Modular Solvers for Constrained Image Restoration Problems Using the Discrepancy Principle

Peter Blomgren
Stanford University
Department of Mathematics
Stanford, CA 94305-2125
blomgren@math.Stanford.EDU

and

Tony F. Chan
University of California, Los Angeles
Department of Mathematics
405 Hilgard Ave
Los Angeles, CA 90095-1555
chan@math.ucla.edu

This work is supported by the ONR under contract ONR N00014-96-1-0277, and by NSF grant DMS-9626755

Abstract

Many problems in image restoration can be formulated as either an unconstrained nonlinear minimization problem, usually with a *Tikhonov*-like regularization, where the regularization parameter has to be determined; or as a fully constrained problem, where an estimate of the noise level, either the variance or the signal-to-noise ratio, is available.

The formulations are mathematically equivalent. However, in practice it is much easier to develop algorithms for the unconstrained problem, and not always obvious how to adapt such methods to solve the corresponding constrained problem.

In this paper, we present a new method which can make use of *any* existing convergent method for the unconstrained problem to solve the constrained one. The new method is based on a Newton iteration applied to an extended system of nonlinear equations, which couples the constraint and the regularized problem, but it does not require knowledge of the Jacobian of the irregularity functional. The existing solver is only used as a black box solver, which for a fixed regularization parameter returns an improved solution to the unconstrained minimization problem given an initial guess.

The new modular solver enables us to easily solve the constrained image restoration problem; the solver automatically identifies the regularization parameter, during the iterative solution process.

We present some numerical results. The results indicate that even in the worst case the constrained solver requires only about twice as much work as the unconstrained one, and in some instances the constrained solver can be even faster.

I. Introduction

The regularized image restoration problem consists of minimizing a regularity functional, $R(u)$, over a function space $X(\Omega)$, subject to constraints relating the resulting image u, to the measured image z. The constraints depend on the noise level, expressed as either the variance, σ^2, of the noise, or the signal-to-noise-ratio, *SNR*; and the blurring of the image, usually expressed as a convolution $\mathbb{K}u$:

$$\arg \min_{u \in X(\Omega)} R(u) \quad \text{subject to} \quad \frac{1}{2}\|\mathbb{K}u - z\|_2^2 = \frac{1}{2}|\Omega|\sigma^2, \tag{1}$$

here $|\Omega|$ is the appropriate measure of the domain over which the images u, and z are defined. This approach is sometimes referred to as applying the *discrepancy principle* to the inverse (restoration) problem, see *e.g.* Engl-Hanke-Neubauer [1]. It is reasonable to consider the constrained problem since estimates for the noise level are obtainable in many applications. The Lagrange function associated with (1) is

$$\mathcal{L}(u, \lambda) = R(u) + \frac{\lambda}{2}\left(\|\mathbb{K}u - z\|_2^2 - |\Omega|\sigma^2\right), \tag{2}$$

and the first order KKT necessary conditions for optimality are

$$\begin{aligned} \nabla R(u) + \lambda \mathbb{K}^* (\mathbb{K}u - z) &= 0 \\ \|\mathbb{K}u - z\|^2 - |\Omega|\sigma^2 &= 0. \end{aligned} \tag{3}$$

There is a wealth of techniques for solving problems of this type:

Newton-based solvers, like the primal-dual solver introduced by Chan-Golub-Mulet [2] solve the system (3) of nonlinear algebraic equations directly. Implementation of Newton-based solvers is difficult, since the linear algebra structures of the functional, $\nabla R(u)$, and the constraints are strongly intertwined in the solver. These solvers are locally quadratically convergent, but global convergence properties are usually not known.

The projected gradient method of Rosen [3] adapted by Rudin-Osher-Fatemi [4] allows the tradeoff between regularization and fit to measured data, $\lambda(t)$, to be dynamically controlled in the setting of an explicit time-marching algorithm:

$$\begin{aligned} u_t = -\nabla R(u) - \lambda(t)\mathbb{K}^* (\mathbb{K}u - z) \\ u(0) = z. \end{aligned} \tag{4}$$

Due to a time-step restriction, this approach converges quite slowly. However, recently Marquina-Osher [5] have found a way to speed up the convergence significantly. In addition, Rudin-Osher [6] used a semi-implicit method to overcome the step restriction.

The motivation for the modular approach is based on the observation that in many cases it is much more straightforward to design schemes for solving the corresponding unconstrained — Tikhonov-like regularized [7] — problem:

$$\arg \min_{u \in X(\Omega)} R(u) + \frac{\hat{\lambda}}{2} \|\mathbb{K}u - z\|_2^2, \tag{5}$$

where $\hat{\lambda}$ is a fixed, given regularization parameter, balancing the trade-off between minimizing the functional and staying true to the measured image. If chosen correctly, $\hat{\lambda}$ is the *Lagrange multiplier*, λ, corresponding to the constraint in the constrained formulation (1). This regularized least squares approach is very common for solving inverse problems.

Most papers in the literature deal with the unconstrained formulation. A variety of methods are used to obtain a good $\hat{\lambda}$, *e.g.* L-curves [8], [9], [10], [11], generalized cross validation [12], the discrepancy principle [1, Section 4.3], etc. This paper does not try to address under what circumstances each of the preceding choices yield the "best" solution.

There exist powerful techniques for solving (5), *e.g.* fixed point iteration [13], unconstrained primal-dual* [2], etc. In most cases the global convergence properties of these algorithms are quite well understood, *e.g.* see Vogel-Dobson [14], Chan-Mulet [15], Chambolle-Lions [16] and Aubert *et al* [17]. In addition they are much easier to implement than solvers for the constrained problem; in some cases the solver may very well be available off-the-shelf in a software library.

