

Unmatched Preconditioning of the Proximal Gradient Algorithm

Marion Savanier, Emilie Chouzenoux, Jean-Christophe Pesquet, and Cyril Riddell

Abstract

This work addresses the resolution of penalized least-squares problems using the proximal gradient algorithm (PGA). PGA can be accelerated by preconditioning strategies. However, typical effective choices of preconditioners may correspond to intricate matrices that are not easily inverted, leading to increased complexity in the computation of the proximity step. To relax these requirements, we propose an unmatched preconditioning approach where two metrics are used in the gradient step and the proximity step. We provide convergence conditions for this new iterative scheme and characterize its limit point. Simulations for tomographic image reconstruction from undersampled measurements show the benefits of our approach for various simple choices of metrics.

Index Terms

Computed tomography, convergence analysis, image reconstruction, matrix approximation, proximal methods.

I. INTRODUCTION

First-order optimization methods are attractive tools for solving a broad class of (possibly nonsmooth) optimization problems due to their scalability and flexibility. Among them is the proximal gradient algorithm (PGA) [13], [16], also called forward-backward splitting method. However, first-order optimization methods often converge slowly [6], bringing into question their relevance for high-dimensional applications. Inertial extensions [1], [25] improve the theoretical convergence rate of these methods, but they may only lead to moderate practical accelerations. In an attempt to reach greater accelerations, several authors have explored preconditioning strategies [7], [26] to improve the conditioning of the gradient step in PGA [2], [5], [8], [15]. Preconditioning PGA consists of performing a change of metric to reduce the condition number of the linear operator (typically, Hessian) involved in updating the smooth part of the cost function. Preconditioning does not affect the fixed points of PGA. The main issue with this strategy is that the preconditioning metric used in the gradient term must theoretically also be included, via its inverse, in the proximity operator, thus increasing the computational cost of PGA iterations. Computationally cheap metrics resulting from rough approximations to the Hessian operator can be conveniently inverted (e.g., diagonal matrices), but they may fail to achieve significant acceleration. Moreover, the proximity operator might no longer have a closed form in the chosen metric. In this case, the proximity operator must be computed through inner iterations so that the extra computations needed outweigh the benefit in terms of convergence rate. Designing both effective and efficient preconditioners is then especially challenging for proximal methods [12].

This work investigates the use of unmatched preconditioners in PGA for penalized least-squares problems. Our contributions focus on (i) showing that convergence of PGA can be established when different preconditioners are used for the gradient and proximity steps, (ii) providing associated error bounds, and (iii) illustrating the good performance of our approach for image reconstruction in Computed Tomography (CT). Section II of this work introduces the considered objective function as well as the general form of the preconditioned PGA. Section III reformulates an iteration of this algorithm with decoupled metrics and investigates its new stability conditions. Numerical simulations are presented in Section IV, and Section V summarizes our findings.

II. PROXIMAL GRADIENT ALGORITHM

a) Notation: Along this work, the underlying signal space is the N -dimensional Euclidean space \mathbb{R}^N endowed with the standard scalar product $\langle \cdot, \cdot \rangle$ and norm $\| \cdot \|$. Moreover, $\|L\|$ denotes the spectral norm of a square matrix L associated to the vector norm $\| \cdot \|$, $\text{Ker } L$ designates its nullspace, L^\dagger is its pseudoinverse, and $\|L\|_F$ its Frobenius norm. Let $Q \in S_N^+$ (resp. \mathcal{D}_N^+) be a symmetric positive-definite (resp. diagonal matrix with positive elements) in $\mathbb{R}^{N \times N}$. We define the proximity operator of a proper, lower semicontinuous, convex function $\psi : \mathbb{R}^N \mapsto \mathbb{R} \cup \{+\infty\}$ (i.e., $\psi \in \Gamma_0(\mathbb{R}^N)$) at $z \in \mathbb{R}^N$, relative to the metric induced by Q , and denoted by $\text{prox}_\psi^Q(z)$ [21], as the unique minimizer of $\psi + \frac{1}{2} \| \cdot - z \|_Q^2$. Here, $\| \cdot \|_Q$ denotes the Q -weighted norm, i.e., for every $z \in \mathbb{R}^N$, $\|z\|_Q = \sqrt{\langle z | Qz \rangle}$. $\text{prox}_\psi^{\text{Id}}$ is the original form of the proximity operator [22].

This work was supported by the European Research Council Starting Grant MAJORIS ERC-2019-STG-850925, the ANRT CIFRE Convention 2018/1587, and the ANR Research and Teaching Chair in Artificial Intelligence BRIDGEABLE.

E. Chouzenoux, J.-C. Pesquet and M. Savanier are with Univ. Paris-Saclay, CentraleSupélec, CVN, Inria, Gif-sur-Yvette, France (e-mail: first-name.lastname@centralesupelec.fr).

