

# Using Randomized Nyström Preconditioners to Accelerate Variational Image Reconstruction

Tao Hong , *Member, IEEE*, Zhaoyi Xu , Jason Hu , *Graduate Student Member, IEEE*,  
and Jeffrey A. Fessler , *Fellow, IEEE*

**Abstract**—Model-based iterative reconstruction plays a key role in solving inverse problems. However, the associated minimization problems are generally large-scale, nonsmooth, and sometimes even nonconvex, which present challenges in designing efficient iterative solvers. Preconditioning methods can significantly accelerate the convergence of iterative methods. In some applications, computing preconditioners on-the-fly is beneficial. Moreover, forward models in image reconstruction are typically represented as operators, and the corresponding explicit matrices are often unavailable, which brings additional challenges in designing preconditioners. Therefore, for practical use, computing and applying preconditioners should be computationally inexpensive. This paper adapts the randomized Nyström approximation to compute effective preconditioners that accelerate image reconstruction without requiring an explicit matrix for the forward model. We leverage modern GPU computational platforms to compute the preconditioner on-the-fly. Moreover, we propose efficient approaches for applying the preconditioners to problems with classical nonsmooth regularizers, i.e., wavelet, total variation, and Hessian Schatten-norm. Our numerical results on image deblurring, super-resolution with impulsive noise, and 2D computed tomography reconstruction illustrate the efficiency and effectiveness of the proposed preconditioner.

**Index Terms**—Image deblur, super-resolution, CT reconstruction, Nyström preconditioner, wavelet, total variation, Hessian Schatten-norm.

## I. INTRODUCTION

THE task of image reconstruction is to recover a clean image from degraded measurements. Model-based iterative reconstruction recovers a clean image  $\mathbf{x}$  by solving the following type of minimization problem:

$$\min_{\mathbf{x} \in \mathbb{R}^N} f(\mathbf{x}) + \lambda g(\mathbf{x}), \quad (1)$$

Received 12 November 2024; revised 18 March 2025 and 4 September 2025; accepted 8 October 2025. Date of publication 20 October 2025; date of current version 20 November 2025. The work of Tao Hong was supported by NIH under Grant R01NS112233. The work of Jeffrey A. Fessler was supported by NIH under Grant R01EB035618 and Grant R21EB034344. The associate editor coordinating the review of this article and approving it for publication was Prof. Yang Chen. (*Corresponding author: Tao Hong.*)

Tao Hong was with the Department of Radiology, University of Michigan, Ann Arbor, MI 48109 USA. He is now with the Oden Institute for Computational Engineering and Sciences, University of Texas at Austin, Austin, TX 78712 USA (e-mail: tao.hong@austin.utexas.edu).

Zhaoyi Xu is with the Department of Mechanical Engineering, University of Michigan, Ann Arbor, MI 48109 USA (e-mail: zhaoyix@umich.edu).

Jason Hu and Jeffrey A. Fessler are with the Department of Electrical and Computer Engineering, University of Michigan, Ann Arbor, MI 48109 USA (e-mail: jashu@umich.edu; fessler@umich.edu).

This article has supplementary downloadable material available at <https://doi.org/10.1109/TCI.2025.3622903>, provided by the authors.

Digital Object Identifier 10.1109/TCI.2025.3622903

where  $f(\mathbf{x})$  denotes a data-fidelity specifying the discrepancy between the model's predictions and the degraded measurements,  $g(\mathbf{x})$  denotes a regularizer imposing prior information on  $\mathbf{x}$ , and  $\lambda > 0$  is a trade-off parameter to balance  $f(\mathbf{x})$  and  $g(\mathbf{x})$ . Both  $f(\mathbf{x})$  and  $g(\mathbf{x})$  can be nonsmooth and nonconvex. This paper focuses on linear inverse problems of the form

$$\min_{\mathbf{x} \in \mathbb{R}^N} \frac{1}{p} \underbrace{\|\mathbf{A}\mathbf{x} - \mathbf{y}\|_p^p}_{f(\mathbf{x})} + \lambda g(\mathbf{x}), \quad (2)$$

where  $\mathbf{A} \in \mathbb{R}^{M \times N}$  with  $M \leq N$  denotes the forward operator, which maps the unknown image  $\mathbf{x}$  to the measurements  $\mathbf{y}$ . In this paper, we assume  $p \in (0, 2]$ .

In practice, the matrix  $\mathbf{A}$  is under-determined or ill-conditioned and the measurements are noisy, so one uses a regularizer  $g(\mathbf{x})$  to stabilize the solution. Typical regularizer choices that are effective include wavelet sparsity [1], total variation (TV) [2], dictionary learning [3], and low-rank models [4]. When using such hand-crafted regularizers, we refer to (1) as variational image reconstruction. Over the past decades, deep learning (DL) [5] has received significant attention in the field of image reconstruction. Instead of using hand-crafted priors, DL-based methods learn implicit priors from massive training data. This includes end-to-end learning [6] and physics-informed deep unrolling [7], [8]. Alternatively, one can use plug-and-play (PnP) [9], [10], [11] or regularization by denoising (RED) [12], [13] where one trains a deep denoiser [14] first and then uses it as an implicit prior for general inverse problems. Recently, building priors with generative models, such as diffusion models, has also attracted much attention. We refer the reader to [15] for a survey of those directions.

While deep learning methods often outperform variational methods in many applications, variational image reconstruction [16], [17], [18], [19] offers significant advantages in terms of interpretability, theoretical guarantees, and robustness to noise or distribution shifts. Variational methods do not require large training datasets and provide predictable, stable behavior. Moreover, they allow for easier customization of priors based on domain knowledge, making them particularly suitable for applications where interpretability and reliability are critical, such as medical imaging. Recent work has also shown that novel neural network structures can be derived from variational methods, yielding state-of-the-art performance [20], [21], [22], [23], [24]. Therefore, this paper considers hand-crafted regularizers for (2). Specifically, we consider  $g(\mathbf{x}) = \Psi(\mathbf{L}\mathbf{x})$  where  $\mathbf{L} \in \mathbb{R}^{L \times N}$

refers to the regularization operator and  $\Psi(\cdot)$  denotes an energy function. A typical choice of  $\Psi(\cdot)$  is a norm or quasi-norm.

#### A. $\ell_p - \ell_q$ Image Reconstruction

For  $p \in (0, 2]$ , we consider  $\Psi(\cdot) = \frac{1}{q} \|\cdot\|_q^q$  with  $q \in (0, 2]$  and (2) becomes

$$\min_{\mathbf{x} \in \mathbb{R}^N} \underbrace{\frac{1}{p} \|\mathbf{Ax} - \mathbf{y}\|_p^p}_{f(\mathbf{x})} + \lambda \underbrace{\frac{1}{q} \|\mathbf{Lx}\|_q^q}_{g(\mathbf{x})}, \quad (3)$$

which is known as the  $\ell_p - \ell_q$  problem [16]. For  $p, q \leq 1$ , both  $f(\mathbf{x})$  and  $g(\mathbf{x})$  are nonsmooth. To address the nonsmooth challenge, Clason et al. considered (3) in its dual formulation [25], where the objective function becomes differentiable, at the cost of additional box and linear equality constraints. For general  $p, q \in (0, 2)$ , Huang et al. [26] first smoothed  $f$  and  $g$  and then applied the majorization-minimization for the smoothed problem, leading to an inner least-squares problem. Then, [26] used the Krylov subspace to represent the image, thereby reducing the dimension of the associated least-squares problem to that of the subspace. Consequently, it can be solved efficiently using a direct method. Following the same strategy, Buccini et al. [27] applied modulus-based iterative methods [28] to the smoothed problem to address the nonnegativity constraint. Alternatively, Chan et al. [29] solved (3) with a half-quadratic algorithm. Lanza et al. [30] proposed a generalized Krylov subspace (GKS) method to reduce computation of the method in [29] by assuming that the signal can be represented in a subspace. Gazzola et al. [31] presented two flexible Krylov subspace methods for sparse reconstruction.

The solvers discussed above for (3) can be considered as iteratively reweighted methods (IRMs) [16] because we eventually need to solve a linear equation (19) with updating of the associated weightings at each iteration. For completeness, Section III-A describes one type IRM for (3). Typically the appealing conjugate gradient (CG) method is chosen to address the associated linear equation. However, the convergence of CG can be extremely slow for some problems, especially for  $p, q < 1$ , as the problem becomes nonconvex. Preconditioning methods [32], [33] are widely used to accelerate the convergence of CG. Because the weightings  $\mathbf{W}_f^k$  and  $\mathbf{W}_g^f$  (defined in (17)) change at each iteration, an effective preconditioner should change at each iteration, so computing such a preconditioner must be computationally inexpensive.

#### B. $\ell_2$ With Mixed Norm Image Reconstruction

When  $p = 2$ , we consider  $\Psi(\cdot)$  to be an  $\ell_{1,\phi}$  mixed norm (i.e.,  $\|\mathcal{V}\|_{1,\phi} = \sum_{l=1}^G \|\mathbf{v}_l\|_\phi$  with  $\phi \geq 1$  and  $\mathbf{v}_l$  denoting the  $l$ th group of  $\mathcal{V} = \{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_G\}$ ), for which we can develop more efficient approaches than IRM to solve (3) and a broader range of choices for  $\mathbf{L}$ . Moreover, we can also address convex constraints  $\mathcal{C}$  easily. Dividing the components of  $\mathcal{V}$  into  $G$  groups, the mixed norm  $\ell_{1,\phi}$  is defined by applying an  $\ell_\phi$  norm to each group and then applying an  $\ell_1$  norm to the entire set of groups. Specifically,

we address

$$\min_{\mathbf{x} \in \mathcal{C}} \underbrace{\frac{1}{2} \|\mathbf{Ax} - \mathbf{y}\|_2^2}_{f(\mathbf{x})} + \lambda \underbrace{\|\mathbf{Lx}\|_{1,\phi}}_{g(\mathbf{x})}. \quad (4)$$

When  $\mathbf{L}$  is a first-order differential operator and  $\phi = 1$  (respectively,  $\phi = 2$ ), then (4) represents the anisotropic (respectively, isotropic) TV-based reconstruction. When  $\mathbf{L}$  denotes the second-order finite-difference operator, then (4) becomes the Hessian Schatten (HS) norm based reconstruction [34]. Appendix A presents an explicit form for  $\|\mathbf{Lx}\|_{1,\phi}$  for various  $\mathbf{L}$  to clarify the definition of mixed norm.

Since (4) is a standard composite minimization problem, an appealing algorithm for solving it is the accelerated proximal gradient (APG) method [35]. Alternatively, one could use a primal-dual method [36], [37] or ADMM [38] etc. To accelerate the convergence of solving (4), one can adopt a quasi-Newton proximal method that uses second-order information [39], [40], [41]. However, the efficiency of using a quasi-Newton proximal method depends highly on the accuracy of approximating the Hessian matrix of  $f(\mathbf{x})$  and the efficiency of solving the associated weighted proximal mapping (WPM, defined in (5)). For  $\phi = 2$ ,  $g(\mathbf{x})$  reduces to the group sparsity case, for which one can adapt a flexible Krylov subspace method [42] to solve (4).

