Supplementary Information for "SAVER: Stochastic Adaptive Variance-Driven Exploration and Reconstruction for Low-Dose Computed Tomography"

By Shunta Nonaga, Koji Tabata, Junya Honda, Hiroyuki Kudo, Wataru Yashiro, Tamiki Komatsuzaki

Supplementary Information Overview

This Supplementary Information provides additional theoretical discussions and experimental results supporting the main conclusions of the present study. First, we present a mathematical interpretation of the axis-prior initialization strategy employed in AIRS, SAVER-A, and SAVER-AO. In particular, we discuss how the initial acquisition of orthogonal projection directions contributes to stabilization of the posterior covariance matrix, improvement of the conditioning of the sequential inverse problem, and reduction of variance estimation instability during the earliest acquisition rounds. We additionally discuss the numerical characteristics of the system matrix constructed using Joseph's reprojection method[1]. Particular attention is paid to the relationship between projection orientation, reprojection-induced smoothing effects, and the stability of early-stage sequential reconstruction. The supplementary figures provide detailed convergence behaviors and statistical comparisons under a high-noise condition. Specifically, we report:

- SSIM convergence curves as a function of acquisition rounds under a relatively higher noise condition ($\sigma_B = 10^{-2}$),
- distributions of normalized AUC values across independent trials.

These additional analyses further support the stability and effectiveness of the proposed SAVER framework and its axis-prior extensions within the investigated parameter ranges.

1 Why Axis-Prior Initialization is Effective in AIRS, SAVER-A, and SAVER-AO

In this supplementary section, we provide a mathematical interpretation of why the axis-prior initialization employed in AIRS, SAVER-A, and SAVER-AO contributes to efficient early-stage reconstruction in the proposed sequential CT acquisition framework.

Role of Early Measurements in Sequential Reconstruction

Recall that the reconstruction problem is formulated as the following Tikhonov-regularized inverse problem[2]: $\hat{\mathbf{x}}_t = \left(A_t^\top A_t + \frac{\sigma_B^2}{\xi^2} I\right)^{-1} A_t^\top \mathbf{y}_t$, where $A_t \in \mathbb{R}^{t\times d}$ denotes the measurement matrix composed of sequentially acquired projection rays. The corresponding posterior covariance matrix is $B_t^{-1} = \left(A_t^\top A_t + \frac{\sigma_B^2}{\xi^2} I\right)^{-1}$. The quality of the reconstruction is therefore strongly governed by the conditioning and rank structure of $A_t^\top A_t$. In the early acquisition stage ($t \ll d$), the matrix $A_t^\top A_t$ is highly rank-deficient, and thus the reconstruction quality depends critically on how efficiently the initial measurements span the image space.

In AIRS, SAVER-A, and SAVER-AO, all detector positions corresponding to 0° and 90° are first exhausted before entering the adaptive scheduling stage. Denote by $A_{\text{axis}} = \begin{pmatrix} A_{0^\circ} \\ A_{90^\circ} \end{pmatrix}$ the submatrix consisting only of projection rays from the horizontal and vertical directions. For a Cartesian image grid, the rows of A_{0° approximately correspond to horizontal line sums, while the rows of A_{90° correspond to vertical line sums. Consequently, $A_{\text{axis}}^\top A_{\text{axis}}$ provides early constraints on both horizontal and vertical spatial frequencies of the image. This can be interpreted as constructing an initial low-frequency basis for the posterior estimate. In contrast, purely random initialization may produce highly anisotropic sampling patterns especially during early rounds, resulting in unstable posterior covariance structures.

Reduction of Posterior Uncertainty

The recursive Woodbury update[3] employed in this study is

$$\begin{cases} D_t = 1 + \mathbf{a}_{i_t,r_t}^{\top} B_{t-1}^{-1} \mathbf{a}_{i_t,r_t} \\ \mathbf{k}_t = (B_{t-1}^{-1} \mathbf{a}_{i_t,r_t})/D_t \\ \hat{\mathbf{x}}_t = \hat{\mathbf{x}}_{t-1} + \mathbf{k}_t (y_t - \mathbf{a}_{i_t,r_t}^{\top} \hat{\mathbf{x}}_{t-1}) \\ B_t^{-1} = B_{t-1}^{-1} - \mathbf{k}_t (\mathbf{a}_{i_t,r_t}^{\top} B_{t-1}^{-1}) \end{cases} .$$

