

Enforcing Integrability by Error Correction Using l_1 -Minimization

Dikpal Reddy, Amit Agrawal, Rama Chellappa

TR2009-044 September 2009

Abstract

Surface reconstruction from gradient fields is an important final step in several applications involving gradient manipulations and estimations. Typically, the resulting gradient field is non-integrable due to linear/non-linear gradient manipulations, or due to presence of noise/outliers in gradient estimation. In this paper, we analyze integrability as error correction, inspired from recent work in compressed sensing, particularly l_0 - l_1 . We propose to obtain the surface by finding the gradient field which best fits the corrupted gradient field in l_1 sense. We present an exhaustive analysis of the properties of l_1 solution for gradient field integration using linear algebra and graph analogy. We consider three cases: (a) noise, but no outliers (b) no-noise but outliers and (c) presence of both noise and outliers in the given gradient field. We show that l_1 solution performs as well as least squares in the absence of outliers. While previous l_0 - l_1 equivalence work has focused on the number of errors (outliers), we show that the location of errors is equally important for gradient field integration. We characterize the l_1 solution both in terms of location and number of outliers, and outline scenarios where l_1 solution is equivalent to l_0 solution. We also show that when l_1 solution is not able to remove outliers, the property of local error confinement holds: i.e., the errors do not propagate to the entire surface as in least squares. We compare with previous techniques and show that l_1 solution performs well across all scenarios without the need for any tunable parameter adjustments.

CVPR 2009

This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

Enforcing Integrability by Error Correction using ℓ_1 -minimization

Dikpal Reddy
University of Maryland
College Park, MD, 20742
dikpal@umd.edu

Amit Agrawal
Mitsubishi Electric Research Labs
Cambridge, MA, 02139
agrawal@merl.com

Rama Chellappa
University of Maryland
College Park, MD, 20742
rama@umiacs.umd.edu

Abstract

Surface reconstruction from gradient fields is an important final step in several applications involving gradient manipulations and estimation. Typically, the resulting gradient field is non-integrable due to linear/non-linear gradient manipulations, or due to presence of noise/outliers in gradient estimation. In this paper, we analyze integrability as error correction, inspired from recent work in compressed sensing, particularly $\ell_0 - \ell_1$ equivalence. We propose to obtain the surface by finding the gradient field which best fits the corrupted gradient field in ℓ_1 sense. We present an exhaustive analysis of the properties of ℓ_1 solution for gradient field integration using linear algebra and graph analogy.

We consider three cases: (a) noise, but no outliers (b) no-noise but outliers and (c) presence of both noise and outliers in the given gradient field. We show that ℓ_1 solution performs as well as least squares in the absence of outliers. While previous $\ell_0 - \ell_1$ equivalence work has focused on the number of errors (outliers), we show that the location of errors is equally important for gradient field integration. We characterize the ℓ_1 solution both in terms of location and number of outliers, and outline scenarios where ℓ_1 solution is equivalent to ℓ_0 solution. We also show that when ℓ_1 solution is not able to remove outliers, the property of local error confinement holds: i.e., the errors do not propagate to the entire surface as in least squares. We compare with previous techniques and show that ℓ_1 solution performs well across all scenarios without the need for any tunable parameter adjustments.

1. Introduction

Surface reconstruction from gradient fields is an important final step in many vision and graphics applications involving gradient domain processing. These can be broadly classified as (a) manipulating gradients and/or (b) estimating gradients before integration. Methods such as Photometric Stereo (PS) [25] and Shape from Shading (SfS) [13] estimate the gradient field of the surface from captured

images. Applications such as image editing [19], stitching [18], HDR compression [9] etc. first apply local/global manipulations to the gradient field of single/multiple images. The final image is then reconstructed from the modified gradient field. The reader is referred to the recent course [2] for more details.

Typically, the resulting gradient field is non-integrable due to linear/non-linear gradient manipulations, or due to presence of noise/outliers in gradient estimation (figure 1). For a reconstruction algorithm, two important considerations are (a) robustness or ability to handle outliers and (b) local error confinement (LEC) [1]. Robustness means that surface features are well reconstructed in presence of outliers. A related property is LEC, which ensures that distortions in the integrated surface are confined spatially close to the errors in the underlying gradient field.

It is well-known that least squares estimate is not robust in presence of outliers. While several integration techniques have been proposed before, we analyze robust surface integration as an error correction problem. We are inspired from recent work in compressed sensing [5], particularly $\ell_0 - \ell_1$ equivalence. We propose to obtain the surface by finding the gradient field which best fits the corrupted gradient field in the ℓ_1 -norm sense. While minimizing the ℓ_1 -norm is not new as a robust statistic, we analyze the properties of ℓ_1 solution and provide new insights using linear algebra and graph analogy. We compare with existing techniques and show that ℓ_1 solution performs well across all scenarios without the need for any tunable parameter adjustments.

1.1. Contributions

- We analyze robust gradient integration as error correction by utilizing ideas from sparse signal recovery literature.
- We show that the location of errors is as important as the number of errors for gradient integration, which is not typically explored when considering $\ell_0 - \ell_1$ equivalence.
- We exhaustively analyze the properties of ℓ_1 solution in terms of robustness and LEC for various outlier patterns and noise in given gradient field.

1.2. Related work

Enforcing integrability: The simplest approach is to find an *integrable* gradient field (or the surface) which best fits the given gradient field, by minimizing the least squares cost function. This amounts to solving the Poisson equation [24]. Frankot & Chellappa [11] project the given gradient field onto integrable functions using Fourier basis to enforce integrability. Cosine basis functions were proposed in [12], while Kovese [17] proposed *shapelets* as a redundant set of basis functions. Petrovic *et al.* [21] used a loopy belief propagation algorithm to find the corresponding integrable field. Methods based on ℓ_2 -norm cannot handle large outliers in the gradient field.

