + \frac{\beta}{2} \|R_y^j D_y s^{(k)} - w_y^j\|_2^2 \\ (p_4) : s^{(k+1)} &\leftarrow \underset{s}{\operatorname{argmin}} \|(\nabla s - \vec{v}) - w_s^{(k+1)}\|_2^2 + \\ &\frac{\lambda}{g} \sum_{j=1}^g \left\{ \|R_x^j D_x s - w_x^{j,(k+1)}\|_2^2 + \|R_y^j D_y s - w_y^{j,(k+1)}\|_2^2 \right\}. \end{aligned} \quad (3)$$

As problems (p_1) , (p_2) and (p_3) are in the proximal form, they can be efficiently solved via pixel-wise shrinkage. Problem (p_4) is quadratic and admits a closed-form.

Solving problem (p_1)

Problem (p_1) takes the following proximal form

$$\underset{x}{\operatorname{argmin}} \frac{1}{\beta} \|x\|_p^p + \frac{1}{2} \|x - y\|_2^2. \quad (4)$$

This problem admits a closed-form solution in the convex case $p = 1$ via soft-thresholding. The non-convex case $p < 1$ admits a first-order approximation [32], [18] (see Appendix)

$$x = \operatorname{shrink}_{l_p}\left(y, \frac{1}{\beta}\right) = \max\left(0, |y| - \frac{1}{\beta}|y + \epsilon|^{p-1}\right) \operatorname{sign}(y). \quad (5)$$

Applying this result to problem (p_1) gives the following solution

$$w_s^{(k+1)} = \begin{cases} w_{s,x}^{(k+1)} = \operatorname{shrink}_{l_{p_1}}\left(D_x s^{(k)} - v_x, \frac{1}{\beta}\right) \\ w_{s,y}^{(k+1)} = \operatorname{shrink}_{l_{p_1}}\left(D_y s^{(k)} - v_y, \frac{1}{\beta}\right) \end{cases} \quad (6)$$

Note that this first-order proximal approach is an approximation which is as fast as the standard nuclear norm minimization. The method in [36] presents a more advanced optimization approach for more accurate results. However, this approach is slower as it requires solving a weighted nuclear norm minimization problem for various iterations.

Solving problems (p₂) and (p₃)

Problems (p₂) and (p₃) take the following general form

$$\operatorname{argmin}_X \frac{1}{\beta} \|X\|_{*,p} + \frac{1}{2} \|X - Y\|_2^2. \quad (7)$$

Consider the SVD of $X = U \operatorname{diag}(\sigma) V^T$, where σ are the singular values. The l_p - nuclear norm is expressed in terms of the singular values only $\|X\|_{*,p} = \|\sigma\|_p^p$. As a result, a first-order approximation of (7) can be expressed using the l_p -thresholding operator (5) as follows

$$X = \operatorname{shrink}_{*,p}(y, \frac{1}{\beta}) = U \operatorname{diag}(\operatorname{shrink}_{l_p}(\sigma, \frac{1}{\beta})) V^T. \quad (8)$$

Applying this result to problems (p₂) and (p₃) gives the following solutions

$$\begin{aligned} w_x^{j,(k+1)} &= \operatorname{shrink}_{*,p_2}(R_x^j D_x s^{(k)}, \frac{1}{\beta}) \\ w_y^{j,(k+1)} &= \operatorname{shrink}_{*,p_2}(R_y^j D_y s^{(k)}, \frac{1}{\beta}). \end{aligned} \quad (9)$$

Solving problem (p₄)

Problem (p₄) is quadratic, hence it admits a closed-form. However, it is not straightforward to get the closed-form due to the presence of the non-local matrices R_x^j and R_y^j . For this reason, we need a reformulation of this form

$$\frac{1}{g} \sum_{j=1}^g \|R^j D s - w^j\|_2^2 \equiv \|D s - w\|_2^2. \quad (10)$$

Once this formulation established, the solution $s^{(k+1)}$ can be efficiently calculated via Fourier transform by considering periodic boundary conditions. Consider the general problem

$$\operatorname{argmin}_x \frac{1}{g} \sum_j \|A^j x - b^j\|_2^2, \quad (11)$$

which admits the following closed-form

$$\frac{1}{g} \left(\sum_j A^{jT} A^j \right) x = \frac{1}{g} \sum_j A^{jT} b^j. \quad (12)$$

Applying this result to problem (11) results in the following solution

$$\left(\frac{1}{g} \sum_j R^{jT} R^j \right) D x = \frac{1}{g} \sum_j R^{jT} w^j. \quad (13)$$

As the non-local groups are not overlapping, we have $\frac{1}{g} \sum_j R^{jT} R^j = I$. The equivalence is thus given as follows

$$\frac{1}{g} \sum_{j=1}^g \|R^j D s - w^j\|_2^2 \equiv \|D s - \frac{1}{g} \sum_j R^{jT} w^j\|_2^2. \quad (14)$$

The $R^{jT} w^j$ simply consists in placing the patches of w^j in their corresponding position. As each pixel can have g estimates, division by g permits to aggregate all the patches. Applying this result to problem (p₄) results in the following subproblem

$$\begin{aligned} (p_4) : s^{(k+1)} &\leftarrow \operatorname{argmin} \|(\nabla s - \vec{v}) - w_s^{(k+1)}\|_2^2 \\ &+ \lambda \|D_x s - \frac{1}{g} \sum_{j=1}^g R_x^{jT} w_x^{j,(k+1)}\|_2^2 \\ &+ \lambda \|D_y s - \frac{1}{g} \sum_{j=1}^g R_y^{jT} w_y^{j,(k+1)}\|_2^2. \end{aligned} \quad (15)$$

Considering periodic boundary conditions, the differential operator can be replaced with convolutions. The problem in this case can be efficiently solved via Fourier Transform \mathcal{F} as follows

$$\begin{aligned} s^{(k+1)} &= \mathcal{F}^{-1} \left\{ \frac{\mathcal{F}(\operatorname{div}(d_x, d_y))}{(1+\lambda)\widehat{\operatorname{lap}}} \right\} \\ d_x &= (v_x + w_{s,x}^{(k+1)}) + \lambda \frac{1}{g} \sum_{j=1}^g R_x^{jT} w_x^{j,(k+1)} \\ d_y &= (v_y + w_{s,y}^{(k+1)}) + \lambda \frac{1}{g} \sum_{j=1}^g R_y^{jT} w_y^{j,(k+1)}, \end{aligned} \quad (16)$$

where div is the discrete divergence operator and $\widehat{\operatorname{lap}}$ is the optical transfer function (OTF) of the discrete Laplacian filter.