The reduction of posterior uncertainty caused by a newly-selected ray $\mathbf{a}_{i_t,r_t}$ is governed by $\mathbf{a}_{i_t,r_t}^{\top} B_{t-1}^{-1} \mathbf{a}_{i_t,r_t}$, which represents the uncertainty along the measurement direction. When the initial rays are concentrated only on a limited subset of directions, the posterior covariance remains highly elongated in unobserved spatial directions. Axis-prior initialization mitigates this effect by rapidly reducing uncertainty along the two principal Cartesian directions. As a result, the posterior covariance becomes more isotropic during early rounds, stabilizing subsequent adaptive scheduling.

Stabilization of Variance Estimation

The proposed SAVER framework selects projection angles using the sample variance $\hat{\sigma}_i^2(t) = \frac{1}{n_i(t)} \sum_{y \in Y_i(t)} y^2 - \left(\frac{1}{n_i(t)} \sum_{y \in Y_i(t)} y \right)^2$. When $n_i(t)$ is very small, the variance estimator becomes highly unstable: $\mathrm{Var}\left[\hat{\sigma}_i^2(t)\right] = O\left(\frac{1}{n_i(t)}\right)$. Therefore, during the earliest acquisition rounds, inaccurate variance estimates may induce erroneous Softmax probabilities: $P_i(t) = \frac{\exp(\tilde{\sigma}_i(t)/T_t)}{\sum_j \exp(\tilde{\sigma}_j(t)/T_t)}$. Axis-prior initialization functions as a deterministic warm-start procedure that stabilizes the initial posterior estimate before the adaptive variance-driven phase begins. This can reduce the probability that the scheduling policy becomes trapped in suboptimal projection directions caused by noisy early-stage variance estimates.

Relation Between AIRS, SAVER-A, and SAVER-AO

The three methods differ from each other on how the method selects the irradiation angles after the axis-prior initialization. After the initialization, AIRS performs uniform random sampling while SAVER-A utilizes adaptive variance-driven scheduling that prioritizes exploration at earlier rounds by estimating underlying variance of projection data values along each angle by the sample variances and gradually shifts exploitation as the number of rounds passes. SAVER-AO uses axis-prior initialization followed by oracle-guided Softmax scheduling using the true variances. Thus, AIRS isolates the contribution of the deterministic initialization itself, while SAVER-A and SAVER-AO additionally exploit adaptive concentration toward informative projection angles. The empirical gap between AIRS and SAVER-A therefore directly quantifies the utility of real-time variance estimation beyond simple axis-prior stabilization, and also a higher performance of SAVER-AO than SAVER-A observed for phantoms especially *3. Triangle* and *7. Gradient* in Fig. 5 indicates that tighter estimation of the projection data values' variance could result in better reconstruction performance.

2 Reconstruction Performance under Relatively Higher Noise

This section evaluates the behavior of the proposed SAVER framework on a higher noise level σ_B (10^4 times larger) than that presented in the main text. Here, we present corresponding results for the representative setting ($\sigma_B = 10^{-2}, \eta = 0.01$) in Figures S1 to S3.

As shown in Figures S1 and S2, SAVER-based methods outperform its non-adaptive baseline (c.f., SAVER and SAVER-O against Random, and SAVER-A and SAVER-AO against AIRS) for phantoms with high structural anisotropy (e.g., *1. Rectangle*, *4. Cross*, and *6. Stripes*) or show comparable performance for most of the other phantoms with more homogeneous information distributions as seen in the lower noise case. For latter phantoms, SAVER-A exhibits performance comparable to AIRS, indicating that the deterministic axis-prior provides the primary contribution to reconstruction stability in such cases. A higher performance of SAVER-AO than SAVER-A was also observed for *3. Triangle* and *7. Gradient*, which suggests some improvement if a tighter estimation of the projection data values' variance is possible as seen in a lower noise level.

