

GRADIENT PROJECTION FOR LINEARLY CONSTRAINED CONVEX OPTIMIZATION IN SPARSE SIGNAL RECOVERY

Zachary Harmany[†], Daniel Thompson^{*}, Rebecca Willett[†], and Roummel F. Marcia^{*}

[†]Department of Electrical and Computer Engineering, Duke University, Durham, NC 27708, USA

^{*}School of Natural Sciences, University of California, Merced, Merced, CA 95343, USA

ABSTRACT

The ℓ_2 - ℓ_1 compressed sensing minimization problem can be solved efficiently by gradient projection. In imaging applications, the signal of interest corresponds to nonnegative pixel intensities; thus, with additional nonnegativity constraints on the reconstruction, the resulting constrained minimization problem becomes more challenging to solve. In this paper, we propose a gradient projection approach for sparse signal recovery where the reconstruction is subject to nonnegativity constraints. Numerical results are presented to demonstrate the effectiveness of this approach.

Index Terms— Gradient projection, sparsity, compressed sensing, convex optimization, Lagrange multipliers, wavelets

1. INTRODUCTION

Modern advances in signal reconstruction algorithms make use of the assumption that the underlying signal is sparse or compressible (admits a sparse approximation) within some basis or representation. Although this is widely known to hold for a wide range of natural signals, the recent trend of compressed sensing [1, 2] offers strong theoretical justification for utilizing sparse models in signal reconstruction algorithms (e.g., [3]). While many numerical experiments support using such sparsity-promoting reconstruction algorithms, the practical application of such methods to real-world imaging problems lags far behind the theory. Often, important limitations and constraints in practical imaging scenarios are overlooked. In this paper, we address two such practical concerns.

First, since the images we wish to accurately estimate correspond to light intensities, they are naturally nonnegative quantities. Therefore we propose an algorithm that incorporates this prior information by constraining the solution to be nonnegative. We will show that incorporating additional information about the signals of interest – in this case, nonnegativity – will lead to more accurate reconstructions with negligible changes in the overall computational cost. In particular, we will develop a projected gradient method that is able to quickly and accurately solve the constrained image reconstruction problem.

Second, many image processing experiments assume that the observations we collect are of very high precision, typically single- or double-floating point precision (32 and 64 bits respectively). However, in practice, the observations taken from a focal plane array are quantized to an accuracy of b bits per pixel, with $b \leq 8$. By limiting the accuracy of our observations in our numerical experiments, we examine the robustness of this class of ℓ_2 - ℓ_1 reconstruction algorithms to low-precision observations.

This work was supported by NSF CAREER Award No. CCF-06-43947, DARPA Grant No. HR0011-07-1-003, and NSF Grant DMS-08-11062.

2. PROBLEM FORMULATION

In our observation model, we measure $y = Af^* + \eta$, where $y \in \mathbb{R}^m$ corresponds to the observation, $A \in \mathbb{R}^{m \times n}$ is the projection matrix, $f^* \in \mathbb{R}^n$ is the signal of interest, and $\eta \in \mathbb{R}^m$ is a vector of errors corresponding to sensor noise, quantization errors, or other measurement inaccuracies. Most methods in current literature solve the following convex ℓ_2 - ℓ_1 optimization problem for estimating f^* :

$$\hat{f} \equiv \arg \min_{f \in \mathbb{R}^n} \frac{1}{2} \|y - Af\|_2^2 + \tau \|f\|_1,$$

for some regularization parameter $\tau > 0$. In imaging applications, f^* corresponds to nonnegative pixel intensities. Thus, additional nonnegative constraints must be imposed on the reconstruction, i.e., \hat{f} now must solve

$$\hat{f} \equiv \arg \min_{f \in \mathbb{R}^n} \frac{1}{2} \|y - Af\|_2^2 + \tau \|f\|_1 \quad \text{subject to } f \geq 0. \quad (1)$$

Often, the signal f^* is not sparse in the canonical basis, but rather in some other (orthonormal) basis W , i.e., $f^* = W\theta^*$, where θ^* is mostly zeros. Therefore, we are interested in solving the more general problem

$$\begin{aligned} \hat{\theta} &\equiv \arg \min_{\theta \in \mathbb{R}^n} \frac{1}{2} \|y - AW\theta\|_2^2 + \tau \|\theta\|_1 \\ &\text{subject to } W\theta \geq 0 \\ \hat{f} &\equiv W\hat{\theta}. \end{aligned} \quad (2)$$

The non-differentiability of the ℓ_1 -penalty term in (2) coupled with a nontrivial linear constraint makes this optimization particularly difficult, whereas the minimization problem (1) has a simple closed form solution (via soft thresholding). In this paper, we propose solving the constrained image reconstruction problem using a well-known method called gradient projection.