M. Savanier and C. Riddell are with GE Healthcare, Buc, France (e-mail: first-name.lastname@ge.com).

b) *Model*: We consider the reconstruction of a vector $\bar{x} \in \mathbb{R}^N$ from an observation $y \in \mathbb{R}^K$ according to

$$y = H\bar{x} + b, \quad (1)$$

where $H \in \mathbb{R}^{K \times N}$ is a linear operator, and $b \in \mathbb{R}^K$ is a noise term assumed to be additive zero-mean, and Gaussian with precision matrix $W \in S_N^+$. To estimate \bar{x} , we rely on the well-known penalized least-squares criterion:

$$\underset{x \in \mathbb{R}^N}{\text{minimize}} \quad \frac{1}{2} \|y - Hx\|_W^2 + \frac{\kappa}{2} \|x\|^2 + g(x), \quad (2)$$

where $\kappa \in [0, +\infty[$ and $g \in \Gamma_0(\mathbb{R}^N)$ is a suitable, possibly non-smooth, regularization function.

c) *Preconditioned PGA*: For solving Problem (2), PGA preconditioned by $Q \in S_N^+$ [5] reads, for every $n \in \mathbb{N}$,

$$x_{n+1} = (1 - \theta_n)x_n + \theta_n \text{prox}_{\gamma g}^Q(x_n - \gamma Q^{-1}(Mx_n - H^\top W y)), \quad (3)$$

where M is the Hessian of the smooth part of the cost function:

$$M = H^\top W H + \kappa \text{Id}. \quad (4)$$

If $\theta_n \in [\epsilon, 1]$ with $\epsilon \in]0, 1[$ and $\gamma \in]0, 2/\alpha[$, where α is the Lipschitz constant of the preconditioned gradient operator $Q^{-1}(M \cdot - H^\top W y)$, i.e. $\alpha = \|\|Q^{-1/2} M Q^{-1/2}\|\|$, then the sequence $(x_n)_{n \in \mathbb{N}}$ generated by Algorithm (3) converges to a solution to Problem (2) when such a solution exists [11], [17]. The basic form of PGA is recovered when $Q = \text{Id}$ [16].

We see that the convergence of PGA depends on the norm of the Hessian of the quadratic part of the objective function, and more particularly on its maximum eigenvalue. A large value results in small gradient steps leading to a slow convergence.

The criteria for identifying a good preconditioner Q are twofolds: (i) the action of Q and Q^{-1} on an element $x \in \mathbb{R}^N$ should be easily computed; (ii) the conditioning number of $Q^{-1}M$ should be small. These criteria often limit the use of preconditioning to a simple form of regularizer g , e.g. a positivity constraint [4], an ℓ_1 or quadratic penalty. In CT reconstruction, most preconditioners were proposed in the case when no proximity operator is involved [18]–[20], [27], [29]. More details can be found in Appendix A.

III. PRECONDITIONING WITH UNMATCHED METRICS

The standard form of preconditioned PGA with matrix Q is appealing because it converges in a reduced number of iterations for a suitably chosen Q and simple regularization choice. However, as we pointed out, such a choice may lead to severe limitations. To include a more generic prior, we introduce a second matrix $P \in \mathbb{R}^{N \times N}$ such that $PM \neq 0$, that leads to a relaxed version of Iteration (3) which reads

$$x_{n+1} = (1 - \theta_n)x_n + \theta_n \text{prox}_{\gamma g}^Q(x_n - \gamma P(Mx_n - H^\top W y)). \quad (5)$$

We say that P and Q are matched when $PQ = \text{Id}$, hence we recover (3). Note that Iteration (5) can be viewed as an inexact version of (3) where the error involved in the argument of the proximity operator is $e(x_n)$ with

$$(\forall x \in \mathbb{R}^N) \quad e(x) = \gamma(Q^{-1} - P)(Mx - H^\top W y). \quad (6)$$

However, the assumption of summability of the error often adopted in the literature [16], [28] is generally not satisfied in our context.

A. Cocoercivity conditions

By introducing the nonlinear operator

$$T_\gamma: \mathbb{R}^N \rightarrow \mathbb{R}^N: x \mapsto \text{prox}_{\gamma g}^Q((\text{Id} - \gamma L)x + \gamma P H^\top W y), \quad (7)$$

Iteration (5) can be expressed more concisely as

$$x_{n+1} = (1 - \theta_n)x_n + \theta_n T_\gamma x_n,$$

$$\text{with} \quad L = P(H^\top W H + \kappa \text{Id}) = PM \in \mathbb{R}^{N \times N}. \quad (8)$$

We also define $L_Q = Q^{1/2} L Q^{-1/2}$. It results from simple algebra that

$$\|\|L\|\|_Q = \sup_{x \in \mathbb{R}^N} \frac{\|Lx\|_Q}{\|x\|_Q} = \|\|L_Q\|\| \quad (9)$$

and the adjoint of L in metric $\|\cdot\|_Q$ is

$$L^* = Q^{-1} L^\top Q = Q^{-1/2} L_Q^\top Q^{1/2}. \quad (10)$$

Let us introduce the following quantities which will play a prominent role in our analysis:

$$\lambda_{\min} = \inf_{\substack{x \in \mathbb{R}^N \\ \|x\|_Q=1}} \langle x | Lx \rangle_Q = \inf_{\substack{x \in \mathbb{R}^N \\ \|x\|=1}} \langle x | L_Q x \rangle, \quad (11)$$

$$\lambda_{\min}^+ = \inf_{\substack{x \in (\text{Ker } L_Q)^\perp \\ \|x\|=1}} \langle x | L_Q x \rangle, \quad \lambda_{\max} = \sup_{\substack{x \in \mathbb{R}^N \\ \|x\|=1}} \langle x | L_Q x \rangle, \quad (12)$$

$$\beta = \frac{1}{2} \|L - L^*\|_Q = \frac{1}{2} \|L_Q - L_Q^\top\|. \quad (13)$$

Note that λ_{\min} (resp. λ_{\max}) is the minimum (resp. maximum) spectral value of $(L_Q + L_Q^\top)/2$. Recall also that cocoercivity, as defined below, is a key concept in fixed point theory [14]:

Definition III.1 Let $\eta \in]0, +\infty[$. L is η -cocoercive in $(\mathbb{R}^N, \|\cdot\|_Q)$ if

$$(\forall x \in \mathbb{R}^N) \quad \langle x | Lx \rangle_Q \geq \eta \|Lx\|_Q^2. \quad (14)$$

We now provide conditions for this property to be satisfied.