The question that begs to be asked is if the easy-to-implement unconstrained solvers can be adapted to solve the fully constrained KKT problem?

The *L-curve* approach of P.C. Hansen [9], [10] is such an approach where a sequence of unconstrained problems are solved in order to identify a $\hat{\lambda}$ that is optimal in some sense. However, the optimality criterion in the L-curve approach does not necessarily correspond to solving the full KKT equations.

The purpose of this paper is to develop a general and easy-to-implement technique whereby any solver for (5) can be adapted to solve (3), without using special knowledge of the functional, the constraints, or the existing solver.

Our approach is conceptually similar to the one of L-curves, in that it solves a sequence of unconstrained problems, while adjusting the regularization parameter. However, the stopping

*Chan-Golub-Mulet used two completely separate implementations for the constrained and unconstrained computations in [2].

criteria for the parameter is quite different: the modular criteria is derived from the constrained formulation (discrepancy principle), whereas the L-curve criteria comes from the shape of the L-curve. Whereas the upper limit for the work required to solve the constrained problem using the modular approach is about two times the cost of solving the unconstrained problem, in section III-D we show that *in practice* the modular approach can solve the constrained problem at a cost not much more than the corresponding unconstrained problem.

We access the unconstrained solver in the form of a black box iteration: $u^{n+1} = S(u^n, \lambda^n)$, which given the estimate u^n to the solution of (5) and a value λ^n approximating λ, returns an improved iterate u^{n+1}. In general it is not necessary for the solver to return an exact solution to (5); it is sufficient that the intermediate solves perform a few iterations of the unconstrained solver.

The modularity facilitates next to effortless replacement of the unconstrained solver, this is conducive to experimentation and performance evaluation of different regularity functionals, discretization schemes, types of constraints, etc...

The modular solver is an adaptation of a general approximate Newton method for coupled nonlinear systems, introduced by Chan [18]. Since the solver is based on a Newton scheme, the overall solver converges quadratically, provided the supplied unconstrained solver has quadratic convergence properties. Local convergence of the method is guaranteed by the theory in [18]. However, this general theory does not make any efficiency predictions. It turns out that in practice the modular approach is very efficient for the image restoration problems considered in this paper.

We present three examples of applications of the modular solver, which demonstrates the versatility of our approach:

(i) Restoration of gray scale images, $u : \Omega \to \mathbb{R}$, where the regularity functional is the total variation norm of Rudin-Osher-Fatemi [4]. Further, we assume Gaussian "white" noise, with variance σ^2, and no blurring, *i.e.* $\mathbb{K} \equiv I$. We base the solver on the quadratically convergent primal-dual solver of Chan-Golub-Mulet [2], which is a sophisticated and powerful technique for the unconstrained problem. The solver is easily adapted to solve the constrained formulation, and retain the local quadratic convergence properties.

(ii) Restoration of vector valued (RGB color) images, $u : \Omega \to \mathbb{R}^3$, where the regularity functional, $R(u) = \mathrm{TV}_{\mathrm{n,m}}(u)$, is the color-TV norm introduced in Blomgren-Chan [19]. Noise and blur are as in *(i)*. For this problem, we base the constrained solver on a fixed-point lagged diffusivity scheme for the unconstrained problem; an adaption of the scheme introduced by Vogel-Oman [13] for the intensity image case. At this point in time, this is the state-of-the art solver for total variation reconstruction of the vector valued images, since the implementation of the primal-dual approach for this case is quite complicated.

(iii) Restoration of a gray scale image degraded by both noise and blur. We use the Total Variation regularization functional, a fixed-point solver using a cosine transform based preconditioner to solver the resulting linear systems [20].

Although the algorithm is presented in the context of total variation functionals, adaptation to other regularizations, *e.g.* H^1, is straight-forward. In addition, it should be quite easy to accommodate for different types of constraints.

II. The Modular Solver

As in Chan [18], we introduce the following notation:

$$G^i(u,\lambda) = -\nabla_i R(u) + \lambda \mathbb{K}^*(\mathbb{K}u_i - z_i), \quad i = 1,2,\ldots,m$$

$$G(u,\lambda) = \begin{bmatrix} G^1(u,\lambda) \\ G^2(u,\lambda) \\ \vdots \\ G^m(u,\lambda) \end{bmatrix}$$

$$N(u,\lambda) = \|\mathbb{K}u - z\|_2^2 - |\Omega|\sigma^2.$$

The KKT first order necessary condition for optimality of the constrained problem (1) — see for instance Nash-Sofer [21, chapter 14]) — can be written as the coupled nonlinear system:

$$\begin{bmatrix} G(u,\lambda) \\ N(u,\lambda) \end{bmatrix} = 0. \tag{6}$$

Our modular solver is based on a Newton iteration applied to this system. Assuming that the solution exists, and is regular enough that the Jacobian

$$J(u,\lambda) = \begin{bmatrix} G_u & G_\lambda \\ N_u & N_\lambda \end{bmatrix}$$

is nonsingular at the solution, at each step of the Newton iteration, we are faced with solving the following *linear* system:

$$\begin{bmatrix} G_u & G_\lambda \\ N_u & N_\lambda \end{bmatrix} \begin{bmatrix} \delta u \\ \delta\lambda \end{bmatrix} = - \begin{bmatrix} G \\ N \end{bmatrix}, \tag{7}$$

yielding the changes $(\delta u, \delta\lambda)$ in the Newton iterates.