3. GRADIENT PROJECTION

To apply gradient-based optimization methods to solve (2), its objective function must be differentiable. To this end, we decompose θ into its positive and negative components, $\theta = u - v$, with $u, v \geq 0$ so that $\|\theta\|_1 = \mathbb{1}_n^T(u + v)$, where $\mathbb{1}_n \in \mathbb{R}^n$ is the n -vector of ones:

$$\begin{aligned} (\hat{u}, \hat{v}) &\equiv \arg \min_{u, v \in \mathbb{R}^n} \frac{1}{2} \|y - AW(u - v)\|_2^2 + \tau \mathbb{1}^T(u + v) \\ &\text{subject to } u, v \geq 0, \quad W(u - v) \geq 0 \\ \hat{f} &\equiv W(\hat{u} - \hat{v}). \end{aligned} \quad (3)$$

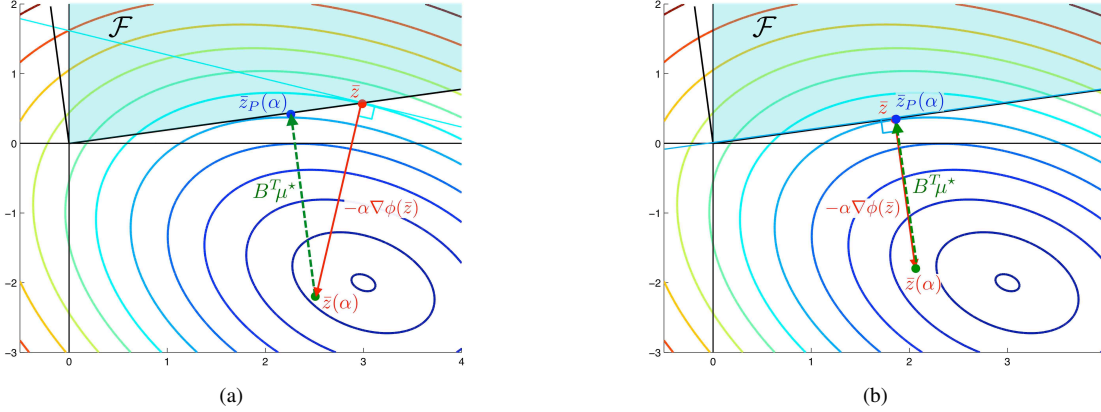


Fig. 1. A two-dimensional illustration of gradient projection onto a linearly constrained feasible set \mathcal{F} (in light green). (a) If \bar{z} is not a stationary point, Prop. 1 implies that there exists a scalar $\bar{\alpha}$ such that if $0 < \alpha < \bar{\alpha}$, $\phi(\bar{z}_P(\alpha)) < \phi(\bar{z})$. (b) For a stationary point \bar{z} , the projection $\bar{z}_P(\alpha)$ of $\bar{z}(\alpha) = \bar{z} - \alpha \nabla \phi(\bar{z})$ for any $\alpha \geq 0$ yields the original point \bar{z} . The matrix-vector product $B^T \mu^*$ is defined in Sec. 4.

The minimization problem (3) can be simplified notationally by letting $z = [u; v] \in \mathbb{R}^{2n}$, $\widetilde{W} = [W \quad -W] \in \mathbb{R}^{n \times 2n}$, and $B = [I_{2n}; \widetilde{W}] \in \mathbb{R}^{3n \times 2n}$:

$$\begin{aligned} \hat{z} &\equiv \arg \min_{z \in \mathbb{R}^{2n}} \phi(z) \equiv \frac{1}{2} \|y - A\widetilde{W}z\|_2^2 + \tau \mathbb{1}^T z \\ &\text{subject to } Bz \geq 0 \\ \hat{f} &\equiv \widetilde{W}\hat{z}. \end{aligned} \quad (4)$$

