arXiv:2602.09527v1 [math.OC] 10 Feb 2026

SPLIT, SKIP AND PLAY: VARIANCE-REDUCED PROXSKIP FOR TOMOGRAPHY RECONSTRUCTION IS EXTREMELY FAST

Evangelos Papoutsellis[1] *Zeljko Kereta*[2] *Kostas Papafitsoros*[3]
[1]Finden Ltd [2]University College London [3]Queen Mary University of London

ABSTRACT

Many modern iterative solvers for large-scale tomographic reconstruction incur two major computational costs per iteration: expensive forward/adjoint projections to update the data fidelity term and costly proximal computations for the regulariser, often done via inner iterations. This paper studies for the first time the application of methods that couple randomised skipping of the proximal with variance-reduced subset-based optimisation of data-fit term, to simultaneously reduce both costs in challenging tomographic reconstruction tasks. We provide a series of experiments using both synthetic and real data, demonstrating striking speed-ups of the order $5\times$–$20\times$ compared to the non-skipped counterparts which have been so far the standard approach for efficiently solving these problems. Our work lays the groundwork for broader adoption of these methods in inverse problems.

***Index Terms*—** Stochastic Optimisation, Tomography Reconstruction, Inverse Problems, Proximal Operator

1. INTRODUCTION

Tomographic modalities like X-ray Computed Tomography (CT), and Positron Emission Tomography (PET) reconstruct an unknown image or volume $\boldsymbol{x}^\dagger \in \mathbb{R}^m$ from indirect and noisy measurements $\boldsymbol{b} \in \mathbb{R}^m$ produced by a (typically linear but ill-posed) forward operator $\boldsymbol{A} : \mathbb{R}^n \to \mathbb{R}^m$. a standard approach is to use variational regularisation and solve

$$\arg\min_{\boldsymbol{x}\in\mathbb{R}^n} \mathcal{D}(\boldsymbol{A}\boldsymbol{x}, \boldsymbol{b}) + \alpha\mathcal{R}(\boldsymbol{x}). \tag{1}$$

The data fidelity term $\mathcal{D}$ penalises the mismatch between the predicted measurements $\boldsymbol{A}\boldsymbol{x}$ and the acquired data $\boldsymbol{b}$ and is determined by the noise statistics, e.g. squared ℓ_2 under Gaussian noise. The regulariser $\mathcal{R}$ encodes prior information and promotes solution desirable properties, such as smoothness, sparsity, or low-rank structure. Commonly used choices for $\mathcal{R}$ include model based priors such as Total Variation (TV), and its extensions [1]. Alternatively, one can use learned implicit priors, e.g. plug-and-play (PnP) approaches [2], where classical proximal regularisation is replaced by a denoiser, such as BM3D [3] or a pretrained neural network.

Problem (1) falls under the general framework of minimising composite, and potentially non-smooth, objectives

$$\min_{\boldsymbol{x}\in\mathbb{R}^n} f(\boldsymbol{x}) + g(\boldsymbol{x}), \tag{2}$$

where f and g correspond to $\mathcal{D}$ and $\mathcal{R}$, respectively. Proximal Gradient Descent, also known as the Iterative Shrinkage-Thresholding Algorithm (ISTA), and its accelerated variant FISTA [4, 5], are widely used for convex f with L-Lipschitz gradient, and proper, convex g, see Algorithms1.

Each iteration of ISTA and FISTA requires evaluating both the gradient of f, and the proximal operator of g. The latter is defined by

$$\operatorname{prox}_{\tau g}(\boldsymbol{x}) := \arg\min_{\boldsymbol{z}\in\mathbb{R}^n} \left\{ \frac{1}{2}\|\boldsymbol{z}-\boldsymbol{x}\|_2^2 + \tau g(\boldsymbol{z}) \right\}, \quad \tau > 0 \tag{3}$$

and has to be practically tractable. While these algorithms are simple and flexible, their practical efficiency is often limited by two dominant per-iteration costs. The first is the gradient evaluation of the data fidelity term which requires repeated application of the forward operator and its adjoint and is costly, especially for large-scale 3D tomography. The second is the proximal step, which can also be expensive, e.g., for TV-type proximals computed via iterative solvers. Moreover, in PnP methods the proximal step is interpreted as a denoising operation which can also be costly e.g., for BM3D.

