MIAT-RPCA: TRAINING-FREE ROBUST PCA VIA MANIFOLD-INSPIRED INTERPOLATION AND ADAPTIVE TRUNCATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Thresholding-based robust principal component analysis (RPCA) separates sparse corruption from low-rank structure. In basic CUR updates, this process does not explicitly exploit local coefficient neighborhoods to refine corrupted observations. Yet observations affected by different corruption levels may retain shared structure that can support mutual refinement. Using this information requires considering both reconstruction inconsistency and neighborhood agreement. Motivated by this observation, we propose Manifold-Inspired Interpolation and Adaptive Truncation for RPCA (MIAT-RPCA), a training-free method that incorporates neighborhood information into sampled CUR recovery. Cosine similarities define local neighborhoods in coefficient space. A quantile-normalized residual score controls interpolation strength, allowing columns with greater reconstruction inconsistency to draw more heavily on neighboring coefficients. A complementary attenuation gate combines this score with neighborhood confidence, applying stronger suppression when inconsistency is high and neighborhood agreement is low. Both operations precede reconstruction and keep the corrected coefficients within their original span. We establish that the correction preserves the target rank bound and satisfies a neighborhood norm bound, characterize the gates' monotonic responses, and analyze the computational cost per iteration. All neighborhood weights and gates are computed from the current observation, with a prescribed threshold schedule and no training data or parameter learning.

Keywords: robust principal component analysis; CUR decomposition; manifold-inspired interpolation; adaptive truncation; image restoration; background separation.

1 INTRODUCTION

Robust principal component analysis (RPCA) separates an observation matrix into a low-rank component and a sparse component, supporting applications such as image restoration and video background separation. Convex formulations and nonconvex iterative methods offer different approaches to robust recovery (Candès et al., 2011; Xu et al., 2012; Netrapalli et al., 2014; Tong et al., 2021). In sampled CUR iterations, residual thresholding estimates sparse corruption, while selected rows and columns provide the low-rank reconstruction (Cai et al., 2021a). This basic update does not include an explicit interpolation step among neighboring coefficient vectors, leaving an opportunity to use local relationships directly in the correction.

Our starting point is that observations affected by different corruption levels may still retain shared structure. Their coefficient neighborhoods can therefore provide information for refinement, even when individual observations are inconsistent with the current reconstruction. A residual score describes this inconsistency, while neighborhood agreement indicates how closely a column aligns with related coefficient vectors. Non-local averaging, graph regularization, and local-information RPCA already exploit relationships between observations in different ways (Buades et al., 2011; Yin et al., 2016; Gao et al., 2024). Within sampled CUR recovery, we investigate how to use such information directly to refine coefficients before reconstruction. This leads to the question: how should a coefficient update combine reconstruction inconsistency with neighborhood confidence?

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

We address this question with MIAT-RPCA. Cosine affinities define CUR coefficient neighborhoods and a weighted estimate for each column. A residual score sets the interpolation strength, while neighborhood confidence modulates attenuation. This coupling draws more heavily on neighbors for inconsistent columns and applies stronger attenuation when neighborhood agreement is also low.

The correction acts before reconstruction, using directly computed weights and gates and a prescribed threshold schedule. It requires no training data or parameter learning. Our analysis establishes rank and norm properties and characterizes computational cost; Figure 1 summarizes the workflow.

Our contributions are twofold:

- **Adaptive coefficient correction.** We incorporate explicit neighborhood interpolation into sampled CUR recovery and couple its strength with confidence-modulated attenuation through a shared residual score.
- **Structural and computational analysis.** We prove that a correction step preserves the target rank bound and satisfies a neighborhood norm bound. We also characterize the gates' monotonic responses and the computational cost per iteration.