Robust estimation: There has been large body of work on robust parametric estimation using RANSAC [10], which becomes combinatorial in nature as the number of parameters increases. For gradient integration on $N \times N$ grid, there are N^2 unknowns (pixel values) and $2N^2$ observations (x and y gradients). Thus, RANSAC is computationally prohibitive [3]. M-estimators modify the least squares cost function to reduce the influence of outliers. Several such influence functions such as Huber, Tukey, etc. have been proposed [14, 22].

Agrawal *et al.* [3] proposed a general framework for robust gradient integration by gradient transformations, such as anisotropic weighting and affine transformations. The diffusion algorithm in [3] solves a modified Poisson equation by applying edge preserving affine transformations to the gradient field. To calculate the local surface edge direction, the algorithm uses gradient values in a neighborhood. We show that it performs poorly when the neighborhood of an edge is corrupted by outliers.

Our approach instead minimizes the ℓ_1 -norm of gradient errors. Minimizing the ℓ_1 -norm has been shown to be effective in correcting outlier errors and recovering sparse signals [8, 15, 18]. Traditionally, ℓ_1 -norm is not preferred since the cost function is not analytically differentiable and minimization is computationally expensive. However, there has been a renewed interest in using ℓ_1 cost functions due to $\ell_0 - \ell_1$ equivalence for sparse reconstructions under the *restricted isometry property* (RIP). We use RIP to show that for gradient integration, the location of outliers is as important as their number. In addition, we use the expander graph structure of gradient-curl pairs to understand the distribution of outliers which can be corrected.

Graph based approach: To avoid the combinatorial nature of RANSAC, a greedy graph based technique was proposed in [1]. This approach treats the underlying 2D grid as a graph, gradients as edges and unknown surface values as nodes. The outlier gradients are heuristically determined by thresholding the curl values over the graph and the corresponding edges are removed. If the graph remains connected, surface could be integrated using the remaining

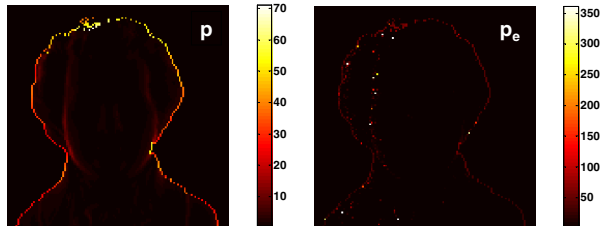


Figure 1. (Left) Ground truth p for Mozart (Right) Outliers along possible shadow regions in the gradient field obtained from PS. The magnitude of the outliers is 5 times the largest ground truth gradient values.

edges (gradients). Else, a minimal set of edges are chosen to connect the graph by assigning edge weights using gradient magnitude or curl values. However, the underlying heuristic of using curl values as a ‘goodness’ measure often fails in presence of noise. We show that [1] performs poorly in presence of noise and that minimizing the ℓ_1 -norm effectively handles noise as well as corrects sparsely distributed outliers in the gradient field. In addition, when the outliers are concentrated, LEC property is maintained.

Denosing and TV regularization: Image denoising is a classical problem and several approaches for feature preserving denoising have been successfully demonstrated. Anisotropic filtering [20] takes into account the local edge direction in a PDE based framework. Rudin *et al.* [23] proposed total variation (TV) regularization, which penalizes the ℓ_1 -norm of the gradients of the estimated (denoised) image. Note that our approach is different: we minimize the ℓ_1 -norm of *gradient errors*, not gradients themselves. Thus, we do not employ any assumptions on the underlying surface such as natural image statistics (distribution of gradients is peaked at zero).

2. Gradient integration as error correction

We use terminology from [1]. Let $S(y, x)$ be the desired surface over a rectangular grid of size $H \times W$. In vector form, we denote it by \mathbf{s} . Let (p, q) denote the given non-integrable gradient field, possibly corrupted by noise and outliers. The goal is to estimate S from (p, q) . The integrable gradient field of S is given by the forward difference equations

$$\begin{aligned} p^0(y, x) &= S(y, x + 1) - S(y, x) \\ q^0(y, x) &= S(y + 1, x) - S(y, x). \end{aligned} \quad (1)$$

In vector form (1) can be written as

$$\mathbf{D}\mathbf{s} = \begin{bmatrix} p^0 \\ q^0 \end{bmatrix} = \mathbf{g}^0, \quad (2)$$

where \mathbf{g}^0 denotes the stacked gradients and \mathbf{D} denotes the gradient operator matrix. Each row of \mathbf{D} has two non-zero entries: ± 1 in pixel positions corresponding to that particular gradient. The curl of the gradient field can be defined as

loop integrals around a box of four pixels [1]

$$\text{curl}(y, x) = p(y+1, x) - p(y, x) + q(y, x) - q(y, x+1)$$

which can be written as

$$\mathbf{d} = \mathbf{C} \begin{bmatrix} p \\ q \end{bmatrix} = \mathbf{C}\mathbf{g}. \quad (3)$$

Here, \mathbf{d} denotes the vector of stacked curl values and \mathbf{C} denotes the curl operator matrix. Each row of \mathbf{C} has only four non-zero entries (± 1) corresponding to the gradients associated with the loop integral.

Since the gradient field \mathbf{g}^0 is integrable, $\mathbf{C}\mathbf{g}^0 = 0$. However, for the given non-integrable gradient field \mathbf{g} , $\mathbf{C}\mathbf{g} \neq 0$. Decomposing \mathbf{g} as the sum of \mathbf{g}^0 and a *gradient error field* \mathbf{e} , we get

$$\mathbf{g} = \mathbf{g}^0 + \mathbf{e} = \mathbf{D}\mathbf{s} + \mathbf{e}. \quad (4)$$

Applying the curl operator on both sides, we obtain

$$\mathbf{d} = \mathbf{C}\mathbf{e} \quad (5)$$

Thus, integrability can also be defined as error correction: the goal is to estimate the gradient error field \mathbf{e} given the curl \mathbf{d} of the corrupted gradient field. Note that in this formulation, there are $M = HW$ knowns (curl values) and $N = 2HW$ unknowns (error gradients), leading to an under-determined system of linear equations. We use $\|\cdot\|_p$ to denote the ℓ_p -norm. $\|\mathbf{e}\|_0$ simply counts the nonzero elements of \mathbf{e} .