1.1. ProxSkip

To reduce the computational cost of the proximal operator, a natural strategy is to avoid computing it at every iteration. This idea is formalised by the *ProxSkip* algorithm [6], which applies the proximal operator only at randomly selected iterations with probability p, while employing a control variate to stabilise the resulting skipping. If $p = 1$, the ProxSkip algorithm coincides with ISTA, see Algorithm 1. Convergence in expectation is established for any $p \in (0, 1]$, in both the strongly convex and merely convex regimes (ergodic convergence) for the function f [6, 7]. In practice, ProxSkip is most beneficial when the skipping probability is relatively low. making the proximal evaluations sufficiently infrequent. However, if p is chosen too small, proximal updates are rare and progress can be slow. Convergence rate can be optimised by selecting the probability parameter as $p = \sqrt{\frac{\mu}{L}}$. This, however, relies on strong convexity, a rather restrictive assumption in imaging inverse problems. Moreover, this choice

is not necessarily optimal in terms of computational time to reach a target accuracy.

ProxSkip was originally proposed in the context of federated/distributed optimisation. It has only recently been explored for imaging inverse problems [8], where it was shown to substantially reduce the computational burden and runtimes of expensive proximal evaluations, while achieving the same image quality. Nevertheless, the cost of forward/adjoint evaluations in tomography remains a significant burden.

1.2. Data Spliting and Variance Reduced Estimators

In cases where the forward model can be decomposed into blocks, e.g. as subsets of projection angles in CT/PET [9], temporal frames in dynamic imaging [10], or groups of measurements in computational photography [11], the data fidelity can be decomposed across measurements. Here one can express the data fidelity term in (1) in a finite-sum form as $\mathcal{D}(\boldsymbol{A}\boldsymbol{x}, \boldsymbol{b}) = \sum_{i=1}^{N} D(\boldsymbol{A}_i\boldsymbol{x}, \boldsymbol{b}_i)$ where $\{(\boldsymbol{A}_i, \boldsymbol{b}_i)\}_{i=1}^{N}$ represent subsets of the forward operator and data. In tomography, each $\boldsymbol{b}_i$ corresponds to a subset of projection angles, see Figure 1.

Evaluating $\nabla D(\boldsymbol{A}_i\boldsymbol{x}, \boldsymbol{b}_i) = \nabla f_i(\boldsymbol{x})$ typically only requires applying $\boldsymbol{A}_i$ and $\boldsymbol{A}_i^\top$, i.e. forward and backward projections for this subset.

For example, when applied to CT reconstruction, SGD uses $G_k(\boldsymbol{x}) = N\boldsymbol{A}_i^\top(\boldsymbol{A}_i\boldsymbol{x} - \boldsymbol{b}_i)$ as an estimator of the full gradient $\boldsymbol{A}^\top(\boldsymbol{A}\boldsymbol{x} - \boldsymbol{b})$.

Variance introduced around the true gradient may prevent convergence to high accuracy solutions under constant step sizes, causing oscillations or stagnation in ill-conditioned inverse problems. These issues can be addressed with *Variance-Reduction* (VR) techniques, which combine cheap stochastic gradient evaluations while controling the variance.

For example, in Stochastic Average Gradient Amelioré (SAGA) [12] this is achieved by storing a table of subset gradients $\{v_k^i\}_{i=1}^{N}$ and defining $\bar{v}_k = \sum_{i=1}^{N} v_k^i$. By sampling i_k the SAGA estimator for the gradient at iteration k is

$$\boldsymbol{G}_k(\boldsymbol{x}) := N(\nabla f_{i_k}(\boldsymbol{x}) - v_k^{i_k}) + \bar{v}_k, \tag{4}$$

and the table is updated only at index $j = i_k$ with $\nabla f_j(\boldsymbol{x})$. Stochastic Variance Reduced Gradient (SVRG) [13, 14] instead uses a snapshot point $\tilde{\boldsymbol{x}}$ and the full gradient $\nabla f(\tilde{\boldsymbol{x}}) := \sum_{i=1}^{N} \nabla f_i(\tilde{\boldsymbol{x}})$, which are stored and updated every N iterations. The gradient estimator is computed as

$$\boldsymbol{G}_k(\boldsymbol{x}) := N(\nabla f_{i_k}(\boldsymbol{x}) - \nabla f_{i_k}(\tilde{\boldsymbol{x}})) + \nabla f(\tilde{\boldsymbol{x}}),$$

Loopless SVRG [15] replaces the deterministic schedule for updating the full gradient with a probabilistic one; in each iteration full gradient is updated with probability $1/N$, cf. [16].