1) *Previous Work on Preconditioning Methods in Image Reconstruction:* Preconditioning methods are used widely to accelerate the convergence of iterative methods [43] in scientific computing and have been extended to accelerate the convergence of image reconstruction. Clinthorne et al. [44] applied a spatially invariant preconditioner to accelerate the convergence of reconstructing single-photon emission computed tomography images with quadratic  $g(\mathbf{x})$ . Fessler et al. [45] proposed a shift-variant preconditioner for accelerating positron emission tomography (PET) reconstruction with non-quadratic smooth  $g(\mathbf{x})$ . Lin et al. [46] derived cosine transform and incomplete factorization based preconditioners for image super-resolution task with smooth  $g(\mathbf{x})$ . Pelletier et al. [47] developed a block circulant with circulant block preconditioner for edge-preserving image super-resolution. The authors in [48], [49] considered diagonal preconditioners for accelerating PET and magnetic resonance imaging reconstruction. To reduce the computation of computing WPM, [50] discussed an unmatched preconditioner method for image reconstruction by computing WPM approximately. To address the nonsmoothness of  $g(\mathbf{x})$ , Koolstra et al. [51] adapted the preconditioner to ADMM to accelerate the convergence of solving the associated least-squares problem. By leveraging deep learning, [52] proposed an approach to learn preconditioners to accelerate the convergence of image reconstruction.

#### C. Our Contribution

The preconditioners discussed in Section I-B1 were either used for smooth  $g(\mathbf{x})$  [44], [45], [46], [48], [49] or applied to a nonsmooth  $g(\mathbf{x})$  after smoothing [47]. Moreover, the previous methods require domain knowledge because they were developed for specific applications, making them difficult to

generalize to other applications. Although the preconditioners proposed in [51], [52] can be used to address nonsmooth  $g(\mathbf{x})$ , the convergence rate of ADMM is only  $\mathcal{O}(1/k)$ . To address these challenges, we adapt the randomized Nyström approximation (RNA) [53], [54] to design an effective preconditioner, called randomized Nyström preconditioner (RNP), to accelerate variational image reconstruction. Our main contributions are summarized as follows.

- To the best of our knowledge, this is the first work to adapt RNP to accelerate image reconstruction with nonsmooth regularizers. Computing RNP requires only the evaluation of  $\mathbf{A}\mathbf{x}$  and  $\mathbf{A}^T\mathbf{x}$  ( $T$  denotes the transpose operator), and does not require to access the explicit matrix  $\mathbf{A}$  or knowledge of any specific structure of  $\mathbf{A}$ . Therefore, RNP can generalize to a wide range of applications.
- By leveraging modern high-performance GPU computational platforms, we compute RNP on-the-fly, allowing it to accelerate IRM for (3) in a computationally efficient way. Our numerical experiments on image deblurring and super-resolution with impulsive noise showed that, by using RNP, we can reduce wall time by more than 90%. To the best of our knowledge, this is the first work to consider the use of hardware platform in conjunction with computing preconditioners on-the-fly in image reconstruction.
- We show how RNP can be adapted to directly address nonsmooth  $g(\mathbf{x})$ . Specifically, we propose efficient approaches for integrating RNP with APG to solve (4). Our experiments on computed tomography (CT) reconstruction with wavelet, TV, and HS norm regularizers demonstrate that RNP can significantly accelerate the convergence of APG.

#### D. Preliminaries and Roadmap

This subsection introduces some definitions and a theorem that will be used frequently in the subsequent discussion, and then presents a roadmap of this paper.

**Definition 1 (Weighted proximal mapping (WPM)):** Given a proper closed convex function  $h(\mathbf{x})$ , a symmetric positive-definite matrix  $\mathbf{W} \succ 0 \in \mathbb{R}^{N \times N}$ , and  $\lambda > 0$ , the WPM associated with  $h$  is defined as

$$\text{prox}_{\lambda h}^{\mathbf{W}}(\mathbf{x}) = \arg \min_{\mathbf{u} \in \mathbb{R}^N} \left( \lambda h(\mathbf{u}) + \frac{1}{2} \|\mathbf{u} - \mathbf{x}\|_{\mathbf{W}}^2 \right), \quad (5)$$

where  $\|\mathbf{x}\|_{\mathbf{W}} \triangleq \sqrt{\mathbf{x}^T \mathbf{W} \mathbf{x}}$  denotes the  $\mathbf{W}$ -norm.

For  $\mathbf{W} = \mathbf{I}$ , (5) simplifies to the proximal mapping:

$$\text{prox}_{\lambda h}(\mathbf{x}) = \arg \min_{\mathbf{u} \in \mathbb{R}^N} \lambda h(\mathbf{u}) + \frac{1}{2} \|\mathbf{u} - \mathbf{x}\|_2^2. \quad (6)$$

If  $h(\mathbf{x})$  is a characteristic function

$$\delta_{\mathcal{C}}(\mathbf{x}) = \begin{cases} 0 & \mathbf{x} \in \mathcal{C} \\ \infty & \mathbf{x} \notin \mathcal{C} \end{cases},$$

then the proximal mapping becomes projection onto the set  $\mathcal{C}$ .

In general, computing a WPM (5) is computationally expensive. However, when  $\mathbf{W}$  is represented by a diagonal matrix plus a rank- $r$  correction, then we can use Theorem 1 to compute WPM relatively cheaply.

**Theorem 1 ([40], Theorem 3.4):** Let  $\mathbf{W} = \mathbf{D} \pm \mathbf{U}\mathbf{U}^T$ ,  $\mathbf{W} \succ 0 \in \mathbb{R}^{N \times N}$ , and  $\mathbf{U} \in \mathbb{R}^{N \times r}$ . Then, it holds that

$$\text{prox}_{\lambda h}^{\mathbf{W}}(\mathbf{x}) = \text{prox}_{\lambda h}^{\mathbf{D}}(\mathbf{x} \mp \mathbf{D}^{-1}\mathbf{U}\boldsymbol{\gamma}^*), \quad (7)$$

where  $\boldsymbol{\gamma}^* \in \mathbb{R}^r$  is the unique solution of

$$\mathbf{U}^T (\mathbf{x} - \text{prox}_{\lambda h}^{\mathbf{D}}(\mathbf{x} \mp \mathbf{D}^{-1}\mathbf{U}\boldsymbol{\gamma})) + \boldsymbol{\gamma} = \mathbf{0}. \quad (8)$$

Semi-smooth Newton methods can solve (8) efficiently [40], [55], [56].

The rest of this paper is organized as follows. Section II describes how RNA can be a preconditioner for linear equations. Section III shows how to adapt RNP to accelerate IRM and APG for solving (3) and (4), respectively. Section IV presents numerical experiments on image deblurring and super-resolution with impulsive noise and computed tomography reconstruction to demonstrate the effectiveness and efficiency of using RNP for acceleration. Section V presents the conclusion and future work. Code to reproduce the results in the paper is available at <https://github.com/hongtao-argmin/RNP-AccImageRecon>.

## II. RANDOMIZED NYSTRÖM APPROXIMATION

The Nyström method provides a way to build a low-rank approximation for a symmetric positive semidefinite (SPSD) matrix  $\Phi \succeq 0 \in \mathbb{R}^{N \times N}$ . Let  $\Omega \in \mathbb{R}^{N \times K}$  be an i.i.d. Gaussian random test matrix [53], [54] with sketch size  $K \geq 1$ . The Nyström method uses the following formula as a low-rank approximation of  $\Phi$ , called the Nyström approximation

$$\Phi \langle \Omega \rangle = (\Phi \Omega)(\Omega^T \Phi \Omega)^\dagger (\Phi \Omega)^T, \quad (9)$$

where  $^\dagger$  denotes the pseudo-inverse. Clearly  $\Phi \langle \Omega \rangle$  is also a SPSP matrix and its rank is at most  $K$  [57, Lemma A.1]. Because a direct implementation of (9) is numerically unstable, we adopt the robust scheme proposed in [53], which is summarized in Algorithm 1. This algorithm computes  $\mathbf{U}$  and  $\hat{\mathbf{S}}$  such that  $\Phi \langle \Omega \rangle = \mathbf{U} \hat{\mathbf{S}} \mathbf{U}^T$ , although this product is not formed explicitly in practice.

Many scientific computing applications [58] require solving the following linear equation:

$$(\Phi + \mu \mathbf{I})\mathbf{x} = \mathbf{b}, \quad \mu > 0. \quad (10)$$

Using RNA to formulate a preconditioner to address (10), the associated formulations of RNP and its inverse [54] are as follows:

$$\begin{aligned} \mathbf{P} &= \frac{1}{\hat{s}_K + \mu} \mathbf{U}(\hat{\mathbf{S}} + \mu \mathbf{I})\mathbf{U}^T + (\mathbf{I} - \mathbf{U}\mathbf{U}^T), \\ \mathbf{P}^{-1} &= (\hat{s}_K + \mu)\mathbf{U}(\hat{\mathbf{S}} + \mu \mathbf{I})^{-1}\mathbf{U}^T + (\mathbf{I} - \mathbf{U}\mathbf{U}^T), \end{aligned} \quad (11)$$

where  $\hat{s}_K$  is the  $K$ th eigenvalue of  $\Phi \langle \Omega \rangle$ .  $\mathbf{U}$  and  $\hat{\mathbf{S}}$  are defined in Algorithm 1. Let  $d_{\text{eff}}(\mu) = \text{tr}(\Phi(\Phi + \mu \mathbf{I})^{-1})$  be the effective dimension. Then, we have the following theorem to describe the condition number of  $(\Phi + \mu \mathbf{I})$  after using RNP, which reveals the effectiveness of RNP.

**Theorem 2 ([54], Theorem 5.1):** Suppose we build the Nyström preconditioner through (11) with sketch size  $K = 2\lceil 1.5d_{\text{eff}}(\mu) + 1 \rceil$ . Then using  $\mathbf{P}$  as the preconditioner for



**Algorithm 1:** Nyström Approximation.

---

**Initialization:** Sketch size  $K$ , machine accuracy  $\epsilon$ , and symmetric positive semidefinite operator  $\Phi \in \mathbb{R}^{N \times N}$ .  
**Output:**  $\mathbf{U} \in \mathbb{R}^{N \times K}$ ,  $\hat{\mathbf{S}} \in \mathbb{R}^{K \times K}$ .  
1: Generate a random matrix  $\Omega \in \mathbb{R}^{N \times K}$   
2:  $\mathbf{Y} = \Phi\Omega \in \mathbb{R}^{N \times K}$   
3:  $v = \epsilon \|\Omega\|_F$   
4:  $\mathbf{Y}_v = \mathbf{Y} + v\Omega$   
5:  $\mathbf{C} = \text{chol}(\Omega^T \mathbf{Y}_v) \in \mathbb{R}^{K \times K}$  % Cholesky decomp.  
6:  $\mathbf{B} = \mathbf{Y}_v \mathbf{C}^{-T} \in \mathbb{R}^{N \times K}$   
7:  $[\mathbf{U}, \mathbf{S}, \sim] = \text{SVD}(\mathbf{B})$   
8:  $\hat{\mathbf{S}} = \max\{\mathbf{0}, \mathbf{S}^2 - v\mathbf{I}\}$

---

$(\Phi + \mu\mathbf{I})$  results in the expectation of condition number bound

$$\mathbb{E}[\kappa(\mathbf{P}^{-1/2}(\Phi + \mu\mathbf{I})\mathbf{P}^{-1/2})] < 28,$$

where  $\kappa(\cdot)$  denotes the condition number function.