Poisson solver finds a least squares fit to the gradients by solving

$$\hat{\mathbf{e}} = \arg \min \|\mathbf{e}\|_2 \quad \text{s.t.} \quad \mathbf{d} = \mathbf{C}\mathbf{e}. \quad (6)$$

The least squares estimate is optimal when the gradient errors obey a Gaussian distribution. If the errors contains outliers, then the estimate is skewed leading to severe artifacts in the reconstructed surface or image. Outliers in the gradient field can be understood as arbitrarily large errors and could obey any distribution. An example of the errors in gradients obtained from PS is shown in figure 1.

ℓ_0 -minimization: An approach to handle outliers is to combinatorially search for the possible locations of outliers, estimate them subject to the curl constraint (5) and pick the combination which satisfies the constraints the best. This can be written mathematically as

$$\hat{\mathbf{e}} = \arg \min \|\mathbf{e}\|_0 \quad \text{s.t.} \quad \mathbf{d} = \mathbf{C}\mathbf{e}. \quad (7)$$

This problem is NP-hard and hence computationally infeasible.

ℓ_1 -minimization: Instead, we solve a convex relaxation of (7) by replacing the ℓ_0 -norm of the gradient error \mathbf{e} with the ℓ_1 -norm. The conditions under which this equivalence holds true is described in detail in Sec. 4.

$$\hat{\mathbf{e}} = \arg \min \|\mathbf{e}\|_1 \quad \text{s.t.} \quad \mathbf{d} = \mathbf{C}\mathbf{e}. \quad (8)$$

(8) can be solved using convex optimization algorithms in polynomial time.



Figure 2. (a) Graph with pixels as nodes & gradients as edges. Curl is calculated along 2×2 loops (b) Spanning tree edges in black (c) Gradient errors in dashed lines (d) Solid black lines and red curl loops have expander graph structure

3. Graph based interpretation

In [1], a graph-based interpretation is provided for integrating the gradient field corrupted by outliers. We discuss this method and borrow its framework to explain our approach. [1] treats the pixel grid as a graph (G, E) , where the pixels are the nodes of the graph and gradients correspond to the edges of the graph (figure 2(a)). Let us first assume that the location of outliers (bad gradients) are known. [1] proposes to remove the corresponding edges from the graph. If the resulting sub-graph remains connected, then integration could be done using the remaining edges/gradients. Else, the graph is connected using a minimal set of edges, by assigning edge weights based on gradient magnitude or curl values. Since in practice, the location of outlier gradients are not known, [1] thresholds the curl values as a heuristic.

Relationship with ℓ_0 -minimization: Note that the key idea is that for integration to be possible, the resulting graph has to be connected. Thus, the minimal set of gradients required for integration should correspond to a spanning tree (ST) [3] as shown in figure 2(b). First, let us assume that the graph remains connected after removing the edges corresponding to outlier gradients. Then, it is easy to see that [1] is a greedy algorithm for ℓ_0 -minimization. This is because the resulting sub-graph trivially minimizes the ℓ_0 -norm of gradient errors.

However, the important point is that even if we know the location of outliers, it does not guarantee error-free reconstruction, since the resulting sub-graph needs to be connected. For example, it is easy to see that if all 4 edges of a node are removed, graph does not remain connected (figure 3, clique-5). On other hand, if the errors are distributed as shown in figure 3 (right), perfect reconstruction can be achieved. Thus, even ℓ_0 -minimization does not guarantee perfect reconstruction. It can handle up to 25% outliers¹, but can fail for as low as 4 outliers. While recent work in compressed sensing [5] has focused on the *number* of errors (outliers), the location of outliers is equally important for gradient reconstruction problem. Since ℓ_0 -minimization can fail depending on spatial distribution of errors, it is important to consider it while analyzing $\ell_0 - \ell_1$ equivalence.

RANSAC: In gradient integration, RANSAC would

¹In general, ℓ_0 -minimization can handle up to 50% outliers. For gradient integration, a unique solution can be obtained only for maximum of 25% outliers

search over different realizations of ST and pick the one which rejects most outliers. As shown in [3], since the number of parameters are large, RANSAC is computationally prohibitive.

3.1. Performance under noise

Note that a robust algorithm should also be able to work well in presence of noise. In [1], a heuristic is used to estimate outlier errors, assuming that the non-zero curl values are related to outlier gradients. However, this assumption is well suited only when the gradient field is corrupted by outliers and fails in presence of noise. Under noise, the algorithm in [1] confuses correct gradients as outliers and performs poorly as shown in figure 5.

In presence of noise, gradient error \mathbf{e} is non-sparse with the largest components corresponding to outliers. To handle noise, the cost function is modified to

$$\hat{\mathbf{e}} = \arg \min \|\mathbf{e}\|_1 \quad \text{s.t.} \quad \|\mathbf{d} - \mathbf{C}\mathbf{e}\|_2 \leq \epsilon \quad (9)$$

for an appropriate ϵ .

