

OPEN ACCESS

Filter factor analysis of scaled gradient methods for linear least squares

To cite this article: Federica Porta *et al* 2013 *J. Phys.: Conf. Ser.* **464** 012006

View the [article online](#) for updates and enhancements.

You may also like

- [A spatially adaptive total variation regularization method for electrical resistance tomography](#)
Xizi Song, Yanbin Xu and Feng Dong
- [A new interpretation of \(Tikhonov\) regularization](#)
Daniel Gerth
- [INVERSE PROBLEMS NEWSLETTER](#)



PRIME
PACIFIC RIM MEETING
ON ELECTROCHEMICAL
AND SOLID STATE SCIENCE

HONOLULU, HI
Oct 6–11, 2024

Abstract submission deadline:
April 12, 2024

Learn more and submit!

Joint Meeting of

The Electrochemical Society
•
The Electrochemical Society of Japan
•
Korea Electrochemical Society

Filter factor analysis of scaled gradient methods for linear least squares

Federica Porta, Anastasia Cornelio, Luca Zanni and Marco Prato

Department of Physics, Computer Science and Mathematics, University of Modena and Reggio Emilia

E-mail: marco.prato@unimore.it

Abstract. A typical way to compute a meaningful solution of a linear least squares problem involves the introduction of a filter factors array, whose aim is to avoid noise amplification due to the presence of small singular values. Beyond the classical direct regularization approaches, iterative gradient methods can be thought as filtering methods, due to their typical capability to recover the desired components of the true solution at the first iterations. For an iterative method, regularization is achieved by stopping the procedure before the noise introduces artifacts, making the iteration number playing the role of the regularization parameter. In this paper we want to investigate the filtering and regularizing effects of some first-order algorithms, showing in particular which benefits can be gained in recovering the filters of the true solution by means of a suitable scaling matrix.

1. Introduction

An ill-posed discrete inverse problem can be modeled with a linear system

$$\mathbf{b} = \mathbf{A}\mathbf{x}_{true} + \boldsymbol{\eta}, \quad (1)$$

where \mathbf{A} is an ill-conditioned full-rank $m \times n$ matrix, with $m \geq n$, $\mathbf{b} \in \mathbb{R}^m$ is the observed data, $\mathbf{x}_{true} \in \mathbb{R}^n$ is the unknown object we want to recover and $\boldsymbol{\eta} \in \mathbb{R}^m$ is the noise corrupting the data. A classical example of the mathematical model (1) can be found in image restoration [2, 11]: in this case the matrix \mathbf{A} represents the blurring effect that the acquisition process introduces on the unknown image \mathbf{x} .

Since discrete ill-posed problems have the property that the singular values decay to zero, a simple inversion of the data would amplify the noise providing a solution without any physical meaning. Especially when dealing with large scale problems, as the image deblurring one, a typical way to overcome this situation consists of iteratively building up a sequence of arrays that converges to the solution of the following least squares problem

$$\min_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x}) \equiv \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2, \quad (2)$$

and stopping the iterations before the noise distorts the reconstructions.

Due to the large diffusion of the least squares approach to address many real-world applications, the development of efficient methods to solve (2) has been widely treated in literature. In this paper, we consider the family of the scaled gradient methods, and we analyze the role played by



well-known steplengths and scaling matrices in reconstructing the true solution. The starting point of our work is a paper of Nagy & Palmer [14], in which the authors investigate how the solution achieved by some classical gradient methods can be written as a linear combination of the singular vectors of \mathbf{A} through appropriate filter factors [11]. Here we extend this analysis to the presence of a scaling matrix multiplying the gradient direction, showing its regularizing effect in recovering the true filter coefficients.

The paper is organized as follows: in Section 2 we introduce the optimization methods for the solution of (2) we decided to analyze, while in Section 3 the state-of-the-art on these algorithms as filtering methods is summarized, and the extension to the scaled case is provided. Section 4 will be devoted to some numerical experiments we carried out on a simulated image deblurring problem, while our conclusions are offered in Section 5.

2. Gradient methods

A gradient method for the solution of (2) is an iterative algorithm whose $(k+1)$ -th element is defined by

$$\mathbf{x}_{k+1} = \mathbf{x}_k - \alpha_k \mathbf{M}_k \mathbf{g}_k, \quad (3)$$

where $\alpha_k > 0$ is the steplength, \mathbf{M}_k is a symmetric and positive definite scaling matrix and $\mathbf{g}_k = \nabla f(\mathbf{x}_k) = \mathbf{A}^T(\mathbf{A}\mathbf{x}_k - \mathbf{b})$ is the gradient vector. The idea of finding a descent direction by scaling the gradient of the objective function is rather classical in numerical optimization. The most famous example of scaled gradient method is provided by the Newton algorithm, in which the scaling matrix is given by the inverse of the Hessian $\nabla^2 f(\mathbf{x}_k)$.