Proposition III.2 L is cocoercive in $(\mathbb{R}^N, \|\cdot\|_Q)$ with $\eta \in]0, +\infty[$ if and only if

$$\begin{cases} \lambda_{\min} \geq 0 \\ \text{Ker}(L_Q + L_Q^\top) = \text{Ker } L_Q. \end{cases} \quad (15)$$

$$(16)$$

Then, the maximum cocoercity constant of L is

$$\bar{\eta} = \frac{2}{\|(\text{Id} + (L_Q - L_Q^\top)(L_Q + L_Q^\top)^\dagger)(L_Q + L_Q^\top)^{1/2}\|^2}. \quad (17)$$

In particular, L is cocoercive in $(\mathbb{R}^N, \|\cdot\|_Q)$ with constant $\underline{\eta} = 1/(\sqrt{\lambda_{\max}} + \beta/\sqrt{\lambda_{\min}^+})^2 \leq \bar{\eta}$.

Proof. For $z \in \mathbb{R}^N$, by setting $x = Q^{-1/2}z$, it follows that Condition (14) is equivalent to

$$(\forall z \in \mathbb{R}^N) \quad \langle z | L_Q z \rangle \geq \eta \|L_Q z\|^2. \quad (18)$$

In other words, L is η -cocoercive in $(\mathbb{R}^N, \|\cdot\|_Q)$ if and only if L_Q is η -cocoercive in \mathbb{R}^N . The result then follows from [9, Proposition 3.4], which provides cocoercivity conditions for L_Q . \square

From (8) and (10), some special cases are worth being considered:

- (i) If $P = Q^{-1}$ (matched preconditioning) then $L_Q = Q^{-1/2}MQ^{-1/2}$ and $\beta = 0$, which leads to $\bar{\eta} = \underline{\eta} = 1/\|L_Q\|$.
- (ii) If M is invertible and $P = M^{-1}$ (Newton preconditioning) then $L = \text{Id}$, $L_Q = Q^{1/2}Q^{-1/2} = \text{Id}$, and $\beta = 0$, which leads to $\bar{\eta} = \underline{\eta} = 1$.

B. Fixed points of T_γ

Let \hat{x} be a solution to Problem (2). Then, \hat{x} satisfies the following first-order optimality condition:

$$0 \in Q^{-1}\partial g(\hat{x}) + Q^{-1}M\hat{x} - Q^{-1}H^\top W y, \quad (19)$$

where $\partial g(\hat{x})$ denotes the subdifferential of g at \hat{x} . Similarly, for every $\gamma \in]0, +\infty[$, the fixed point set of operator T_γ is

$$\mathcal{F} = \{\tilde{x} \in \mathbb{R}^N \mid 0 \in Q^{-1}\partial g(\tilde{x}) + L\tilde{x} - PH^\top W y\}. \quad (20)$$

One can notice that \mathcal{F} is no longer the set of minimizers of the objective function in (2). We first characterize the existence and uniqueness of a fixed point of T_γ .

Proposition III.3 Assume that Conditions (15)-(16) hold.

- (i) \mathcal{F} is nonempty if $\text{dom } \partial g = \mathbb{R}^N$ and the function defined as

$$(\forall x \in \mathbb{R}^N) \quad h(x) = \frac{1}{2} \langle x | Lx \rangle_Q + g(x) \quad (21)$$

is coercive, i.e., $\lim_{\|x\| \rightarrow +\infty} h(x) = +\infty$.

- (ii) \mathcal{F} is a singleton if g is strongly convex or $\lambda_{\min} > 0$.

Proof.

(i) According to (8) and (20), $\tilde{x} \in \mathcal{F}$ if and only if $\tilde{x} = Q^{-1/2}\tilde{z}$ and

$$Q^{1/2}PH^\top W y \in Q^{-1/2}\partial g(Q^{-1/2}\tilde{z}) + L_Q\tilde{z}. \quad (22)$$

To prove the existence of a solution to this inclusion, let us define the auxiliary function \tilde{h} as

$$\tilde{h}: x \mapsto h(Q^{-1/2}x) = \frac{1}{2}\langle x \mid L_Q x \rangle + g(Q^{-1/2}x). \quad (23)$$

If $\text{dom } \partial g = \mathbb{R}^N$ and h (hence \tilde{h} , since $Q^{-1/2}$ is invertible) is coercive, the cocoercivity of L_Q allows us to apply [9, Proposition 3.10(iv)(a-b)] that establishes the surjectivity of $\partial(g \circ Q^{-1/2}) + L_Q$. Given that $Q^{-1/2} \circ \partial g \circ Q^{-1/2} + L_Q = \partial(g \circ Q^{-1/2}) + L_Q$, this shows the existence of \tilde{z} in (22).

(ii) According to (22), $\tilde{x} \in \mathcal{F}$ is uniquely defined if and only if \tilde{z} is. From the assumption of convexity made on g , and the cocoercivity of L_Q , it follows from [3, Corollary 23.37] that \tilde{x} is unique if either $\partial(g \circ Q^{-1/2})$ or L_Q is strongly monotone [3]. This holds in particular if g is strongly convex or $\lambda_{\min} > 0$.