4. $\ell_0 - \ell_1$ equivalence

One of the earliest methods in sparse signal recovery by minimizing the ℓ_1 -norm is Basis Pursuit [8] but it is recently that conditions for equivalence between minimizing ℓ_0 and ℓ_1 -norm have been provided in the compressed sensing literature [5, 7, 6]. In fact, the gradient error correction problem is similar to the classical error correction problem analyzed in [7], but the location of errors is equally important as discussed in section 3. Continuing the notation, we present sufficient conditions for $\ell_0 - \ell_1$ equivalence as described in [6]. They are

- \mathbf{e} is k -sparse ($\|\mathbf{e}\|_0 = k$)
- The matrix \mathbf{C} obeys RIP with isometry constant δ_{2k}

RIP (with δ_{2k}) is a sufficient condition on a matrix (\mathbf{C}) which guarantees recovery of *all* k -sparse vectors (\mathbf{e}) from its projection (\mathbf{d}) using ℓ_0 -minimization (if $\delta_{2k} \leq 1$) or ℓ_1 -minimization (if $\delta_{2k} < \sqrt{2} - 1$). This implies that ℓ_1 -minimization can recover a k -sparse vector as well as ℓ_0 -minimization when $\delta_{2k} < \sqrt{2} - 1$. \mathbf{C} is said to satisfy RIP with isometry constant δ_{2k} , if the eigenvalues of $\mathbf{C}_T^* \mathbf{C}_T$ ² lie between $(1 - \delta_{2k})$ and $(1 + \delta_{2k})$ for every submatrix \mathbf{C}_T , formed by choosing $2k$ columns with index set T . Note that the condition to recover k -sparse \mathbf{e} is actually on $2k$ columns of \mathbf{C} . This is to ensure that the true k -sparse vector is not confused with any other k -sparse vector with the same projection \mathbf{d} , thereby ensuring a *unique* solution. Typically, dense matrices such as *i.i.d.* Gaussian or partial Fourier matrices [5] satisfy RIP for large k .

² \mathbf{C}^* is the transpose of \mathbf{C}

As discussed in section 3, if all 4 edges of a node are in error, they can't be corrected even if we knew their locations. It implies that the recovery of 4-sparse gradient error vector \mathbf{e} using either ℓ_0 or ℓ_1 -minimization is impossible. Thus, RIP doesn't hold for $k = 4$ and hence for all $k > 4$. But, the constant δ_{2k} corresponding to a $2k$ edge set T does inform us whether any k gradient errors in T can be corrected using either ℓ_0 or ℓ_1 -minimization

For a $2k$ edge set T , $\delta_{2k} < 1$ means that $\mathbf{C}_T^* \mathbf{C}_T$ is non-singular. This implies that the 2D graph remains connected after removing the corresponding $2k$ edges T . Conversely, in figure 3 clique-5, $\delta_{2k} = 1$ since the graph does not remain connected when all the four edges are in error.

4.1. Spatial distribution of errors

Figure 3 lists several spatial distribution of errors in a isolated neighborhood. We qualitatively analyze which of these can be corrected with the help of isometry constant δ_{2k} . Few of them are described in detail below. For example, in clique-2 ($2k = 2$), $\delta_{2k} = 0.5$ implying clique-1 ($k = 1$) can be corrected perfectly by ℓ_0 -minimization. However, in practice ℓ_1 -minimization can also correct the single outlier although $\delta_{2k} > \sqrt{2} - 1$. Likewise, $\delta_{2k} = 0.5$ in clique-8 ($2k = 4$) implies clique-6 ($k = 2$) can be corrected perfectly by both ℓ_0 & ℓ_1 -minimization. This confirms that conditions on δ_{2k} are just sufficient. Nevertheless, the conditions provide insight into the error locations that can be corrected.

Since the condition for ℓ_1 recovery is stronger than ℓ_0 recovery, there exist outlier distributions which ℓ_0 -minimization can correct but ℓ_1 cannot. For example, since $\delta_{2k} = 0.5$ in clique-8, ℓ_0 -minimization can correct clique-3 but ℓ_1 cannot always. Conversely, if ℓ_0 -minimization cannot correct a gradient error \mathbf{e} then neither can ℓ_1 . In other words, ℓ_1 -minimization corrects less errors compared to ℓ_0 .

We generalize to other outlier spatial distributions. Let T denote the indices of some $2k$ edge locations and T^c the complement edges. If T^c is not a connected subgraph, the matrix $\mathbf{C}_T^* \mathbf{C}_T$ is singular and $\delta_{2k} = 1$. This implies that there exist k error locations in T , which ℓ_0 -minimization cannot correct uniquely. If T^c is a connected subgraph, then the matrix $\mathbf{C}_T^* \mathbf{C}_T$ is non-singular and $\delta_{2k} < 1$ suggesting ℓ_0 -minimization can correct any k error locations in T . For sufficiently small k we will have $\delta_{2k} < \sqrt{2} - 1$ and ℓ_1 -minimization corrects all of them. For example, ℓ_1 -minimization can correct outliers distributed as shown in figure 3 (right).

4.2. Expander graph structure

Unlike typical dense matrices in compressed sensing, curl matrix \mathbf{C} is sparse and hence doesn't satisfy RIP for even few edges in error. Each curl value carries information about four gradients and each gradient contributes to two curl values. In the graph obtained by removing the bor-