Among the large variety of selection rules for the steplength, classical examples are the Steepest Descent (SD) [5] and the Minimal Gradient (MG) [6, 15] parameters, which minimize $f(\mathbf{x}_k - \alpha \mathbf{M}_k \mathbf{g}_k)$ and $\|\nabla f(\mathbf{x}_k - \alpha \mathbf{M}_k \mathbf{g}_k)\|_2$, respectively:

$$\alpha_k^{SD} = \frac{\mathbf{g}_k^T \mathbf{M}_k \mathbf{g}_k}{\|\mathbf{A} \mathbf{M}_k \mathbf{g}_k\|_2^2} \quad ; \quad \alpha_k^{MG} = \frac{\mathbf{g}_k^T \mathbf{M}_k \mathbf{A}^T \mathbf{A} \mathbf{g}_k}{\|\mathbf{A}^T \mathbf{A} \mathbf{M}_k \mathbf{g}_k\|_2^2}. \quad (4)$$

In order to accelerate the slow convergence exhibited in most cases by the standard formulas (4), many other strategies for the steplength selection have been proposed, as the two Barzilai and Borwein rules [1] adapted to account for the scaling matrix [4], obtained by regarding the matrix $\mathbf{B}(\alpha_k) = (\alpha_k \mathbf{M}_k)^{-1}$ as an approximation of the Hessian $\nabla^2 f(\mathbf{x}_k)$ and forcing a quasi-Newton property on $\mathbf{B}(\alpha_k)$. The resulting steplengths are given by

$$\alpha_k^{BB1} = \frac{\mathbf{s}_{k-1}^T \mathbf{M}_k^{-1} \mathbf{M}_k^{-1} \mathbf{s}_{k-1}}{\mathbf{s}_{k-1}^T \mathbf{M}_k^{-1} \mathbf{y}_{k-1}} \quad ; \quad \alpha_k^{BB2} = \frac{\mathbf{s}_{k-1}^T \mathbf{M}_k \mathbf{y}_{k-1}}{\mathbf{y}_{k-1}^T \mathbf{M}_k \mathbf{M}_k \mathbf{y}_{k-1}},$$

where $\mathbf{s}_{k-1} = \mathbf{x}_k - \mathbf{x}_{k-1}$, $\mathbf{y}_{k-1} = \mathbf{g}_k - \mathbf{g}_{k-1}$, and their several alternated versions [9, 15].

As for the scaling matrices, in this paper we consider the following two examples arising from the constrained optimization:

- the one provided by the iterative space reconstruction algorithm (ISRA) [8], whose explicit expression of the scaling matrix is

$$\mathbf{M}_k^{ISRA} = \text{diag} \left(\frac{\mathbf{x}_k}{\mathbf{A}^T \mathbf{A} \mathbf{x}_k} \right),$$

where the quotient is intended in the Hadamard sense;

- the one proposed by Hager, Mair and Zhang (HMZ) [10], that exploits a gradient splitting strategy [12, 13] with a resulting scaling matrix given by

$$\mathbf{M}_k^{HMZ} = \text{diag} \left(\frac{\alpha_k^{CBB1} \mathbf{x}_k}{\mathbf{x}_k + \alpha_k^{CBB1} (\mathbf{A}^T \mathbf{A} \mathbf{x}_k + \mathbf{A}^T \mathbf{b})^+} \right),$$

where $t^+ = \max\{0, t\}$ and α_k^{CBB1} is a cyclic version of the first Barzilai and Borwein steplength rule computed by reusing α_k^{BB1} for p consecutive iterations [7].

For both matrices, the positive definiteness is ensured by thresholding the diagonal elements in a prefixed interval $0 < L_{\min} < L_{\max}$. We remark that here we are *not* considering the ISRA and HMZ algorithms, but we only borrow the scaling matrices defined in the algorithms themselves and use them in our scaled gradient scheme.