Algorithm 1 Family of all considered algorithms

Parameters: $\gamma > 0$, probability $p > 0$, data subsets N
Initialize: $\boldsymbol{x}_0, \boldsymbol{h}_0 \in \mathbb{R}^n$
for $k = 0, 1, \ldots, K-1$ **do**
  Compute $\boldsymbol{G}_k$ (unbiased estimator of $\nabla f(\boldsymbol{x}_k)$)
  $\hat{\boldsymbol{x}}_{k+1} = \boldsymbol{x}_k - \gamma(\boldsymbol{G}_k(\boldsymbol{x}_k) - \boldsymbol{h}_k)$
  Sample $\theta_k \sim \mathrm{Bernoulli}(p)$, $\theta_k \in \{0, 1\}$
  if $\theta_k = 1$ **then**
    $\boldsymbol{x}_{k+1} = \mathrm{prox}_{\frac{\gamma}{p}g}\left(\hat{\boldsymbol{x}}_{k+1} - \frac{\gamma}{p}\boldsymbol{h}_k\right)$
  else $\boldsymbol{x}_{k+1} = \hat{\boldsymbol{x}}_{k+1}$
  end if
  $\boldsymbol{h}_{k+1} = \boldsymbol{h}_k + \frac{p}{\gamma}(\boldsymbol{x}_{k+1} - \hat{\boldsymbol{x}}_{k+1})$
end for

$p = 1$	$0 < p < 1$	$\boldsymbol{G}_k$
ISTA	ProxSkip	$\nabla f(\boldsymbol{x}_k)$
ProxSGD	ProxSGDSkip	$N\nabla f_{i_k}(\boldsymbol{x})$
ProxSAGA	ProxSAGASkip	$N(\nabla f_{i_k}(\boldsymbol{x}) - v_k^{i_k}) + \bar{v}_k$
ProxSVRG	ProxSVRGSkip	$N(\nabla f_{i_k}(\boldsymbol{x}) - \nabla f_{i_k}(\tilde{\boldsymbol{x}})) + \nabla f(\tilde{\boldsymbol{x}})$
ProxLSVRG	ProxLSVRGSkip	as above, updated with $p = 1/N$

FISTA is ISTA with an acceleration step [5]

Combining proximal skipping with stochastic VR gradients yields a family of *ProxSkip-VR* methods (ProxSAGASkip, ProxSVRGSkip, ProxLSVRGSkip), all outlined in Algorithms 1 that reduce both dominant per-iteration costs in terms of the forward model and the proximal operator.

Hence, ProxSkip-VR provides two complementary computational savings, *data splitting* for cheaper gradient steps and *proximal skipping* for cheaper regularisation steps, while maintaining stable convergence behaviour comparable to full-gradient proximal methods. ProxSkip-VR algorithms were only recently introduced for the problem of federated learning [17].

1.3. Our contribution

Here, we show for the first time the feasibility of ProxSkip-VR algorithms in challenging large scale inverse problems. We use the proposed versions of ProxSkip-VR algorithms on tomographic reconstruction problems (ProxSGDSkip, ProxSAGASkip, ProxSVRGSkip, ProxLSVRGSkip), and compare them against (i) non-skipped counterparts (ProxSGD, ProxSVRG, ProxLSVRG) (ii) stochastic methods without variance reduction (ProxSGDSkip), and (iii) deterministic full-gradient algorithms, skipped and non-skipped, (ISTA, FISTA, ProxSkip). We demonstrate a speed-up of at least $5\times$ of ProxSkip-VR algorithms over their non-skipped counterparts which so far have been considered among the most efficient methods with respect to computational time for variational regularisation problems in tomography.

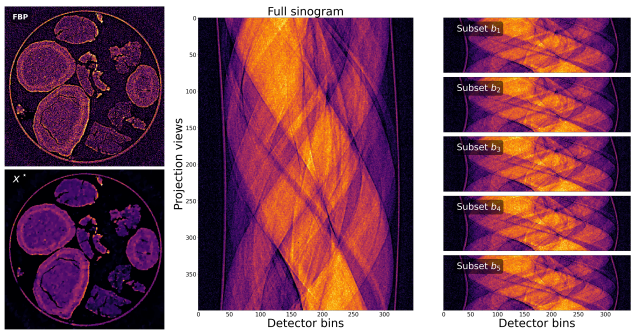

Fig. 1: Left: FBP reconstruction of the post–partial oxidation of methane reaction [18] and high-accuracy reference (TV-regularised) solution, computed with 2×10^5 PDHG iterations using diagonal preconditioning [19]. Right: Example of data splitting ($N = 5$) using staggered indexing for the projection angles: $b_k = \{k,\ k+5,\ k+10,\ \ldots\}$ for $k = 0, \ldots, 4$.