We denote the feasible set by $\mathcal{F} \equiv \{z \in \mathbb{R}^{2n} : Bz \geq 0\}$. Now the two-step gradient projection method [4] defines its iterates $z^{(k+1)}$ from the previous iterate $z^{(k)}$ by first projecting onto the feasible set the resulting vector defined by a steepest descent method:

$$z_P^{(k)}(\alpha^{(k)}) = P \left(z^{(k)} - \alpha^{(k)} \nabla \phi(z^{(k)}) \right), \quad (5)$$

where P is the projection operator onto the feasible set \mathcal{F} and $\alpha^{(k)} > 0$. Then a linesearch is performed along this direction to obtain a suitable steplength $v^{(k)}$:

$$z^{(k+1)} = z^{(k)} + v^{(k)} \left(z_P^{(k)}(\alpha^{(k)}) - z^{(k)} \right).$$

For ease of notation, we drop the superscripts corresponding to the iterates and denote the current iterate by \bar{z} . We define $\bar{z}(\alpha) = \bar{z} - \alpha \nabla \phi(\bar{z})$ and $\bar{z}_P(\alpha) \equiv P(\bar{z} - \alpha \nabla \phi(\bar{z}))$. Then the following proposition holds [5].

Proposition 1: *Let \bar{z} be a feasible point, i.e., $B\bar{z} \geq 0$. (a) If \bar{z} is not a stationary point, then there exists a scalar $\bar{\alpha} > 0$ with $\phi(\bar{z}_P(\alpha)) < \phi(\bar{z})$ for all $\alpha \in (0, \bar{\alpha}]$. (b) The point \bar{z} is stationary if and only if $\bar{z}_P(\alpha) = \bar{z}$ for all $\alpha \geq 0$.*

Prop. 1 implies that unless a feasible point is a local minimum, there always exists a nonnegative step along the negative gradient direction such that its projection onto the feasible set results in a decrease in the objective function $\phi(z)$. These results are illustrated in Fig. 1.

4. PROJECTING ONTO THE FEASIBLE SET

The projection $\bar{z}_P(\alpha)$ of $\bar{z}(\alpha)$ onto the feasible set \mathcal{F} is the closest point in \mathcal{F} in Euclidean distance, which is given by the solution to the optimization problem

$$\begin{aligned} \bar{z}_P(\alpha) &\equiv \arg \min_{z \in \mathbb{R}^{2n}} \pi(z) \equiv \frac{1}{2} \|z - \bar{z}(\alpha)\|_2^2 \\ &\text{subject to } Bz \geq 0. \end{aligned} \quad (6)$$

The projection must be computed at each iterate, which means that (6) must be solved easily and efficiently. Unfortunately, solving this minimization problem can be challenging because of the linear constraints. We propose solving its *dual* problem instead since, as we will demonstrate, it will lead to an easier optimization problem.

The Lagrangian function $\mathcal{L} : \mathbb{R}^{2n} \times \mathbb{R}^{3n} \rightarrow \mathbb{R}$ associated with (6) is given by $\mathcal{L}(z, \mu) = \frac{1}{2} \|z - \bar{z}(\alpha)\|_2^2 - \mu^T Bz$, with the Lagrange multipliers $\mu \in \mathbb{R}^{3n}$ and $\mu \geq 0$. The Lagrange dual function $g : \mathbb{R}^{3n} \rightarrow \mathbb{R}$ is given by $g(\mu) = \inf_z \mathcal{L}(z, \mu)$. Taking the partial derivative of $\mathcal{L}(z, \mu)$ with respect to z and setting it equal to zero yields

$$z = \bar{z}(\alpha) + B^T \mu. \quad (7)$$

Thus, the dual associated with (6) is

$$\begin{aligned} \mu^* &\equiv \text{maximize } g(\mu) = -\frac{1}{2} \mu^T B B^T \mu - \mu^T B \bar{z}(\alpha) \\ &\text{subject to } \mu \geq 0. \end{aligned} \quad (8)$$

We note here that (8) is easier to solve than (6) since the constraints in (8) are simple bound constraints whereas those in (6) are (the more general) linear constraints.

Strong duality. Standard primal-dual optimization theory indicates that for a primal feasible z and a nonnegative (dual feasible) μ , $g(\mu) \leq \pi(z)$. Now, the objective function $\pi(z)$ in the primal problem (6) is convex and the constraints $Bz \geq 0$ are affine. The feasible set \mathcal{F} is non-empty since it contains the origin. Thus, a weaker version of Slater's condition (see [6]) implies that the duality gap is zero, i.e., $g(\mu^*) = \pi(\bar{z}_P(\alpha))$. Thus, the solution to (6) can then be defined from (7) as $\bar{z}_P(\alpha) \equiv \bar{z}(\alpha) + B^T \mu^*$ by solving (8).

