A proximal approach to the inversion of ill-conditioned matrices

Pierre Maréchal, Aude Rondepierre

A proximal approach to the inversion of ill-conditioned matrices

Pierre Maréchal[*]   Aude Rondepierre[†]

Abstract

We propose a general proximal algorithm for the inversion of ill-conditioned matrices. This algorithm is based on a variational characterization of pseudo-inverses. We show that a particular instance of it (with constant *regularization parameter*) belongs to the class of *fixed point* methods. Convergence of the algorithm is also discussed.

1 Introduction

Inverting ill-conditioned large matrices is a challenging problem involved in a wide range of applications, including inverse problems (image reconstruction, signal analysis, etc.) and partial differential equations (computational fluid dynamics, mechanics, etc.). There are two classes of methods: the first one involves factorization of the matrix (SVD, QR, LU, LQUP); the second one involves iterative schemes (fixed point methods, projection onto increasing sequences of subspaces).

The main purpose of this note is to show that a particular instance of the Proximal Point Algorithm provides a fixed point method for the problem of matrix inversion. This fact is based on the observation that the pseudo-inverse $M^\dagger$ of a matrix $M \in \mathbb{R}^{m \times n}$ satisfies the *fixed point equation*:

$$\Phi = \varphi(\Phi) := B\Phi + C, \quad \text{with} \quad B := (I + \mu M^\top M)^{-1} \quad \text{and} \quad C := (M^\top M + \mu^{-1} I)^{-1} M^\top$$

where $\mu > 0$. The corresponding *fixed point iteration* $\Phi_{k+1} = B\Phi_k + C$ is nothing but a proximal iteration. We see that φ is a contraction and that, if $M^\top M$ is positive definite, then φ is a strict contraction. It is worth noticing that, in the proximal algorithm, μ may depend on k, allowing for large (but inaccurate) steps for early iteration and small (but accurate) steps when approaching the solution.

---

[*]Institut de Mathématiques, Université Paul Sabatier, 31062 Toulouse, France. *pr.marechal@gmail.com*

[†]Institut de Mathématiques, INSA de Toulouse, Département GMM, 31077 Toulouse, France. *aude.rondepierre@math.univ-toulouse.fr*

The Proximal Point Algorithm (PPA) was introduced in 1970 by Martinet [5], in the context of the regularization of variational inequalities. A few years later, Rockafellar [6] generalized the PPA to the computation of zeros of a maximal monotone operator. Under suitable assumptions, it can be used to efficiently minimize a given function, by finding iteratively a zero in its Clarke subdifferential.

Throughout, we denote by $\|M\|_F$ the Frobenius norm of a matrix $M \in \mathbb{R}^{m \times n}$ and by $\langle M, N \rangle_F$ the Frobenius inner product of $M, N \in \mathbb{R}^{m \times n}$ (which is given by $\langle M, N \rangle_F = \operatorname{tr}(MN^\top) = \operatorname{tr}(N^\top M)$). In $\mathbb{R}^{m \times n}$, we denote by $\operatorname{dist}_F(M, \mathcal{S})$ the distance between a matrix M and a set $\mathcal{S}$:

$$\operatorname{dist}_F(M, \mathcal{S}) := \inf\{\|M - M'\|_F | M' \in \mathcal{S}\}.$$

The identity matrix will be denoted by I, its dimension being always clear from the context.

The next theorem, whose proof may be found *e.g.* in [1], provides a variational characterization of $M^\dagger$.