IV. ANALYSIS

l_1 vs. $l_{p<1}$

As explained before, we use non-convex sparsity via the $l_{p<1}$ penalty instead of the popular l_1 -norm. The reasons why we do so is that the $l_{p_1<1}$ penalty reflects better the empirical residual gradient distribution that is highly kurtotic. For the low-rank regularization, the $l_{p_2<1}$ penalty would respect better the nature of the singular values. As the singular values are sorted in decreasing order and that the first singular values contain more energy than the rest, the shrinkage operator should act accordingly. As one can observe, the shrinkage operator given by the l_1 -norm consists in shrinking all the values with the same value¹

$$\operatorname{shrink}_{l_1}(\sigma, \frac{1}{\beta}) = \max\left(0, \sigma - \frac{1}{\beta}\right). \quad (17)$$

As a result, all the singular values would be shrunk the same way and this would severely affect the first singular values. In contrast, the shrinkage operator given by the $l_{p<1}$ penalty performs a weighted shrinkage of this form

$$\operatorname{shrink}_{l_p}(\sigma, \frac{1}{\beta}) = \max\left(0, \sigma - \frac{1}{\beta}(\sigma + \epsilon)^{p-1}\right), \quad (18)$$

where the weights $(\sigma + \epsilon)^{p-1}$ guide the shrinkage operator. Note that, for $p \rightarrow 0$ and $\epsilon \rightarrow 0$, we have

$$\lim_{\sigma \rightarrow +\infty} (\sigma + \epsilon)^{p-1} = 0. \quad (19)$$

As a result, the amount of penalization decreases as the magnitude of the elements increases, which is exactly the nature of singular values. To demonstrate this in our case, we run surface reconstruction experiments on three popular datasets with various values of $p = 1, 0.5, 0.15$. The results are presented in Table I. As can be seen, the non-convex

¹The sign is being omitted as the singular values are positive.

	5% Outliers	15% Outliers	30% Outliers	5% Outliers + Noise	15% Outliers + Noise	30% Outliers + Noise
Ramp Peaks						
$p_1 = 1, p_2 = 1$	68.95 dB	45.67 dB	32.88 dB	40.59 dB	35.27 dB	31.88 dB
$p_1 = 0.5, p_2 = 0.5$	147.64 dB	65.10 dB	39.58 dB	42.55 dB	38.11 dB	35.14 dB
$p_1 = 0.15, p_2 = 0.15$	153.53 dB	107.10 dB	42.98 dB	43.24 dB	40.17dB	37.41 dB
Vase						
$p_1 = 1, p_2 = 1$	67.62 dB	43.14 dB	31.78 dB	37.33 dB	33.21 dB	25.09 dB
$p_1 = 0.5, p_2 = 0.5$	77.44 dB	62.11 dB	39.29 dB	38.46 dB	35.14 dB	26.56 dB
$p_1 = 0.15, p_2 = 0.15$	142.71 dB	66.96 dB	44.75 dB	40.35 dB	36.55 dB	27.84 dB
Mozart						
$p_1 = 1, p_2 = 1$	62.16 dB	42.45 dB	34.51 dB	45.61 dB	40.24 dB	33.10 dB
$p_1 = 0.5, p_2 = 0.5$	68.99 dB	44.19 dB	35.14 dB	47.32 dB	41.03 dB	34.02 dB
$p_1 = 0.15, p_2 = 0.15$	71.20 dB	52.94 dB	38.38 dB	49.33 dB	41.40 dB	35.50 dB

TABLE I: Mean PSNR comparison between the convex case ($p_1, p_2 = 1$) and the non-convex case ($p_1, p_2 < 1$). As can be seen, non-convexity improves the reconstruction results as it promotes more sparsity.

	5% Outliers	15% Outliers	30% Outliers	5% Outliers + Noise	15% Outliers + Noise	30% Outliers + Noise
Ramp Peaks						
Ground-Truth Block-Matching	153.72 dB	109.25 dB	48.04 dB	44.07 dB	42.19 dB	41.99 dB
Iterative Block-Matching	153.53 dB	107.10 dB	42.98 dB	43.24 dB	40.17dB	37.41 dB
Vase						
Ground-Truth Block-Matching	144.41 dB	72.75 dB	50.68 dB	41.11 dB	38.17 dB	35.88 dB
Iterative Block-Matching	142.71 dB	66.96 dB	44.75 dB	40.35 dB	36.55 dB	27.84 dB
Mozart						
Ground-Truth Block-Matching	71.70 dB	53.30 dB	38.84 dB	49.39 dB	41.99 dB	36.31 dB
Iterative Block-Matching	71.20 dB	52.94 dB	38.38 dB	49.33 dB	41.40 dB	35.50 dB

TABLE II: Mean PSNR comparison between ground-truth block-matching calculated on clean ground-truth gradients and iterative block-matching in our method. As can be seen, a simple l_2 -distance clustering block-matching produces good quality results even in the presence of both outliers and noise.