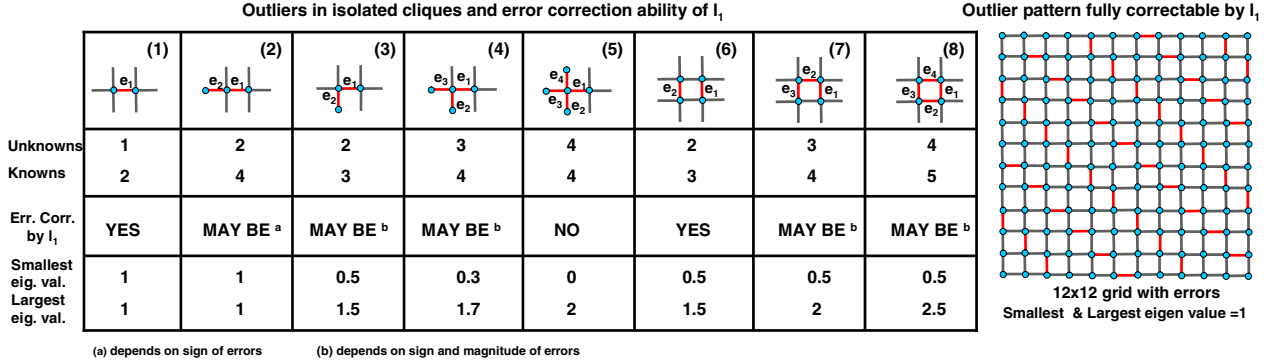


Figure 3. (Left) Top row: Isolated cliques (set T) in error (red). Second row: Number of unknown (gradients) and known (curls) variables. Third row: Shows whether ℓ_1 -minimization can correct these errors. Fourth row: Smallest & largest eigenvalue (λ_{min} & λ_{max}) of $C_T^* C_T$. (Right) Distribution of outliers (red) on 12×12 grid which can be corrected perfectly by ℓ_0 and ℓ_1 . Note that the errors are distributed apart and follow clique-1 structure.

der edges of the grid, the gradients and curl values have an *expander graph* relationship where every gradient contributes to two curl values and every curl value has contribution from four gradients. The truncated curl matrix C_{int} corresponding to gradients in the interior has the structure of an adjacency matrix of an expander graph, where the gradients are the left nodes U and the curl values, the right nodes V (figure 2(d)). But, each column of C_{int} has both $+1$ and -1 entries unlike the adjacency matrix which has only $+1$ as both entries.

In the compressed sensing literature, the concept of RIP has been extended to sparse matrices such as the adjacency matrix of an expander graph [4]. Theorem 1 in [4] states that if any matrix C_{ex} of size $M \times N'$ is the adjacency matrix of an (k, α) expander $G = (U, V, E)$ with left degree d such that $1/\alpha, d$ are smaller than N' , then the scaled matrix $C_{ex}/d^{1/p}$ satisfies the $RIP_{p,k,\delta}$ property, for $1 \leq p \leq 1 + 1/\log n$ and $\delta = \beta\alpha$ for some absolute constant $\beta > 1$.

Although C_{int} is not truly an adjacency matrix, it follows the proof of Theorem 1 in [4] for the case $p = 1$ in a straightforward way with parameters $d = 2$ and $\alpha \sim 3/4$. $\alpha \sim 3/4$ implies a poor expander and hence ℓ_1 -minimization fails to correct the errors for even simple outlier distributions. Nevertheless, the expander graph structure of the problem provides a nice framework to analyze the error distributions which can be corrected completely (such as figure 3 (right)) and also opens the door for greedy algorithms which can correct such error distributions. For example, the standard decoding algorithm for expander codes with $d = 2$ and $\alpha \sim 3/4$ would first look for two neighboring curl values which have been affected by a corrupt edge and then account for that edge in the curl values and iterate this search. This procedure indicates that for a decoding algorithm to be successful on the 2D graph (poor expander), the gradient errors should be distributed apart as shown in figure 3 (right)).

5. Experiments and Results

We compare the performance of our algorithm with the Least squares [24], Shapelets [17], Algebraic approach [1] and the Diffusion algorithm [3]. For shapelets, we use the default parameters (nscales=6, minradius=1, mult=2). Shapelets produce a scaled surface with unknown scale, which is fixed by setting the surface mean to the mean of the ground truth surface (for synthetic experiments). We assume Neumann boundary conditions for integration, which results in an unknown additive constant of integration. This needs to be set for meaningful comparisons among approaches for which we align the median of the reconstructed surface values³. Note that although mean square error (MSE) values in table 1 are indicative of the algorithm performance, it may not be related to the visual performance.

To solve (9), we use the regularized formulation: $\arg \min \mu \|e\|_1 + 1/2 \|d - Ce\|_2^2$, since faster software [16] exists for the latter. μ is the only parameter which we need to set and to enforce sparsity in the gradient error e , we found that $\mu = 10^{-3}$ works over a wide range of problems and outlier distributions.

Effect of noise without outliers: First, we compare the performance of the algorithms when the the gradient field is corrupted only by noise. We added Gaussian noise with $\sigma = 10\%$ of the maximum gradient value to the gradients of **Ramp-peaks** synthetic dataset shown in figure 5. ℓ_1 -minimization performs as well as Least squares in presence of noise in the gradient field. Algebraic method performs poorly due to the simplifying assumptions it makes about the relationship between curl and gradient error. MSE numbers are reported in table 1.

Effect of outliers without noise: To analyze the effect of outliers, we added outliers to 10% of the ground truth gradient field. The outliers are salt and pepper noise with a range five times that of the original gradient field. The re-

³Effective as long as 50% of the surface values remain uncorrupted after surface reconstruction.

	ℓ_1 -minimization	Least Squares	Diffusion	Shapelets	Algebraic
Ramp-peaks Noise only	0.5581	0.2299	0.3980	0.7221	4.5894
Ramp-peaks Outliers only	0.3136	9.9691	2.0221	24.7759	0.2430
Ramp-peaks Noise & Outliers	0.5064	6.8096	1.8382	16.8603	3.1849
Mozart PS	550.1	575.8	521.4	1179.1	708.7

Table 1. MSE of reconstructed surfaces using different methods on Ramp-peaks dataset and PS experiment on Mozart dataset.

constructed surfaces are shown in figure 6. ℓ_1 -minimization performs as well as the Algebraic approach as shown in table 1. Note that ℓ_1 -minimization corrects most of the outliers and preserves the surface edges and details. It also confines the errors locally when it fails to correct them.