Theorem 1.1 *The pseudo-inverse of a matrix $M \in \mathbb{R}^{m \times n}$ is the solution of minimum Frobenius norm of the optimization problem*

$$(\mathcal{P}) \quad \text{Minimize } f(\Phi) := \frac{1}{2}\|M\Phi - I\|_F^2 \quad \text{over } \mathbb{R}^{n \times m}.$$

2 The proximal point algorithm

The proximal point algorithm is a general algorithm for computing zeros of maximal monotone operators. A well-known application is the minimization of a convex function f by finding a zero in its subdifferential. In our setting, it consists in the following steps:

1. Choose an initial matrix $\Phi_0 \in \mathbb{R}^{m \times n}$;

2. Generate a sequence $(\Phi_k)_{k \geq 0}$ according to the formula

$$\Phi_{k+1} = \underset{\Phi \in \mathbb{R}^{m \times n}}{\operatorname{argmin}} \left\{ f(\Phi) + \frac{1}{2\mu_k}\|\Phi - \Phi_k\|_F^2 \right\}, \tag{1}$$

   in which $(\mu_k)_{k \geq 0}$ is a sequence of positive numbers, until some stopping criterion is satisfied.

Equation (1) will be subsequently referred to as the *proximal iteration* of Problem $(\mathcal{P})$. The stopping criterion may combine, as usual, conditions such as

$$\|\nabla f(\Phi_k)\|_F \leq \varepsilon_1 \quad \text{and} \quad \|\Phi_k - \Phi_{k-1}\|_F \leq \varepsilon_2,$$

where the parameters ε_1 and ε_2 control the precision of the algorithm.

Clearly, the function $f : \Phi \mapsto \|M\Phi - I\|_F^2/2$ is convex and indefinitely differentiable. Therefore, solutions of the proximal iteration (1) are characterized by the relationship $\nabla f(\Phi_{k+1}) + \mu_k^{-1}(\Phi_{k+1} - \Phi_k) = 0$ i.e.,

$$(I + \mu_k M^\top M)\Phi_{k+1} = \Phi_k + \mu_k M^\top. \tag{2}$$

Since $M^\top M$ is positive semi-definite and μ_k is chosen to be positive for all k, the matrix $(I + \mu_k M^\top M)$ is nonsingular and the proximal iteration also reads:

$$\Phi_{k+1} = \left(I + \mu_k M^\top M\right)^{-1} \left(\Phi_k + \mu_k M^\top\right). \tag{3}$$

The iterates Φ_k could be computed either exactly (in the ideal case), or approximately, using *e.g.* any efficient minimization algorithm to solve the proximal iteration (1). In that case, we need another stopping criterion and we here choose the following one suggested in [4]:

$$\|\Phi_{k+1} - A(\Phi_k + \mu_k M^\top)\|_F \leq \epsilon_k \min\{1, \|\Phi_{k+1} - \Phi_k\|_F^r\}, \; r > 1 \tag{4}$$

where $\epsilon_k > 0$ and the series $\sum \epsilon_k$ is convergent. Notice that the larger r, the more accurate the computation of Φ_{k+1}. Notice also that, in the case where $\mu_k = \mu$ for all k, each proximal iteration involves the multiplication by the same matrix $A := (I + \mu M^\top M)^{-1}$, and that the latter inverse may be easy to compute numerically, if the matrix $I + \mu M^\top M$ is well-conditioned.

We now turn to convergence issues. Recall that our objective function f is a quadratic function whose Hessian $M^\top M$ is positive semi-definite. Nevertheless, unless $M^\top M$ is positive definite, the matrix $I - (I + \mu M^\top M)^{-1}$ is in general singular and the classical convergence theorem for iterative methods (see e.g. [2]) is not helpful here to prove the convergence of our proximal scheme. The following proposition is a consequence of Theorem 2.1 in [4]. For clarity, we shall denote by $\mathcal{M}$ the linear mapping $\Phi \mapsto M\Phi$, by $\mathcal{L}$ the linear mapping $\Phi \mapsto M^\top M\Phi$ and by $\mathcal{A}$ the linear mapping $\Phi \mapsto A\Phi = (I + \mu M^\top M)^{-1}\Phi$.