2. PROXSKIP-VR FOR TOMOGRAPHY

We consider two experimental setups to assess performance under different regularisations and evaluation criteria.

2.1. TV tomography reconstruction on real dataset

Our goal is to compare instances of Algorithm 1 in terms of the CPU time needed to reach a prescribed tolerance $\varepsilon > 0$ to a high-accuracy reference solution $\boldsymbol{x}^*$. We solve

$$\min_{\boldsymbol{x}\in\mathbb{R}^n} \ \frac{1}{2}\|\boldsymbol{A}\boldsymbol{x} - \boldsymbol{b}\|_2^2 + \alpha\mathrm{TV}(\boldsymbol{x}) + \mathbb{I}_{\{\boldsymbol{x}\geq 0\}}(\boldsymbol{x}), \qquad (5)$$

where $\boldsymbol{A}$ is the discrete Radon transform, on a real tomographic dataset, where ground truth is not available. We use the isotropic TV as the regularisation term weighted with a regularisation parameter $\alpha > 0$, optimised by visual inspection in the absence of ground truth. Here $\mathrm{prox}_{\alpha\mathrm{TV}}(\cdot)$, line 8 of Algorithm 1, is solved using the Fast Gradient Projection (FGP) algorithm [20] with a fixed number of iterations.

We use a fixed number of warm-started FGP inner iterations to emulate cheap (10 iterations) and expensive (100 iterations) proximal steps, enabling controlled comparisons.

Dataset: We use a chemical imaging tomography dataset, see Figure 1. The initial dataset was acquired for 800 projection angles with 695×695 detector size and 700 vertical slices. For demonstration purposes we consider one 2D vertical slice with half the projections and 2× rebinned detector size, i.e. 400 angles and 347 detector points.

Reference solution and target accuracy: The goal in this experiment is to identify which method achieves the desired reconstruction accuracy the fastest. The methods are evaluated in terms of the wall-clock CPU time required to either reach an error tolerance $\frac{\|\boldsymbol{x}-\boldsymbol{x}^*\|^2}{\|\boldsymbol{x}^*\|^2} < 10^{-5}$ or to reach 200 iterations for deterministic methods and 200 data passes for stochastic methods. Here, a data pass denotes one iteration for deterministic algorithms and when using N subset gradients for stochastic algorithms. The high-accuracy reference solution $\boldsymbol{x}^*$ is computed using a separate primal–dual proximal method to avoid bias, see Figure 1.

Controlling gradient and proximal costs: Our proposed algorithms allow independent control of the two dominant per-iteration costs. (i) *Proximal cost:* We control the expected proximal workload through the skipping probability $p \in \{0.01, 0.05, 0.1, 0.3, 0.5\}$ which determines how often the (cheap or expensive) TV proximal operators are performed. (ii) *Gradient cost:* For stochastic methods we vary the number of subsets N in the decomposition of the data fidelity term. Larger N corresponds to smaller-size subsets and therefore cheaper (but noisier) stochastic gradients. Decreasing N yields more expensive gradients that better approximate the full gradient ∇f and have lower variance.

2.2. PnP-BM3D on a simulated dataset

We perform the tomographic reconstruction within a plug-and-play (PnP) framework using the BM3D denoiser and combining variance-reduced and skipping techniques [21].

Dataset: We use simulated data, enabling quantitative evaluation of the iterates against a known ground truth. Specifically, we generate a cylindrical foam phantom with non-overlapping bubble microstructures [22], see Figure 5. The image dimensions match those of the real dataset in Section 2.1. We then compute analytical projection data from 400 uniformly spaced view angles and corrupt the measurements with additional simulated noise.

BM3D denoiser:

We replace step 8 in Algorithms 1 with a BM3D denoiser $D_\sigma(\cdot)$, with $\sigma > 0$ separately optimised for all the algorithms.

Stopping criterion and evaluation: In this experiment, algorithms are compared with respect to the CPU time required to reach a given image-quality level, measured by PSNR and SSIM against the ground truth. We impose a fixed time budget and stop all methods after 3 minutes of CPU time, enabling a direct time-to-quality comparison.