3. Filter factors analysis

Let $\mathbf{A} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^T$ be the singular value decomposition (SVD) of \mathbf{A} , where $\mathbf{U} = [\mathbf{u}_1 \mathbf{u}_2 \dots \mathbf{u}_m]$ and $\mathbf{V} = [\mathbf{v}_1 \mathbf{v}_2 \dots \mathbf{v}_n]$ are unitary matrices and $\mathbf{\Sigma} \in \mathbb{R}^{m \times n}$ is a diagonal matrix with entries $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_n > 0$. When solving the linear ill-posed inverse problem (1), the generalized solution

$$\mathbf{x}^\dagger = \sum_{i=1}^n \frac{\mathbf{u}_i^T \mathbf{b}}{\sigma_i} \mathbf{v}_i = \mathbf{x}_{true} + \sum_{i=1}^n \frac{\mathbf{u}_i^T \boldsymbol{\eta}}{\sigma_i} \mathbf{v}_i$$

results to be totally useless due to the division by the small singular values and the consequent magnification of the corresponding noise components. A typical regularization strategy consists in simultaneously preserving the highest σ_i while filtering out the smaller ones by means of an array of positive weights φ_i . The corresponding regularized solution is given by

$$\mathbf{x}_{reg} = \sum_{i=1}^n \varphi_i \frac{\mathbf{u}_i^T \mathbf{b}}{\sigma_i} \mathbf{v}_i, \quad (5)$$

where φ_i are the so-called filter factors. The truncated singular value decomposition (TSVD) and the Tikhonov method are examples of this “direct” regularization approach, in which the filter factors are defined by

$$\varphi_i = \begin{cases} 1 & \text{if } i \leq r \\ 0 & \text{if } i > r \end{cases} \quad ; \quad \varphi_i = \frac{\sigma_i^2}{\sigma_i^2 + \lambda},$$

where $r \in \{1, \dots, n\}$ and $\lambda > 0$.

The general iteration of any gradient method can also be interpreted as a filtered regularized solution, since it can be written as

$$\mathbf{x}_{k+1} = \sum_{i=1}^n \varphi_i^{k+1} \frac{\mathbf{u}_i^T \mathbf{b}}{\sigma_i} \mathbf{v}_i,$$

where φ_i^{k+1} are opportune filter factors, depending on α_k and \mathbf{M}_k , automatically defined during the iterative procedure. In particular, if we assume $\mathbf{x}_0 = \mathbf{0}$, then:

- when $\mathbf{M}_k = \mathbf{I}$ for each iteration k , the iterative filter factors can be written as

$$\varphi_i^{k+1} = 1 - \prod_{\ell=0}^k (1 - \alpha_\ell \sigma_i^2), \quad k = 0, 1, \dots \quad (6)$$

- when \mathbf{M}_k is not trivial, the expression for φ_i^{k+1} is more complicated. In particular, once defined the polynomial P_k , acting on a $n \times n$ matrix $\mathbf{\Omega}$ as

$$P_k(\mathbf{\Omega}) = P_{k-1}(\mathbf{\Omega}) + \alpha_k \mathbf{M}_k(\mathbf{I} - \mathbf{\Omega} P_{k-1}(\mathbf{\Omega})) , \quad P_{-1}(\mathbf{\Omega}) = \mathbf{0},$$

we have

$$\varphi_i^{k+1} = \sigma_i^2 \left(\frac{\sum_{j=1}^n (\mathbf{Q}_k)_{ij} \sigma_j \mathbf{u}_j^T \mathbf{b}}{\sigma_i \mathbf{u}_i^T \mathbf{b}} \right) , \quad k = 0, 1, \dots \quad (7)$$

where $\mathbf{Q}_k = \mathbf{V}^T P_k(\mathbf{A}^T \mathbf{A}) \mathbf{V}$.

From equation (7) we can see that the presence of \mathbf{M}_k on the filters expression makes any factor φ_i^{k+1} explicitly related to the whole singular system of \mathbf{A} , that acts in a different way on each i -th component of the filter vector. For the nonscaled case, this dependence is hidden in the α_ℓ coefficients, that are the same for all filters $\varphi_1^{k+1}, \dots, \varphi_n^{k+1}$, while in the analytical expression only the i -th singular value σ_i is explicitly present. In Section 4, we will show the positive effect of such more complicated dependence in reconstructing the actual values of the true solution filter factors

$$\varphi_i^{true} = \sigma_i \frac{\mathbf{v}_i^T \mathbf{x}_{true}}{\mathbf{u}_i^T \mathbf{b}} . \quad (8)$$

4. Numerical experiments

The numerical experiments are carried out on a 2-dimensional image restoration test problem. We used the *phantom* test image from Matlab's Image Processing toolbox, artificially blurred by a Gaussian point spread function (PSF) with variance 1, in order to simulate the degrading effect on a real image due to the action of a general acquisition system. The resulting blurred data has been corrupted with 1% white Gaussian noise. We set $n = m = 4096$, thus obtaining \mathbf{x}_{true} and \mathbf{b} of 64×64 pixels. The first row of Figure 1 contains the true image, the PSF modeling the blurring effect and the measured data. When performing the scaled methods, we adopt the BB2 rule for the steplength selection, with a further Armijo reduction procedure to ensure the convergence.