Proposition 2.1 *Let α_1 be the smallest nonzero eigenvalue of $\mathcal{L}$ and let E_1 be the corresponding eigenspace. Assume that $\mu_k = \mu$ for all k and that Φ_0 is not $\langle \cdot, \cdot \rangle_F$-orthogonal to the eigenspace E_1. Then,*

$$\frac{\|A(\Phi_{k+1} - \Phi_k)\|_F}{\|\Phi_{k+1} - \Phi_k\|_F} \to \frac{1}{1 + \alpha_1 \mu} \quad and \quad \frac{\Phi_{k+1} - \Phi_k}{\|\Phi_{k+1} - \Phi_k\|_F} \to \Psi_1 \quad as \quad k \to \infty,$$

in which Ψ_1 is a unit eigenvector in E_1. Moreover the sequence (Φ_k) generated by the proximal algorithm, either with infinite precision or using the stopping criterion (4) for the inner loop, converges linearly to the orthogonal projection of Φ_0 onto the solution set $\mathcal{S} := \operatorname{argmin} f = M^\dagger + \ker \mathcal{M}$.

PROOF. *Step 1.* The error $\Delta_{k+1} := \Phi_{k+1} - \Phi_k$ (at iterate $k + 1$) satisfies: $\Delta_{k+1} = (I + \mu M^\top M)^{-1}\Delta_k$. The latter iteration is that of the power method for the linear mapping $\mathcal{A}$.

Clearly, $\mathcal{A}$ is symmetric and positive definite. Consequently, its eigenvalues are strictly positive and there exists an unique eigenvalue of largest modulus (not necessarily simple).

Notice now that the eigenspace E_1 associated to the eigenvalue α_1 is nothing but the eigenspace of $\mathcal{A}$ for its largest eigenvalue strictly smaller than 1, namely, $1/(1+\alpha_1\mu)$. We proceed as in [7, Theorem 1] to obtain the desired convergence rate *via* that of the iterated power method.

Step 2. We now establish the linear convergence of the sequence (Φ_k). First, the solution set $\mathcal{S}$ is clearly nonempty since it contains $M^\dagger$. Moreover, let (Φ_k) be a sequence generated by the PPA algorithm using the stopping criterion (4). Let us prove that

$$\exists a > 0, \ \exists \delta > 0, \ \forall \Phi \in \mathbb{R}^{m \times n}, \quad \left[\|\nabla f(\Phi)\|_F < \delta \Rightarrow \mathrm{dist}_F(\Phi, \mathcal{S}) \leq a\|\nabla f(\Phi)\|_F\right], \qquad (5)$$

which is nothing but Condition (2.1) in [4, Theorem 2.1], in our context. Let $\Phi \in \mathbb{R}^{m \times n}$ and let $\bar{\Phi}$ be the orthogonal projection of ϕ over $(\ker \mathcal{M})^\perp$. It results from the classical theory of linear least squares that $\mathrm{dist}_F(\Phi, \mathcal{S}) = \|\bar{\Phi} - M^\dagger\|_F$. Since $M^\top M \bar{\Phi} - M^\top = \nabla f(\bar{\Phi})$ and $M^\top M M^\dagger - M^\top = 0$, we also have: $\nabla f(\Phi) = M^\top M(\bar{\Phi} - M^\dagger)$. Moreover, $\bar{\Phi} - M^\dagger \in (\ker \mathcal{M})^\perp = (\ker \mathcal{L})^\perp$, so that

$$\|\nabla f(\Phi)\|_F = \|M^\top M(\bar{\Phi} - M^\dagger)\|_F \geq \alpha_1\|\bar{\Phi} - M^\dagger\|_F.$$

It follows that (5) is satisfied with $a = 1/\alpha_1$. The conclusion then follows from [4, Theorem 2.1]: the sequence (Φ_k) converges linearly with a rate bounded by $a/\sqrt{a^2 + \mu^2} = 1/\sqrt{1 + \mu^2\alpha_1^2} < 1$.