Theorem 2 shows that, by choosing an appropriate sketch size  $K$ , RNP can significantly reduce the condition number of  $(\Phi + \mu\mathbf{I})$ . Thus, one can significantly accelerate the convergence of iterative methods for solving (10) with RNP. However, in practice, computing  $d_{\text{eff}}$  is computationally expensive. Moreover, the value of  $K = 2\lceil 1.5d_{\text{eff}}(\mu) + 1 \rceil$  can be very large, which will significantly increase the computation of computing and applying RNP. In our experimental settings, we found that using a moderate value for  $K$  is sufficient. We refer the reader to [53], [54], where the authors discussed more theoretical properties of using RNP for considering (10). The main purpose of this paper is to study how to adapt RNP to accelerate variational image reconstruction, focusing on evaluating its effectiveness and efficiency. We also refer the reader to [54], [57], [59] and the references therein that discuss using RNP for other applications. Although Theorem 2 presents the split preconditioner  $\mathbf{P}^{-1/2}$ , we use  $\mathbf{P}^{-1}$  in practice for simplicity.

### III. ON ADAPTING NYSTRÖM PRECONDITIONERS IN IMAGE RECONSTRUCTION

For impulsive noise, a natural data-fidelity term is the  $\ell_p$  quasi-norm with  $p \leq 1$ , allowing image reconstruction by solving the  $\ell_p - \ell_q$  problem. When the noise is assumed to be i.i.d. Gaussian, we use the  $\ell_2$  norm as the metric in the data-fidelity term, resulting in the  $\ell_2 - \ell_{1,\phi}$  problem. This section presents two baseline algorithms for solving  $\ell_p - \ell_q$  and  $\ell_2 - \ell_{1,\phi}$ , and then study how they can be accelerated using RNP.

Following we first describe IRM in Section III-A and then show how RNP can be adapted to accelerate IRM for solving (3) efficiently. Because the weightings in IRM change at each iteration, computing RNP must be computationally inexpensive. We leverage the modern GPU computational platform and the PyTorch library to present an on-the-fly implementation for computing RNP at the end of this section. Additionally, we show how RNP can be incorporated with APG [60] to efficiently handle (4). Furthermore, we propose efficient approaches to solve the associated WPM to reduce the computational cost

involved in applying RNP. We consider separately the  $\ell_p - \ell_q$  and  $\ell_2 - \ell_{1,\phi}$  cases.

#### A. $\ell_p - \ell_q$ Image Reconstruction Acceleration

Following [29, Lemma 1], we represent the  $p$ th power of  $r \in \mathbb{R} \setminus \{0\}$  with  $0 < p < 2$  as

$$|r|^p = \min_{\beta > 0} \left\{ \beta r^2 + \frac{1}{b_p \beta^{a_p}} \right\}, \quad (12)$$

where  $a_p = \frac{p}{2-p}$  and  $b_p = \frac{2^{2-p}}{(2-p) \cdot p^{2-p}}$  are two positive scalars. The minimizer of (12) is

$$\beta^* = \frac{p}{2} |r|^{p-2}. \quad (13)$$

For the singular case,  $r = 0$ , we simply set  $|r|_\epsilon = \sqrt{r^2 + \epsilon}$  with  $\epsilon > 0$ , following [29].

Representing (3) in a component-wise form, we get

$$\min_{\mathbf{x}} \frac{1}{p} \sum_{m=1}^M |(\mathbf{A}\mathbf{x})_m - y_m|^p + \frac{\lambda}{q} \sum_{l=1}^L \sum_{\zeta=1}^{\pi} |(\mathbf{L}\mathbf{x})_l^\zeta|^q, \quad (14)$$

where  $(\mathbf{A}\mathbf{x})_m$  and  $y_m$  represent the  $m$ th element of  $\mathbf{A}\mathbf{x}$ ,  $\mathbf{y}$ , respectively.  $(\mathbf{L}\mathbf{x})_l^\zeta$  denotes the  $\zeta$ th element in the  $l$ th group of  $\mathbf{L}\mathbf{x}$ . Here, we assume each group has the same number of elements and consider cases where each group represents a scalar or vector. Here,  $\pi$  denotes the number of elements in each group. Using (12), we rewrite (14) as

$$\begin{aligned} \min_{\substack{\mathbf{x}, \\ \mathbf{v} > \mathbf{0}, \\ \mathbf{z} > \mathbf{0}}} F(\mathbf{x}, \mathbf{v}, \mathbf{z}) &\equiv \frac{1}{p} \sum_{m=1}^M (v_m |(\mathbf{A}\mathbf{x})_m - y_m|^2 \\ &+ \frac{1}{b_p v_m^{a_p}}) + \frac{\lambda}{q} \sum_{l=1}^L \sum_{\zeta=1}^{\pi} \left( z_l^\zeta |(\mathbf{L}\mathbf{x})_l^\zeta|^2 + \frac{1}{b_q (z_l^\zeta)^{a_q}} \right), \end{aligned} \quad (15)$$

where  $v_m$  and  $z_l^\zeta$  denote the  $m$  and  $(l-1)\pi + \zeta$ th element of  $\mathbf{v}$  and  $\mathbf{z}$ , respectively. By using an alternating minimization algorithm, IRM [16] solves (15) by successively updating  $\mathbf{v}$ ,  $\mathbf{z}$ , and  $\mathbf{x}$ . Algorithm 2 summarizes the IRM steps. Using (13), we derive closed-form solutions for  $\mathbf{v}^k$  and  $\mathbf{z}^k$  at steps 2 and 3 of Algorithm 2, i.e.,

$$\mathbf{v}^k = \frac{p}{2} |(\mathbf{A}\mathbf{x}^k) - \mathbf{y}|^{p-2} \text{ and } \mathbf{z}^k = \frac{q}{2} |\mathbf{L}\mathbf{x}^k|^{q-2}, \quad (16)$$

where the powers  $(p-2)$  and  $(q-2)$  are element-wise.

By defining

$$\mathbf{W}_f^k = \text{diag} \left( \frac{2}{p} \mathbf{v}^k \right) \text{ and } \mathbf{W}_g^k = \text{diag} \left( \frac{2}{q} \mathbf{z}^k \right), \quad (17)$$

the minimization problem at step 4 of Algorithm 2 becomes

$$\mathbf{x}^k = \arg \min_{\mathbf{x} \in \mathbb{R}^N} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_{\mathbf{W}_f^k}^2 + \lambda \|\mathbf{L}\mathbf{x}\|_{\mathbf{W}_g^k}^2. \quad (18)$$

The first-order optimality condition of (18) is given by

$$\Phi_k \mathbf{x} = \mathbf{A}^T \mathbf{W}_f^k \mathbf{y}, \quad \Phi_k = \mathbf{A}^T \mathbf{W}_f^k \mathbf{A} + \lambda \mathbf{L}^T \mathbf{W}_g^k \mathbf{L}. \quad (19)$$



---

**Algorithm 2:** Alternating Minimization Method for Solving (15).

---

**Initialization:** Initialization  $\mathbf{x}^0$ , tolerance  $\varepsilon$ , and maximum number of iterations  $\text{Max\_Iter}$ .

**Output:**  $\mathbf{x}^*$ .

```

1: for  $k = 1, 2, \dots, \text{Max\_Iter}$  do
2:    $\mathbf{v}^k = \arg \min_{\mathbf{v} > 0} F(\mathbf{x}^{k-1}, \mathbf{v}, \mathbf{z}^{k-1})$ 
3:    $\mathbf{z}^k = \arg \min_{\mathbf{z} > 0} F(\mathbf{x}^{k-1}, \mathbf{v}^k, \mathbf{z})$ 
4:    $\mathbf{x}^k = \arg \min_{\mathbf{x}} F(\mathbf{x}, \mathbf{v}^k, \mathbf{z}^k)$ 
5:   if  $\frac{\|\mathbf{x}^k - \mathbf{x}^{k-1}\|}{\|\mathbf{x}^k\|} < \varepsilon$  then
6:     break
7:   end if
8: end for
9:  $\mathbf{x}^* = \mathbf{x}^k$ 

```

---

Since the Hessian  $\Phi_k$  is a SPSD matrix, we use CG to solve (19). However, solving (19) is time-consuming when  $\Phi_k$  is ill-conditioned. We use (11) to construct a preconditioner for  $\Phi_k$  and apply the preconditioned CG (PCG) method to solve (19). Because  $\mathbf{W}_f^k$  and  $\mathbf{W}_g^k$  change at each iteration<sup>1</sup>, computing  $\mathbf{P}$  in (11) must be computationally inexpensive. Section III-C discusses an on-the-fly implementation of RNP to address this computational issue. If  $p, q = 2$ , then (3) is a quadratic minimization problem, allowing us to directly apply PCG with RNP. When  $p = 2$  and  $q \neq 1$ , IRM is still applicable, but APG is more efficient, as discussed in the following subsection.

### B. $\ell_2 - \ell_{1,\phi}$ Image Reconstruction Acceleration

APG, which only uses first-order information and achieves the optimal worst-case convergence rate of  $\mathcal{O}(1/k^2)$  [61], is an appealing algorithm for addressing (4). To further accelerate the convergence of APG, we adapt RNP to APG, resulting in the weighted accelerated proximal gradient (WAPG) algorithm that is summarized in Algorithm 3. For any  $\mathbf{x}_1, \mathbf{x}_2 \in \mathbb{R}^N$ , the Lipschitz constant  $L_f^{\mathbf{P}}$  of  $f$  with respect to the  $\mathbf{P}$ -norm satisfies the following inequality

$$\|\mathbf{P}^{-1}(\nabla f(\mathbf{x}_1) - \nabla f(\mathbf{x}_2))\|_{\mathbf{P}} \leq L_f^{\mathbf{P}} \|\mathbf{x}_1 - \mathbf{x}_2\|_{\mathbf{P}}. \quad (20)$$

If  $f, g$  are convex, the worst-case convergence rate of the cost function sequence generated by Algorithm 3 is  $\frac{2L_f^{\mathbf{P}} \|\mathbf{x}^1 - \mathbf{x}^*\|_{\mathbf{P}}^2}{(k+1)^2}$  [62, Chapter 10.7.5] where  $\mathbf{x}^1$  and  $\mathbf{x}^*$  are the initial and optimal values, respectively. Clearly, for  $\mathbf{P} = \mathbf{I}$ , WAPG simplifies to APG. By choosing a suitable  $\mathbf{P}$ , one can expect a lower bound than  $\mathbf{P} = \mathbf{I}$  yielding faster convergence. Indeed, our experimental results in CT reconstruction illustrate the faster convergence of WAPG with RNP, i.e., experiments in Section IV-B and IV-C.

For noninvertible  $\mathbf{L}$ , the associated WPM in Algorithm 3 generally does not have a closed-form solution, even with  $\mathbf{P} = \mathbf{I}$ , so WPM requires iterative methods. Following [56], we solve WPM using its dual problem to address the nonsmoothness of  $g(\mathbf{x})$ . Moreover, compared with computing WPM with  $\mathbf{P} = \mathbf{I}$ ,

<sup>1</sup>Note that  $\mathbf{W}_f^k$  and  $\mathbf{W}_g^k$  are changed at each outer iteration but remain unchanged within CG iterations when solving (19).

we show that solving WPM with RNP negligibly increases computation. Whereas [56] only discussed constrained TV regularization, we generalize their derivation to an abstract  $\mathbf{L}$  with mixed norm in this work.