We apply a block elimination algorithm to the system (*see* Keller [22]):

First we notice that $N_\lambda = 0$ and pre-multiplying the equation for δu by G_u^{-1} yields:

$$\begin{bmatrix} I & G_u^{-1}G_\lambda \\ N_u & 0 \end{bmatrix} \begin{bmatrix} \delta u \\ \delta\lambda \end{bmatrix} = - \begin{bmatrix} G_u^{-1}G \\ N \end{bmatrix}. \tag{8}$$

Now the key idea in the modular algorithm is to approximate the terms $w = G_u^{-1}G$, and $v = G_u^{-1}G_\lambda$ by calls to the unconstrained solver S. First, observe that w can be approximated by one call to the existing constrained solver, $w = S(u,\lambda) - u$. The reason this works is that $G_u^{-1}G$ corresponds to the Newton correction to u for the unconstrained problem (where λ is fixed).

Moreover, we can approximate v by another call to the solver: Differentiating $G(u,\lambda) = 0$ with respect to λ yields $G_u u_\lambda + G_\lambda = 0$. Hence $v = G_u^{-1}G_\lambda = -u_\lambda$. At convergence $-u_\lambda = -S(u,\lambda)_\lambda = -S_u u_\lambda - S_\lambda$. Provided S is sufficiently contractive, *i.e.* $\|S_u\| \ll 1$, it is reasonable to approximate $v \approx -S_\lambda$. In particular, if the solver S is a Newton solver, this approximation is exact, since $S_u = 0$. In practice, all reasonably fast solvers are contractive enough to make this approximation work. Finally, we use a finite difference approximation to S_λ, so that $v = \frac{S(u,\lambda)-S(u,\lambda+\epsilon)}{\epsilon}$, where $|\epsilon| \ll 1$. Note that due to issues of finite precision, while $|\epsilon| \ll \lambda$ it cannot be too small in practice, see *e.g.* Dennis-Schnabel [23, page 105].

Using the approximation v of $G_u^{-1}G_\lambda$ we rewrite (8):

$$\begin{bmatrix} I & v \\ N_u & 0 \end{bmatrix} \begin{bmatrix} \delta u \\ \delta\lambda \end{bmatrix} = - \begin{bmatrix} -w \\ N \end{bmatrix}. \tag{9}$$

Now, elimination of the $(2,1)$-block yields

$$\begin{bmatrix} I & v \\ 0 & -N_u v \end{bmatrix} \begin{bmatrix} \delta u \\ \delta\lambda \end{bmatrix} = - \begin{bmatrix} -w \\ N + N_u w \end{bmatrix}. \tag{10}$$

Thus we can summarize the algorithm as follows:

ALGORITHM: *Modular Solver*

PROBLEM:

$\arg\min R(u)$, subject to $\frac{1}{2}\|\mathbb{K}u - z\| = \frac{1}{2}|\Omega|\sigma^2$.

ASSUMPTION:

$u \leftarrow S(u, \lambda)$ is a convergent solver for

$\nabla R(u) + \lambda \mathbb{K}^* (\mathbb{K}u - z) = 0.$

1. Compute $\quad w = S(u^n, \lambda^n) - u^n.$
2. Compute $\quad v = [\epsilon]^{-1} [S(u^n, \lambda^n) - S(u^n, \lambda^n + \epsilon)].$
3. Compute $\quad \delta\lambda = [N_u v]^{-1}(N + N_u w)$, where
   - $N = \frac{1}{2}\left((\mathbb{K}u - z)^2 - |\Omega|\sigma^2\right)$,
   - $N_u = \mathbb{K}^* (\mathbb{K}u - z)$.
4. Compute $\quad \delta u = w - v\, \delta\lambda.$
5. Update $\quad \begin{cases} u^{n+1} &= u^n + \delta u \\ \lambda^{n+1} &= \lambda^n + \delta\lambda. \end{cases}$

Hence, each iteration requires two calls to the unconstrained solver, S. The solution obtained from the first call can be used as an initial guess for the second call, thus speeding up this call considerably: for example, when S is the fixed point algorithm of Vogel-Oman [13], the first call may use 5–7 iterations, and the second 2 to reach the same precision.

III. NUMERICAL RESULTS

We present three applications of the modular solver. For stability reasons we damp the update $\delta\lambda$ so that λ never changes by more than one order of magnitude between iterations.

The total variation functionals used in the examples are non-smooth at the origin. The primal-dual solver used in the first example handles this non-linearity by introducing auxiliary (dual) variables, and the solvers in the latter two examples use a small degree of smoothing near the origin. The details are not essential to the discussion of the modular approach, but can be found in [2], [13], [19].

A. Gray Scale Reconstruction

In the first example, we use the total variation norm of Rudin-Osher-Fatemi [4] to restore a gray scale image. The regularity functional is defined by:

$$\mathrm{TV}_{\mathrm{n},1}(\Phi) \equiv \int_\Omega |\nabla \Phi| \, dx,$$

and the associated *Euler-Lagrange* equations are:

$$-\nabla \circ \left(\frac{\nabla u}{|\nabla u|} \right) + \lambda \mathbb{K}^* (\mathbb{K} u - z) = 0.$$

In the primal dual approach of Chan-Golub-Mulet [2] an auxiliary variable $w = \frac{\nabla u}{|\nabla u|}$, $\|w\| = 1$ is introduced. This leads to an enlarged system of equations which is less non-linear than the original one. Thus, the convergence properties of the expanded problem are more favorable.

The result of using the primal-dual algorithm on a simple test image can be seen in Figure 1. As can be seen in Figure 2, the correct *Lagrange multiplier* λ is identified in 5–6 iterations. After that, convergence is quadratic, as expected.