Figure S3 highlights the specific performance gap between random sampling, deterministic axis-prior initialization, and our adaptive scheduling. These results confirm that the performance trends observed in the low-noise environment (main text) are preserved even when the measurement noise is increased by four orders of magnitude ($\sigma_B = 10^{-2}$). The synergy between orthogonal initial measurements and variance-driven exploration continues to yield efficient reconstruction, particularly for structured objects, demonstrating the stability of the SAVER-A framework against stochastic perturbations.

Sensitivity Analysis on Annealing Hyperparameter

We also conducted extensive evaluations by varying the annealing parameter $\eta = 0.1, 1.0$, and 10. In these settings, SAVER-A maintained its superiority over the other methods that do not use the information of true variance. Axis-prior initialization during

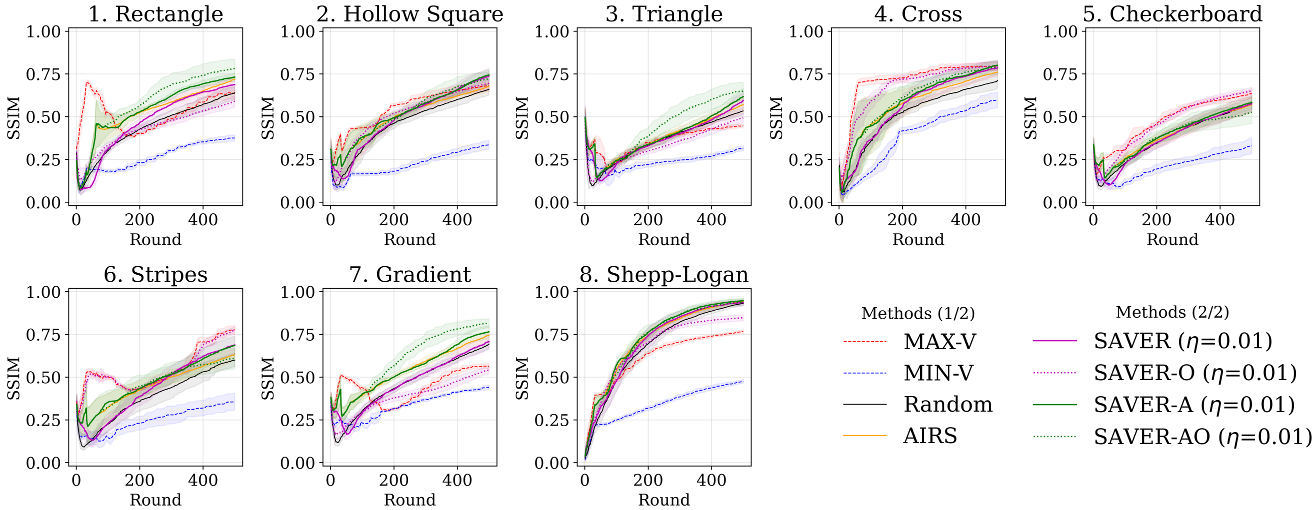

Supplementary Figure S1. SSIM convergence for all methods ($\sigma_B = 10^{-2}, \eta = 0.01$)**.** Shaded regions indicate standard deviation.