Table 1 shows the best reconstruction errors $\|\mathbf{x}_k - \mathbf{x}_{true}\|_2 / \|\mathbf{x}_{true}\|_2$ for the methods we considered. For the HMZ matrix, we chose the value $p = 4$ for the cyclic version of the BB1 steplength rule defining \mathbf{M}_k . Moreover, the diagonal elements of the scaling matrices provided by the ISRA and HMZ approaches have been thresholded in the range $[10^{-3}, 10^8]$. From the results of Table 1, we can notice that the presence of the scaling matrix provides clear improvements in the minimum error reached by the corresponding methods, even if higher number of iterations are required. We have to remark that, for the scaled algorithms, the error behaviour rapidly decrease with the iterations and then stands in a very flat region, therefore values very close to the minimum can be obtained in both cases by stopping the procedure hundreds of iterations before the optimal value. In order to appreciate the positive effects of the scaling matrix on the reconstructions, the second row of Figure 1 reports the restored images by the different algorithms (since the reconstructions of the nonscaled methods were very similar, we reported only the BB2 case). In particular, we can observe that the presence of a scaling matrix helps to remove some artifacts due to the noise on the data and better recover some details of the solution. The better performances of the scaled algorithms are also visible in the plots of the filter coefficients; in particular, in Figure 2, we compare the filter factors generated at the iteration corresponding to the minimum error for the different nonscaled and scaled gradient methods and the "optimal" coefficients (8) associated to the true solution. The filter factors generated by the nonscaled method with different steplength exhibit smooth trends, which are not able to follow the (somewhat surprising) irregular profile of the filters related to \mathbf{x}_{true} , due

to the presence of noise in the vector \mathbf{b} at the denominator of the filters expression (8). On the contrary, the presence of the same denominator in the analytical form (7) allows a more faithful reconstruction of the filters by the scaled algorithms, whose behaviour results to be in general as erratic as that of the true ones.

Table 1. Numbers of iterations and minimum error reached for solving the image deblurring problem.

	Nonscaled gradient				Scaled gradient	
	MG	SD	BB1	BB2	ISRA+BB2	HMZ+BB2
Iter	1339	1331	195	204	1495	570
Min Err	0.2636	0.2636	0.2636	0.2636	0.1820	0.2258