This example shows how we can adapt a complicated, sophisticated, and powerful solver for the unconstrained problem and adapt it to solve the constrained problem with the same efficiency.

B. Color Reconstruction

The second example shows a reconstruction of a color image, where the regularizing functional is the Color-TV norm of Blomgren-Chan [19]:

$$\mathrm{TV}_{\mathrm{n,m}}(u) \equiv \left[\sum_{i=1}^{3} \left[\mathrm{TV}_{\mathrm{n},1}(u_i) \right]^2 \right]^{1/2},$$

and the unconstrained module is a lagged diffusivity fixed point scheme, which is globally and linearly convergent, *see* Vogel-Oman [13]:

$$-\nabla \circ \left(\frac{\nabla u^{n+1}}{|\nabla u^n|} \right) + \lambda \mathbb{K}^* (\mathbb{K} u^{n+1} - z) = 0.$$

For a detailed discussion of the $\mathrm{TV}_{\mathrm{n,m}}$ norm, including comparisons to other possible extensions, see Blomgren-Chan [19]. Sochen-Kimmel-Malladi [24] relate the $\mathrm{TV}_{\mathrm{n,m}}$ norm to a general framework of flows.

The result of the reconstruction can be seen in figure 3. Figure 4 shows the evolution of the Lagrange multiplier, λ, the residual of the constraint, as well as the constraint of the coupled nonlinear system. The convergence is linear as expected, since the fixed point solver is only linearly convergent. Since the tolerance levels for the unconstrained solver were set so that "convergence" meant "below perceptible change in the image," the method stagnates in this example; the recovered image does not change perceptibly after 12 iterations.

Figure 5 shows the evolution of λ for several initial guesses. Even though the modular method is based on a Newton iteration, which is not guaranteed to be globally convergent, we notice that the method is quite robust — typically it converges for initial guesses off by as much as 3 orders of magnitude.

This example shows the state-of-the-art implementation for total variation reconstruction of vector valued images. Implementation of the fixed-point algorithm is quite straight forward, whereas the implementation of the corresponding fully constrained primal-dual algorithm is quite complicated.

C. Reconstruction with Gaussian Blur

The final example shows reconstruction of an image degraded by Gaussian blur, and a small amount of additive noise (SNR $= 25.2\,\mathrm{dB}$). Figure 6 shows the true, degraded, and recovered images, and figure 7 shows convergence statistics. The underlying unconstrained solver, based on a fixed-point algorithm [13] with cosine transform based preconditioning [20], converges linearly, a property which is inherited by the modular solver.

This example shows a setting where the constraint, and pre-conditioning issues are more complicated. Still, the resulting solver is very stable.

D. Work Comparison

Figure 8 shows a convergence comparison of the constrained and unconstrained methods (the data is from example 2, color restoration). The nonlinear residual is plotted against the number of calls to the unconstrained solver. This is a good measure of the total work under the assumption that each call involves the same amount of work. However, if the second call in each iteration, $S(u^n, \lambda^n + \epsilon)$, is given the result of the first call, $S(u^n, \lambda^n)$, as the initial state it will require *less* work since the perturbation is small. Each modular iteration requires three vector-vector inner products, and two matrix-vector products (if $\mathbb{K}$ is non-trivial). This additional work is of lower order, and is neglected in the comparison.

We notice that after the initial search for the correct value of the Lagrange multiplier λ, the constrained solver converges faster than the unconstrained one (which was given the correct multiplier λ).

In practice, the modular solver produces a solution using about the same amount of work compared with the unconstrained solver (which is given the correct multiplier).

IV. Conclusion

As indicated by the three applications presented, the modular solver is easy to implement. The experimental settings are quite different, nonetheless the same solver is successful in all three.

The modularity facilitates next to effortless replacement of the unconstrained solver, this is conducive to experimentation and evaluation of different regularization methods. Identifying the correct Lagrange multiplier is the only way to make fair comparisons between different models.

The approach is very robust. In most cases the modular solver converges for a wide range of initial guesses for the Lagrange multiplier.

The modular solver is efficient. In practice we can compute a solution to the constrained solution in the same amount of time it takes to compute the unconstrained solution.

References

[1] Hienz W. Engl, Martin Hanke, and Andreas Neubauer, *Regularization of Inverse Problems*, *Mathematics and Its Applications* **375**, Kluwer Academic Publishers, (1996).

[2] Tony F. Chan, Gene Golub, and Pep Mulet, "A Nonlinear Primal-Dual Method for TV-Based Image Restoration," in *ICAOS '96, 12th International Conference on Analysis and Optimization of Systems: Images, Wavelets, and PDEs, Paris, June 26–28, 1996*, M. Berger, R. Deriche, I. Herlin, J. Jaffre, and J. Morel, Eds., (1996), number 219 in Lecture Notes in Control and Information Sciences, 241–252.

[3] J. G. Rosen, "The Gradient Projection Methods for Nonlinear Programming, Part II, Nonlinear Constraints," *J. Soc. Indust. Appl. Math* **9**, 514 (1961).

[4] Leonid I. Rudin, Stanley Osher, and Emad Fatemi, "Nonlinear Total Variation Based Noise Removal Algorithms," *Physica D* **60**, 259–268 (1992).

[5] Antonio Marquina and Stanley Osher, "Explicit Algorithms for a New Time Dependent Model Based on Level Set Motion for Nonlinear Deblurring and Noise Removal," Tech. Rep. CAM 99-5, UCLA Department of Mathematics, (January 1999).