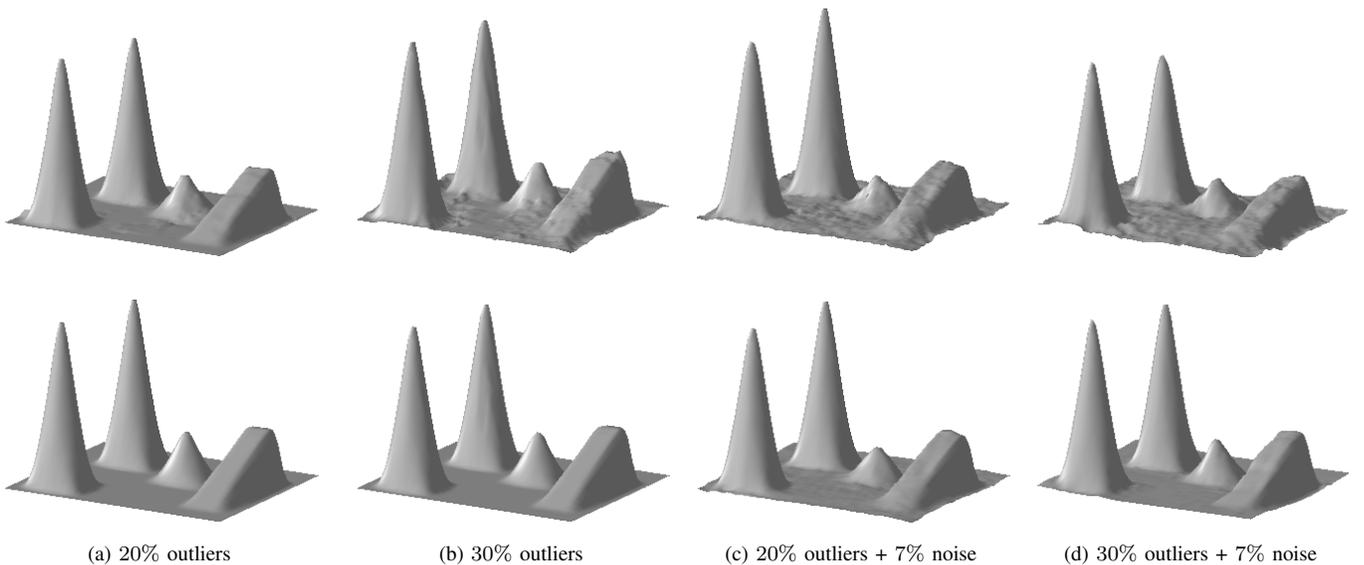


Fig. 3: Proposed reconstruction method in the convex case ($p_1 = p_2 = 1.0$) (top) and non-convex case ($p_1 = p_2 = 0.15$) (bottom). The non-convex approach leads to better results in both outliers only case (a-b) and the mixed corruption case (c-d).

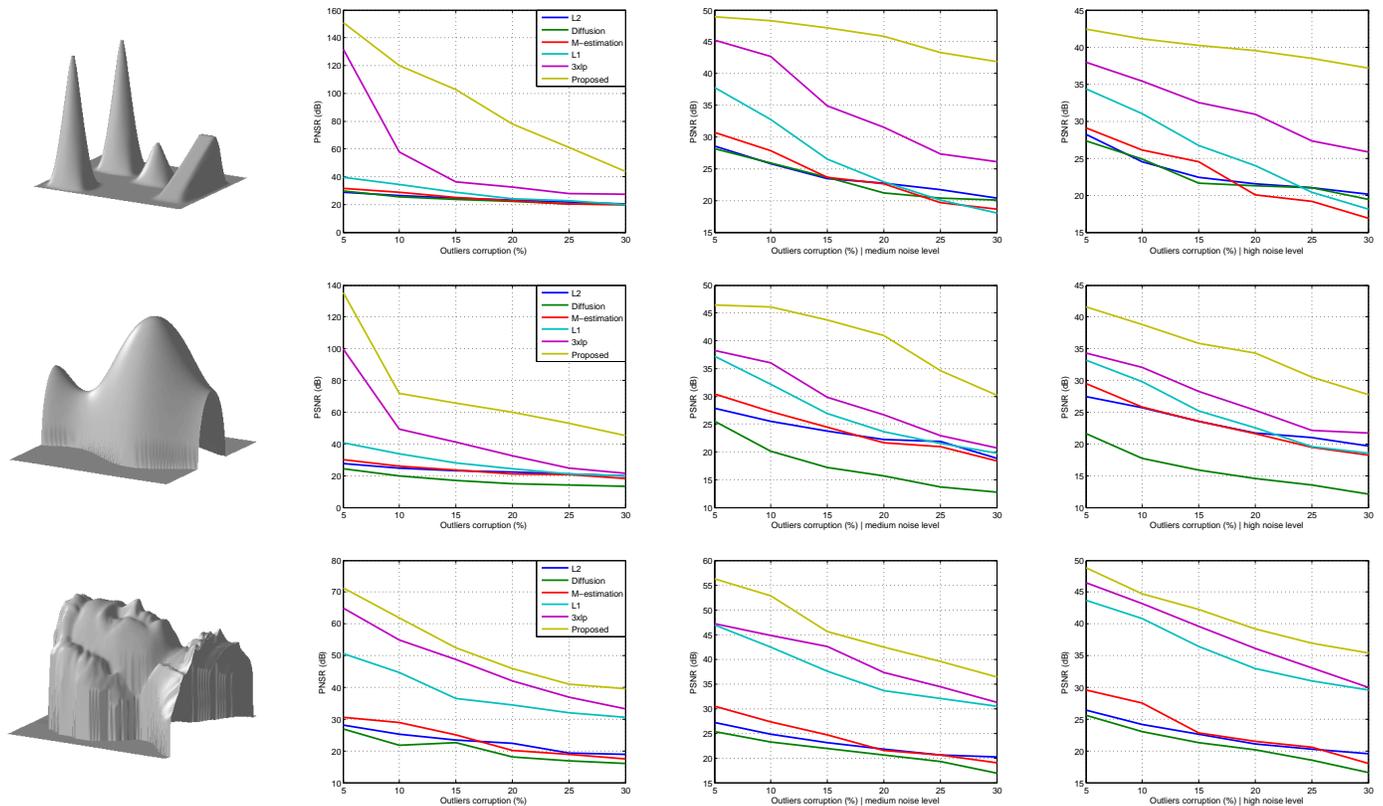


Fig. 4: Performance of the proposed method under various levels of outliers corruption when there is no noise, medium amount of noise and high amount of noise. The proposed method outperforms previous optimization methods in all the cases (visual results are reported in 6).

penalty with $p < 1$ outperforms the l_1 -norm, for both the residual and the low-rank term. For visual comparison, we show reconstruction results of the Ramp-Peaks dataset in the convex and non-convex cases under high outliers corruption (20%, 30% of the gradient points are corrupted with high-magnitude), and high mixed corruption (20%, 30% of the gradient points are corrupted with high-magnitude + high Gaussian noise corruption with standard deviation 7% of the maximum intensity). As can be seen in Figure 3, the non-convex approach leads to better results. It is able to better smooth the surface in both the outliers case only (a-b) and the mixed noise situation (c-d). Throughout the paper, we consider only the non-convex case.