C. Convergence results

The unmatched preconditioned scheme (5) benefits from convergence results similar to those existing for the standard preconditioned PGA. However, there usually exists a discrepancy between the limit point \tilde{x} of (5) and any minimizer \hat{x} of the objective function in (2).

Proposition III.4¹ *Suppose that $\mathcal{F} \neq \emptyset$. Assume that Conditions (15) and (16) hold. Let $\bar{\eta}$ be defined by (17), let $\gamma \in]0, 2\bar{\eta}[$, let $\delta = 2 - \gamma/(2\bar{\eta})$, and let $(\theta_n)_{n \in \mathbb{N}}$ be a sequence in $[0, \delta]$ such that $\sum_{n \in \mathbb{N}} \theta_n(\delta - \theta_n) = +\infty$. Then any sequence $(x_n)_{n \in \mathbb{N}}$ generated by Iteration (5) converges to a point $\tilde{x} \in \mathcal{F}$.*

In addition, let $\nu \in [0, +\infty[$ be such that $g = h + \nu/2\|\cdot\|_Q^2$ where $h \in \Gamma_0(\mathbb{R}^N)$.² If $\nu > 0$ or $\lambda_{\min} > 0$, then

$$\|\tilde{x} - \hat{x}\|_Q \leq \inf_{\gamma \in]0, 2\bar{\eta}[} \frac{\|e(\hat{x})\|_Q}{1 + \gamma\nu - \|\|\text{Id} - \gamma L_Q\|\|}. \quad (24)$$

It is worthy to note that the symmetry and positiveness of P are not mandatory to ensure the convergence of (5).

Usually, the regularization function g depends on some parameter vector ω . The previous analysis assumes that ω is set to the same value for the exact and unmatched preconditioned PGA. A better strategy may be to adjust ω in the unmatched case to reduce the discrepancy between \tilde{x} and \hat{x} .

IV. NUMERICAL EXPERIMENTS

Hereafter, in Model (1), y represents limited tomographic measurements, H is the projector in parallel-beam geometry, and \bar{x} is the ground truth phantom. The projector and the backprojector H^\top are the most computation-intensive operations. To limit the number of PGA iterations and thus of multiplications with these operators, we use an effective preconditioner P for the gradient step. The numerical experiments were carried out in MATLAB with the ASTRA Toolbox [32].

a) Tomographic geometry: Reference image \bar{x} is a slice of size $N = 256 \times 256$ extracted from a computerized tomographic scan of an abdomen with values belonging to $[1000, 2100]$. The image contains simulated small structures of comparatively high intensity (3000). We simulate 60 projections at uniformly spaced angular positions within the interval $[0^\circ, 180^\circ]$. The detector has $363 \simeq \sqrt{2} \times 256$ bins of the same size as the pixels, so that the data is not truncated and $K = 60 \times 363$. Operator H is based on the line-length model, which corresponds to ASTRA GPU implementation [31]. For the measurement y , we simulate a noise-free sinogram and add some low level of noise b drawn from the Gaussian distribution $\mathcal{N}(0, 100 \times \text{Id})$, so that $W = 10^{-2} \times \text{Id}$.

b) Regularization: For such an underdetermined problem, we adopt an Ivanov variational formulation [23] combined with a total variation (TV) bound [10]: an estimate $\rho > 0$ of the value range for the TV bound of our target image is supposed to be known. Thus, in (2), $g = \iota_{B_{1,2}(0, \rho)}(\nabla \cdot)$, where ∇ is the 2D discrete spatial gradient operator and $\iota_{B_{1,2}}$ is the indicator function of the ball of radius ρ associated with the $\ell_{1,2}$ norm $\|\cdot\|_{1,2}$ as defined in [48]. More precisely, $B_{1,2}(0, \rho) = \{u \in \mathbb{R}^{2N} \mid \|u\|_{1,2} \leq \rho\}$. Even without any metric, the proximity operator of g does not have a closed form and is computed by the dual forward-backward (DFB) algorithm [13] with a maximum of 500 iterations and a warm restart procedure.

¹The proof of this proposition is given in Appendix B.

² ν is the strong convexity modulus of g in $(\mathbb{R}^N, \|\cdot\|_Q)$.

c) *Parameter selection:* Except otherwise stated, we choose ρ equal to $\bar{\rho} = \|\nabla\bar{x}\|_{1,2}$. Parameter κ in (4) is set to 10^{-5} , which implies that M is invertible and we set $P = M^{-1}$. In order to compute products between P and any vector, we use the conjugate gradient algorithm with a tolerance of 10^{-4} initialized with the previous iterate (warm restart). Preconditioning the gradient term with P might require extra projections and backprojections per iteration due to the CG inner loop. However, we demonstrate in the sequel that this inversion of the Hessian effectively accelerates the method, especially when combined with a simpler preconditioner for the proximity operator. Following Proposition III.2, the cocoercivity constant of L equals 1. In addition, $\lambda_{\min} = 1$ and the conditions in Proposition III.3(ii) are satisfied; \mathcal{F} is thus a singleton. In PGA, the step size γ is set to 0.9 which, according to Proposition III.4, guarantees the convergence of Algorithm (5). In the DFB iterations, the metric Q weighting the proximity operator is inverted. Setting $Q = P^{-1} = M$ would require using CG again in the DFB step, which would be computationally heavy. Several other choices of proximity metrics are considered instead, as described hereafter.