Computing WPM at step 2 of Algorithm 3 at  $k$ th iteration is equivalent to solving

$$\min_{\mathbf{x} \in \mathcal{C}} \frac{1}{2} \|\mathbf{x} - \mathbf{s}^k\|_{\mathbf{P}}^2 + \bar{\lambda} \|\mathbf{L}\mathbf{x}\|_{1,\phi}, \quad (21)$$

where  $\mathbf{s}^k = \mathbf{u}^{k-1} - \alpha \mathbf{P}^{-1} \nabla f(\mathbf{u}^{k-1})$  and  $\bar{\lambda} = \alpha \lambda$ . By using the dual representation

$$\|\mathbf{L}\mathbf{x}\|_{1,\phi} = \max_{\mathbf{Q} \in B_{\infty,\psi}} \langle \mathbf{Q}, \mathbf{L}\mathbf{x} \rangle, \quad (22)$$

where  $B_{\infty,\psi}$  denotes the  $\ell_{\infty,\psi}$  mixed unit-norm ball, i.e.,

$$B_{\infty,\psi} = \{\mathbf{Q} = \{\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_G\} : \|\mathbf{q}_l\|_{\psi} \leq 1, \forall l\}, \quad (23)$$

and  $\psi \geq 1$  satisfies  $\frac{1}{\phi} + \frac{1}{\psi} = 1$ , we rewrite (21) as

$$\min_{\mathbf{x} \in \mathcal{C}} \max_{\mathbf{Q} \in B_{\infty,\psi}} \frac{1}{2} \|\mathbf{x} - \mathbf{s}^k\|_{\mathbf{P}}^2 + \bar{\lambda} \langle \mathbf{L}^T \mathbf{Q}, \mathbf{x} \rangle, \quad (24)$$

where  $\mathbf{L}^T$  denotes the adjoint of  $\mathbf{L}$ . Denote by  $\mathcal{V} = \mathbf{L}\mathbf{x}$  with  $\mathcal{V} = \{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_G\}$ . The inner product in (22) is defined as  $\sum_{l=1}^G \langle \mathbf{q}_l, \mathbf{v}_l \rangle$ . Reorganizing (24), we get

$$\begin{aligned} \max_{\mathbf{Q} \in B_{\infty,\psi}} \min_{\mathbf{x} \in \mathcal{C}} & \|\mathbf{x} - (\mathbf{s}^k - \bar{\lambda} \mathbf{P}^{-1} \mathbf{L}^T \mathbf{Q})\|_{\mathbf{P}}^2 \\ & - \|\mathbf{s}^k - \bar{\lambda} \mathbf{P}^{-1} \mathbf{L}^T \mathbf{Q}\|_{\mathbf{P}}^2. \end{aligned} \quad (25)$$

Because  $\mathbf{x}$  appears only in the first term of (25), the minimizer of  $\mathbf{x}$  is exactly the WPM associated with  $\delta_{\mathcal{C}}$ :

$$\bar{\mathbf{x}}^* = \text{prox}_{\delta_{\mathcal{C}}}^{\mathbf{P}}(\mathbf{s}^k - \bar{\lambda} \mathbf{P}^{-1} \mathbf{L}^T \mathbf{Q}). \quad (26)$$

Substituting (26) into (25), we get

$$\begin{aligned} \mathcal{Q}^* &= \arg \min_{\mathbf{Q} \in B_{\infty,\psi}} \|\mathbf{s}^k - \bar{\lambda} \mathbf{P}^{-1} \mathbf{L}^T \mathbf{Q}\|_{\mathbf{P}}^2 \\ & - \left\| \text{prox}_{\delta_{\mathcal{C}}}^{\mathbf{P}}(\mathbf{s}^k - \bar{\lambda} \mathbf{P}^{-1} \mathbf{L}^T \mathbf{Q}) \right. \\ & \left. - (\mathbf{s}^k - \bar{\lambda} \mathbf{P}^{-1} \mathbf{L}^T \mathbf{Q}) \right\|_{\mathbf{P}}^2. \end{aligned} \quad (27)$$

Following [56], the gradient and the corresponding Lipschitz constant of the objective function in (27) are

$$-2\bar{\lambda} \mathbf{L} \text{prox}_{\delta_{\mathcal{C}}}^{\mathbf{P}}(\mathbf{s}^k - \bar{\lambda} \mathbf{P}^{-1} \mathbf{L}^T \mathbf{Q}), \quad (28)$$

and  $2\sigma_{\min}^{\mathbf{P}} \bar{\lambda}^2 \|\mathbf{L}\|^2$ , respectively.  $\sigma_{\min}^{\mathbf{P}}$  is the smallest eigenvalue of  $\mathbf{P}$ . Since the objective function in (27) is differentiable, we apply APG to solve (27). After getting  $\mathcal{Q}^*$ , we compute

$$\mathbf{x}^{k+1} = \text{prox}_{\delta_{\mathcal{C}}}^{\mathbf{P}}(\mathbf{s}^k - \bar{\lambda} \mathbf{P}^{-1} \mathbf{L}^T \mathbf{Q}^*). \quad (29)$$

Both (28) and (29) require computing  $\text{prox}_{\delta_{\mathcal{C}}}^{\mathbf{P}}(\cdot)$  and  $\mathbf{P}^{-1}$ , which can be computationally expensive for a general  $\mathbf{P}$ . A natural way to choose  $\mathbf{P}$  is to approximate the Hessian  $\mathbf{A}^T \mathbf{A}$ . By using RNP,  $\mathbf{P}$  (see (11)) can be rewritten as  $\mathbf{I} + \bar{\mathbf{U}} \bar{\mathbf{U}}^T$  with  $\bar{\mathbf{U}} = \mathbf{U} \sqrt{\frac{1}{s_K + \mu} (\hat{\mathbf{S}} + \mu \mathbf{I})} - \mathbf{I}$ , where the square root here is applied elementwise. Clearly,  $\mathbf{P}$  has the same structure as  $\mathbf{W}$  proposed in Theorem 1, so getting  $\text{prox}_{\delta_{\mathcal{C}}}^{\mathbf{P}}(\cdot)$  reduces to addressing a nonlinear (8). Thus, the total computation for  $\text{prox}_{\delta_{\mathcal{C}}}^{\mathbf{P}}(\cdot)$

---

**Algorithm 3:** Weighted Accelerated Proximal Gradient (WAPG) Method.

---

**Initialization:** Initialization  $\mathbf{u}^1$ ,  $\mathbf{x}^1$ , stepsize  $\alpha$ ,  $t^0 = 1$ ,  $\mathbf{P}$ , and maximum number of iterations  $\text{Max\_Iter}$ .

**Output:**  $\mathbf{x}^{\text{Max\_Iter}}$ .

```

1: for  $k = 1, 2, \dots, \text{Max\_Iter}$  do
2:    $\mathbf{x}^{k+1} = \text{prox}_{(\alpha\lambda)g+\delta_C}^{\mathbf{P}}(\mathbf{u}^k - \alpha\mathbf{P}^{-1}\nabla f(\mathbf{u}^k))$ 
3:    $t^k = \frac{1+\sqrt{1+4(t^{k-1})^2}}{2}$ 
4:    $\mathbf{u}^{k+1} = \mathbf{x}^{k+1} + \frac{t^{k-1}-1}{t^k}(\mathbf{x}^{k+1} - \mathbf{x}^k)$ 
5: end for

```

---

can be dramatically reduced, as demonstrated by our numerical experiments in Section IV-B. Note that  $\hat{s}_K$  is the  $K$ th eigenvalue of  $\Phi \langle \Omega \rangle$ , so the last column in  $\bar{\mathbf{U}}$  will be a zero vector. In practice, we found that using  $\sqrt{\hat{s}_K}$  instead of  $\hat{s}_K$  resulted in faster convergence. So, in this paper, we use  $\hat{s}_K \leftarrow \sqrt{\hat{s}_K}$  in (11). Since  $\mathbf{U}$  is an orthogonal matrix, we have  $\sigma_{\min}^{\mathbf{P}} = 1$ , so the Lipschitz constant of the cost function in (27) becomes  $2\bar{\lambda}^2 \|\mathbf{L}\|^2$ .

When  $\mathbf{L}$  is invertible and  $\mathcal{C} = \mathbb{R}^N$ , we can rewrite (4) as

$$\min_{\bar{\mathbf{x}} \in \mathbb{R}^N} \frac{1}{2} \|\mathbf{A}\mathbf{L}^{-1}\bar{\mathbf{x}} - \mathbf{y}\|_2^2 + \lambda \|\bar{\mathbf{x}}\|_{1,\phi}, \quad (30)$$

where  $\mathbf{x} = \mathbf{L}^{-1}\bar{\mathbf{x}}$ . If  $\phi = 1$  and  $\mathbf{L}$  refers to a wavelet transform, (30) represents wavelet-regularized reconstruction. In such cases, we do not need to compute WPM iteratively, saving further computation. By setting  $\mathbf{P}$  to approximate  $(\mathbf{L}^{-1})^T \mathbf{A}^T \mathbf{A} \mathbf{L}^{-1}$  with RNP, we only need to solve a small nonlinear equation to compute the associated WPM. We did not observe any degradation in our subsequent CT reconstruction when using (30), even though it does not account for the constraint  $\mathcal{C}$ .

### C. Implementation Details of On-the-Fly RNP Method

Sections III-A and III-B show how to adapt RNP to accelerate IRM and APG for (3) and (4), respectively. However, if computing RNP is expensive, then using RNP may lose its practical utility. Indeed, in (19),  $\mathbf{W}_f^k$  and  $\mathbf{W}_g^k$  change at each iteration, so the preconditioner should be recomputed. Therefore, the entire algorithm may be slower in terms of wall time than one without using RNP if obtaining RNP is computationally expensive.

Algorithm 1 shows that, to compute RNP, one must evaluate  $\Phi\Omega$ . Each column of  $\Omega$  represents an image. In image reconstruction  $\mathbf{A}$  and  $\mathbf{L}$  are represented as operators, so the associated  $\Phi$  is also an operator, such as  $\Phi = \mathbf{A}^T \mathbf{A}$  in (4) and  $\Phi = (\mathbf{L}^{-1})^T \mathbf{A}^T \mathbf{A} \mathbf{L}^{-1}$  in (30). For some  $\mathbf{A}$  and  $\mathbf{L}$ , evaluating  $\mathbf{A}\mathbf{x}$  and  $\mathbf{L}\mathbf{x}$  many times could be computationally expensive. Moreover, in practice, one computes  $\Phi\Omega$  by applying  $\Phi$  to each column of  $\Omega$  sequentially, which could be slow when  $K$  is moderately large. So step 2 in Algorithm 1 will dominate the overall computation, making it crucial to significantly reduce this computational cost.

On classical CPU computational platforms, a parallel scheme may be used to evaluate  $\Phi\Omega$ . However, a central server still must

collect the results to formulate  $\mathbf{Y}$ , making communication time a potential bottleneck. Modern machine learning libraries (e.g., PyTorch [63] and TensorFlow [64]) support a batch mode, allowing efficient parallel computation of  $\Phi\Omega$ . Additionally, these libraries are highly optimized for modern GPU computational platforms, and the batch mode is easy to implement with basic Python knowledge. These features enable us to compute RNP on-the-fly. Table II presents the wall time of computing RNP for CT reconstruction with different acquisitions on CPU and GPU computational platforms, clearly showing the advantage of using GPU. However, RNP requires storing an additional  $K$  images. As a result, it may encounter memory limitations when handling high-dimensional images on a GPU platform. Therefore, developing a memory-efficient approach to apply RNP to high-dimensional images would be an interesting future direction.