Block-Matching

Due to the high-amount of corruption that may occur in the estimated gradient field, we analyze the robustness of block-matching. We use a standard approach (same code as [7]) that consists in clustering patches in a neighborhood that minimize the l_2 -distance. In all the experiments, each non-local cluster contains 20 patches of size 6×6 . This may not be the optimal setting as one should increase the patch size with respect to the amount of noise, but this setup already works very well in our experiments. We start with an initial solution calculated via the Poisson equation, calculate the gradients of the solution and iteratively improve the result. Block-matching is performed

on each gradient component separately. As this operation is quite slow, we do not perform it at each iteration. We rather run block-matching each 20 iterations to offer a good balance between processing time and quality of reconstruction. We compare with a ground-truth matching calculated on clean ground-truth gradients. The results are presented in Table II. As can be seen, it turns out that a simple l_2 -distance matching works very well even in severe corruption situations.

Parameters

The parameter β typically takes a small value, it is fixed to $\beta = 10^{-4}$ for all the experiments. β is increased at each iteration, we do this by multiplying its value by 1.20 at each iteration $\beta^{(k+1)} \rightarrow 1.20\beta^{(k)}$. The remaining parameters are : the regularization parameter λ , the patch size $s \times s$ for the non-local regularization and the number of similar patches per cluster b and the choice of the sparsity $p_{1,2}$. As mentioned before, we fix parameters to $s = 6$, $b = 20$ for all experiments. Parameter λ is proportional to the amount of corruption as in any regularization (low λ for low corruption and higher for more corruption). We fix the penalty parameters to around $p_1 = p_2 = 0.15$, which seems to work well in practice as mentioned before.

Complexity

The method consists in shrinkage operations that are fast to calculate as they correspond to pixel-wise operations, a Fourier-based reconstruction that can be performed efficiently and a non-local sparsity part. The block-matching operation in our case is a straightforward MATLAB code. This operation can be accelerated using fast dedicated methods such as PatchMatch [21] or parallel processing. For the non-local sparsity part, as processing each group is independent (because the groups are non-overlapping), the low-rank processing step can be performed in parallel as well. Patch manipulation (extraction and reconstruction) are easy operations that are efficiently implemented in C (MEX). On a laptop with a i7-2670QM CPU using 4 cores, the block-matching reduces to 0.40 second and the low-rank estimation subproblem to 0.25 second per iteration for a heavily corrupted 128×128 gradient field.

V. EXPERIMENTS

We run various experiments to demonstrate the effectiveness of the proposed method on synthetic and real data.

Quantitative Experiments

The first part of the experiments consists in quantitative experiments on standard synthetic surfaces. Given a ground-truth gradient field for each surface, we generate sparse outliers of density $\tau\%$ and magnitude 5 times the maximum gradient intensity value. Dense Gaussian noise is then added in three cases : no noise, medium noise and high noise². We then evaluate how the proposed method is able to recover the true surface for various corruption situations. We compare with previous optimization-based methods l_2 [9], M-estimation [12], l_1 [16], [15], $3 \times l_p$ [18]. The results are given in Figure 4. As can be seen, the proposed method outperforms the other methods for all outliers corruption levels (from medium (5%) to high (30%)) for the noise free case, as well as in the presence of medium and high levels of dense noise. To compare the visual quality, we present the results in Figure 6 in the case of mixed outliers/noise (high corruption 30% with a high amount of noise). For this comparison, we consider other methods such as [10], [14], [17]. As can be seen, our method preserves better important structures and correctly smooths the surface. While the sparse regularization method [18] leads to an important improvement compared to the other methods, it tends to over-smooth the surface. This is because the local prior prefers a piecewise constant solution. In contrast, the proposed non-local solution promotes non-local smoothness of the patches, which does not produce over-smoothing. It is worth noting also that the method in [18] needs various parameters to set in the case of mixed outliers/noise. The proposed technique requires setting only the regularization parameter λ , it is thus easier to use.

²Because the surfaces do not have the same gradient dynamic range, the medium noise level is determined visually so it matches the three surfaces. The high noise level is simply twice the medium noise level.

Real Photometric Stereo

We apply the proposed method to a real Photometric Stereo (PS) problem³. In this case, the corruptions in the estimated gradient field are caused by imperfections in the input images of the PS method. To make the experiment more challenging, we use a simple least-squares solution to estimate the normals. This leads to much more corruption in the estimated gradient field. We add two types of corruption to the input images : additive Gaussian dense noise with standard deviation equals to 5% of the maximum intensity, and random sparse outliers with density 1% and magnitude three times the maximum intensity. Note that this is a very challenging reconstruction case as the input images are heavily corrupted. The results are given in Figure 5. As can be seen, the proposed method is able to reconstruct a high-quality smooth surface even though the input images for the PS method are severely corrupted.

VI. CONCLUSION

We present a new approach to surface-from-gradients using a non-local low-rank regularization. The proposed method iteratively gathers non-local patches of the corrupted vector field and applies low-rank estimation to reduce perturbations such as outliers and dense noise. This regularization is used along with a sparse data-fitting term that imposes non-convex sparsity on the residual gradient and ensures better outliers handling. We propose an efficient alternate minimization solution to the problem using a Half-Quadratic approach. Experiments conducted on synthetic and real data show that the proposed method is able to recover a high-quality surface from severely corrupted vector fields in mixed dense/sparse corruption situations.