[6] Leonid Rudin and Stanley Osher, "Total Variation Based Image Restoration with Free Local Constraints," in *Proceedings of the IEEE International Conference on Image Processing*, (1994), 31–35.

[7] A. N. Tikhonov and V. Y. Arsenin, *Solutions of Ill-Posed Problems*, John Wiley, (1977).

[8] Per Christian Hansen, "Analysis of Discrete Ill-posed Problems by Means of the L-curve," *SIAM Review* (561–580) (1992).

[9] Per Christian Hansen and D.P. O'Leary, "The Use of the L-curve in the Regularization of Discrete Ill-posed probles," *SIAM Journal on Scientific Computing* **14**, 1487–1503 (1993).

[10] Per Christian Hansen, *Rank-Deficient and Discrete Ill-Posed Problems*, Ph.D. thesis, Technical University of Denmark, (1996).

[11] Per Christian Hansen, *Rank-Deficient and Discrete Ill-Posed Problems: Numerical Aspects of Linear Inversion*, SIAM, (1998).

[12] D. A. Girard, "The Fast Monte-Carlo Cross-validation and C_L Procedures: Comments, New Resultsm and Application to Image Recover Problems," *Computational Statistics* **10**, 205–231 (1995).

[13] C. R. Vogel and M. E. Oman, "Fast Total Variation-Based Image Reconstruction," in *Proceedings of the ASME Symposium on Inverse Problems*, (1995).

[14] C.R. Vogel and D. C. Dobson, "Convergence of an Iterative Method for Total Variation Denoising," *SIAM Journal on Numerical Analysis* (submitted).

[15] Tony F. Chan and Pep Mulet, "On the Convergence of the Lagged Diffusivity Fixed Point Method in Total Variation Image Restoration," Tech. Rep. CAM 97-46, UCLA Department of Mathematics, (August 1997), to appear in SNUM.

[16] A. Chambolle and Pierre-Louis Lions, "Image Recovery via Total Variation Minimization and Related Problems," *Numerische Mathematik* **76**, 167–188 (1997).

[17] G. Aubert, M. Barlaud, L. Blanc-Féraud, and P. Charbonnier, "Deterministic Edge-Preserving Regularization in Computed Imaging," Tech. Rep., Université de Nice-Sophia Antipolis, (January 1994), submitted to IEEE Image Processing.

[18] Tony F. Chan, "An Approximate Newton Method for Coupled Nonlinear Systems," *SIAM Journal on Numerical Analysis* **22**(5), 904–913 (October 1985).

[19] Peter Blomgren and Tony F. Chan, "Color TV: Total Variation Methods for Restoration of Vector Valued Images," *IEEE Transactions on Image Processing* **7**(3), 304–309 (March 1998).

[20] R. Chan, T. F. Chan, and C. K. Wong, "Cosine Transform Based Preconditioners for Total Variation Minimization Problems in Image Processing," in *Iterative Methods in Linear Algebra*, S. Margenov and P. Vassilevski, Eds., *IMACS Series in Computational and Applied Math.* **3**, 311–329. IMACS, (1996).

[21] Stephen G. Nash and Ariela Sofer, *Linear and Nonlinear Programming*, McGraw-Hill, (1996).

[22] H. B. Keller, "Numerical Solution of Bifurcation and Nonlinear Eigenvalue Problems," in *Applications of Bifurcation Theory*, P. Rabinowitz, Ed., 359–384. Academic Press, (1977).

[23] J. E. Dennis and Robert B. Schnabel, *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*, *Classics in Applied Mathematics* **16**, SIAM, (1996).

[24] Nir Sochen, Ron Kimmel, and Ravikanth Malladi, "From High Energy Physics to Low Level Vision," Tech. Rep. LBNL-39243 / UC-405, Ernest Orlando Lawrence Berkeley National Laboratory, Physics Division, Mathematics Department, (August 1996).

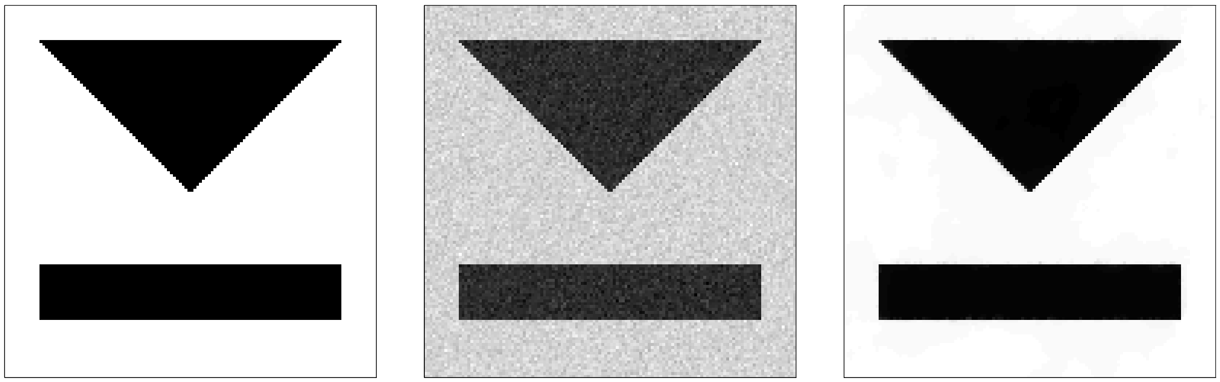

Fig. 1. The true, noisy (SNR = 16.4 dB), and the recovered (SNR = 30.2 dB) images.