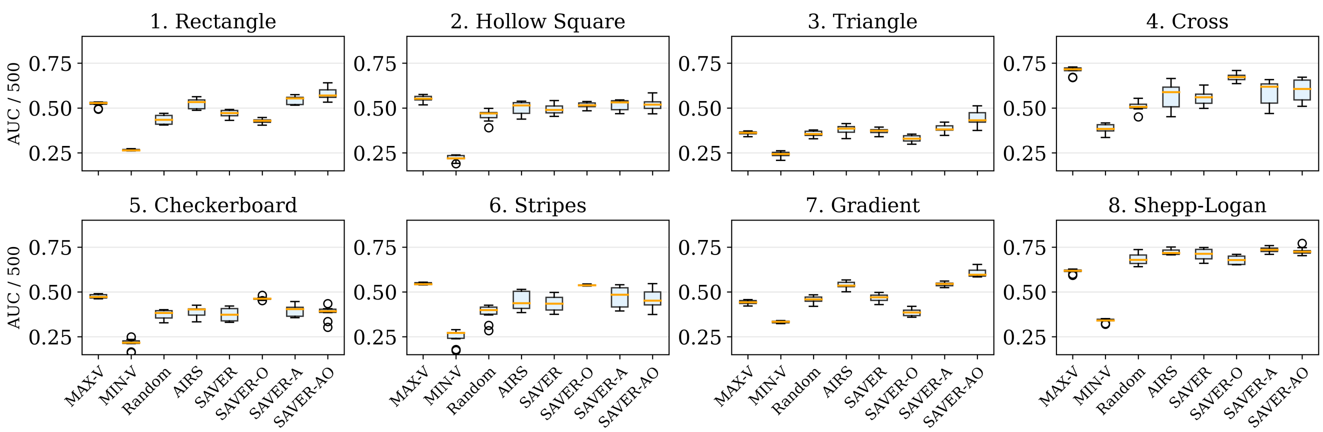

Supplementary Figure S2. Distribution of AUC/500 values ($\sigma_B = 10^{-2}, \eta = 0.01$)**.** Boxplots represent 10 independent trials for each method.

the earliest acquisition rounds ($t < 100$) was seen against non-adaptive baseline Random. This suggests that the axis-prior serves as a stability warm-start mechanism across the tested signal-to-noise ratios and hyperparameter settings. Figures S4, S5, and S6 exemplified the SSIM convergence with the noise level $\sigma_B = 10^{-6}$ for four methods, Random, AIRS, SAVER-A, and SAVER-AO. The performance gap among SAVER-A, SAVER-AO, and AIRS across different η values indicates that variance-driven exploration provides a meaningful advantage for anisotropic samples, although this gain is naturally moderated in isotropic cases: e.g., compare the isotropic phantoms such as *5. Checkerboard* and *Shepp-Logan* and the other aisotropic phantoms over ηs. For dependency on annealing parameter η of SAVER-A performance, it should be noted that SAVER-A gets closer to the oracle SAVER-AO for some of the phantoms such as *3. Triangle* and *7. Gradient* as η increases from $\eta = 0.01$ to $\eta = 0.1$ and 1.0 in latter rounds > 300 although extremely high annealing rate $\eta = 10$ suppresses this improvement. This may indicate that while most phantoms in this paper have less dependency on the annealing rates, there exists room to optimize the annealing schedule in practical applications. This is one of the forthcoming subjects to be unresolved.

References

1. Joseph, P. M. An improved algorithm for reprojecting rays through pixel images. *IEEE transactions on medical imaging* **1**, 192–196 (2007).