*Remark 1:* In this section, we introduce IRM (respectively, WAPG) for solving (3) (respectively, (4)) and discuss how RNP can be incorporated for acceleration. However, in some special cases, e.g.,  $p = 2$  and  $\mathcal{C} = \mathbb{R}^N$ , (3) is identical to (4). This naturally raises the question of how to choose between IRM and WAPG. From our experiments, when both IRM and WAPG are applicable, we find that IRM converges faster than WAPG in terms of iterations. However, since IRM requires more computation per iteration than WAPG, we consistently observe that WAPG converges faster than IRM in terms of wall time. We validate this observation in Section IV-A4 and Section IV-B1. In this paper, we only consider linear inverse problems. However, we believe that RNP also has the potential to accelerate numerical solvers for nonlinear inverse problems. One example is the full-waveform inversion [65] in geophysics, where truncated Newton methods are among the appealing algorithms. However, one needs to solve a linear equation at each iteration, and the related coefficient matrix (i.e., the Hessian matrix) changes at every iteration. So using effective on-the-fly preconditioners could significantly accelerate the numerical efficiency. We leave the investigation of this direction for future work, as it is beyond the scope of this paper.

## IV. NUMERICAL EXPERIMENTS

This section studies the performance of using RNP on three different image reconstruction tasks: image deblurring and super-resolution with impulsive noise, and CT. Following [16], we use a TV regularizer for image deblurring and super-resolution tasks. For CT reconstruction, we consider wavelet, TV, and HS norm regularizers. Table I presents the associated minimization problems and regularizers for different image tasks in this part. The trade-off parameter  $\lambda$  is selected through a grid search to maximize the PSNR. All experiments were implemented with PyTorch and run on an NVIDIA GeForce RTX 3090.

### A. Image Deblurring and Super-Resolution

For the image deblurring task, we first degraded the test images by convolving them with two different  $9 \times 9$  kernels: a uniform blur and a Gaussian blur with a standard deviation

TABLE I  
THE ASSOCIATED MINIMIZATION PROBLEMS AND THE USED PRIORS FOR DIFFERENT IMAGING TASKS.  $\nabla$  AND  $\nabla^2$  REFER TO THE FIRST- AND SECOND-DIFFERENCE OPERATORS, RESPECTIVELY.

Image Task	Problem	$q/\phi$	$\mathbf{L}$
Deblurring/Super-Resolution	(3) with $p \in (0, 1]$	$q = 1$	$\nabla, \text{TV}$
Computed Tomography	(30)	$\phi = 1$	wavelet
	(4)	$\phi = 1$	$\nabla, \text{TV}$
	(4)	$\phi = 1, 2, \infty$	$\nabla^2, \text{HS}$

TABLE II  
COMPARISON OF WALL TIME FOR OBTAINING RNP IN COMPUTED TOMOGRAPHY RECONSTRUCTION ON CPU AND GPU.

	Parallel-beam			Fan-beam		
	Wavelet	TV	HS	Wavelet	TV	HS
$K$	20	20	100	20	20	100
GPU	0.11s	0.10s	0.40s	0.11s	0.09s	0.46s
CPU	0.48s	0.40s	4.20s	0.48s	0.52s	4.25s

of 1.6. In the image super-resolution task, we generated the low-resolution image by convolving the high-resolution with a  $7 \times 7$  Gaussian blur with a standard deviation of 1.6, followed by downsampling with a factor of 2. We added 5% salt-and-pepper noise to the degraded images.<sup>2</sup> The supplementary material also reports the performance of using RNP with different noise levels. We set the sketch size  $K = 100$  and tolerance  $\epsilon = 10^{-4}$  in Algorithm 1 and PCG, respectively. Similar to [12], the restoration of an RGB image was conducted by converting it to the YCbCr color-space first. Next, we applied the reconstruction algorithm to the luminance channel only and then transformed the result back to the RGB domain. For the deblurring task, we simply applied a median filter to the chroma channels for denoising. For the super-resolution task, we first used a median filter to denoise the chroma channels and then applied bicubic interpolation to upscale them to the desired resolution. Moreover, the peak signal-to-noise ratio (PSNR) values of RGB images were evaluated on the luminance channel only. We used  $\ell_p - \ell_1$  to denote the method with different  $p$ . For the one using RNP, we added “P-”, i.e.,  $\mathbf{P} - \ell_p - \ell_1$ .

1) *Image Deblurring*: Fig. 1 presents the PSNR values versus iteration and wall time for different  $p$ . The first (respectively, second) row of Fig. 1 was tested on the starfish (respectively, leaves) image with uniform (respectively, Gaussian) blur. Fig. 1(a) and (c) show a small  $p$  yielded a higher PSNR than a large  $p$ . It is not surprising because a small  $p$  is more robust to outliers [66]. However, (3) becomes nonconvex for  $p < 1$ , so solving (3) is more challenging. Indeed, Fig. 1(b) and (d) illustrate that a small  $p$  required much more wall time than a large one. Fig. 1 illustrates that using RNP significantly accelerated the convergence speed in terms of wall time. The time for computing RNP was included in the whole wall time. From Fig. 1(b) and (d), we even saw  $\mathbf{P} - \ell_{0.5} - \ell_1$  converged faster than  $\ell_1 - \ell_1$  in terms of the wall time illustrating the effectiveness of using RNP. Fig. 2 describes the number of iterations within CG at each iteration of IRM for different  $p$ , with the uniform and Gaussian kernels for the starfish and leaves images, respectively. Using RNP reduced the number

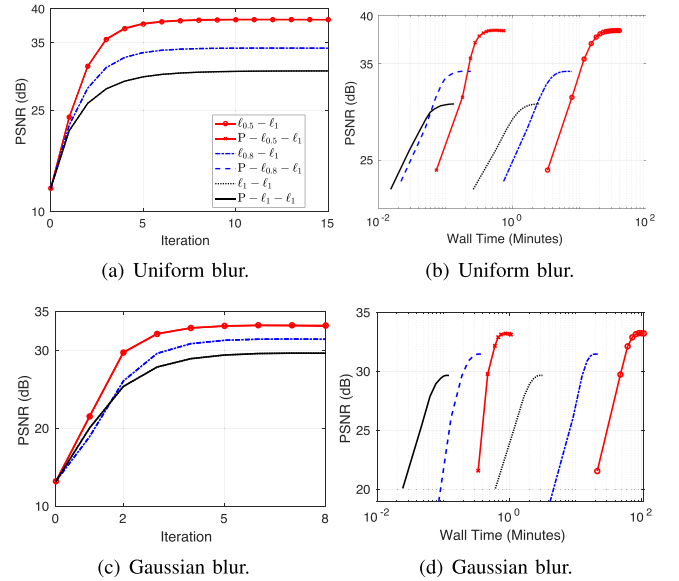


Fig. 1. PSNR values versus iteration and wall time for different  $p$  in the image deblurring task. First (respectively, second) row is tested on the starfish (respectively, leaves) image.




Fig. 2. Number of iterations within CG for different  $p$  in the image deblurring task. The legend is identical to Fig. 1.

of iterations required by CG more than 90%, while achieving the same solution accuracy. The first and second rows of Fig. 3 present the noisy, reconstructed, and ground truth images, where the degradation is reduced significantly.

<sup>2</sup>We first randomly chose 5% of the pixels and set them to 1, and then randomly selected another 5% of the pixels and set them to 0.





Fig. 3. The reconstructed images for different  $p$  on the image deblurring (first and second rows) and super-resolution (third row) tasks. The PSNR value is labeled at the right bottom corner of each image. The first and fifth columns are the noisy measurement and the ground truth, respectively. The first (respectively, second) row is the reconstructed starfish (respectively, leaves) image at 10th (respectively, 6th) iteration for the image deblurring task with uniform (respectively, Gaussian) blur. The third row is the reconstructed bike image at 7th iteration for the image super-resolution task with a downsampling factor 2 and 5% salt-and-pepper noise.

TABLE III

COMPARISON OF PSNR AND ST FOR DIFFERENT  $p$  IN THE IMAGE DEBLURRING TASK WITH UNIFORM AND GAUSSIAN BLURS. WE RAN IRM 20 ITERATIONS ON EACH IMAGE AND PRESENTED THE HIGHEST PSNR VALUES WITHIN THESE 20 ITERATIONS ALONG WITH THE ASSOCIATED ST VALUES. THE HIGHEST PSNR AND ST VALUES FOR EACH IMAGE ARE MARKED IN BOLD.

Image	Butterfly		Boats		C. Man		House		Parrot		Barbara		Starfish		Peppers		Leaves	
Deblurring: Uniform blur, 5% salt-and-pepper noise																		
$p = 1$	PSNR	ST	PSNR	ST	PSNR	ST	PSNR	ST	PSNR	ST	PSNR	ST	PSNR	ST	PSNR	ST	PSNR	ST
	31.6	0.94	32.9	0.94	31.8	0.95	35.9	0.93	32.3	0.95	27.9	0.95	30.8	0.94	32.2	0.94	30.7	0.94
$p = 0.8$	35.3	0.96	35.9	0.95	35.1	0.97	38.7	0.96	35.7	0.97	31.1	0.97	34.2	0.96	35.2	0.96	34.7	0.96
$p = 0.5$	<b>39.9</b>	<b>0.98</b>	<b>39.9</b>	<b>0.97</b>	<b>38.8</b>	<b>0.98</b>	<b>42.5</b>	<b>0.98</b>	<b>39.2</b>	<b>0.98</b>	<b>34.9</b>	<b>0.98</b>	<b>38.4</b>	<b>0.98</b>	<b>38.8</b>	<b>0.98</b>	<b>39.7</b>	<b>0.98</b>
Deblurring: Gaussian blur, 5% salt-and-pepper noise																		
$p = 1$	29.5	0.95	30.9	0.96	29.4	0.97	34.1	0.96	31.9	0.96	28.1	0.97	29.9	0.96	31.3	0.95	29.7	0.96
$p = 0.8$	31.5	0.98	32.9	0.98	31.5	0.98	35.6	0.98	33.5	0.98	30.3	0.98	31.9	0.98	32.5	0.98	31.5	0.98
$p = 0.5$	<b>33.5</b>	<b>0.99</b>	<b>35.2</b>	<b>0.99</b>	<b>33.5</b>	<b>0.99</b>	<b>37.4</b>	<b>0.99</b>	<b>34.9</b>	<b>0.99</b>	<b>31.9</b>	<b>0.99</b>	<b>33.7</b>	<b>0.99</b>	<b>34.1</b>	<b>0.99</b>	<b>33.3</b>	<b>0.99</b>

TABLE IV

COMPARISON OF PSNR AND ST FOR DIFFERENT  $p$  IN THE IMAGE SUPER-RESOLUTION TASK. WE RAN IRM 20 ITERATIONS ON EACH IMAGE AND PRESENTED THE HIGHEST PSNR VALUES WITHIN THESE 20 ITERATIONS ALONG WITH THE ASSOCIATED ST VALUES. THE HIGHEST PSNR AND ST VALUES FOR EACH IMAGE ARE MARKED IN BOLD.