3. NUMERICAL RESULTS

All experiments were run on an Apple M2 Pro with 16 GB RAM and no GPU acceleration. We first examine the average wall-clock time (10 runs) of the regularisers in each setting; see Figure 2. The TV-proximal is relatively inexpensive, taking 0.023sec and 0.189sec on average for 10 and 100 inner iterations respectively. In contrast, BM3D is substantially more costly, requiring more than 1sec per iteration.

We also report the cost of the deterministic full gradient ∇f used in every iteration for ISTA, FISTA and ProxSkip, the stochastic subset gradients ∇f_i used in the data-splitting

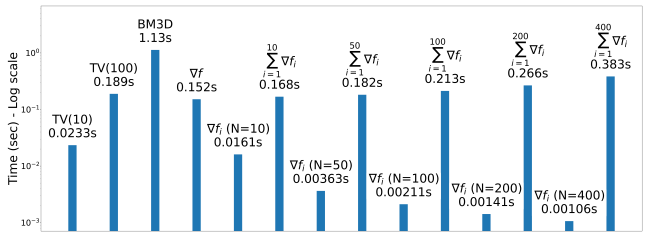

Fig. 2: CPU time for TV proximal (10/100 iterations), BM3D, deterministic/stochastic gradients, and full stochastic gradients for different N.

cases, and the full-gradient computations required by ProxSVRG/ProxLSVRG for different numbers of subsets N. For the stochastic gradients, we observe that increasing N (using smaller subsets) substantially reduces the per-update cost of ∇f_i. In contrast, the cost of a full gradient pass constructed from N subset evaluations increases with N, mainly due to the overhead of looping over more subsets. However, such full-gradient computations are performed infrequently in ProxSVRG (periodically) and randomly in ProxLSVRG, so they do not dominate runtime. Finally, the timings are consistent so that the cost of approximately N stochastic gradients evaluations is roughly the same as full-gradient evaluation.

TV-Real dataset: We first compare ProxSVRG and ProxSVRGSkip with $N = 100$ and $p \in \{0.01, 0.05, 0.1, 0.3, 0.5\}$, Figure 3 (top). The step size is fixed, $\gamma = \frac{1}{L}$. Similarly to observations in [8], the error trajectories in terms of the data passes are identical for all probabilities except the extreme case of $p = 0.01$ in which the proximal is rarely applied. For all the other probabilities, we observe a faster convergence in terms of computational times, with a $3.13\times$ acceleration for $p = 0.05$ compared to the non-skipped version and only 232 $\mathrm{prox}_{\alpha \mathrm{TV}}$ evaluations with 10 inner iterations.

In Figure 3 (bottom), we perform the same comparison as above by keeping the probability fixed, $p = 0.3$ and varying the number of subsets N. When N is small we observe that both skip and non-skip versions have slow convergence in terms of data passes and fail to reach high accuracy within 200 data passes. As N, and hence the variance, increase, both skip and non-skip versions converge in less than 200 data passes. Moreover, ProxSVRGSkip consistently reaches the target accuracy faster than ProxSVRG, with a $1.65\times$–$2.33\times$ speed-up in wall-clock time. The error trajectories plotted against data passes are identical for $N = 50, 100, 200$ and begin to differ for $N = 400$ due to the high stochastic variance. Interestingly, measured in data passes, ProxSVRG with $N = 400$ reaches the target accuracy in the fewest data passes, but is the slowest in terms of computation time. This highlights that data passes do not fully capture the overall computational effort, particularly for large N, where each update is cheap but the stochastic gradient variance is high.

We finally compare all the methods: ISTA ($\gamma = 1.99/L$),

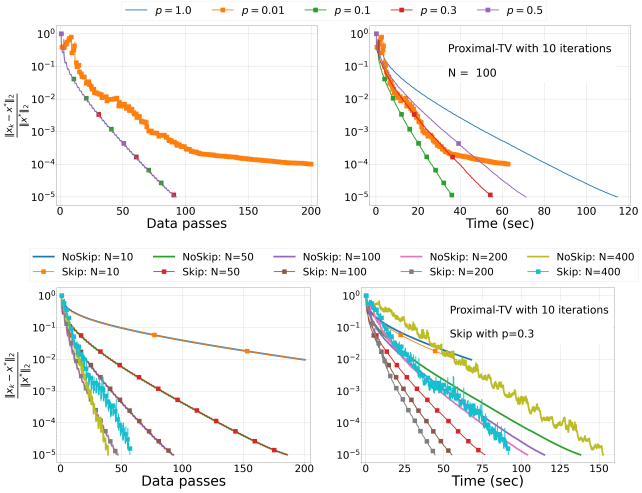

ProxSVRG / p	0.01	0.1	0.3	0.5	1
	–	**36.6 sec**	55.3 sec	71.1 sec	115 sec

ProxSVRG / N	10	50	100	200	400
NoSkip	–	138 sec	115 sec	104 sec	153 sec
Skip	–	**76.6 sec**	**55.3 sec**	**44.6 sec**	**91.8 sec**