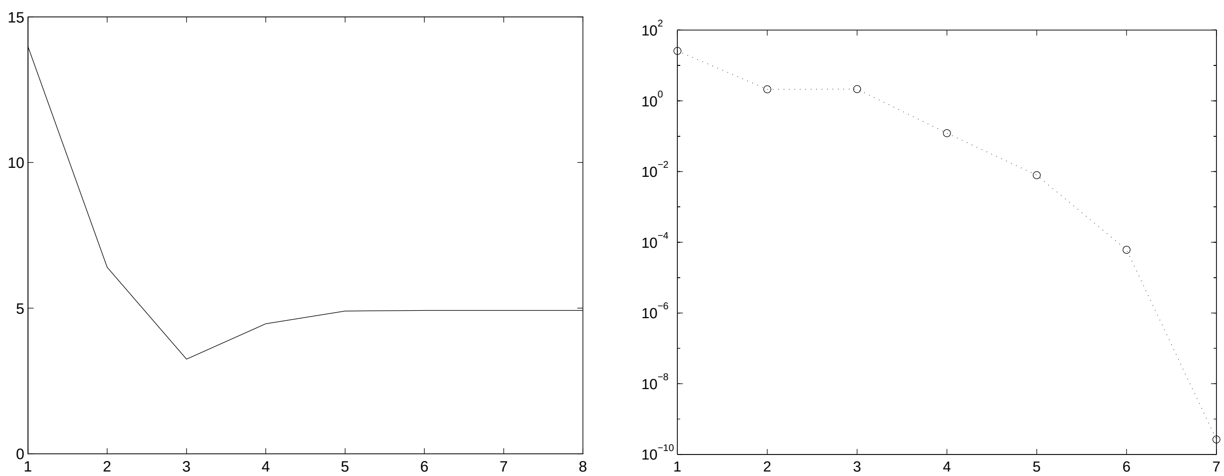

Fig. 2. **Left:** The regularization parameter, λ, as a function of the number of iterations; and **Right:** The residual of the constraint, *e.g.* $N(u, \lambda)$ as a function of the iterations. We notice that the correct λ is found within a few iterations, and convergence is quadratic.

Fig. 3. The true, noisy, and the recovered images.

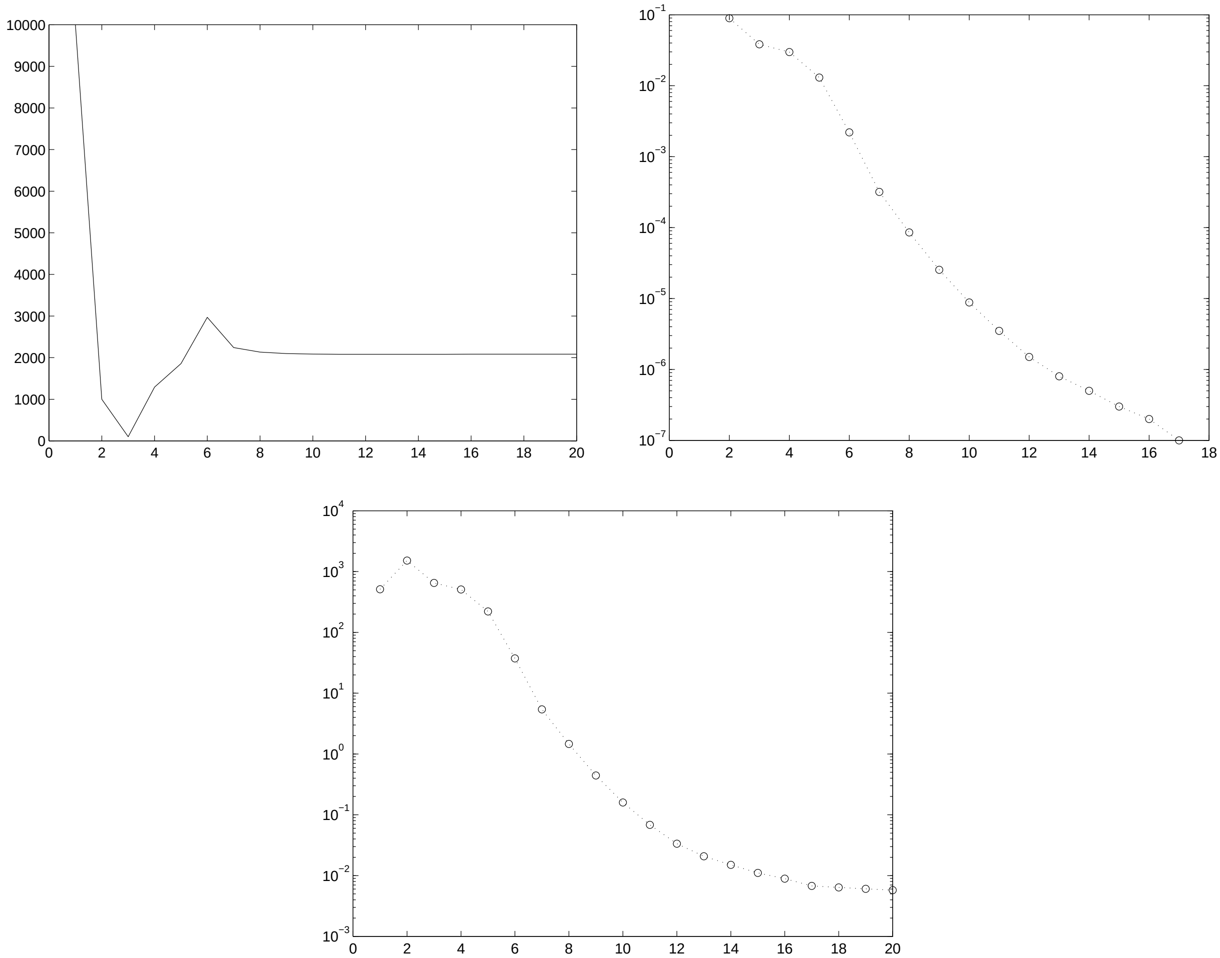

Fig. 4. **Upper-Left:** The regularization parameter, λ, as a function of the number of iterations. **Upper-Right:** The residual of the constraint, *e.g.* $N(u, \lambda)$ as a function of the iterations. **Center:** The residual of the coupled (G, N) system. Again, λ is quickly identified, and the convergence is linear as expected.