2. Tikhonov, A. N. *et al.* On the stability of inverse problems. In *Dokl. akad. nauk sssr*, vol. 39, 195–198 (1943).

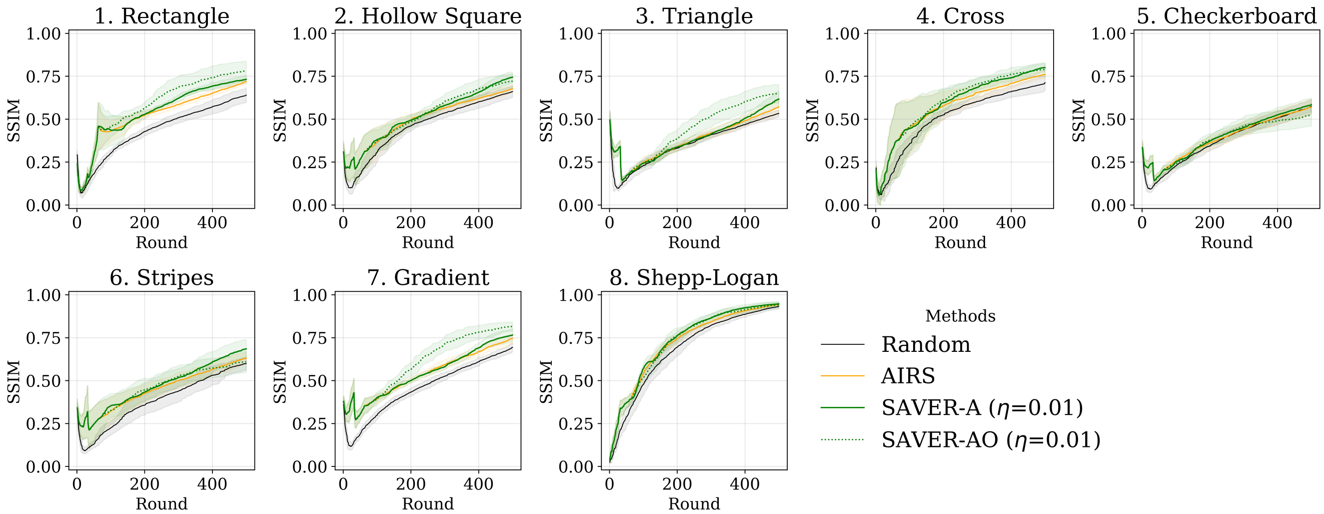

Supplementary Figure S3. Detailed SSIM convergence comparison of representative methods ($\sigma_B = 10^{-2}, \eta = 0.01$). This figure highlights the performance gap between purely random sampling (Random), axis-prior initialization (AIRS), and the proposed adaptive variance-driven scheduling (SAVER-A/AO). Shaded regions indicate the standard deviation across 10 independent trials.

3. Sherman, J. & Morrison, W. J. Adjustment of an inverse matrix corresponding to a change in one element of a given matrix. *The Annals Math. Stat.* **21**, 124–127 (1950).

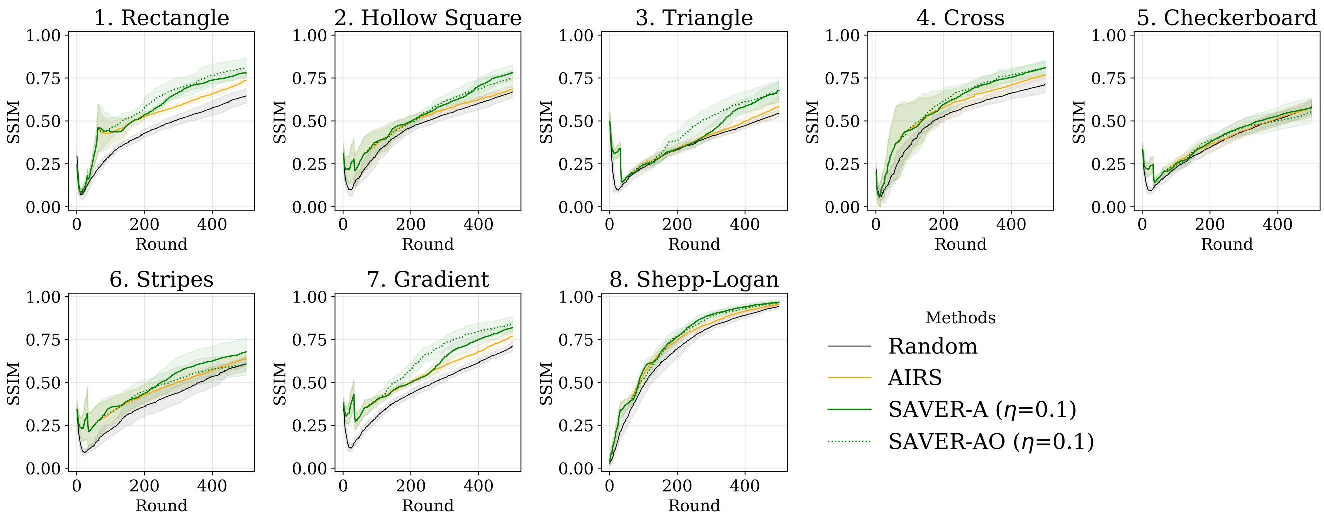

Supplementary Figure S4. Detailed SSIM convergence comparison of representative methods ($\sigma_B = 10^{-6}, \eta = 0.1$).