Super-resolution: scaling = 2, 5% salt-and-pepper noise																		
Image	Butterfly		Flower		Girl		Parth.		Parrot		Raccoon		Bike		Hat		Plants	
	PSNR	ST	PSNR	ST	PSNR	ST	PSNR	ST	PSNR	ST	PSNR	ST	PSNR	ST	PSNR	ST	PSNR	ST
$p = 1$	25.6	0.74	27.4	0.74	31.6	0.78	24.5	0.43	27.5	0.76	27.3	0.68	23.2	0.76	29.4	0.76	31.9	0.76
$p = 0.8$	27.1	0.94	28.7	0.93	31.5	0.91	25.4	0.91	28.6	0.94	28.4	0.90	24.3	0.93	30.3	0.94	33.3	0.92
$p = 0.5$	<b>28.1</b>	<b>0.96</b>	<b>29.6</b>	<b>0.97</b>	<b>33.1</b>	<b>0.96</b>	<b>25.9</b>	<b>0.96</b>	<b>29.7</b>	<b>0.97</b>	<b>29.4</b>	<b>0.97</b>	<b>25.3</b>	<b>0.97</b>	<b>31.4</b>	<b>0.96</b>	<b>34.4</b>	<b>0.96</b>

To quantitatively measure the performance of using RNP, we define a “saved time” (ST) criterion:

$$ST = \frac{\text{Time}_{w/o} - \text{Time}_w}{\text{Time}_{w/o}}, \quad (31)$$

where  $\text{Time}_w$  and  $\text{Time}_{w/o}$  denote the wall time with and without using RNP, respectively. If ST is close to 1, then using RNP significantly reduced the wall time. To study the performance of RNP further, we tested on additional 8 images. We ran IRM 20 iterations for different  $p$ . Table III presents the highest PSNR and the associated ST value for each image within 20

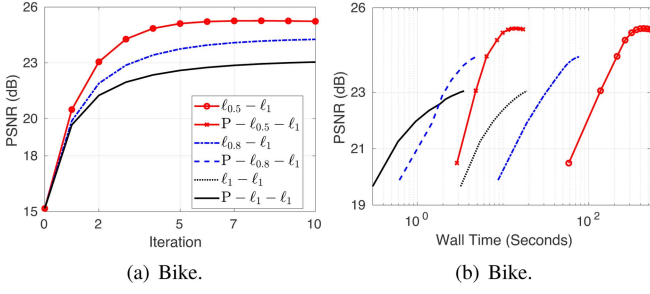


Fig. 4. PSNR values versus iteration and wall time for different  $p$  in the image super-resolution task.




Fig. 5. Number of iterations within CG for different  $p$  in the super-resolution task. The downsampling factor is 2. The legend is identical to Fig. 4.

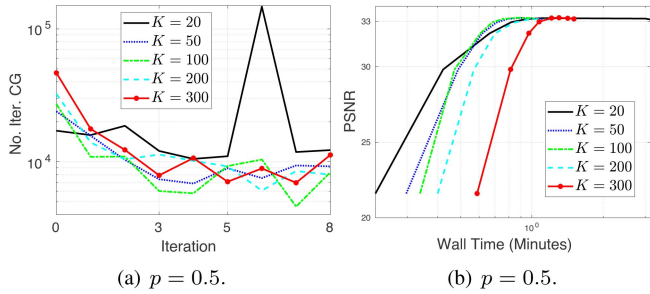


Fig. 6. Comparison of different  $K$  in the image deblurring task with Gaussian kernel on the leaves image for  $p = 0.5$ . (a): the number of iterations within CG in each iteration; (b): PSNR values versus wall time.

iterations. Using RNP saved more than 95% wall time than without RNP, illustrating the effectiveness and efficiency of using RNP. Moreover, Table III also shows that a smaller  $p$  yielded a higher PSNR.

2) *Image Super-Resolution*: Fig. 4 shows the PSNR values for different  $p$  in the image super-resolution of the bike image. The results are consistent with the trends in Section IV-A1, where a small  $p$  yielded a higher PSNR than a larger one and using RNP significantly accelerated the convergence speed in terms of wall time. Moreover, Fig. 5 presents the number of CG iterations for different  $p$ ; the preconditioned one required fewer iterations illustrating the effectiveness of using RNP. The third row of Fig. 3 presents the reconstructed images, where we observed that a small  $p$  yielded a higher quality image. We

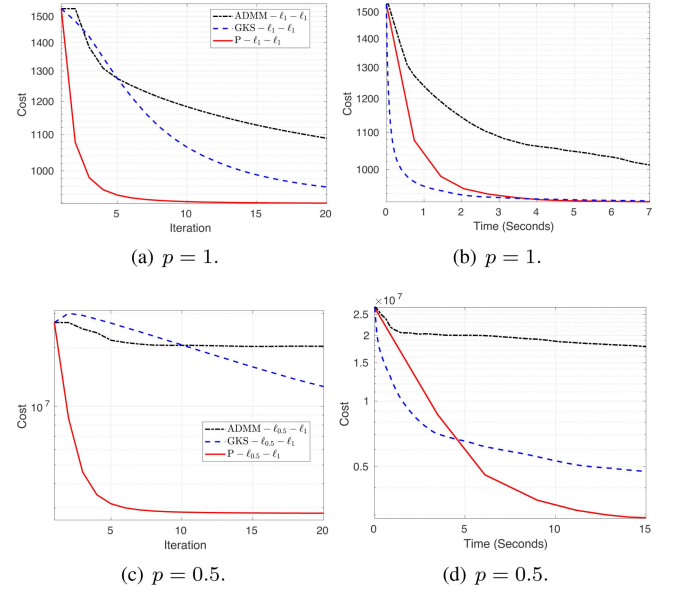


Fig. 7. Comparison with GKS and ADMM on the image super-resolution task on the bike image with  $p = 0.5, 1$ .




Fig. 8.  $L_f^P \|x^1 - x^*\|_P^2$  values versus different  $K$  for fan-beam and parallel-beam acquisitions.




Fig. 9. Ground truth CT images.






Fig. 10. The parallel-beam CT reconstructed images with wavelet, TV, and  $HS_1$  based regularization at iterations 10, 20, 40, and 60. Columns 1, 3, 5, and 7 (respectively, 2, 4, 6, and 8) show the reconstructions without (respectively, with) RNP. The associated PSNR values are listed at the right corner of each image.

tested on additional 8 images to study the performance of using RNP further. Table IV shows using RNP saved almost 70% (respectively, 95%) time for  $p = 1$  (respectively,  $p = 0.5$ ) in line with the observation in Section IV-A1.

3) *The Choice of Sketch Size*: From Theorem 2, one option for choosing the sketch size is to set  $K = 2\lceil 1.5d_{\text{eff}}(\mu) + 1 \rceil$ . However,  $d_{\text{eff}}(\mu)$  is difficult to compute in practice. Moreover,  $d_{\text{eff}}(\mu)$  can be extremely large, making computing RNP time- and memory-intensive, even with an on-the-fly implementation. Indeed, for an extremely large  $K$ , steps 5-7 in Algorithm 1 will dominate the computation instead of step 2. Furthermore, applying RNP can be also expensive for an extremely large  $K$ . Murray et al. [67] presented an adaptive strategy to update  $K$  at each iteration, avoiding the need for  $d_{\text{eff}}(\mu)$ . However, the adaptive strategy requires calling Algorithm 1 multiple times and executing many iterations of the power method to estimate the error between  $\Phi$  and  $\hat{U}\hat{U}^T$ . Therefore, it is unsuitable for our experimental settings, where RNP must be computed on-the-fly. We found that using a fixed, moderate  $K$  works well in practice.

Fig. 6(a) presents the number of iterations within CG in the image deblurring task with Gaussian blur and  $p = 0.5$  on the leaves image for different  $K$ . For  $K > 20$ , the number of iterations within CG did not change significantly. Notice that at the 7th iteration,  $K = 20$  requires significantly more CG iterations than other iterations, where the preconditioner is not very effective. However, the required number of CG iterations is still much lower than in the case without a preconditioner. In general, a larger  $K$  would yield a more effective preconditioner, but it also increases the computational cost to compute and apply. Fig. 6(b) shows the reconstructed PSNR values versus wall time, where clearly  $K = 300$  required more wall time than the others. Moreover,  $K = 50$  and  $K = 100$  performed better than the others. Since  $K = 100$  performed slightly better than  $K = 50$ , we simply set  $K = 100$  in this subsection.

4) *Comparison of GKS and ADMM*: We compared the performance of IRM with RNP (w-20), GKS [30], and ADMM for  $p = 0.5, 1$  on the super-resolution task with the bike image.

Fig. 7 presents the cost values in terms of number of iterations and wall time. Fig. 7(a) and (c) show that w-20 is the fastest algorithm in terms of iterations. Compared to w-20, GKS requires significantly less computation per iteration. Indeed, Fig. 7(b) indicates GKS converged fastest than the other methods in terms of wall time for  $p = 1$ . However, for  $p = 0.5$ , we observed that w-20 converged faster than GKS in terms of wall time illustrating the effectiveness of using RNP for acceleration when the problem becomes more challenging.

### B. Computed Tomography Reconstruction

We studied the performance of using RNP for CT reconstruction with three different regularizers: wavelet, TV with  $\phi = 1$ , and  $HS_\phi$  norm. Moreover, we also compared our method with the ADMM algorithm [38]. For  $HS_\phi$ , we mainly studied  $\phi = 1$ . We used the “daub4” wavelet with 4 levels in the “pywt” [68] toolbox. The Operator Discretization Library (ODL) [69] was used for the CT forward model. We investigated both parallel-beam and fan-beam acquisition geometries with 100 projections. For both geometries, the reconstruction space was uniformly discretized on  $[-20, 20] \times [-20, 20]$  cm<sup>2</sup> with a resolution of  $512 \times 512$ . In the parallel-beam case, projection data were collected over uniformly spaced views in  $[0, 180^\circ]$ , using a linear detector with 1024 bins spanning  $[-40, 40]$  cm. In the fan-beam case, projection data were acquired over uniformly spaced views in  $[0, 360^\circ]$ , with a linear detector of 1024 bins covering  $[-60, 60]$  cm. We scaled two slices from subject “067” in the AAPM CT Grand Challenge data [70] to  $[0, 1]$  as the test images; see Fig. 9. The projections were generated by applying the forward model and then adding Gaussian noise with zero mean and variance  $10^{-4}$ . Due to the page limit, the supplement shows the reconstruction results for Fig. 9(b) and for the fan-beam geometry. Again we use “w/o” to denote results without using RNP, while “w- $K$ ” denotes RNP with sketch size  $K$ .

Table II summarizes the wall time for computing RNP for different acquisitions with various regularizers on both CPU and GPU computational platforms, clearly highlighting the






Fig. 11. Comparison of using RNP and ADMM for wavelet based CT reconstruction. w/o denotes the one without RNP.




Fig. 12. Comparison of using RNP and ADMM for TV based CT reconstruction. w/o denotes the one without RNP.

advantage of using the batch mode and the GPU computational platforms. In general, a larger  $K$  would yield a faster convergence in terms of iterations than a small  $K$ . Fig. 8 presents the values of  $L_f^P \|x^1 - x^*\|_P^2$  along the change of  $K$  with fan-beam and parallel-beam acquisitions. Fig. 8 shows that by using RNP, we can significantly reduce  $L_f^P \|x^1 - x^*\|_P^2$  that one can expect an acceleration. In this case, we observe that RNP becomes less effective at reducing  $L_f^P \|x^1 - x^*\|_P^2$  when  $K > 100$ . Moreover, a larger  $K$  may slightly increase  $L_f^P \|x^1 - x^*\|_P^2$  but still remains effective at reducing it. In practice, large  $K$  would also increase the computation of applying RNP. Here we set  $K = 20$  for wavelet and TV-based reconstruction, and  $K = 100$




Fig. 13. Comparison of using RNP and ADMM for  $HS_1$  based CT reconstruction. w/o denotes reconstruction without RNP.