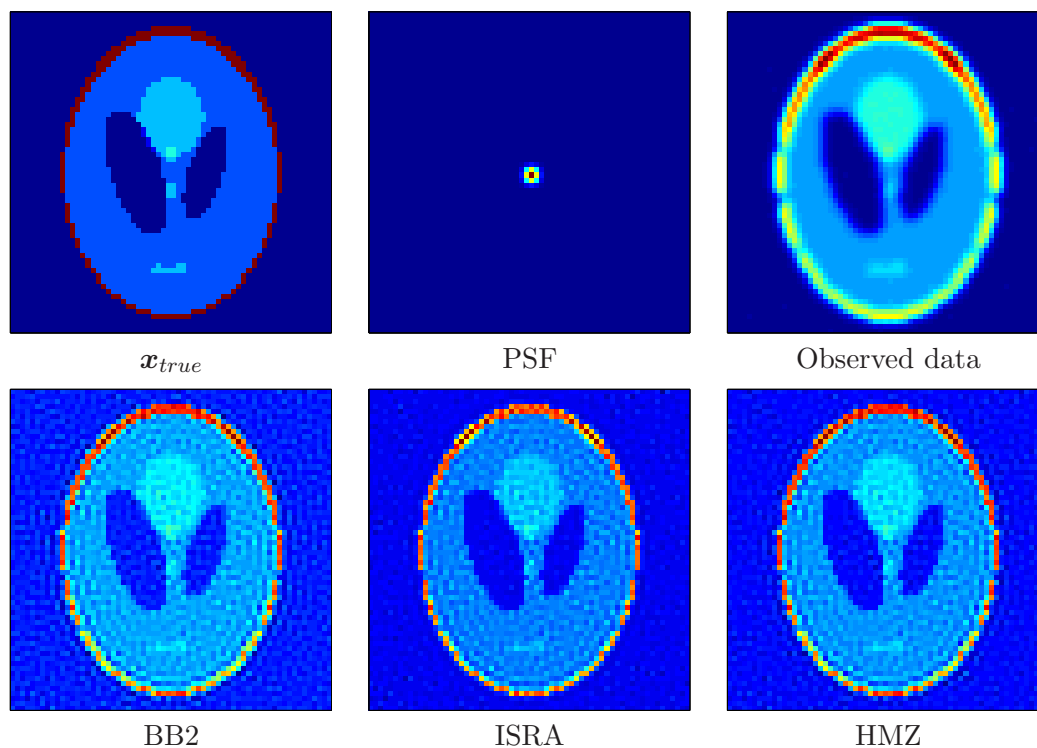


Figure 1. First row: true image, Gaussian PSF and measured data. Second row: reconstructions obtained with BB2 steplength and, from the left, $\mathbf{M}_k = \mathbf{I}$, $\mathbf{M}_k = \mathbf{M}_k^{ISRA}$, $\mathbf{M}_k = \mathbf{M}_k^{HMZ}$.

5. Conclusions

In this paper we conducted a comparative study between several iterative gradient methods for linear least squares problems, with the aim of investigating their ability in providing a regularized solution, i.e., a useful and stable solution not corrupted by the presence of noise on the data. In particular, the analysis we carried out regards the ability of a given scheme to reproduce faithfully the filter factors of the true solution. The filtering properties of nonscaled gradient methods had already been described by Nagy & Palmer: in our paper we generalized their considerations to the presence of a nontrivial scaling matrix in the descent direction. In particular, the formal




Figure 2. Comparison of the \mathbf{x}_{true} filter factors (red dashed) with the ones generated by the gradient method with SD, MG, BB1, BB2 and the scaled (ISRA, HMZ) versions (black solid).

expression of the filter factors has been given and their numerical evaluation has been performed. We showed that the nonscaled methods are not able to generate filter factors that behave as erratically as those associated to the true solution. On the other hand, the scaled methods not only are able to recover the irregular profile of the exact filters, but also provide a better approximation of the unknown solution in terms of accuracy and goodness of the results (as remarked e.g. in [3]).

Acknowledgments

This work has been partially supported by the Italian Spinner2013 PhD Project “High-complexity inverse problems in biomedical applications and social systems” and by a grant of the Italian Gruppo Nazionale per il Calcolo Scientifico (GNCS) - Istituto Nazionale di Alta Matematica (INdAM).

References

- [1] Barzilai J and Borwein J M 1988 *IMA J. Numer. Anal.* **8** 141–8
- [2] Bertero M and Boccacci P 1998 *Introduction to inverse problems in imaging* (Bristol: Institute of Physics)
- [3] Bonettini S Landi G Loli Piccolomini E and Zanni L 2013 *Int. J. Comput. Math.* **90** 9–29
- [4] Bonettini S Zanella R and Zanni L 2009 *Inverse Probl.* **25** 015002
- [5] Cauchy A 1847 *C. R. Acad. Sci. Paris* **25** 536–8
- [6] Dai Y H and Yuan Y X 2003 *IMA J. Numer. Anal.* **23** 377–93
- [7] Dai Y H Hager W W Schittkowski K and Zhang H 2006 *IMA J. Numer. Anal.* **26** 604–27
- [8] Daube-Witherspoon M E and Muehlener G 1986 *IEEE T. Med. Imaging* **5** 61–6
- [9] Frassoldati G Zanni L and Zanghirati G 2008 *J. Indust. Manag. Optim.* **4** 299–312
- [10] Hager W W Mair B A and Zhang H 2009 *Math. Program.* **119** 1–32
- [11] Hansen P C Nagy J G O’Leary D P 2006 *Deblurring images: matrices, spectra and filtering* (Philadelphia, PA: SIAM)
- [12] Lantéri H Roche M Cuevas O and Aime C 2001 *Signal Process.* **81** 945–74
- [13] Lantéri H Roche M and Aime C 2002 *Inverse Probl.* **18** 1397–419
- [14] Nagy J and Palmer K 2003 *BIT* **43** 1003–17
- [15] Zhou B Gao L and Dai Y H 2006 *Comput. Optim. Appl.* **35** 69–86