5. SOLVING THE DUAL PROBLEM

If we partition μ as $\mu = [v; \zeta]$, where $v \in \mathbb{R}^{2n}$ and $\zeta \in \mathbb{R}^n$, then the dual problem (8) can be equivalently written as

$$\begin{aligned} \text{minimize } h(\mu) &\equiv \frac{1}{2} \|v + \widetilde{W}^T \zeta + \bar{z}(\alpha)\|_2^2 - \frac{1}{2} \|\bar{z}(\alpha)\|_2^2 \\ &\text{subject to } v, \zeta \geq 0. \end{aligned} \quad (9)$$

Note that the objective function $h(\mu)$ is almost separable with respect to v and ζ , with only the term $\zeta^T \widetilde{W} v$ coupling the variables, i.e., $h(\mu)$ is given by

$$h(\mu) = \left(\frac{1}{2} \|v\|_2^2 + v^T \bar{z}(\alpha) \right) + \zeta^T \widetilde{W} v + \left(\|\zeta\|_2^2 + \zeta^T \widetilde{W} \bar{z}(\alpha) \right).$$

We can solve this minimization problem by component-wise minimization, first by fixing ζ and minimizing with respect to v , and then fixing the computed v and optimizing for ζ . Each step can easily be done since the each component-wise minimization has an analytic solution. This process is repeated until convergence. Multiplication by W can be done efficiently for many bases, such as the wavelet transform, which can be performed in $O(n)$ flops. Although solving the constrained problem requires many such multiplications, it is clear from the results below that this computational expense is worth paying since it leads to better performance. However, for most practical imaging situations the multiplication by the sensing matrix A is the primary computational expense.

We now describe each step more explicitly.

Step 1. Given ζ^{j-1} from the previous iterate, solve

$$v^j = \arg \min_{v \in \mathbb{R}^{2n}} \frac{1}{2} \|v\|_2^2 + v^T \left(\bar{z}(\alpha) + \widetilde{W}^T \zeta^{j-1} \right) \quad \text{subject to } v \geq 0,$$

the solution to which is obtained using soft thresholding:

$$v^j = \left[- \left(\bar{z}(\alpha) + \widetilde{W}^T \zeta^{j-1} \right) \right]_+. \quad (10)$$

Step 2. Given v^j , solve

$$\zeta^j = \arg \min_{\zeta \in \mathbb{R}^n} \|\zeta\|_2^2 + \zeta^T \widetilde{W} \left(\bar{z}(\alpha) + v^j \right) \quad \text{subject to } \zeta \geq 0,$$

the solution to which is also obtained using soft thresholding:

$$\zeta^j = \frac{1}{2} \left[-\widetilde{W} \left(\bar{z}(\alpha) + v^j \right) \right]_+. \quad (11)$$

Inexact solution. At each iteration j , the primal variable associated with the dual variable $\mu^j = (v^j, \zeta^j)$, given by $z_j \equiv \bar{z}(\alpha) + B^T \mu^j$ from (7), is feasible with respect to the constraints in (3):

$$\widetilde{W} z_j = \widetilde{W} \left(\bar{z}(\alpha) + B^T \mu^j \right) = \widetilde{W} \bar{z}(\alpha) + \widetilde{W} v^j + 2\zeta^j \geq 0,$$

using (11). This indicates that we can terminate the iterations for the dual problem early and still obtain a feasible point. This feature is particularly useful when the iterates are far from the solution, and the projection onto the feasible set need not be done very accurately.

6. NUMERICAL EXPERIMENTS

To test the effectiveness of our methods, we compare the performance of our proposed algorithms with current gradient projection algorithms. For this experiment, we consider an example of imaging in astronomy, where the scene to be imaged contains regions of low-intensity. In particular, we consider the 256×256 gray-scale image of a dark region on the surface of the planet Mercury [7]. We create the observations using the model in Fig. 2, in which the true intensity f^* (the Mercury image) passes through the optical system described by A . In these experiments A corresponds to a blur operation using the 2D blur kernel $h_{i,j} = 1/(1+i^2+j^2)$, for $i, j = -4, \dots, 4$, which is the same kernel used in Experiment 2 of [8]. This blurred




Fig. 2. Block diagram representation of our observation model. The scene intensity f is passed through an optical system described by the sensing matrix A . We then sense Af on a CCD array, modeled by adding zero-mean Gaussian noise to Af and quantizing the result to an accuracy of b bits. This results in our observations y .

intensity is then captured by a CCD array, which we model as adding zero-mean Gaussian noise of variance $\sigma^2 = 25$, and quantizing the result to an accuracy of $b = 3$ bits per pixel. This model is similar to the one employed in [9].