We also analyze the performance of various algorithms as the percentage of outliers increase. In figure 4(a), we vary the percentage of outliers in the Ramp-peaks gradient field and compute the percentage of surface values in error. We declare a surface value to be in error if it deviates more than 5% from the maximum surface value. The plot shows that the Algebraic approach is the most effective in correcting outliers with similar performance by ℓ_1 -minimization. Note that both Least squares and Shapelets fail to preserve the surface shape even for small percentage of outliers. For this experiment, we averaged the performance over 200 realizations for every percentage of outliers.

Effect of noise and outliers: The true test of a robust algorithm is in presence of both noise and outliers. We test the realistic scenario of both noise and outliers by adding outliers to 7% of the gradients and Gaussian noise with $\sigma = 7\%$ of the maximum gradient value. ℓ_1 -minimization performs better than all the other methods. It captures the characteristic of Least squares to handle noise and that of a combinatorial method such as Algebraic approach to correct outliers.

Photometric stereo (PS): We perform PS experiment on Mozart synthetic dataset to simulate the realistic occurrence of outliers in gradient fields. We first generate images assuming Lambertian reflectance model, distant point source lighting and constant albedo. Then we estimate the surface normals (n_x, n_y, n_z) and albedo through PS on images corrupted by random noise ($\sigma = 5\%$ of the maximum intensity). The estimated gradient field is given by $p = \frac{-n_x}{n_z}$ and $q = \frac{-n_y}{n_z}$ and is corrupted by outliers as shown in figure 1. Figure 8 shows that our method and the diffusion algorithm give the best results. Both these methods correct the outlier errors which corrupt the gradient field during gradient estimation. Although ℓ_1 -minimization is marginally less successful compared to the diffusion algorithm in terms of MSE, note that our method corrects more outliers on the side of the face and also avoids the pinching artifacts near the flatter regions of the surface. It should be noted that the diffusion algorithm introduces artifacts close to sharp edges corrupted by outliers as illustrated in figure 4(c).

Local error confinement: We show that even when ℓ_1 -

minimization fails to correct the outliers, it confines the errors locally. In figure 4(b), we add outliers in a 5×5 region on a 20×20 flat surface. Both ℓ_1 -minimization and Least squares fail to correct the errors. However, Least squares spreads the error globally, while ℓ_1 -minimization confines it locally. By further regularizing the gradients themselves as in TV regularization, these remaining errors could be removed.

6. Discussions

While minimizing ℓ_1 -norm is computationally expensive compared to solving the Poisson equation, researchers in compressed sensing are devising new and better algorithms. Also, the expander graph structure of the problem opens avenue to accurate greedy algorithms. For gradient integration problem, the errors which are not corrected by ℓ_1 -minimization result in spikes in the reconstructed surface. This could be handled by adding small regularization on the gradient magnitude itself, or by median filtering the final result. The benefits of ℓ_1 -minimization over previous approaches is that no parameter tuning is required and it combines the best of least squares and combinatorial search to handle noise and correct outliers respectively. Further, even when it fails to correct the outliers, it has the local error confinement property and preserves sharp edges in the reconstructed surface.

We showed how robust gradient integration can be analyzed as error correction. Although ℓ_1 -minimization is not new, we provide new insight into the gradient integration formulation using graph theory and linear systems, by borrowing ideas from compressed sensing and sparse signal recovery. We showed that the location of outliers is equally important for gradient integration. We hope that our analysis will be useful for several other parametric robust estimation problems.

References

- [1] A. Agrawal, R. Chellappa, and R. Raskar. An algebraic approach to surface reconstruction from gradient fields. In *Proc. Int'l Conf. Computer Vision*, pages 174–181, 2005.
- [2] A. Agrawal and R. Raskar. Gradient domain manipulation techniques in vision and graphics. ICCV short course, 2007.
- [3] A. Agrawal, R. Raskar, and R. Chellappa. What is the range of surface reconstructions from a gradient field? In *Proc. European Conf. Computer Vision*, volume 3951, pages 578–591, 2006.
- [4] R. Berinde, A. C. Gilbert, P. Indyk, H. J. Karloff, and M. J. Strauss. Combining geometry and combinatorics: A unified approach to sparse signal recovery. *CoRR*, abs/0804.4666, 2008.

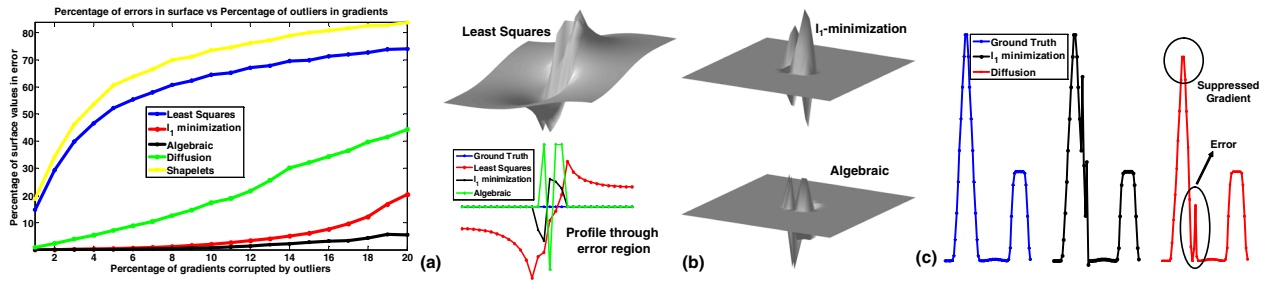


Figure 4. (a) Plot showing the number of surface values in error versus varying outlier percentage. (b) LEC is satisfied by ℓ_1 -minimization and Algebraic approach but fails in Least squares. (c) In presence of outliers near a sharp edge, the Diffusion technique results in artifacts.