Fig. 1: Evolution of the NRMSE along iterations for Algorithms (3) and (5), $\rho = \bar{\rho}$, and $Q = Q_1$.




Fig. 2: From left to right: \bar{x} , reconstructed images for $\rho = \bar{\rho}$ with Alg. (3) using Q_1 , and with Alg. (5) using Q_1 .

d) *Matched versus unmatched preconditioning:* We first set metric Q to the diagonal majorant matrix proposed in [24]: $Q_1 = W \text{diag}(S^T 1_N) + \kappa \text{Id}$, where $1_N = [1, \dots, 1]^T \in \mathbb{R}^N$ and $S = \left(|H_{i,k}| \sum_{k'=1}^K |H_{i,k'}| \right)_{1 \leq i \leq N, 1 \leq k \leq K}$. We also compare PGA schemes (3) and (5) to the Chambolle-Pock (CP) primal-dual algorithm [49]. CP is a competing proximal-based method to minimize our cost function involving neither sub-iterations nor preconditioning. We assess the performance of the three methods in Figure 1 showing the overall normalized reconstruction error $\text{NRMSE} = \|x_n - \bar{x}\| / \|\bar{x}\|$ as a function of the execution time in seconds. We notice that Algorithm (5) converges faster than both Algorithm (3) and CP. The convergence curve of CP oscillates in the first iterations. After only 10 seconds, the reconstruction error associated with Algorithm (5) is lower than the error corresponding to the estimate delivered by Algorithm (3) at convergence. Interestingly, the resulting fixed point was observed to be closer to the ground truth \bar{x} than the minimizer of Problem (2). The reconstructed solutions are displayed in Figure 2. Residual deconvolution artifacts (undershooting) and subsampling streaks are present in the estimate produced by Algorithm (3) with Q_1 but not in the one yielded by our unmatched scheme.

e) *Alternative choices for Q :* Four additional approximations to M in \mathcal{S}_N^+ have been tested for metric Q , namely $Q_2 = \text{Id}$, Q_3 the inverse of 2D Laplacian filter,³ $Q_4 = \text{Diag}((M_{i,i})_{1 \leq i \leq N})$ the Jacobi preconditioner,⁴ $Q_5 = \text{argmin}_{Q \in \mathcal{D}_N^+} \|Q^{1/2} - Q^{-1/2} M\|_F^2$ and Q_6 a tridiagonal approximation to M i.e., the symmetric matrix whose elements on the main diagonal are $(M_{i,i})_{1 \leq i \leq N}$ and those on the next upper / lower diagonals are $\frac{1}{2}(M_{i+1,i} + M_{i,i+1})_{1 \leq i \leq N-1}$.

Table I (first row) contains the NRMSE values after 1000 iterations obtained with Algorithm (5) for the different metrics, and $\rho = \bar{\rho}$. First, we see that all choices are competitive compared to the baseline Algorithm (3) and their NRMSE values are

³The cone filter is decomposed into a local Laplacian operator coupled with a non local logarithmic kernel filtering [30].

⁴The entries of H correspond to line integrals between a ray and a pixel, which are positive. Since $\lambda_{\min} > 0$, the entries of M are strictly positive.

close. Metric Q_1 leads to the best quantitative results. Metric Q_3 provides the poorest reconstruction with the highest NRMSE and a patchy look (see Appendix C).

f) *Sensitivity to ρ* : Table I also shows the sensitivity of the reconstruction to the TV bound when performing sets of trials reconstructions for $\rho \in \{\rho_-, \rho_+\}$, with $\rho_+ = 1.1 \times \bar{\rho}$, $\rho_- = 0.9 \times \bar{\rho}$. We observe that ρ_- generally leads to lower reconstruction errors. Q_1 is consistently associated with the lowest NRMSE. The choice of $Q = Q_2$ was noticeably shown to rank second-best for the lowest bound ρ_- .

	Q_1	Q_1	Q_2	Q_3	Q_4	Q_5	Q_6
$\bar{\rho}$	0.0621	0.0543	0.0562	0.0622	0.0558	0.0559	0.0560
ρ_-	0.559	0.0484	0.0489	0.0544	0.0494	0.0494	0.0490
ρ_+	0.0664	0.0614	0.0622	0.0708	0.0621	0.0628	0.0626

TABLE I: NRMSE after 1000 iterations for Algorithms (3) (first column) and (5) (all other columns), for various choices of Q and ρ .

V. CONCLUSION

In this article, we have addressed the problem of accelerating the convergence of PGA through unmatched preconditioning. By leveraging recent tools of fixed point theory, we have characterized the resulting solutions, calculated error bounds, and provided convergence conditions for our modified PGA. Our experimental results on a CT imaging inverse problem prove that satisfactory solutions can be obtained in reduced time, when the metric in the proximity term has lower complexity than the gradient preconditioner.

APPENDIX A

PRECONDITIONING IN CT RECONSTRUCTION

This section briefly reviews the family of preconditioners used in CT reconstruction. We show that existing accelerated reconstruction algorithms can be reformulated as unmatched preconditioned algorithms. Our framework thus clarifies the convergence properties of these accelerated algorithms.

In CT reconstruction, diagonal preconditioners [35], [39] provide limited acceleration. Fourier preconditioners are used instead [18]–[20], [27], [33], [34], [40]. Preconditioning PGA when the precision matrix $W = \text{Id}$ in Problem (2) derives from the Fourier-slice theorem, which states that the continuous version of the normal operator $M = H^\top H$ is a convolutional circulant operator.