We compare the results of four methods. First, we consider the Gradient Projection for Sparse Reconstruction (GPSR) [8] method of Figueiredo et al. The reconstruction obtained from this method is not necessarily nonnegative. Therefore we threshold the computed reconstruction to obtain a new result, which leads to our second method which we call GPSR-T. The third method is the method we propose which we call *Linearly Constrained Gradient Projection* (LCGP), which is described in detail in Secs. 4 and 5. We initialize LCGP using the standard initialization $\theta^{(0)} = (AW)^T y$. We propose a fourth method wherein we solve (1) using GPSR, threshold the result, and use it to initialize LCGP. We call this GPSR-initialized method I-LCGP.

For this experiment we fixed the total time $t = 10$ sec allotted for each algorithm. For the I-LCGP method, we run GPSR for five seconds and LCGP for the remaining five seconds. Optimal τ values are determined independently for each method by running the algorithms and minimizing the root mean square (RMS) error at the exhaustion of the time budget. The optimal τ value for I-LCGP is determined by first finding the optimal τ for GPSR after five seconds. Then using this result, we independently find the optimal τ for the LCGP step by minimizing the RMS error at the exhaustion of the total ten second time budget. All the methods use the Barzilai-Borwein approach [4] for choosing α in (5), which often results in faster convergence by allowing nonmonotonic decreases in the objective. We repeat the above procedure a total of ten times so we may examine the ten-trial average performance of the reconstruction methods considered, removing the bias associated with presenting results for a single realization of the noise.

The proposed LCGP method solves (2) by a sequence of quadratic programming subproblems (6), which is solved using the dual formulation (8). From our numerical experience, these subproblems need not be solved very accurately. In our experiment, we needed only three inner iterations of the alternating minimization described in Sec. 5 to obtain an accurate approximation to the projection $\bar{z}_P(\alpha)$. Table 1 shows the resulting RMS from each method, and Fig. 3 shows the reconstruction obtained from the GPSR-T and LCGP methods, as well as the true signal f and degraded observations y . We highlight a particular region of the reconstruction by magnifying it and using a contrast-enhancing colormap.

The reconstruction using GPSR-T improves upon the reconstruction by simply using GPSR, which highlights the importance of achieving a nonnegative reconstruction. However, explicitly constraining the solution to be nonnegative within the optimization

Method	Single-Trial RMS (%)	Ten-Trial Average RMS (%)
GPSR	26.0732	26.0897
GPSR-T	24.9383	24.8635
I-LCGP	24.0673	24.0085
LCGP	23.8694	23.8741

Table 1. Reconstruction RMS for single-trial results and results averaged over ten trials. $\text{RMS}(\hat{f}) \equiv \|\hat{f} - f^*\|/\|f^*\|$.

formulation results in a greater increase in performance. This is seen from Table 1, as the proposed LCGP method achieves the lowest RMS of all the methods considered. Interestingly, the I-LCGP method results in an intermediate RMS error between the GPSR-T and LCGP methods, indicating that splitting the time budget between GPSR and LCGP is sub-optimal. These trends continue if we examine the ten-trial performance in Table 1. The differences between the GPSR-T and LCGP methods are subtle, but by considering a particular image location shown in Figs. 3(e) and 3(f), we see the GPSR-T solution suffers from blocking artifacts near boundaries, and spurious artifacts appear in regions of near-zero intensity. In contrast, the LCGP reconstruction demonstrates fewer blocking artifacts and captures regions of low intensity more accurately.

7. CONCLUSIONS

In this work, we formulated nonnegative sparse image recovery as a constrained ℓ_2 - ℓ_1 convex program (1). The incorporation of ℓ_1 sparsity constraints has proven highly successful in accurately reconstructing natural signals, and the additional nonnegativity constraint realistically estimates light intensities. We proposed the LCGP method for solving the constrained ℓ_2 - ℓ_1 minimization problem by solving a sequence of quadratic subproblems. We demonstrated in a numerical experiment that with very few subproblem iterations, we are able to improve upon the performance of state-of-the-art gradient projection methods currently available. This result suggests that with slightly more computational effort, our proposed approach can lead to more accurate reconstruction. We intend to extend this approach to video applications in future work.