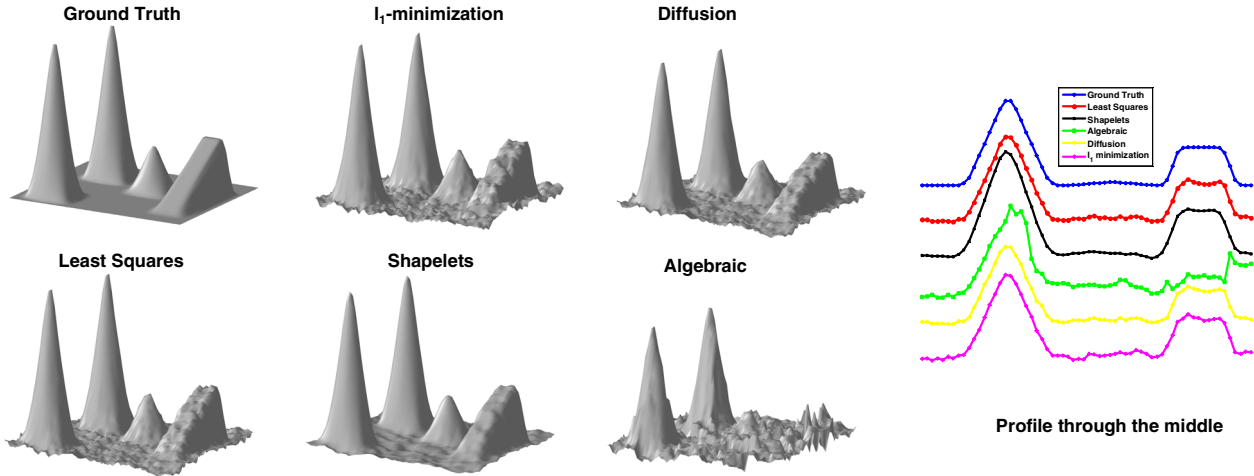


Figure 5. Reconstructed surface when the gradient field is corrupted by only Gaussian noise ($\sigma=10\%$ of the maximum gradient value). Note that the Algebraic approach performs poorly under noise and ℓ_1 -minimization performs as well as Least squares.

[5] E. J. Candès. Compressive sampling. In *Proc. International Congress of Mathematicians*, volume 3, pages 1433–1452, Madrid, Spain, 2006.

[6] E. J. Candès. The restricted isometry property and its implications for compressed sensing. *Compte Rendus de l'Academie des Sciences*, 2008.

[7] E. J. Candès and T. Tao. Decoding by linear programming. *IEEE Trans. Inform. Theory*, 51:4203–4215, Dec. 2005.

[8] S. S. Chen, D. L. Donoho, and M. A. Saunders. Atomic decomposition by basis pursuit. *SIAM J. Sci. Comput.*, 20(1):33–61, 1998.

[9] R. Fattal, D. Lischinski, and M. Werman. Gradient domain high dynamic range compression. In *SIGGRAPH*, pages 249–256, 2002.

[10] M. A. Fischler and R. C. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM*, 24(6):381–395, 1981.

[11] R. T. Frankot and R. Chellappa. A method for enforcing integrability in shape from shading algorithms. *IEEE Trans. Pattern Anal. Machine Intell.*, 10(4):439–451, 1988.

[12] A. S. Georghiadis, P. N. Belhumeur, and D. J. Kriegman. From few to many: Illumination cone models for face recognition under variable lighting and pose. *IEEE Trans. Pattern Anal. Machine Intell.*, 23(6):643–660, 2001.

[13] B. K. P. Horn. Height and gradient from shading. *Int'l J. Computer Vision*, 5(1):37–75, 1990.

[14] P. Huber. *Robust Statistics*. John Wiley and Sons, New York, 1981.

[15] Q. Ke and T. Kanade. Robust ℓ_1 norm factorization in the presence of outliers and missing data by alternative convex programming. In *Proc. Conf. Comp. Vision and Pattern Recognition*, June 2005.

[16] S.-J. Kim, K. Koh, M. Lustig, S. Boyd, and D. Gorinevsky. An interior-point method for large-scale ℓ_1 -regularized least squares. *Selected Topics in Signal Processing, IEEE Journal of*, 1(4):606–617, Dec. 2007.

[17] P. Kovési. Shapelets correlated with surface normals produce surfaces. In *Proc. Int'l Conf. Computer Vision*, pages 994–1001, 2005.

[18] A. Levin, A. Zomet, S. Peleg, and Y. Weiss. Seamless image stitching in the gradient domain. In *Proc. European Conf. Computer Vision*, pages 377–389. Springer-Verlag, 2004.

[19] P. Pérez, M. Gangnet, and A. Blake. Poisson image editing. *ACM Trans. Graph.*, 22(3):313–318, 2003.

[20] P. Perona and J. Malik. Scale-space and edge detection using anisotropic diffusion. In *IEEE Transactions on Pattern Analysis and Machine Intelligence*, volume 12, pages 629–639, 1990.

[21] 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. *Proc. Conf. Comp. Vision and Pattern Recognition*, 1:743, 2001.

[22] H. Poor. *An Introduction to Signal Detection and Estimation*. Springer-Verlag, 1988.

[23] L. I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Phys. D*, 60(1-4):259–268, 1992.

[24] T. Simchony, R. Chellappa, and M. Shao. Direct analytical methods for solving poisson equations in computer vision problems. *IEEE Trans. Pattern Anal. Machine Intell.*, 12(5):435–446, 1990.

[25] R. J. Woodham. Photometric method for determining surface orientation from multiple images. *OptEng*, 19(1):139–144, 1980.