This theorem provides two filters associated with the Caldéron operator [45]. The first is the 2D cone filter, which is applied after the backprojector H^\top , yielding a potential choice of preconditioner P in Algorithm (5). However, it requires infinite support and is considered computationally expensive. Many variants have been proposed to simplify [44] or improve it by taking into account the underlying sampling in $H^\top H$. Improved filters were obtained by approximating the SIRT algorithm [43], or in a learning framework [42], [46].

The second choice is the 1D ramp filter, which is simpler to compute and does not require infinite support. However, since it is applied prior to H^\top , it does not fit exactly the form of Algorithm (5): the product $PM = PH^\top H$ is replaced by the more general structure $H^\dagger H$. Both filters are particularly efficient because they provide a close approximate inversion of H .

Regularized Iterative Filtered Backprojection (i-FBP) [37] (and Iterative Feldkamp-Davis-Kress i-FDK methods for cone-beam geometry [36]) speeds up convergence by using this approximate inversion. This method consists in applying $H^\dagger H$ rather than $H^\top H$ in the gradient step, followed by an unweighted proximity step [38], [41]. Regularization behaves as a filtering process whose behavior and strength are independent of the acquisition model H . The output of the regularization becomes easily predictable [47]. These methods relate to our unmatched preconditioned Algorithm (5) with $Q = \text{Id}$, for which we establish conditions of convergence to a fixed point.

APPENDIX B

PROOF OF PROPOSITION III.4

In the renormed space $(\mathbb{R}^N, \|\cdot\|_Q)$, Algorithm (5) with L defined by (8) takes the same form as the algorithm investigated in [9], although for other purposes than preconditioning. The convergence thus follows from [9, Proposition 3.15].

The existence of a unique point $\tilde{x} \in \mathcal{F}$ is a direct consequence of Proposition III.3(ii). Let \hat{x} be a solution to Problem (2) and let $\gamma \in]0, +\infty[$. We have

$$\begin{cases} \tilde{x} = \text{prox}_{\gamma g}^Q(\tilde{x} - \gamma P(M\tilde{x} - H^\top W y)), \\ \hat{x} = \text{prox}_{\gamma g}^Q(\hat{x} - \gamma Q^{-1}(M\hat{x} - H^\top W y)). \end{cases} \quad (25)$$

$$\hat{x} = \text{prox}_{\gamma g}^Q(\hat{x} - \gamma Q^{-1}(M\hat{x} - H^\top W y)). \quad (26)$$

From properties of the proximity operator [13], since $\nu \geq 0$, for every $x \in \mathbb{R}^N$, $\text{prox}_{\gamma g}^Q(x) = \text{prox}_{\frac{\gamma}{1+\gamma\nu}h}^Q\left(\frac{x}{1+\gamma\nu}\right)$. Since $\text{prox}_{\frac{\gamma}{1+\gamma\nu}h}^Q$ is nonexpansive in $(\mathbb{R}^N, \|\cdot\|_Q)$, we deduce from (25) and (26) that

$$(1 - \tau_\gamma)\|\tilde{x} - \hat{x}\|_Q \leq \frac{\gamma}{1 + \gamma\nu} \|(Q^{-1} - P)(M\hat{x} - H^\top W y)\|_Q$$

with $\tau_\gamma = \frac{\| \text{Id} - \gamma L \|_Q}{1 + \gamma \nu} = \frac{\| \text{Id} - \gamma L_Q \|}{1 + \gamma \nu}$. In addition, according to [9, Proposition 3.9], for $\lambda_{\min} > 0$ or $\nu > 0$, if $\gamma \in]0, 2\bar{\eta}[$, then $\tau_\gamma < 1$.

In summary, if $\gamma < 2\bar{\eta}$ and either $\lambda_{\min} > 0$ or $\nu > 0$,

$$\| \tilde{x} - \hat{x} \|_Q \leq \frac{\| e(\hat{x}) \|_Q}{(1 - \tau_\gamma)(1 + \gamma \nu)} = \frac{\| e(\hat{x}) \|_Q}{1 + \gamma \nu - \| \text{Id} - \gamma L_Q \|},$$

which leads to (24).

APPENDIX C

COMPARISON OF THE EFFECT OF METRICS Q_3 AND Q_1 IN ALGORITHM (5)

Figure 3 displays two reconstructions obtained with our unmatched preconditioned scheme (5) using two choices of $Q \in \{Q_1, Q_3\}$ in the proximity operator. Algorithm (5) using metric Q_3 yields patchy artifacts and an unnatural appearance.




Fig. 3: Reconstructed images for $\rho = \bar{\rho}$ with Alg. (5) using Q_3 (left) and Q_1 (right).