REFERENCES

- [1] B. K. P. Horn, "Height and gradient from shading," *Int. J. Comput. Vision*, vol. 5, no. 1, pp. 37–75, Sep. 1990.
- [2] R. J. Woodham, "Shape from shading," B. K. P. Horn and M. J. Brooks, Eds. Cambridge, MA, USA: MIT Press, 1989, ch. Photometric method for determining surface orientation from multiple images, pp. 513–531.
- [3] F. Roddier, *Adaptive Optics in Astronomy*. Cambridge University Press, 1999.
- [4] P. Pérez, M. Gangnet, and A. Blake, "Poisson image editing," in *ACM SIGGRAPH 2003 Papers*, ser. SIGGRAPH '03. New York, NY, USA: ACM, 2003, pp. 313–318.
- [5] R. Fattal, D. Lischinski, and M. Werman, "Gradient domain high dynamic range compression," *ACM Trans. Graph.*, vol. 21, no. 3, pp. 249–256, Jul. 2002.
- [6] A. Levin, A. Zomet, S. Peleg, and Y. Weiss, "Seamless image stitching in the gradient domain," in *In Proceedings of ECCV*, 2006.
- [7] W. Dong, G. Shi, and X. Li, "Nonlocal image restoration with bilateral variance estimation: A low-rank approach," *IEEE Trans. Image Processing*, vol. 22, no. 2, 2013.
- [8] S. Gu, L. Zhang, W. Zuo, and X. Feng, "Weighted nuclear norm minimization with application to image denoising," in *In Proceedings of CVPR*, 2014.
- [9] T. Simchony, R. Chellappa, and M. Shao, "Direct analytical methods for solving poisson equations in computer vision problems," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 12, no. 5, pp. 435–446, May 1990.
- [10] P. Kovsi, "Shapelets correlated with surface normals produce surfaces," in *In Proceedings of CVPR*, 2005.

³The dataset is publicly available on Neil Alldrin's homepage : vision.ucsd.edu/~nalldrin/research/