2 Related Work

Robust low-rank recovery and sampled CUR. RPCA solvers range from convex formulations and augmented Lagrangian methods (Candès et al., 2011; Lin et al., 2010) to alternating projections and scaled gradient descent (Netrapalli et al., 2014; Tong et al., 2021). Related approaches adapt the rank estimate (Xu et al., 2021) or incorporate prior subspace information (Chen et al., 2025). Sampled CUR reduces the size of the low-rank update: IRCUR uses selected rows and columns (Cai et al., 2021a), and eRPCA combines CUR with block Krylov iteration (Fang et al., 2023). Recovery theory studies sparsity and incoherence conditions (Cai et al., 2021b), while tensor CUR extends sampling to low-Tucker-rank recovery (Cai et al., 2024). MIAT-RPCA augments the sampled matrix representation with explicit neighborhood interpolation before reconstruction.

Neighborhood and geometric information. Local relationships support recovery through similarity-weighted averaging (Buades et al., 2011), graph regularization (Yin et al., 2016), fuzzy local information (Gao et al., 2024), and adaptive sample weighting (Zhang & Tong, 2019). Nonlinear RPCA also models a manifold component and studies tangent-space incoherence (Lyu et al., 2019). These approaches motivate using relationships retained across corrupted observations. In MIAT-RPCA, neighborhood information enters through interpolation of CUR coefficient vectors. Residual scores determine the interpolation strength, and neighborhood confidence modulates attenuation, linking local agreement to reconstruction feedback.

Learned recovery and restoration priors. Learned approaches include unfolded RPCA (Cai et al., 2021c), robust matrix completion (Cai et al., 2026), and neural restoration architectures such as Restormer, NAFNet, and PromptIR (Zamir et al., 2022; Chen et al., 2022; Potlapalli et al., 2023). Plug-and-play methods embed denoisers in iterative reconstruction (Zhang et al., 2022); gradient-step and proximal denoisers support convergence analyses for specified schemes (Hurault et al., 2022a;b). Kernel-PCA interpretations also connect self-attention with robust principal-component attention (Teo & Nguyen, 2024). Deep Image Prior and Zero-Shot Noise2Noise fit networks to individual observations (Ulyanov et al., 2018; Mansour & Heckel, 2023). MIAT-RPCA instead computes its correction directly from the current observation using preset hyperparameters. This places its contribution in the design and analysis of the coefficient update within sampled CUR.

3 Manifold-Inspired Interpolation and Adaptive Truncation

Overview. Figure 1 summarizes how MIAT-RPCA uses local coefficient relationships within sampled CUR recovery. Each iteration forms the representation, computes residual scores and neighborhood information, applies interpolation and attenuation, and reconstructs the low-rank output. Appendix C collects the threshold definition, neighborhood-selection rule, and default settings.

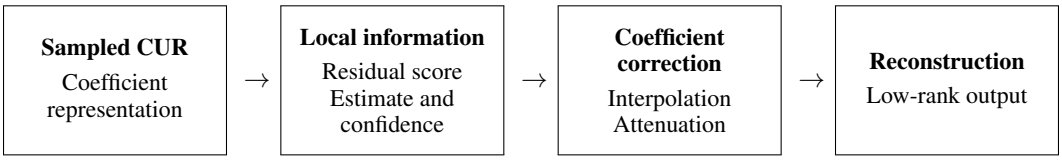

Figure 1: MIAT-RPCA workflow. Sampled CUR supplies a coefficient representation. Residual scores and cosine neighborhoods determine the interpolated estimate and correction strength. Interpolation and attenuation then act on the coefficients before reconstruction.

3.1 Sampled CUR representation

Let $X \in \mathbb{R}^{m\times n}$ be the observation matrix, r the target rank, and H the number of iterations. We seek $X \approx L + S$, where L is low-rank and S contains thresholded residuals. Initialize $L = S = 0$. At iteration $\ell = 0, \ldots, H-1$, sample row and column index sets I, J of sizes

$$n_r = \min(\lfloor cr \log m \rfloor, m), \qquad n_c = \min(\lfloor cr \log n \rfloor, n), \quad c = 4. \tag{1}$$