REFERENCES

- [1] A. Beck and M. Teboulle, “A fast iterative shrinkage-thresholding algorithm for linear inverse problems”, *SIAM J. Imaging Sci.*, vol. 2, pp. 183–202, 2009.
- [2] S. Becker and M.-J. Fadili, “A quasi-Newton proximal splitting method”, in *Proc. NeurIPS*, Lake Tahoe, Nevada, USA, 2012, pp. 2618–2626.
- [3] H. H. Bauschke, and P. L. Combettes, *Convex Analysis and Monotone Operator Theory in Hilbert Spaces* Corrected 2nd ed. New York, NY, USA: Springer, 2019.
- [4] S. Bonettini, F. Porta, V. Ruggiero, and L. Zanni, “Variable metric techniques for forward–backward methods in imaging”, *Comput. Appl. Math.*, vol. 385, no. 1, pp. 113192, 2021.
- [5] J. Bonnans, J.-C. Gilbert, C. Lemaréchal, and C. Sagastizabal, “A family of variable metric proximal methods”, *Math. Program.*, vol. 68, pp. 15–47, 1995.
- [6] G. H. Chen and R. T. Rockafellar, “Convergence rates in forward-backward splitting”, *SIAM J. Optim.*, vol. 7, pp. 421–444, 1997.
- [7] K. Chen, *Matrix Preconditioning Techniques and Applications*, Cambridge University Press, 2005.
- [8] E. Chouzenoux, J.-C. Pesquet and A. Repetti, “Variable Metric Forward-Backward algorithm for minimizing the sum of a differentiable function and a convex function”, *J. Optim. Theory Appl.*, vol. 162, no. 1, pp. 107–132, 2014.
- [9] E. Chouzenoux, J.-C. Pesquet, C. Riddell, M. Savanier, and Y. Troussel, “Convergence of proximal gradient algorithm in the presence of adjoint mismatch”, *Inverse Problems*, vol. 37, no. 6, pp. 5009, 2021.
- [10] P. L. Combettes and J.-C. Pesquet, “Image restoration subject to a total variation constraint”, *IEEE Trans. Image Process.*, vol. 13, pp. 1213–1222, 2004.
- [11] P. L. Combettes and J.-C. Pesquet, “Proximal thresholding algorithm for minimization over orthonormal bases”, *SIAM J. Optim.*, vol. 18, no. 4, pp. 1351–1376, 2007.
- [12] S. Becker, J. M. Fadili, and P. Ochs, “On quasi-Newton forward-backward splitting: proximal calculus and convergence”, *SIAM J. Optim.*, 29:2445–2481, 2019.
- [13] P. L. Combettes and J.-C. Pesquet, “Proximal Splitting Methods in Signal Processing” in *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*, New York, NY, USA: Springer, 2011, pp. 185–212.
- [14] P. L. Combettes and J.-C. Pesquet, “Fixed point strategies in data science”, *IEEE Trans. Signal Process.*, vol. 69, pp. 3878–3905, 2021.
- [15] P. L. Combettes and B. Vũ, “Variable metric forward–backward splitting with applications to monotone inclusions in duality”, *Optimization*, vol. 63, no. 9, pp. 1289 – 1318, 2012.
- [16] P. L. Combettes and V.R. Wajs, “Signal recovery by proximal forward-backward splitting”, *Multiscale Model. Simul.*, vol. 4, no. 4, pp. 1164–1200, 2005.
- [17] I. Daubechies, M. Defrise, and C. De Mol, “An iterative thresholding algorithm for linear inverse problems with a sparsity constraint”, *Commun. Pure Appl. Math.*, vol. 57, pp. 1413–1457, 2004.
- [18] J. A. Fessler and S. D. Booth, “Conjugate-gradient preconditioning methods for shift-variant PET image reconstruction”, *IEEE Trans. Image Process.*, vol. 8, no. 5, pp. 688–699, 1999.
- [19] L. Fu, B. De Man, K. Zeng, T. M. Benson, Z. Yu, G. Cao, and J.-B. Thibault, “A preliminary investigation of 3D preconditioned conjugate gradient reconstruction for cone-beam CT”, in *Proc. of SPIE Medical Imaging: Physics of Medical Imaging*, 2012, San Diego, California, USA, pp. 1051 – 1059.
- [20] H. Guo and X. Cui, “Block-tridiagonal shift-variant preconditioner for iterative cone beam CT reconstruction”, in *Proc. of Fully 3D*, Xián Shaanxi, China, 2017.
- [21] J. B. Hiriart-Urruty and C. Lemaréchal, *Convex Analysis and Minimization Algorithms*, New York, NY, USA: Springer-Verlag, 1993.
- [22] J. J. Moreau, “Fonctions convexes duales et points proximaux dans un espace hilbertien”, *C. R. Acad. Sci. Paris Sér. A Math.*, vol. 93, pp. 273–299, 1962.
- [23] D. A. Lorenz and N. S. Worliczek, “Necessary conditions for variational regularization schemes”, *arXiv: Functional Analysis*, 2012.