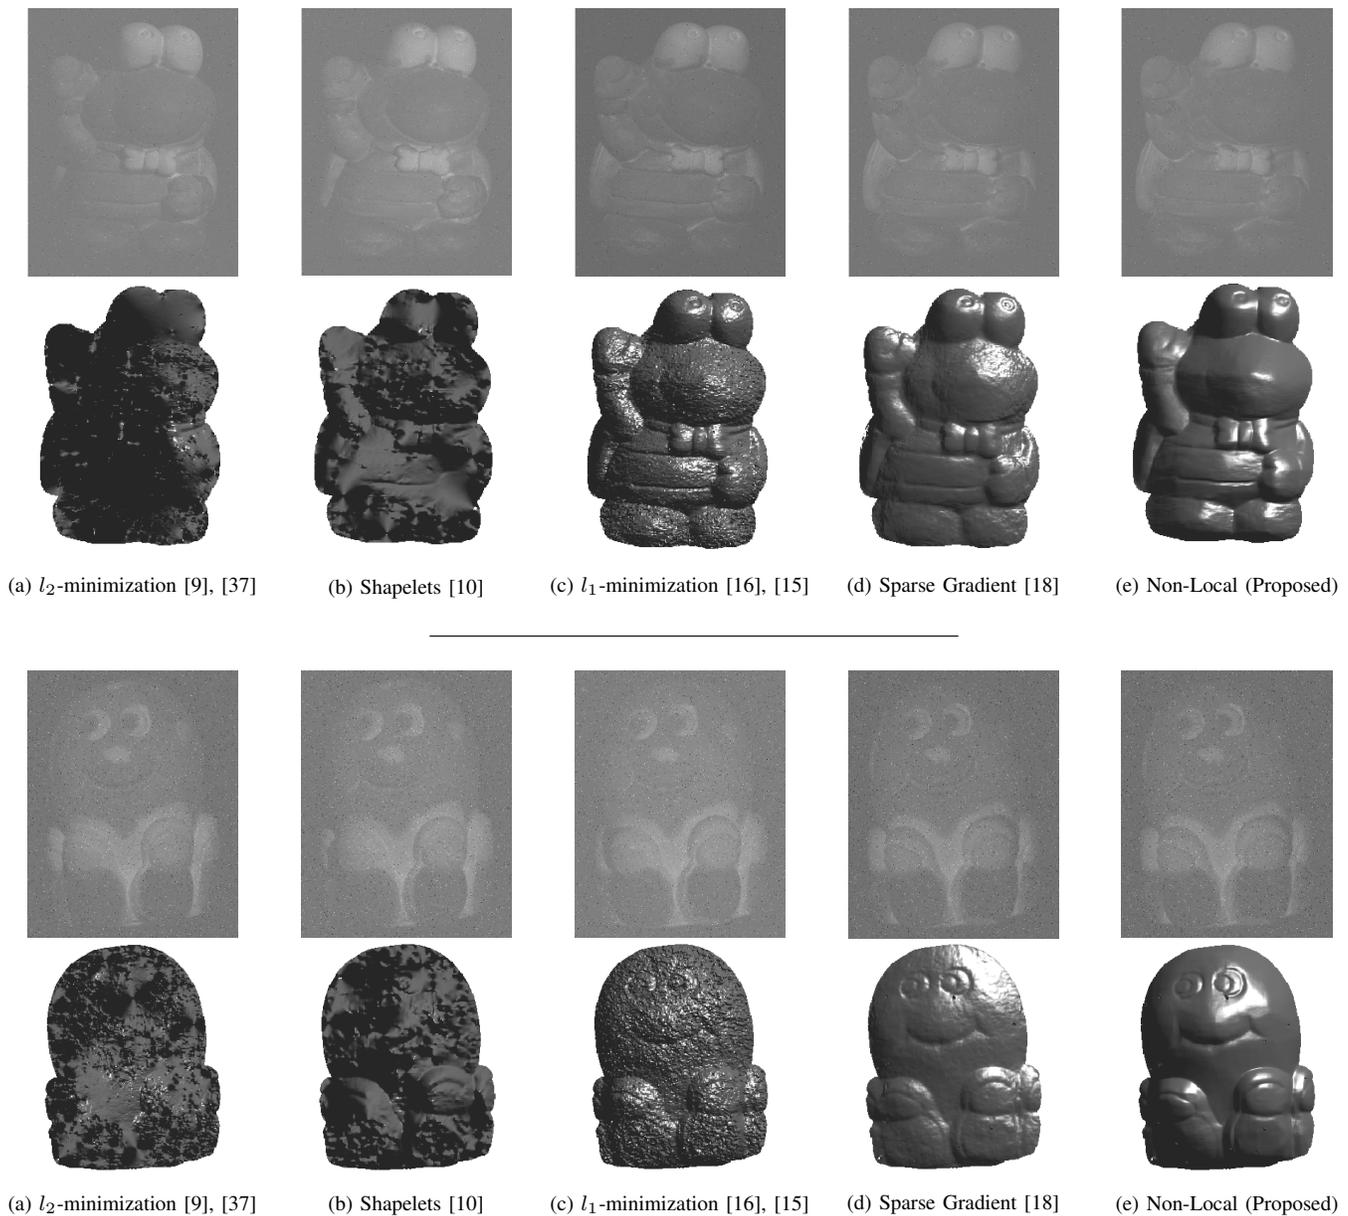


Fig. 5: Photometric stereo on real corrupted images. The proposed method is able to recover a high-quality surface even when the input PS images are heavily corrupted.

- [11] N. Petrovic, I. Cohen, B. J. Frey, R. Koetter, and T. S. Huang, "Enforcing integrability for surface reconstruction algorithms using belief propagation in graphical models." in *In Proceedings of CVPR*, 2001.
- [12] A. Agrawal, R. Raskar, and R. Chellappa, "What is the range of surface reconstructions from a gradient field?" in *In Proceedings of ECCV*, 2006.
- [13] A. Agrawal, R. Chellappa, and R. Raskar, "An algebraic approach to surface reconstruction from gradient fields," in *In Proceedings of ICCV*, 2005.
- [14] M. Harker and P. O'Leary, "Least squares surface reconstruction from gradients: Direct algebraic methods with spectral, tikhonov, and constrained regularization." in *In Proceedings of CVPR*, 2011.
- [15] Z. Du, A. Robles-Kelly, and F. Lu, "Robust surface reconstruction from gradient field using the l_1 norm," ser. DICTA '07, 2007, pp. 203–209.
- [16] D. Reddy, A. K. Agrawal, and R. Chellappa, "Enforcing integrability by error correction using l_1 -minimization," in *In Proceedings of CVPR*, 2009.
- [17] M. Rostami, O. V. Michailovich, and Z. Wang, "Surface reconstruction in gradient-field domain using compressed sensing," *IEEE Trans. Image Processing*, vol. 24, no. 5, pp. 1628–1638, 2015.
- [18] H. Badri, H. Yahia, and D. Aboutajdine, "Robust surface reconstruction via triple sparsity," in *In Proceedings of CVPR*, 2014.
- [19] T.-P. Wu, J. Sun, C.-K. Tang, and H.-Y. Shum, "Interactive normal reconstruction from a single image," in *ACM SIGGRAPH Asia 2008 Papers*, ser. SIGGRAPH Asia '08, 2008, pp. 119:1–119:9.
- [20] W. Xie, Y. Zhang, C. C. L. Whang, and R. C.-K. Chung, "Surface-from-gradients: An approach based on discrete geometry processing," in *In Proceedings of CVPR*, 2014.
- [21] C. Barnes, E. Shechtman, D. B. Goldman, and A. Finkelstein, "The generalized patchmatch correspondence algorithm," in *In Proceedings of ECCV*, 2010.
- [22] A. Buades, B. Coll, and J.-M. Morel, "A non-local algorithm for image denoising," in *In Proceedings of CVPR*, 2005.
- [23] K. Dabov, A. Foi, V. Katkovich, and K. Egiazarian, "Image denoising by sparse 3d transform-domain collaborative filtering," *IEEE Trans. Image*