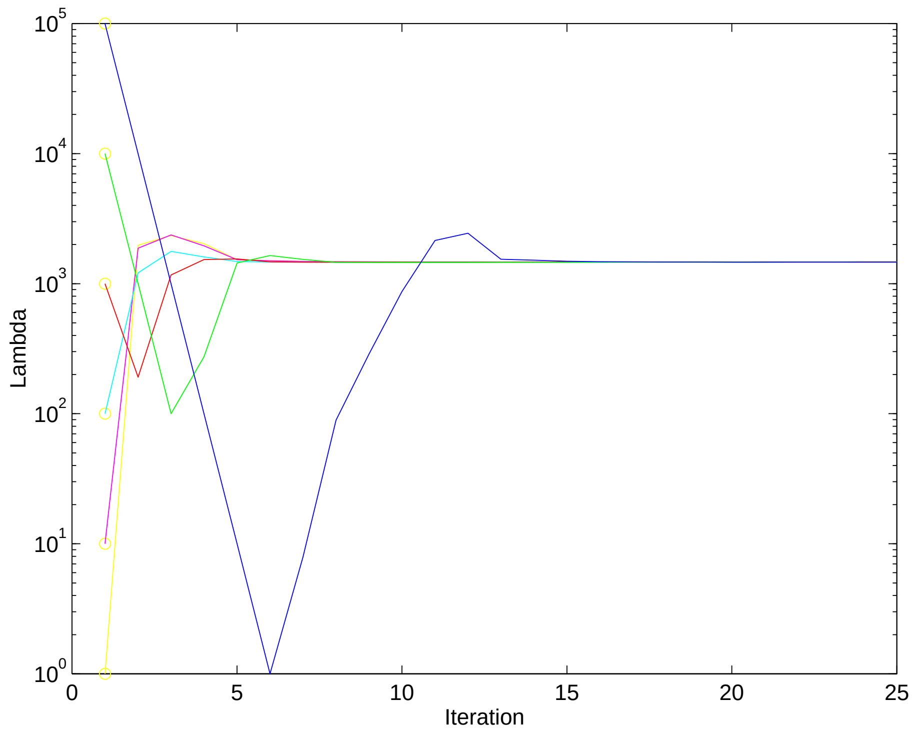

Fig. 5. Convergence of λ for initial guesses ranging five orders of magnitude..

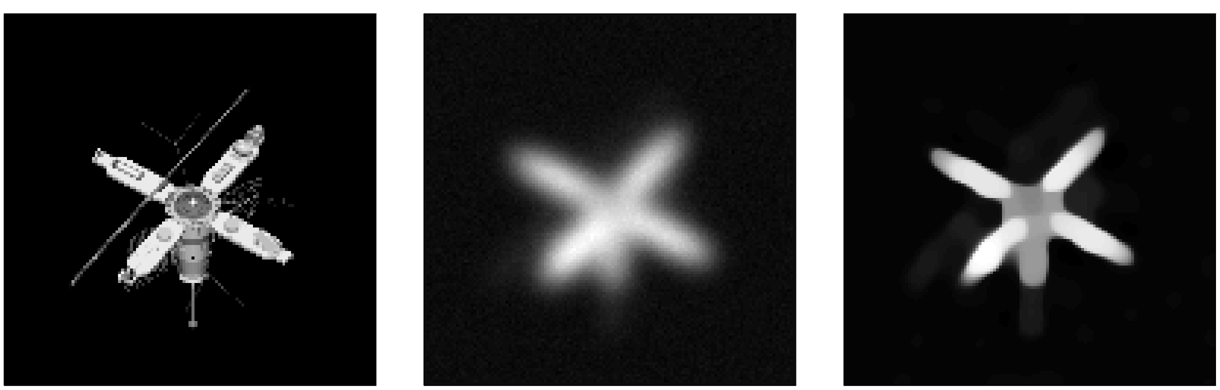

Fig. 6. The true, degraded, and the recovered images.