Fig. 14. Comparison of using RNP, IRM, and GKS for wavelet based CT reconstruction.




Fig. 15. Comparison of using RNP and ADMM for TV based reconstruction on the walnut data. w/o denotes reconstruction without RNP.

for  $HS_1$ -based reconstruction, because these values worked well in practice.

Fig. 11 presents the comparison of using RNP for wavelet based CT reconstruction. Fig. 11(a) and (c) shows that w-20 converged faster than w/o in terms of iterations, illustrating the effectiveness of using RNP. Moreover, Fig. 11(b) and (d) shows that w-20 converged faster than w/o in terms of wall time, demonstrating the wall time for solving the related WPM with Theorem 1 is negligible. Fig. 12 presents the results of TV-based




Fig. 16. CT reconstructions of the walnut data using TV based regularization are shown at iterations 10, 20, 40, and 60. Columns 1, 3, 5, and 7 (respectively, 2, 4, 6, and 8) show the reconstructions without (respectively, with) RNP.

reconstruction that shows similar trends as in wavelet-based regularization. Obviously, ADMM is the slowest algorithm in these comparisons. Fig. 10 illustrates the reconstructed images at iterations 10, 20, 40, and 60, for both w- $K$  and w/o.

Since TV based reconstruction will introduce blocking artifacts [34], we included the study of HS norm based reconstruction. Fig. 13 shows the results of HS<sub>1</sub> based reconstruction, where we observed that using RNP accelerated the convergence compared to not using RNP. Moreover, we observed that HS<sub>1</sub> yielded a higher PSNR than wavelet and TV regularizers. The supplementary material reports the performance of HS<sub>2</sub> and HS<sub>∞</sub> based reconstruction that trended similar to HS<sub>1</sub>.

1) *Comparison of IRM and GKS*: We compared the performance of WAPG using RNP with IRM and GKS for wavelet-based parallel-beam CT reconstruction. Fig. 14 presents the cost values versus the number of iterations and wall time for different methods. For IRM and GKS, the  $\ell_1$  norm is smoothed, but the cost values presented here are still computed with the  $\ell_1$  norm. We observed that IRM is the fastest algorithm in terms of the number of iterations. However, since IRM requires higher computation per iteration, we found that w-20 is the fastest algorithm in terms of wall time. Compared to IRM, GKS significantly reduces the computational cost per iteration, which is also validated in our experiment. However, we observed that w-20 is still faster than GKS in terms of wall time in our setting.

### C. Computed Tomography Reconstruction With Real Data

We used the publicly available walnut CT dataset provided by the Finnish Inverse Problems Society, which was acquired with a fan-beam geometry [71]. The dataset includes a given matrix  $\mathbf{A}$  and the corresponding 120 projection measurements  $\mathbf{y}$ . We evaluated the effectiveness of RNP for image reconstruction by solving (4) with TV-based regularization, with the sketch size  $K = 100$ . In this experiment, we subsampled the data and used only 40 projections. Fig. 15 shows the cost values versus the number of iterations and wall time, demonstrating our approach converged faster than others. Fig. 16 presents the reconstructed images. Using RNP recovered sharper details than APM without RNP at the same number of iterations, illustrating the effectiveness of RNP. Because walnut data did not include the ground truth image, PSNR values are not reported.

## V. CONCLUSION AND FUTURE WORK

In this paper, we introduced RNP to accelerate variational image reconstruction. Additionally, we showed how to efficiently adapt RNP to solve reconstruction problems that involve

nonsmooth priors. By leveraging modern GPU computational platforms and PyTorch, we achieved an on-the-fly implementation, making RNP suitable for applications that require real-time processing. We extensively evaluated the performance of RNP on image deblurring, super-resolution tasks, and computed tomography with various regularizers, demonstrating its effectiveness and efficiency in accelerating variational image reconstruction.

Extending RNP to magnetic resonance imaging (MRI) reconstruction presents an interesting future direction. Modern MRI scanners use multi-coils to acquire k-space data, which causes the forward model to change for each scan, requiring the preconditioner to be computed on-the-fly. Additionally, in functional MRI applications, the blood oxygenation level-dependent (BOLD) contrast, which is sensitive to magnetic field inhomogeneity, is primarily used. In some cases, images are jointly reconstructed with the associated field maps [72], causing the forward model to change at each iteration. Given RNP's on-the-fly implementation, we believe it holds great potential as a preconditioner for acceleration.

In this paper, we considered impulsive noise, and set the data-fidelity term to be the  $\ell_p$  norm with  $p \leq 1$ . Another interesting direction for future work is to incorporate deep learning to address this problem. Developing a novel PnP/RED framework for impulsive noise would be an exciting research direction.

## APPENDIX A

### EXAMPLES OF $\|\mathbf{L}\mathbf{x}\|_{1,\phi}$

Let  $\mathbf{X} \in \mathbb{R}^{N_1 \times N_2}$  be the matrix form of  $\mathbf{x}$ , where  $\mathbf{x}$  is the column stacking of  $\mathbf{X}$ . So if  $\mathbf{L}$  represents the first-order differential operator, we have  $G = N$  and

$$\mathbf{v}_l = \begin{bmatrix} \mathbf{X}_{i,j} - \mathbf{X}_{i-1,j} \\ \mathbf{X}_{i,j} - \mathbf{X}_{i,j-1} \end{bmatrix},$$

with  $l = jN_1 + i$  and  $\mathbf{X}_{i,j}$  denotes the component of  $\mathbf{X}$  at the  $i$ th row and the  $j$ th column. So for  $\phi = 1$  (respectively,  $\phi = 2$ ), we have  $\|\mathbf{L}\mathbf{x}\|_{1,1} = \sum_l |\mathbf{X}_{i,j} - \mathbf{X}_{i-1,j}| + |\mathbf{X}_{i,j} - \mathbf{X}_{i,j-1}|$  (respectively,  $\|\mathbf{L}\mathbf{x}\|_{1,2} = \sum_l \sqrt{(\mathbf{X}_{i,j} - \mathbf{X}_{i-1,j})^2 + (\mathbf{X}_{i,j} - \mathbf{X}_{i,j-1})^2}$ ), which represents the anisotropic (respectively, isotropic) total variation.

If  $\mathbf{L}$  represents the second-order finite-difference operator, we have

$$\mathbf{v}_l = \begin{bmatrix} v_l^{11} & v_l^{12} \\ v_l^{21} & v_l^{22} \end{bmatrix},$$

with  $l = jN_1 + i$  and  $v_l^{11} = \mathbf{X}_{i-1,j} - 2\mathbf{X}_{i,j} + \mathbf{X}_{i+1,j}$ ,  $v_l^{22} = \mathbf{X}_{i,j-1} - 2\mathbf{X}_{i,j} + \mathbf{X}_{i,j+1}$ ,  $v_l^{12} = v_l^{21} = \frac{1}{4}(\mathbf{X}_{i+1,j+1} - \mathbf{X}_{i+1,j-1} - \mathbf{X}_{i-1,j+1} + \mathbf{X}_{i-1,j-1})$ . Thus  $\|\mathbf{Lx}\|_{1,\phi} = \sum_l \|\mathbf{v}_l\|_{\phi}$ .