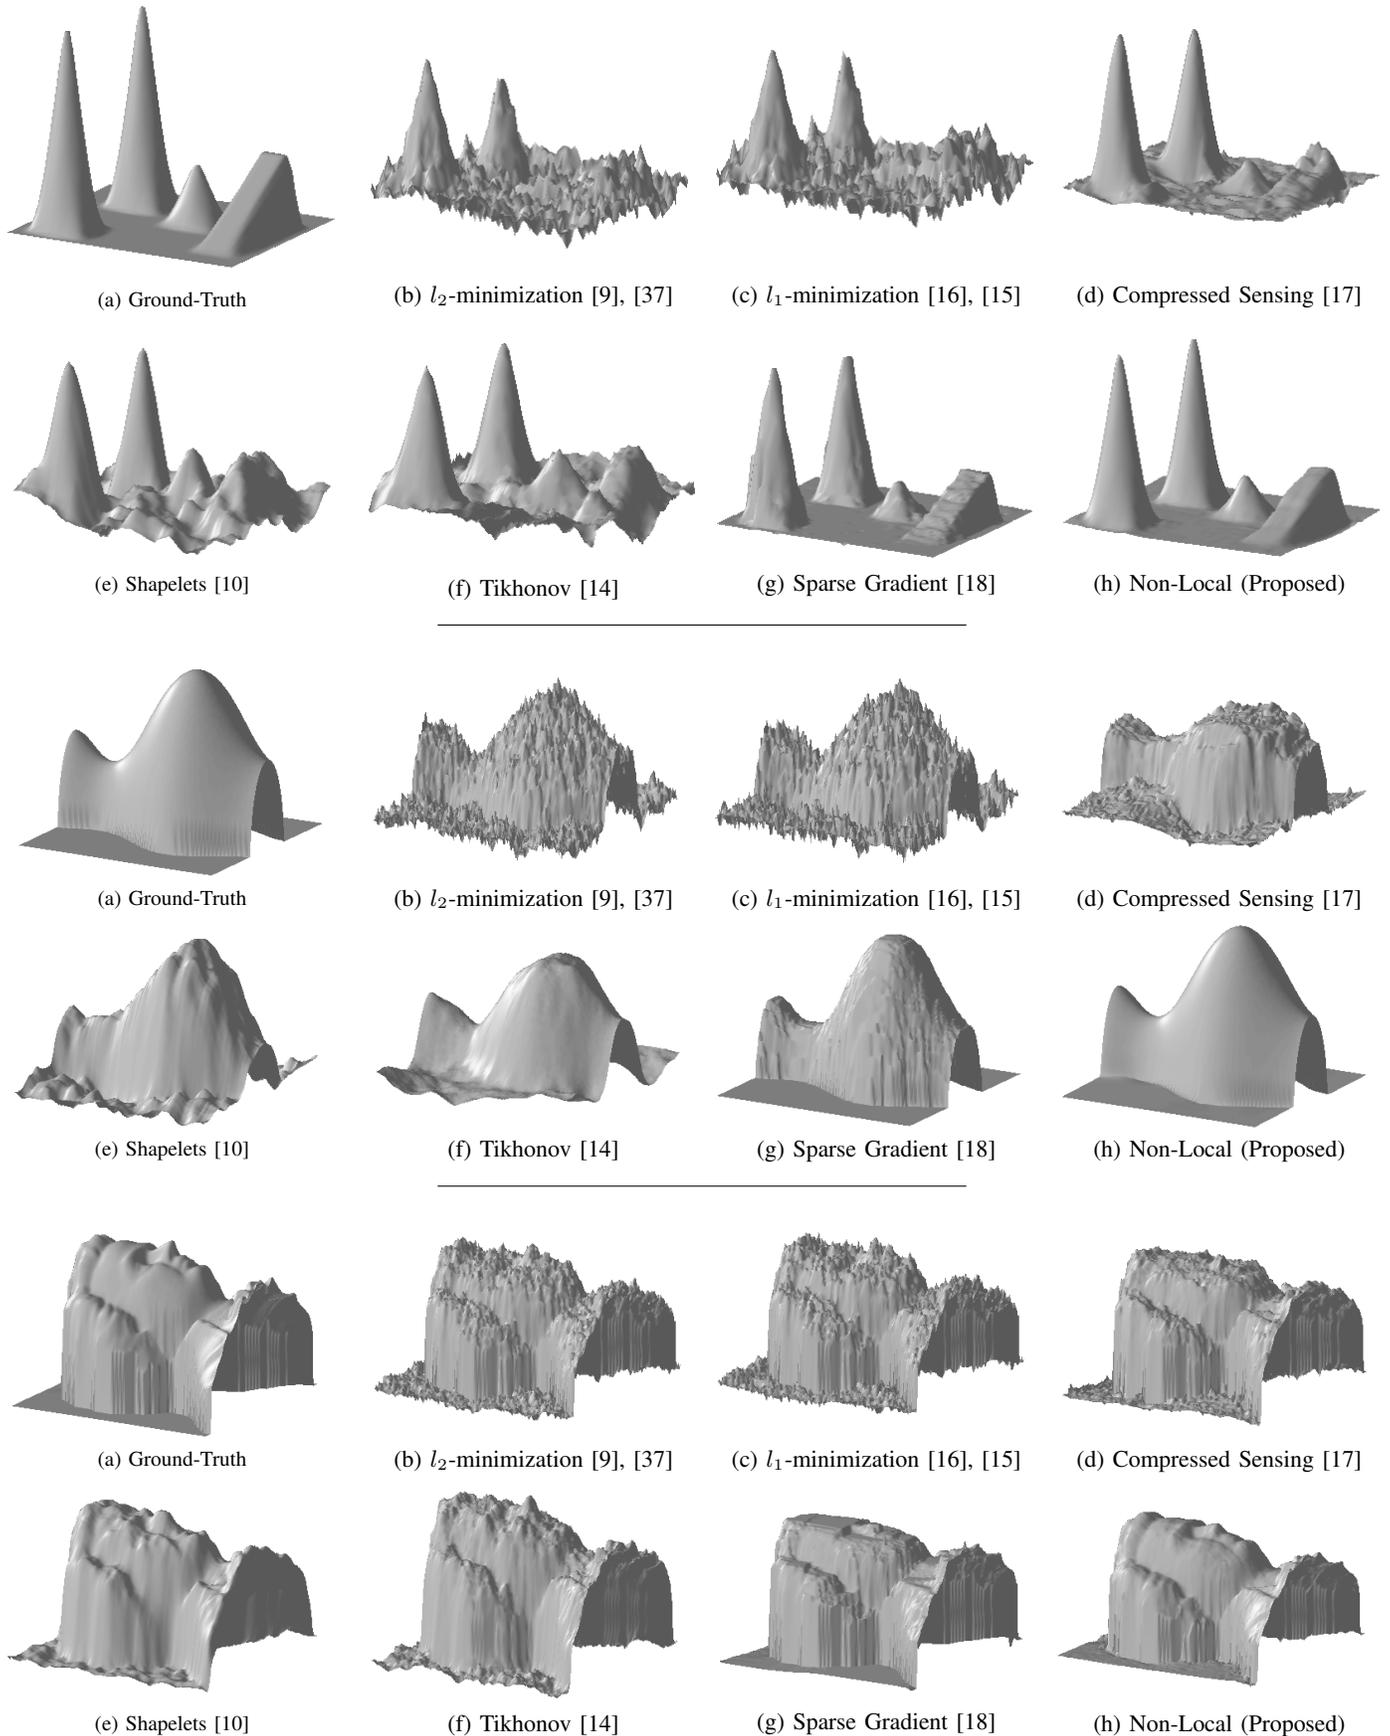


Fig. 6: Reconstruction quality with various methods in the case of high outliers corruption (30%) mixed with a high level noise. The proposed method leads to a better reconstruction quality even in extreme corruption situations.

Processing, vol. 16, no. 8, 2007.