Fig. 3: ProxSVRG vs ProxSVRGSkip comparison in terms of data passes (left) and time (right). Top: fixed number of subsets N and varying probabilities p. Bottom: Fixed probability p and varying number of subsets N.

FISTA ($\gamma = 1/L$), stochastic estimators, ProxSAGA ($\gamma = 1/3L$) and ProxSVRG/ProxLSVRG ($\gamma = 1/L$), and their skipped counterparts ProxSkip ($\gamma = 1.99/L$), ProxSAGASkip ($\gamma = 1/3L$), ProxSVRGSkip/ProxLSVRGSkip ($\gamma = 1/L$). We evaluate all algorithms over the values of N and p specified above (with full-gradient update probability set to $\frac{1}{N}$ where applicable) and for the two cases of inner iterations. In Figure 4, we report the best-performing algorithms, showing only those that reached the target accuracy $\varepsilon = 10^{-5}$ within our computational budget. We also report in Table 1 the best-performing configurations (10 inner iterations setting) at two target tolerances, $\varepsilon = 10^{-3}$ and 10^{-5}, summarising the fastest time-to-accuracy and the corresponding iteration count K achieved by each method.

We observe *striking* speed-ups. ProxSVRGSkip and ProxLSVRGSkip achieve about a $5\times$ speed-up over their non-skipped counterparts and $20\times$ speed-up over deterministic ISTA. For 100 inner iterations, the corresponding speed-ups are about $23\times$ (vs. non-skipped) and $21\times$ (vs. ISTA).

PnP–BM3D-Simulated dataset: In Figure 5, we report the PnP–BM3D reconstructions obtained after a fixed CPU-time budget of 3 minutes, alongside the ground truth and the Filtered Back Projection (FBP) reconstruction. To ensure a fair comparison in this PnP setting, we tune the denoiser

Table 1: Best time-to-accuracy for TV reconstruction (10 inner iterations): Time (T) and number of iterations (K) at error thresholds ε; "–" denotes not reached.

Algorithm	N	p	T/K ($\varepsilon = 10^{-3}$)	T/K ($\varepsilon = 10^{-5}$)
ISTA	1	1	182.1 / 1050	405.72 / 2336
FISTA	1	1	72.9 / 415	–
ProxSkip	1	0.1	162.7 / 1037	373.6 / 2328
ProxSAGA	100	1	128.9 / 6226	–
ProxSVRG	200	1	47.4 / 2113	103.9 / 4651
ProxLSVRG	200	1	52.8 / 2419	109.8 / 4971
ProxSAGASkip	200	0.05	14.9 / 6248	–
ProxSVRGSkip	200	0.05	11.3 / 2447	**20.8 / 4545**
ProxLSVRGSkip	200	0.05	**11.1 / 2267**	21.7 / 4769

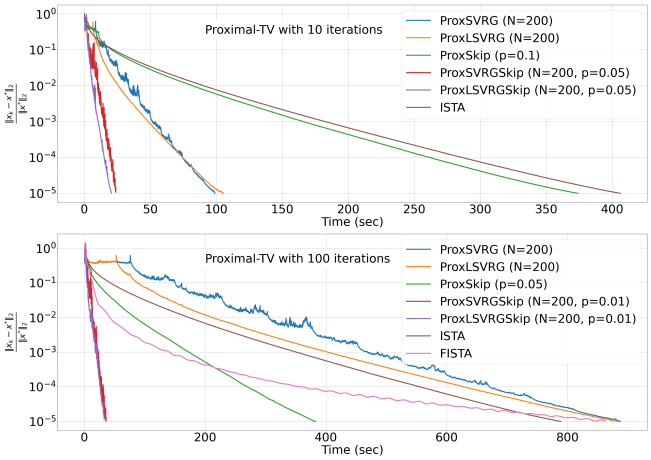

Fig. 4: Convergence (relative error vs. CPU time) for TV tomography reconstruction with 10 (top) and 100 (bottom) inner iterations. Only methods reaching the target accuracy 10^{-5} within our computational budget are shown.