## REFERENCES

- [1] J.-L. Starck, D. L. Donoho, and E. J. Candes, "Very high quality image restoration by combining wavelets and curvelets," *Proc. SPIE*, vol. 4478, pp. 9–19, 2001.
- [2] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Phys. D: Nonlinear Phenomena*, vol. 60, no. 1–4, pp. 259–268, 1992.
- [3] M. Aharon, M. Elad, and A. Bruckstein, "K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation," *IEEE Trans. Signal Process.*, vol. 54, no. 11, pp. 4311–4322, Nov. 2006.
- [4] W. Dong, G. Shi, X. Li, Y. Ma, and F. Huang, "Compressive sensing via nonlocal low-rank regularization," *IEEE Trans. Image Process.*, vol. 23, no. 8, pp. 3618–3632, Aug. 2014.
- [5] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, 2015.
- [6] M. T. McCann, K. H. Jin, and M. Unser, "Convolutional neural networks for inverse problems in imaging: A review," *IEEE Signal Process. Mag.*, vol. 34, no. 6, pp. 85–95, Nov. 2017.
- [7] T. Chen et al., "Learning to optimize: A primer and a benchmark," *J. Mach. Learn. Res.*, vol. 23, no. 189, pp. 1–59, 2022.
- [8] V. Monga, Y. Li, and Y. C. Eldar, "Algorithm unrolling: Interpretable, efficient deep learning for signal and image processing," *IEEE Signal Process. Mag.*, vol. 38, no. 2, pp. 18–44, Mar. 2021.
- [9] S. V. Venkatakrishnan, C. A. Bouman, and B. Wohlberg, "Plug-and-play priors for model based reconstruction," in *Proc. IEEE Glob. Conf. Signal Inf. Process.*, 2013, pp. 945–948.
- [10] Y. Sun, Z. Wu, X. Xu, B. Wohlberg, and U. S. Kamilov, "Scalable plug-and-play ADMM with convergence guarantees," *IEEE Trans. Comput. Imag.*, vol. 7, pp. 849–863, 2021.
- [11] T. Hong, X. Xu, J. Hu, and J. A. Fessler, "Provable preconditioned plug-and-play approach for compressed sensing MRI reconstruction," *IEEE Trans. Comput. Imag.*, vol. 10, pp. 1476–1488, 2024.
- [12] Y. Romano, M. Elad, and P. Milanfar, "The little engine that could: Regularization by denoising (RED)," *SIAM J. Imag. Sci.*, vol. 10, no. 4, pp. 1804–1844, 2017.
- [13] T. Hong, I. Yavneh, and M. Zibulevsky, "Solving RED with weighted proximal methods," *IEEE Signal Process. Lett.*, vol. 27, pp. 501–505, 2020.
- [14] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang, "Beyond a Gaussian denoiser: Residual learning of deep CNN for image denoising," *IEEE Trans. Image Process.*, vol. 26, no. 7, pp. 3142–3155, Jul. 2017.
- [15] G. Daras et al., "A survey on diffusion models for inverse problems," 2024, *arXiv:2410.00083*.
- [16] P. Rodríguez and B. Wohlberg, "Efficient minimization method for a generalized total variation functional," *IEEE Trans. Image Process.*, vol. 18, no. 2, pp. 322–332, Feb. 2009.
- [17] M. Guerquin-Kern, M. Haberland, K. P. Pruessmann, and M. Unser, "A fast wavelet-based reconstruction method for magnetic resonance imaging," *IEEE Trans. Med. Imag.*, vol. 30, no. 9, pp. 1649–1660, Sep. 2011.
- [18] T. Wang, K. Nakamoto, H. Zhang, and H. Liu, "Reweighted anisotropic total variation minimization for limited-angle CT reconstruction," *IEEE Trans. Nucl. Sci.*, vol. 64, no. 10, pp. 2742–2760, Oct. 2017.
- [19] J. A. Fessler and R. R. Nadakuditi, *Linear Algebra For Data Science Machine Learning and Signal Processing*. Cambridge, U.K.: Cambridge Univ. Press, 2024.
- [20] H. Cherkaoui, J. Sulam, and T. Moreau, "Learning to solve TV regularised problems with unrolled algorithms," in *Proc. Adv. Neural Inf. Process. Syst.*, 2020, vol. 33, pp. 11513–11524.
- [21] D. Zheng, X. Zhang, K. Ma, and C. Bao, "Learn from unpaired data for image restoration: A variational Bayes approach," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 45, no. 5, pp. 5889–5903, May 2022.
- [22] H. Gu, B. Yaman, S. Moeller, J. Ellermann, K. Ugurbil, and M. Akçakaya, "Revisiting  $\ell_1$ -wavelet compressed-sensing MRI in the era of deep learning," *Proc. Nat. Acad. Sci.*, vol. 119, no. 33, 2022, Art. no. e2201062119.
- [23] T. T. N. Mai, E. Y. Lam, and C. Lee, "Deep unrolled low-rank tensor completion for high dynamic range imaging," *IEEE Trans. Image Process.*, vol. 31, pp. 5774–5787, 2022.
- [24] J. W. Soh and N. I. Cho, "Variational deep image restoration," *IEEE Trans. Image Process.*, vol. 31, pp. 4363–4376, 2022.
- [25] C. Clason, B. Jin, and K. Kunisch, "A duality-based splitting method for  $\ell_1$ -TV image restoration with automatic regularization Parameter choice," *SIAM J. Sci. Comput.*, vol. 32, no. 3, pp. 1484–1505, 2010.
- [26] G. Huang, A. Lanza, S. Morigi, L. Reichel, and F. Sgallari, "Majorization-minimization generalized Krylov subspace methods for  $\ell_p - \ell_q$  optimization applied to image restoration," *BIT Numer. Math.*, vol. 57, no. 2, pp. 351–378, 2017.
- [27] A. Buccini, M. Pasha, and L. Reichel, "Modulus-based iterative methods for constrained  $\ell_p - \ell_q$  minimization," *Inverse Problems*, vol. 36, no. 8, 2020, Art. no. 084001.
- [28] Z.-Z. Bai, A. Buccini, K. Hayami, L. Reichel, J.-F. Yin, and N. Zheng, "Modulus-based iterative methods for constrained Tikhonov regularization," *J. Comput. Appl. Math.*, vol. 319, pp. 1–13, 2017.
- [29] R. H. Chan and H.-X. Liang, "Half-quadratic algorithm for problems with applications to TV-image restoration and compressive sensing," in *Proc. Efficient Algorithms Glob. Optim. Methods Comput. Vis.: Int. Dagstuhl Seminar, Dagstuhl Castle*, Springer, 2014, pp. 78–103.
- [30] A. Lanza, S. Morigi, L. Reichel, and F. Sgallari, "A generalized Krylov subspace method for  $\ell_p - \ell_q$  minimization," *SIAM J. Sci. Comput.*, vol. 37, no. 5, pp. S30–S50, 2015.
- [31] S. Gazzola, J. G. Nagy, and M. S. Landman, "Iteratively reweighted FGMRES and FLSQR for sparse reconstruction," *SIAM J. Sci. Comput.*, vol. 43, no. 5, pp. S47–S69, 2021.
- [32] M. Benzi, "Preconditioning techniques for large linear systems: A survey," *J. Comput. Phys.*, vol. 182, no. 2, pp. 418–477, 2002.
- [33] D. Bertaccini and F. Durastante, *Iterative Methods and Preconditioning for Large and Sparse Linear Systems With Applications*. London, U.K.: Chapman and Hall/CRC, 2018.
- [34] S. Lefkimmiatis, J. P. Ward, and M. Unser, "Hessian Schatten-norm regularization for linear inverse problems," *IEEE Trans. Image Process.*, vol. 22, no. 5, pp. 1873–1888, May 2013.
- [35] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM J. Imag. Sci.*, vol. 2, no. 1, pp. 183–202, 2009.
- [36] A. Chambolle and T. Pock, "A first-order primal-dual algorithm for convex problems with applications to imaging," *J. Math. Imag. Vis.*, vol. 40, pp. 120–145, 2011.
- [37] N. Komodakis and J.-C. Pesquet, "Playing with duality: An overview of recent primal-dual approaches for solving large-scale optimization problems," *IEEE Signal Process. Mag.*, vol. 32, no. 6, pp. 31–54, Nov. 2015.
- [38] S. Boyd et al., "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Found. Trends Mach. Learn.*, vol. 3, no. 1, pp. 1–122, 2011.
- [39] A. Repetti and Y. Wiaux, "Variable metric forward-backward algorithm for composite minimization problems," *SIAM J. Optim.*, vol. 31, no. 2, pp. 1215–1241, 2021.
- [40] S. Becker, J. Fadili, and P. Ochs, "On quasi-newton forward-backward splitting: Proximal calculus and convergence," *SIAM J. Optim.*, vol. 29, no. 4, pp. 2445–2481, 2019.
- [41] T. Hong, L. Hernandez-Garcia, and J. A. Fessler, "A complex quasi-Newton proximal method for image reconstruction in compressed sensing MRI," *IEEE Trans. Comput. Imag.*, vol. 10, pp. 372–384, 2024.
- [42] J. Chung and M. S. Landman, "Flexible Krylov methods for group sparsity regularization," *Phys. Scripta*, vol. 99, no. 12, 2024, Art. no. 125006.
- [43] K. Chen, *Matrix Preconditioning Techniques and Applications*. Cambridge UK: Cambridge Univ. Press, 2005, no. 19.
- [44] 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. Imag.*, vol. 12, no. 1, pp. 78–83, Mar. 1993.
- [45] 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, May 1999.
- [46] L. Fu et al., "A preliminary investigation of 3D preconditioned conjugate gradient reconstruction for cone-beam CT," *Proc. SPIE*, vol. 8313, SPIE, 2012, pp. 1051–1059.
- [47] S. Pelletier and J. R. Cooperstock, "Preconditioning for edge-preserving image super resolution," *IEEE Trans. Image Process.*, vol. 21, no. 1, pp. 67–79, Jan. 2012.
- [48] Y.-J. Tsai et al., "Fast quasi-newton algorithms for penalized reconstruction in emission tomography and further improvements via preconditioning," *IEEE Trans. Med. Imag.*, vol. 37, no. 4, pp. 1000–1010, Apr. 2018.



- [49] F. Ong, M. Uecker, and M. Lustig, "Accelerating non-cartesian MRI reconstruction convergence using K-space preconditioning," *IEEE Trans. Med. Imag.*, vol. 39, no. 5, pp. 1646–1654, May 2020.
- [50] M. Savanier, E. Chouzenoux, J.-C. Pesquet, and C. Riddell, "Unmatched preconditioning of the proximal gradient algorithm," *IEEE Signal Process. Lett.*, vol. 29, pp. 1122–1126, 2022.
- [51] K. Koolstra, J. van Gemert, P. Börnert, A. Webb, and R. Remis, "Accelerating compressed sensing in parallel imaging reconstructions using an efficient circulant preconditioner for cartesian trajectories," *Magn. Reson. Med.*, vol. 81, no. 1, pp. 670–685, 2019.
- [52] K. Koolstra and R. Remis, "Learning a preconditioner to accelerate compressed sensing reconstructions in MRI," *Magn. Reson. Med.*, vol. 87, no. 4, pp. 2063–2073, 2022.
- [53] P.-G. Martinsson and J. A. Tropp, "Randomized numerical linear algebra: Foundations and algorithms," *Acta Numerica*, vol. 29, pp. 403–572, 2020.
- [54] Z. Frangella, J. A. Tropp, and M. Udell, "Randomized Nyström preconditioning," *SIAM J. Matrix Anal. Appl.*, vol. 44, no. 2, pp. 718–752, 2023.
- [55] A. Eberhard, R. Hill, D. Ralph, B. M. Glover, L. Qi, and D. Sun, "A survey of some nonsmooth equations and smoothing newton methods," in *Progress in Optimization Contributions from Australasia*. Boston, MA, USA: Springer, 1999, pp. 121–146.
- [56] T. Hong, T.-a. Pham, I. Yavneh, and M. Unser, "A mini-batch quasi-newton proximal method for constrained total-variation nonlinear image reconstruction," 2023, *arXiv:2307.02043*.
- [57] S. Zhao, Z. Frangella, and M. Udell, "NysADMM: Faster composite convex optimization via low-rank approximation," in *Proc. Int. Conf. Mach. Learn.*, PMLR, 2022, pp. 26824–26840.
- [58] Y. Saad, *Iterative Methods For Sparse Linear Systems*. Philadelphia, PA, USA: SIAM, 2003.
- [59] Y.-C. Chu, L.-R. Santos, and M. Udell, "Randomized Nyström preconditioned interior point-proximal method of multipliers," *Accepted SIAM J. Scientific Comput.*, 2024.
- [60] N. Parikh et al., "Proximal algorithms," *Found. Trends Optim.*, vol. 1, no. 3, pp. 127–239, 2014.
- [61] A. Beck and M. Teboulle, "Fast gradient-based algorithms for constrained total variation image denoising and deblurring problems," *IEEE Trans. Image Process.*, vol. 18, no. 11, pp. 2419–2434, Nov. 2009.
- [62] A. Beck, *First-Order Methods in Optimization*. Philadelphia, PA, USA: SIAM, 2017, vol. 25.
- [63] A. Paszke et al., "Pytorch: An imperative style, high-performance deep learning library," *Proc. 33rd Int. Conf. Neural Inf. Process. Syst.*, 2019, pp. 8026–8037.
- [64] M. Abadi et al., "TensorFlow: A system for large-scale machine learning," in *Proc. 12th USENIX Symp. Operating Syst. Des. Implementation*, 2016, pp. 265–283.
- [65] L. Métivier, R. Brossier, J. Virieux, and S. Operto, "Full waveform inversion and the truncated Newton method," *SIAM J. Sci. Comput.*, vol. 35, no. 2, pp. B401–B437, 2013.
- [66] E. J. Candes, M. B. Wakin, and S. P. Boyd, "Enhancing sparsity by reweighted  $\ell_1$  minimization," *J. Fourier Anal. Appl.*, vol. 14, pp. 877–905, 2008.
- [67] R. Murray et al., "Randomized numerical linear Algebra: A perspective on the field with an eye to software," 2023, *arXiv:2302.11474*.
- [68] M. Wolter, F. Blanke, J. Garcke, and C. T. Hoyt, "PTWT-The pytorch wavelet toolbox," *J. Mach. Learn. Res.*, vol. 25, no. 80, pp. 1–7, 2024.
- [69] J. Adler et al., "ODL—The operator discretization library," version 0.6.0, 2017. [Online]. Available: <https://github.com/odlgroup/odl>
- [70] C. H. McCollough et al., "Results of the 2016 low dose CT grand challenge," *Med. Phys.*, vol. 44, no. 10, pp. e339–e352, Oct. 2017.
- [71] K. Hämäläinen, L. Harhanen, A. Kallonen, A. Kujanpää, E. Niemi, and S. Siltanen, "Tomographic X-ray data of a walnut," 2015, *arXiv:1502.04064*.
- [72] V. T. Olafsson, D. C. Noll, and J. A. Fessler, "Fast joint reconstruction of dynamic  $R_2^*$  and field maps in functional MRI," *IEEE Trans. Med. Imag.*, vol. 27, no. 9, pp. 1177–1188, Sep. 2008.