- [24] J. Mairal, F. Bach, J. Ponce, G. Sapiro, and A. Zisserman, "Non-local sparse models for image restoration," in *proceedings of ICCV*, 2009.
- [25] A. Danielyan, V. Katkovnik, and K. Egiazarian, "Bm3d frames and variational image deblurring," *IEEE Trans. Image Processing*, vol. 21, no. 4, 2012.
- [26] X. H. Weisheng Dong, Guangming Shi and Y. Ma, "Nonlocal sparse and low-rank regularization for optical flow estimation," *IEEE Trans. Image Processing*, vol. 23, no. 10, 2014.
- [27] H. Talebi and P. Milanfar, "Nonlocal image editing," *IEEE Trans. Image Processing*, vol. 23, no. 10, 2014.
- [28] Z. Zhang, A. Ganesh, X. Liang, and Y. Ma, "Tilt: Transform invariant low-rank textures," *International Journal of Computer Vision*, vol. 99, no. 1, 2012.
- [29] D. Zhang, Y. Hu, J. Ye, and X. Li, "Matrix completion by truncated nuclear norm regularization," in *In Proceedings of CVPR*, 2012.
- [30] H. Ji, C. Liu, Z. Shen, and Y. Xu, "Robust video denoising using low rank matrix completion," in *In Proceedings of CVPR*, 2010.
- [31] L. Wu, A. Ganesh, B. Shi, Y. Matsushita, Y. Wang, and Y. Ma, "Robust photometric stereo via low-rank matrix completion and recovery," in *In Proceedings of Asian Conference on Computer Vision (ACCV)*.
- [32] R. Chartrand, "Fast algorithms for nonconvex compressive sensing: Mri reconstruction from very few data," in *Int. Symp. Biomedical Imaging*, 2009.
- [33] D. Krishnan and R. Fergus, "Fast image deconvolution using hyper-laplacian priors," in *In Proceedings of the Neural Information Processing Systems Conference (NIPS)*, 2009.
- [34] J. Wright, A. Ganesh, S. Rao, and Y. Ma, "Robust principal component analysis: Exact recovery of corrupted low-rank matrices via convex optimization," in *In Proceedings of NIPS*, 2009.
- [35] D. Geman and G. Reynolds, "Constrained restoration and the recovery of discontinuities," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 14, no. 3, pp. 367–383, 1992.
- [36] C. Lu, J. Tang, S. Yan, and Z. Lin, "Generalized nonconvex nonsmooth low-rank minimization," in *In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [37] R. T. Frankot and R. Chellappa, "A method for enforcing integrability in shape from shading algorithms," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 10, no. 4, pp. 439–451, Jul. 1988.
- [38] N. Parikh and S. Boyd, "Proximal algorithms," *Foundations and Trends in Optimization*, 2013.
- [39] H. Badri, H. Yahia, and D. Aboutajdine, "Fast edge-aware processing via first order proximal approximation," *IEEE Trans. Visualization and Computer Graphics*, vol. 21, no. 6, 2015.

APPENDIX

l_p -Shrinkage Solution

The proof of the approximation of the proximal operator associated to the l_p -norm is given here. First, consider the following proximal operator [38]

$$\mathbf{prox}_{\lambda f}(x) = \underset{y}{\operatorname{argmin}} \lambda f(y) + \frac{1}{2} \|x - y\|_2^2. \quad (20)$$

This problem admits an exact solution via an inverse function

$$\mathbf{prox}_{\lambda f}(x) = (I + \lambda \partial f)^{-1}(x). \quad (21)$$

Unfortunately, this inverse function cannot be evaluated directly for the l_p -norm. However, a solution can be approximated via a first-order Taylor expansion [39] :

$$\mathbf{prox}_{\lambda f}(x) \approx x - \lambda \partial f(x). \quad (22)$$

As the solution is a shrinkage operation, the solution does not depend on the sign

$$\mathbf{prox}_{\lambda f}(x) = \max \{0, |x| - \lambda \partial f(|x|)\} \operatorname{sign}(x). \quad (23)$$

The shrinkage operator is also pixel-wise, we thus have

$$\mathbf{prox}_{\lambda f}(x_i) = \max \{0, |x_i| - \lambda \partial f(|x_i|)\} \operatorname{sign}(x_i). \quad (24)$$

For the l_p -norm, by considering $f(x) = \frac{1}{p} \|x\|_p^p$, we have $\partial f(|x_i|) \approx |x_i + \epsilon|^{p-1}$, where ϵ is a small value to prevent division by zero. The solution is thus given as follows

$$\mathbf{prox}_{\lambda f}(x_i) = \max \{0, |x_i| - \lambda |x_i + \epsilon|^{p-1}\} \operatorname{sign}(x_i). \quad (25)$$

This solution is similar to the one proposed by Chartrand [32] calculated using Legendre-Fenchel transform.



Hicham Badri received the M.S. degree in computer science and telecommunications from Mohammed V-Agdal University and the Ph.D. degrees in image processing from both University of Bordeaux (INRIA) and Mohammed V-Agdal University. His research interests include image processing, sparse methods and machine learning.



Hussein Yahia is the head of the INRIA research team GEOSTAT (Geometry and statistics in acquisition data). Dr. H. Yahia received the Doctorat ode 3eme cycle from University Paris 11 (Orsay) and the HDR (Habilitation Diriger des Recherches) from University Paris 13. He specializes in non-linear signal processing and the analysis of complex signals and systems using advanced nonlinear physics. Dr. H. Yahia has also made substantial contributions to computer graphics and image processing. Dr. H. Yahia is developing a strong collaboration with the LEGOS Laboratory in Toulouse (UMR CNRS 55 66), ICM-CSIC in Barcelona, and IIT Roorkee in India (Associated team OPTIC with Prof. Dharmendra Singh). He is also involved in many contracts with the French spatial agency CNES and the European Spatial Agency (ESA). Dr. H. Yahia is the author or co-author of about 80 publications in international peer-reviewed journals and conferences including top-ranked conferences such as ACM SIGGRAPH, CVPR and ECCV. He is a member of the editorial board of journals in signal processing or complex systems and has supervised more than 10 PhD students.