8. REFERENCES

- [1] E. J. Candès and T. Tao, “Decoding by linear programming,” *IEEE Trans. Inform. Theory*, vol. 15, no. 12, pp. 4203–4215, 2005.
- [2] D. L. Donoho, “Compressed sensing,” *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.
- [3] R. Tibshirani, “Regression shrinkage and selection via the lasso,” *J. Roy. Statist. Soc. Ser. B*, vol. 58, no. 1, pp. 267–288, 1996.
- [4] J. Barzilai and J. M. Borwein, “Two-point step size gradient methods,” *IMA J. Numer. Anal.*, vol. 8, no. 1, pp. 141–148, 1988.
- [5] P. H. Calamai and J. J. Moré, “Projected gradient methods for linearly constrained problems,” *Math. Programming*, vol. 39, no. 1, pp. 93–116, 1987.
- [6] S. Boyd and L. Vandenberghe, *Convex optimization*, Cambridge University Press, Cambridge, 2004.
- [7] NASA/Johns Hopkins University Applied Physics Laboratory/ Carnegie Institution of Washington, “PIA12272: Seeing double?,” 2009, <http://photojournal.jpl.nasa.gov/tiff/PIA12272.tif>.
- [8] M. A. T. Figueiredo, R. D. Nowak, and S. J. Wright, “Gradient projection for sparse reconstruction: Application to compressed sensing and other inverse problems,” *IEEE J. Sel. Top. Sign. Proces.: Special Issue on Convex Optimization Methods for Signal Processing*, vol. 1, no. 4, pp. 586–597, 2007.
- [9] Y. Tsin, V. Ramesh, and T. Kanade, “Statistical calibration of CCD imaging process,” in *Proceedings of Eighth IEEE Int. Conf. on Comput. Vision*, 2001, vol. 1, pp. 480–487.

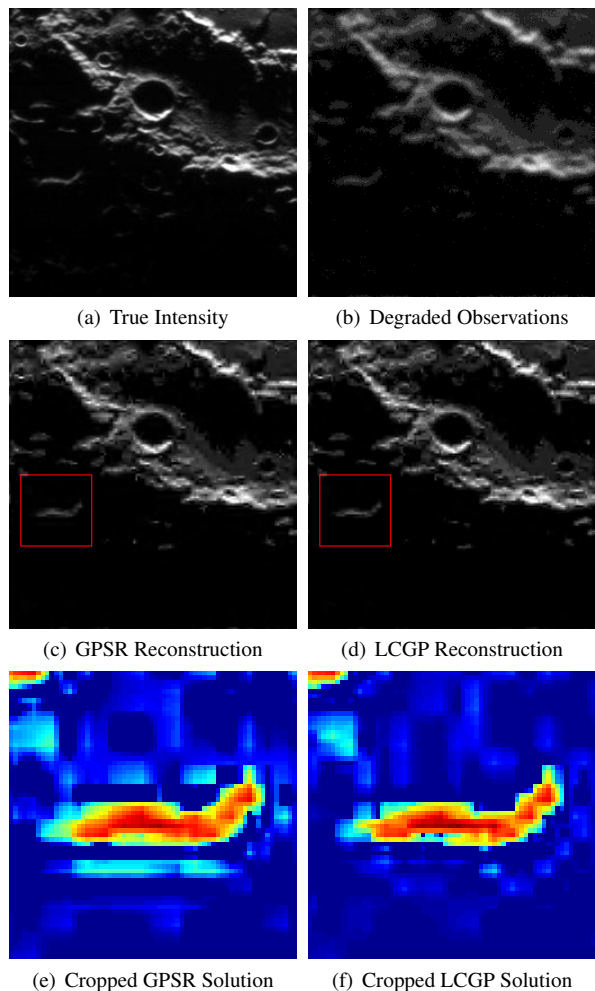


Fig. 3. Results of our numerical experiments. Here we show (a) the true intensity f , (b) the degraded observations y . (c) the reconstruction using GPSR, (d) the reconstruction with the proposed LCGP method. The images (e) and (f) crop to a particular region (in red square) in the reconstructions (c) and (d) using a different colormap on a log-scale to highlight the differences. Note the blocking artifacts near boundaries and in regions of low intensity present in the GPSR solution are less pronounced in the LCGP solution.