Step 3. By rewriting the proximal iteration in an orthonormal basis of eigenvectors of $\mathcal{L}$, we finally prove that the limit of the sequence (Φ_k) is the orthogonal projection of Φ_0 onto $\mathrm{argmin}\, f$. $\blacksquare$

A complete numerical study, which goes beyond the scope of this paper, is currently in progress and will be presented in a forthcoming publication. Let us merely mention that our proximal approach makes it possible to combine features of factorization methods (in the proximal iteration) with features of iterative schemes. In particular, if M if invertible, it shares with iterative methods the absence of error propagation and amplification, since each iterate can be regarded as a new initial point of a sequence which converges to the desired solution.

3 Comments

Tikhonov approximation. A standard approximation of the pseudo-inverse of an ill-conditioned matrix M is $(M^\top M + \varepsilon I)^{-1}M^\top$, where ε is a small positive number. This approximation is nothing but the Tikhonov regularization of $M^\dagger$, with regularization parameter ε. It is worth noticing that the choice $\Phi_0 = 0$ in the proximal algorithm yields the latter approximation for $\varepsilon = 1/\mu$ after one proximal iteration.

Trade-offs. At the k-th proximal iteration, the perturbation of the objective function f is, roughly speaking, proportional to the square of the distance between the current iterate

and the solution set of $(\mathcal{P})$, and inversely proportional to μ_k. In order to speed up the algorithm, it seems reasonable to choose large μ_k for early iterations, yielding large but inaccurate steps, and then smaller μ_k for late iterations, where *proximity* with the solution set makes it suitable to perform small and accurate steps. This is especially true in the case where M is invertible, since the solution set then reduces to $\{M^{-1}\}$. Moreover, numerical accuracy in early proximal iteration may be irrelevant, since the limit of the proximal sequence is what really matters. A trade-off between a rough approximation of the searched proximal point and an accurate and costly solution must be found. As suggested in [3], one may use the following stopping criterion for the proximal iteration:

$$f(\Phi_{k+1}) - f(\Phi_k) \leq \delta \langle \nabla f(\Phi_{k+1}), \Phi_{k+1} - \Phi_k \rangle_F.$$

This criterion is an Armijo-like rule: the algorithm stops when the improvement of the objective function f is at least a given fraction $\delta \in (0,1)$ of its ideal improvement.

Inversion versus linear systems. It is often unnecessary to compute the inverse of a matrix M, in particular when the linear system $M\mathbf{x} = \mathbf{d}$ must be solved for a few data vectors $\mathbf{d}$ only. In such cases, of course, the usual proximal strategy may be used to compute least squares solutions. It is important to realize that, although the regularization properties of the proximal algorithm are effective at every proximal iteration, perturbations of $\mathbf{d}$ may still have dramatic effects on the algorithm if M is ill-conditioned. In applications for which no perturbation of the data must be considered, accurate solutions may be reached by a proximal strategy. We emphasize that, in the minimization of $\Phi \mapsto \|M\Phi - I\|_F$, the *data I* undergoes no perturbation whatsoever.

References

[1] L. Amodei and J.-P. Dedieu, *Analyse Numérique Matricielle*. Collection Sciences Sup, Dunod, 2008.

[2] P.G. Ciarlet. *Introduction to numerical linear algebra and optimisation*. Cambridge Texts in Applied Mathematics, Cambridge University Press, 1989.

[3] R. Correa and C. Lemaréchal. *Convergence of some algorithms for convex minimization*. Math. Programming, vol. 62, pp. 161–275, 1993.

[4] F.J. Luque. Asymptotic convergence analysis of the proximal point algorithm. SIAM Journal on Control and Optimization, vol. 22 (2), pp. 277–293, 1984.

[5] B. Martinet. *Régularisation d'inéquations variationelles par approximations successives*. Revue Française d'Informatique et de Recherche Opérationelle, pp. 154–159, 1970.

[6] R.T. Rockafellar. *Monotone operators and the proximal point algorithm*. SIAM Journal on Control and Optimization, vol. 14 (5), p877–898, 1976.

[7] G. Vige *Proximal-Point Algorithm for Minimizing Quadratic Functions.* INRIA research report RR-2610, 1995.