strength separately for skipped and non-skipped variants so that they converge to similar-quality reconstructions. For the non-skip algorithms, we use $\sigma_{\text{non-skip}} = 2 \times 10^{-4}$. For the skipped variants, we empirically found a simple heuristic that yields comparable image-quality metrics across all tested values of p, i.e. $\sigma_{\text{skip}} = \sigma_{\text{non-skip}}/\sqrt{p}$. Since BM3D is approximately $42\times$ and $5\times$ more expensive than the cheap and expensive versions of prox_{TV}, respectively, the overall per-iteration runtime is dominated by the denoiser rather than the gradient computation. As a result, and to keep the experimental sweep computationally feasible, we restrict the stochastic experiments to a single representative configuration $N = 50, p = 0.05$.

The results show that skipping-based algorithms, e.g., ProxSkip and ProxSVRGSkip reach high PSNR/SSIM substantially faster than their non-skipped counterparts. For ISTA, FISTA, and ProxSVRG, BM3D is applied at every iteration, explaining their slower improvement in wall-clock time despite comparable iteration counts. Note, however, that ProxSVRG completes only three data passes for $N = 100$ within the 3-minute budget, which explains its noticeably lower PSNR/SSIM in this experiment.

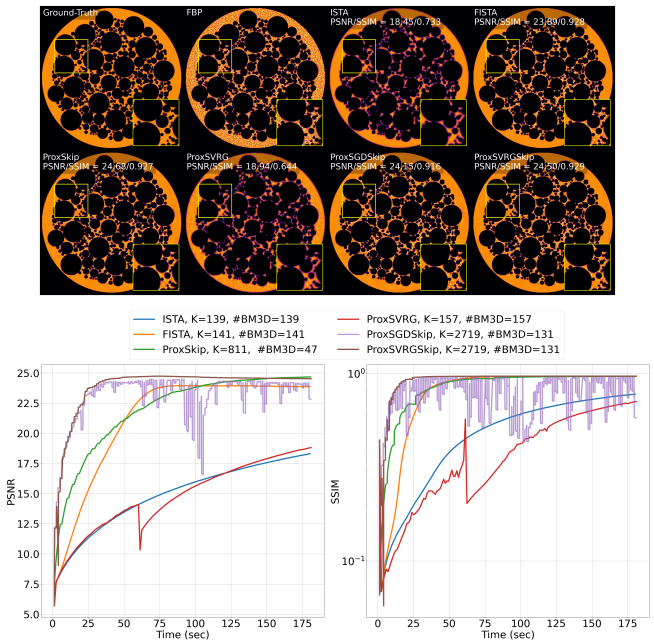

Fig. 5: PnP–BM3D tomographic reconstruction on simulated foam data (3-minute budget). Top: ground truth, FBP, and reconstructions from all algorithms. Middle: PSNR (left) and SSIM (right) versus wall-clock time. The number of subsets is $N = 50$ with $p = 0.05$ and #BM3D denotes the number of calls for the BM3D denoiser.

Code reproducibility: The code and datasets needed to reproduce the results will be available on the corresponding author's webpage.

4. CONCLUSION

Until now, variational regularisation problems in tomographic imaging have been most efficiently addressed using stochastic variance-reduced methods. Here we show for the first time that the performance of these algorithms can be significantly accelerated when the proximal operator that corresponds to regularisation is randomly skipped. This combination of subset-based variance-reduced gradients with random proximal skipping, significantly reduces the overall computational cost while maintaining stable convergence. Our work opens up several directions for future research. Integrating variance reduction and proximal skipping to accelerated (momentum-based) schemes, has the potential to further reduce computational times. Furthermore, the incorporation of these mechanisms to stochastic primal-dual approaches is also worth of investigation. Finally, the adoption of ProxSkip-VR schemes to other challenging tasks, e.g. diffraction tomography and dynamic imaging, is also expected to be largely beneficial.

5. REFERENCES

[1] M. Benning and M. Burger, "Modern regularization methods for inverse problems," *Acta Numer.*, vol. 27, pp. 1–111, 2018.

[2] S. V. Venkatakrishnan, C. A. Bouman, and B. Wohlberg, "Plug-and-play priors for model based reconstruction," in *IEEE GlobalSIP*, 2013, pp. 945–948.

[3] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising by sparse 3-D transform-domain collaborative filtering," *IEEE trans. Image Process.*, vol. 16, no. 8, pp. 2080–2095, Aug. 2007.