- [24] Y. Marnissi, E. Chouzenoux, A. Benazza-Benyahia, and J.-C. Pesquet, “Majorize-Minimize Adapted Metropolis-Hastings Algorithm”, *IEEE Trans. Signal Process.*, vol. 68, pp. 2356 – 2369, 2020.
- [25] Y. Nesterov. “A method for solving the convex programming problem with convergence rate $o(1/k^2)$ ”, *Proc. of the USSR Academy of Sciences*, vol. 269, pp. 543–547, 1983.
- [26] J. Nocedal and S. Wright, *Numerical Optimization*, New York, NY, USA: Springer, 2006.
- [27] S. Ramani and J. A. Fessler, “A splitting-based iterative algorithm for accelerated statistical X-ray CT reconstruction”, *IEEE Trans. Med. Imaging*, vol. 31, no. 3, pp. 677–688, 2012.
- [28] M. Schmidt, N. Roux, and F. Bach, “Convergence rates of inexact proximal-gradient methods for convex optimization”, in *Proc. of NeurIPS*, 2011, Granada, Spain, pp. 1458–1466.
- [29] M. Tivnan, W. Wang, and J. Stayman, “A preconditioned algorithm for model-based iterative CT reconstruction and material decomposition from spectral CT data”, *arXiv: Medical Physics*, 2020.
- [30] Y. Xia, A. Maier, H. Hofmann, F. Dennerlein, K. Mueller, and J. Hornegger, “Reconstruction from truncated projections in cone-beam CT using an efficient 1D filtering”, in *Proc. of SPIE Medical Imaging: Physics of Medical Imaging*, 2013, Lake Buena Vista, Florida, US, pp. 348 – 354.
- [31] W. Palenstijn, K. Batenburg, and J. Sijbers, “Performance improvements for iterative electron tomography reconstruction using graphics processing units (GPUs)”, *J. Struct. Biol.*, vol. 176, no. 2, pp. 250–253, 2011.
- [32] W. van Aarle, W. Palenstijn, J. Cant, E. Janssens, F. Bleichrodt, A. Dabravolski, J. Beenhouwer, K. Batenburg, and J. Sijbers, “Fast and flexible X-ray tomography using the ASTRA toolbox”, *Opt. Express*, vol. 24, no. 22, pp. 25129–25147, 2016.
- [33] N. H. Clinthorne, T. S. Pan, P. C. Chiao, W. L. Rogers, and J. A. Stamos, “Preconditioning methods for improved convergence rates in iterative reconstructions”, *IEEE Trans. Med. Imaging*, vol. 12, no. 1, pp. 78–83, 1993.
- [34] L. Fu, Z. Yu, J.-B. Thibault, B. De Man, M. McGaffin, and J. Fessler, “Space-Variant Channelized Preconditioner Design for 3D Iterative CT reconstruction”, in *Proc. Fully 3D*, Lake Tahoe, California, USA, 2013.
- [35] J. Gregor and J. A. Fessler, “Comparison of SIRT and SQS for regularized weighted least squares image reconstruction”, *IEEE Trans. Comput. Imaging*, vol. 1, no. 1, pp. 44–55, 2015.
- [36] L. A. Feldkamp, L. C. Davis, and J. W. Kress, “Practical cone-beam algorithm”, *J. Opt. Soc. Am. A: Opt. Image Sci. Vis.*, vol. 1, pp. 612–619, 1984.
- [37] D. Lalush and B. Tsui, “Improving the convergence of iterative filtered backprojection algorithms”, *Med Phys.*, vol. 21, no. 8, pp. 1283–1286, 1994.
- [38] H. Langet, C. Riddell, A. Reshef, Y. Trouset, A. Tenenhaus, E. Lahalle, G. Fleury, and N. Paragios, “Compressed-sensing-based content-driven hierarchical reconstruction: Theory and application to C-arm cone-beam tomography”, *Med Phys.*, vol. 42, no. 9, pp. 5222–5237, 2015.
- [39] M. Jiang and G. Wang, “Convergence of the simultaneous algebraic reconstruction technique (SART)”, *IEEE Trans. Image Process.*, vol. 12, no. 8, pp. 957–961, 2003.
- [40] M. McGaffin, S. Ramani, and J. Fessler, “Reduced memory augmented Lagrangian algorithm for 3D iterative x-ray CT image reconstruction”, in *Proc. SPIE Medical Imaging 2012: Physics of Medical Imaging*, San Diego, California, USA, 2012, pp. 639–644.
- [41] M. Nilchian, C. Vonesch, P. Modregger, M. Stampanoni, and M. Unser, “Iterative FBP for improved reconstruction of X-ray differential phase-contrast tomograms”, in *Proc. ISBI*, San Francisco, CA, USA, 2013, pp. 1260–1263.
- [42] D. Pelt, and K. Batenburg, “Improving filtered backprojection reconstruction by data-Dependent filtering”, *IEEE Trans. Image Process.*, vol. 23, pp. 4750–4762, 2014.
- [43] D. Pelt and V. De Andrade. Improved tomographic reconstruction of large-scale real-world data by filter optimization. *Adv. Struct. Chem. Imaging*, vol. 2, 2016.
- [44] G. Poludniowski, N. M. Allinson, and P. M. Evans, “Proton computed tomography reconstruction using a backprojection-then-filtering approach”, *Phys. Med. Biol.*, vol. 59, no. 24, pp. 7905–7918, 2014.
- [45] E. T. Quinto, “An introduction to X-ray tomography and Radon transforms” in *The Radon Transform, Inverse Problems, and Tomography*, Atlanta, Georgia, USA: American Mathematical Society, 2006, pp. 14–37.
- [46] C. Syben, B. Stimpel, K. Breininger, T. Würfl, R. Fahrig, A. Dörfler, and A. Maier, “Precision learning: Reconstruction filter kernel discretization”, *arXiv: Computer Vision and Pattern Recognition*, 2017.
- [47] C. Riddell, H. Benali, and I. Buvat, “Diffusion regularization for iterative reconstruction in emission tomography”, *IEEE Trans. on Nucl. Sci.*, vol. 52, no. 3, pp. 669–675, 2005.
- [48] L. Condat, “Discrete total variation: New definition and minimization”, *SIAM J. Imaging Sci.*, vol. 10, no. 3, pp. 1258–1290, 2017.
- [49] A. Chambolle, T. Pock, “A first-order primal-dual algorithm for convex problems with applications to imaging”, *J. Math. Imaging Vis.* vol. 40, pp. 120–145, 2011.