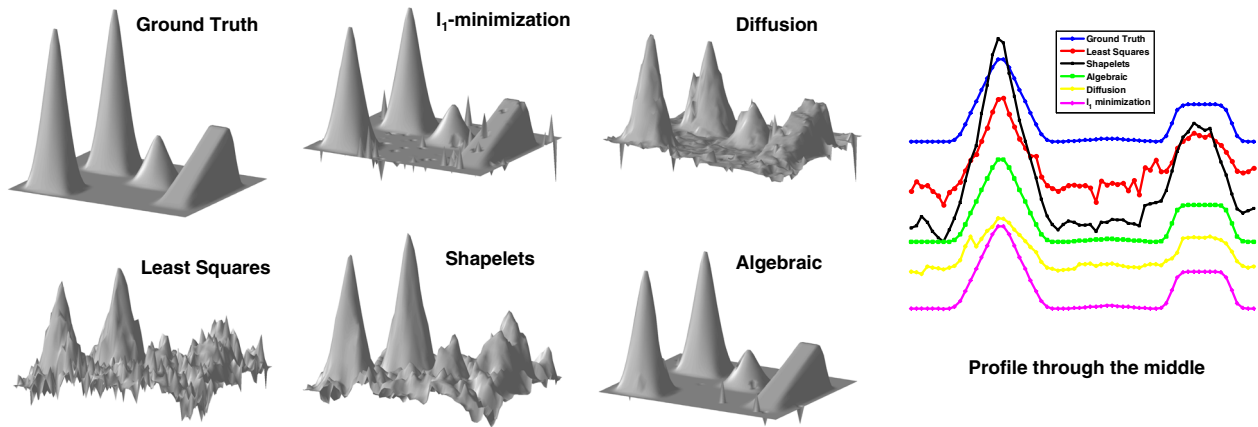


Figure 6. Reconstructed surface when 10% of the gradient field is corrupted by outliers with no noise. Note that because there is no noise, Algebraic approach performs best. ℓ_1 -minimization also reconstructs with high fidelity. Other techniques perform poorly. Even when ℓ_1 -minimization can't correct all the errors, it confines the errors locally and preserves sharp edges in the surface.

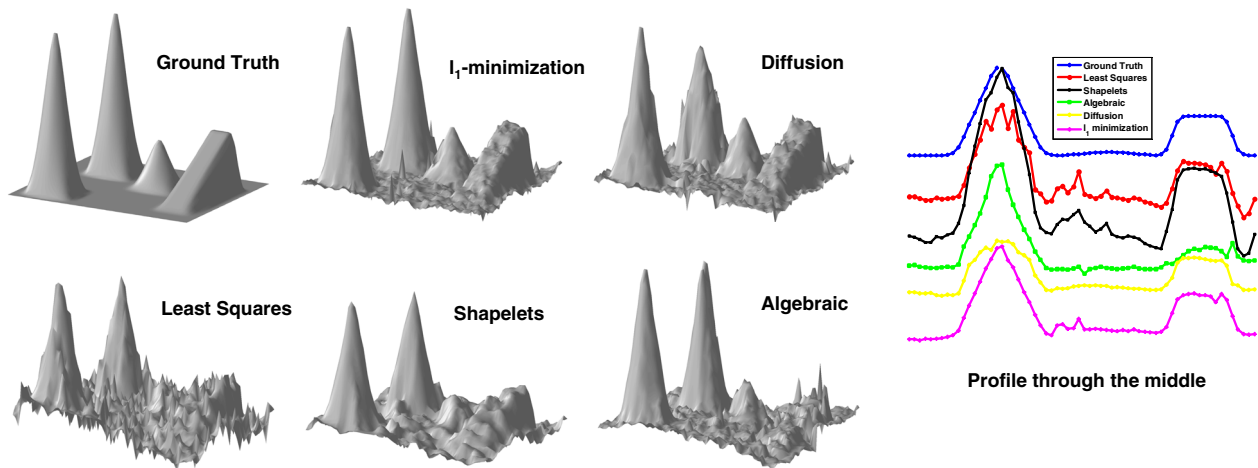


Figure 7. Reconstructed surface when gradient field is corrupted by both outliers (at 7% locations) and noise (Gaussian with $\sigma=7\%$ the maximum gradient value). ℓ_1 -minimization performs significantly better with the best characteristic of Algebraic approach for outliers and Least squares for noise.

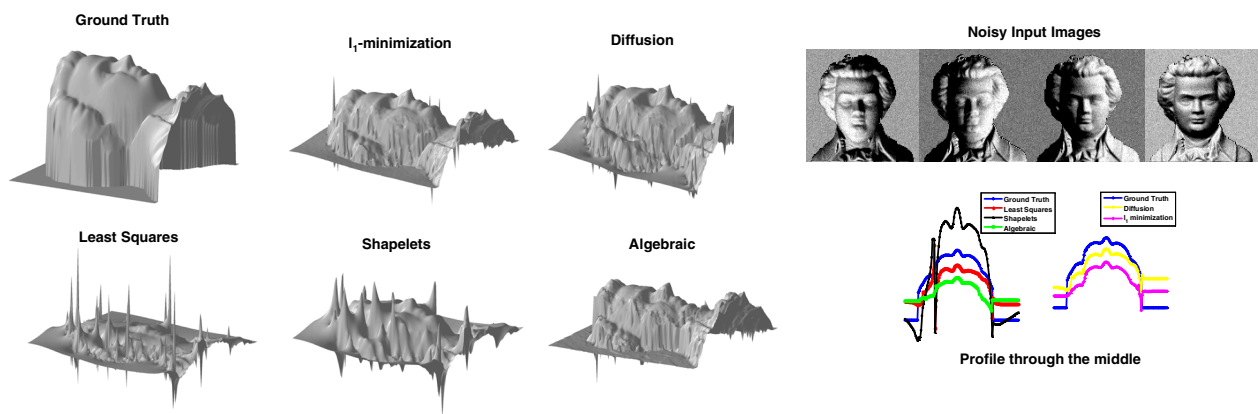


Figure 8. Surface of the Mozart bust reconstructed from gradient field obtained from PS. The gradient errors are shown in figure 1. Both ℓ_1 -minimization and Diffusion perform significantly better compared to other methods but the flatter regions of the surface have artifacts in Diffusion.