[4] P. L. Combettes and V. R. Wajs, "Signal recovery by proximal forward-backward splitting," *Multiscale Model. Sim.*, vol. 4, no. 4, pp. 1168–1200, 2005.

[5] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM J. Imaging Sci.*, vol. 2, no. 1, pp. 183–202, 2009.

[6] K. Mishchenko, G. Malinovsky, S. Stich, and P. Richtárik, "ProxSkip: Yes! Local gradient steps provably lead to communication acceleration! Finally!," in *PMLR 39*, 2022, vol. 162, pp. 15750–15769.

[7] L. Condat and P. Richtárik, "RandProx: Primal-dual optimization algorithms with randomized proximal updates," in *11th ICLR*, 2023.

[8] E. Papoutsellis, Z. Kereta, and K. Papafitsoros, "Why do we regularise in every iteration for imaging inverse problems?," in *SSVM*, 2025, pp. 43–55.

[9] E. Papoutsellis, E. Ametova, C. Delplancke, G. Fardell, J. S. Jørgensen, E. Pasca, M. Turner, R. Warr, W. R. B. Lionheart, and P. J. Withers, "Core Imaging Library - part II: multichannel reconstruction for dynamic and spectral tomography," *Philos. Trans. A Math. Phys. Eng. Sci.*, vol. 379, no. 2204, pp. 20200193, 2021.

[10] L. Protopapa, M. Duff, J. Mayer, J. Schulz-Menger, K. Thielemans, C. Kolbitsch, and E. Pasca, "Efficient motion-corrected image reconstruction for 3d cardiac mri through stochastic optimisation," *Physics in Medicine & Biology*, vol. 70, no. 18, pp. 185012, 2025.

[11] Y. Sun, B. Wohlberg, and U. Kamilov, "An online plug-and-play algorithm for regularized image reconstruction," *IEEE Trans. Comput. Imaging*, vol. 5, no. 3, pp. 395–408, Sept. 2019.

[12] A. Defazio, F. Bach, and S. Lacoste-Julien, "SAGA: A fast incremental gradient method with support for non-strongly convex composite objectives," in *NeurIPS*, 2014, vol. 2, pp. 1646–1654.

[13] R. Johnson and T. Zhang, "Accelerating stochastic gradient descent using predictive variance reduction," in *NeurIPS*, 2013, vol. 26.

[14] L. Xiao and T. Zhang, "A Proximal Stochastic Gradient Method with Progressive Variance Reduction," *SIAM Journal on Optimization*, vol. 24, no. 4, pp. 2057–2075, Jan. 2014.

[15] D. Kovalev, S. Horváth, and P. Richtárik, "Don't jump through hoops and remove those loops: SVRG and Katyusha are better without the outer loop," in *PMLR*. 08 Feb–11 Feb 2020, vol. 117 of *PMLR*, pp. 451–467, PMLR.

[16] R. M. Gower, M. Schmidt, F. Bach, , and P. Richtárik, "Variance-reduced methods for machine learning," *Proc. IEEE*, vol. 108, no. 11, pp. 1968–1983, 2020.

[17] G. Malinovsky, K. Yi, and P. Richtárik, "Variance reduced ProxSkip: Algorithm, theory and application to federated learning," *NeurIPS*, vol. 35, pp. 15176–15189, 2022.

[18] D. Matras, A. Vamvakeros, S. D. M. Jacques, M. di Michiel, V. Middelkoop, I. Z. Ismagilov, E. V. Matus, V. V. Kuznetsov, R. J. Cernik, and A. M. Beale, "Multi-length scale 5D diffraction imaging of Ni-Pd/CeO2-ZrO2/Al2O3 catalyst during partial oxidation of methane," *J. Mater. Chem. A*, vol. 9, no. 18, pp. 11331–11346, 2021.

[19] T. Pock and A. Chambolle, "Diagonal preconditioning for first order primal-dual algorithms in convex optimization," in *ICCV*, 2011, pp. 1762–1769.

[20] 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, 2009.

[21] V. Monardo, A. Iyer, S. Donegan, M. De Graef, and Y. Chi, "Plug-and-play image reconstruction meets stochastic variance-reduced gradient methods," in *ICIP*, 2021, pp. 2868–2872.

[22] D. J. Ching and D. Gürsoy, "XDesign: an open-source software package for designing X-ray imaging phantoms and experiments," *J. Synchrotron Rad.*, vol. 24, no. 2, pp. 537–544, 2017.