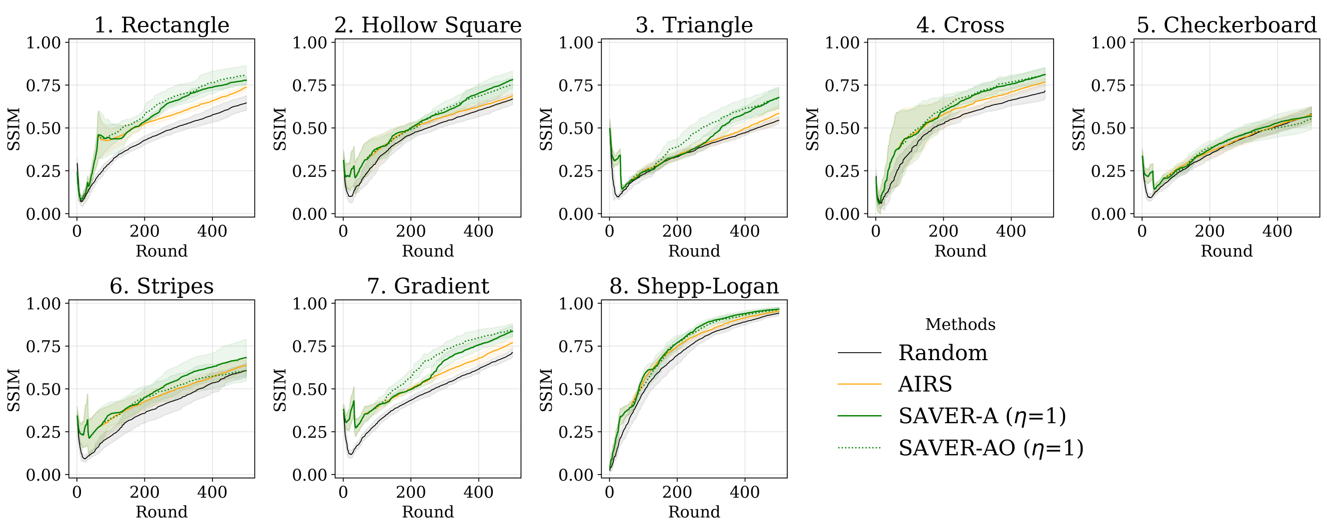

Supplementary Figure S5. Detailed SSIM convergence comparison of representative methods ($\sigma_B = 10^{-6}, \eta = 1$).

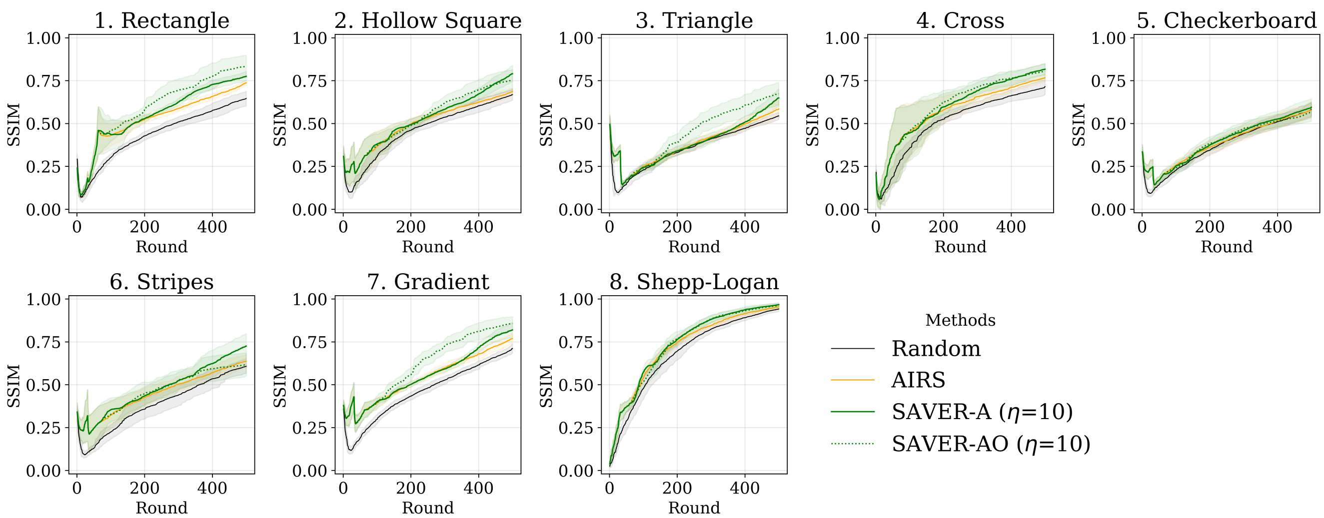

Supplementary Figure S6. Detailed SSIM convergence comparison of representative methods ($\sigma_B = 10^{-6}, \eta = 10$).