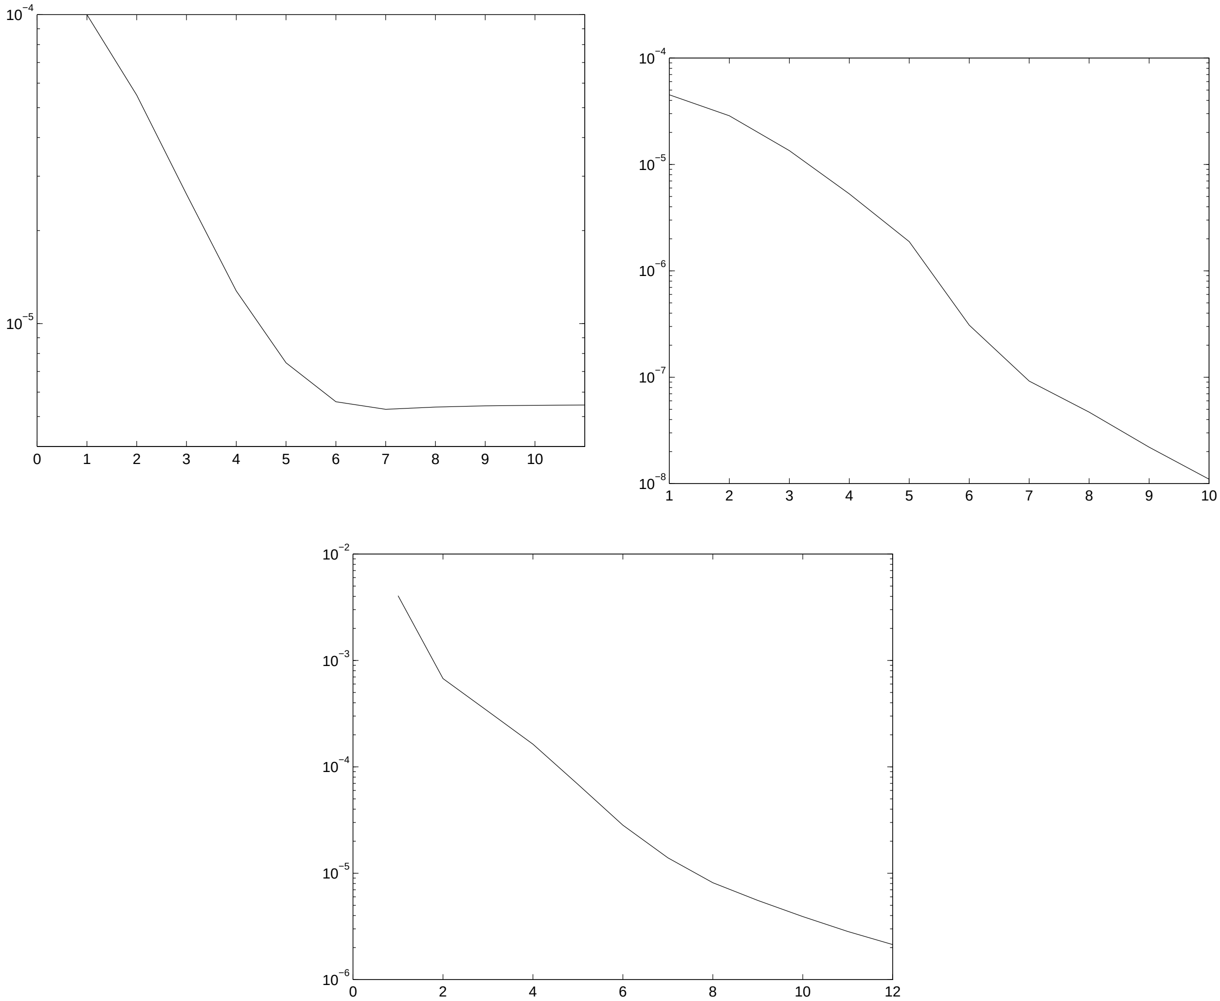

Fig. 7. **Upper-Left:** Convergence of the Lagrange multiplier. **Upper-Right:** The estimated error in the Lagrange multiplier; notice the linear convergence. **Center:** The residual of the nonlinear system; again, the convergence is linear.

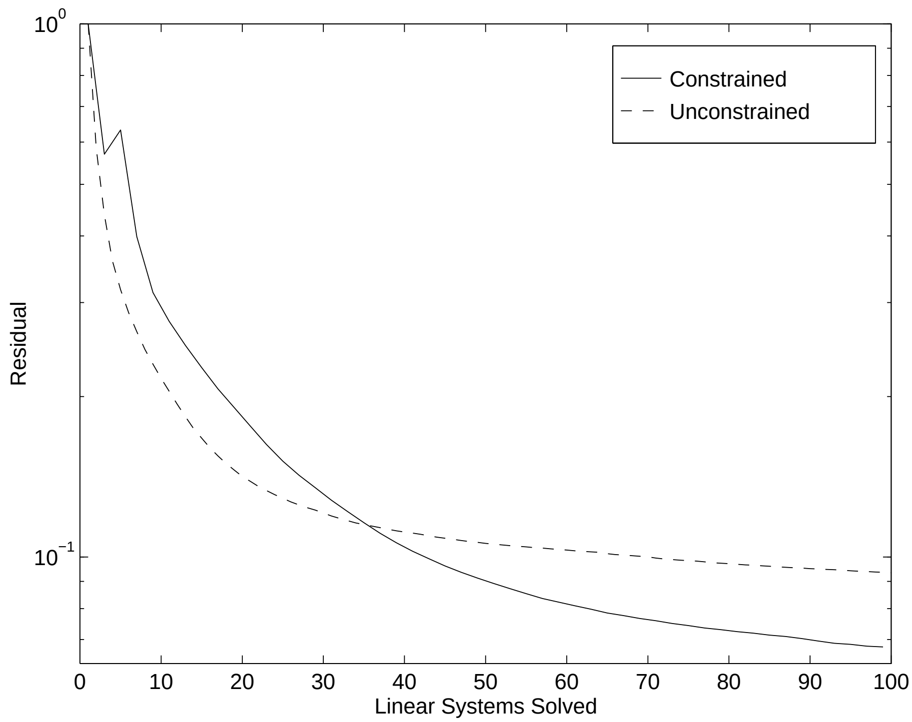

Fig. 8. Comparison of the residual for the constrained, and unconstrained solvers. Notice that the x-axis shows the total number of calls to the unconstrained solver, not the number of iterations.