Sampling uses random permutations without replacement, with new index sets at each iteration. Apply the smooth threshold $h_{\theta,\bar{\theta}}$ defined in Appendix C to the sampled entries of $X - L$, updating $S_{:,J}$ and $S_{I,:}$. Form the sampled matrices and their intersection:

$$C = X_{:,J} - S_{:,J}, \qquad R = X_{I,:} - S_{I,:}, \qquad U = R_{:,J}. \tag{2}$$

For nonzero leading singular values, let the first r singular triplets of U be (P_r, Σ_r, Q_r). Define

$$U_r^\dagger = Q_r \Sigma_r^{-1} P_r^\top, \qquad Z = U_r^\dagger R \in \mathbb{R}^{n_c \times n}. \tag{3}$$

This construction follows sampled CUR recovery (Cai et al., 2021a;b). The factor C maps coefficients back to the observation space, and Z supplies the representation for the neighborhood correction. The target rank satisfies $1 \le r \le \min(n_r, n_c)$.

3.2 Residual score and neighborhood confidence

Residual score. The residual on the sampled rows measures reconstruction inconsistency. For column k, define the residual e_k and normalized residual score o_k as

$$e_k = \frac{1}{n_r} \sum_{i=1}^{n_r} |R_{ik} - (C_{I,:}Z)_{ik}|, \tag{4}$$

$$o_k = \mathrm{clip}_{[0,1]}\left(\frac{e_k - Q_{0.5}(e)}{\max\{Q_{0.9}(e) - Q_{0.5}(e), 10^{-8}\}} \right), \tag{5}$$

where Q_q denotes the empirical q-quantile. The median sets a reference level, and the upper quantile scales larger residuals. The score is recomputed at each iteration and measures inconsistency relative to the current set of columns.

Neighborhood estimate and confidence. Let $z_k = Z_{:,k}$ and $\tilde{z}_k = z_k / \max(\|z_k\|_2, 10^{-12})$. The cosine affinity $a_{kj} = \tilde{z}_k^\top \tilde{z}_j$ measures their agreement. We construct $\mathcal{N}_k$ from the highest-ranked affinities using the selection rule in Appendix C. Define the interpolation weights, interpolated vector, and neighborhood confidence by

$$w_{kj} = \frac{\exp(a_{kj}/t)}{\sum_{q\in\mathcal{N}_k} \exp(a_{kq}/t)}, \qquad \widehat{z}_k = \sum_{j\in\mathcal{N}_k} w_{kj} z_j, \tag{6}$$

$$c_k = \mathrm{sigmoid}\left(10 \left[\frac{1}{|\mathcal{N}_k|} \sum_{j\in\mathcal{N}_k} a_{kj} - \tau_d \right] \right). \tag{7}$$

Temperature t controls the concentration of interpolation weights, and τ_d centers the confidence response. Thus c_k summarizes neighborhood agreement, while o_k measures reconstruction inconsistency. Candidate neighbors span all coefficient columns; for RGB sequences, each column represents a frame–channel observation, as detailed in Appendix B.

3.3 Adaptive interpolation and truncation

The mixing coefficient β_k sets the interpolation strength. Adaptive truncation is implemented as a multiplicative attenuation factor T_k, applied before reconstruction:

$$\beta_k = \beta_l + (\beta_h - \beta_l)o_k, \qquad T_k = 1 - [o_k(1 - c_k)]^\eta, \tag{8}$$

$$z_k^+ = T_k[(1 - \beta_k)z_k + \beta_k \widehat{z}_k], \qquad L^+ = CZ^+, \tag{9}$$

The two operations share the residual score: increasing o_k gives more weight to neighboring coefficients through β_k and strengthens attenuation through T_k when neighborhood confidence is low. High confidence moderates this attenuation. Algorithm 1 gives the full iteration; Appendix C lists the default gate settings.

Algorithm 1 MIAT-RPCA: manifold-inspired interpolation and adaptive truncation

Require: X, rank r, iterations H, initial threshold $\theta_0 > 0$, neighborhood/gate settings.

1: Initialize $L = S = 0$; prescribe $\theta_\ell = \theta_0(0.7)^\ell$, $\bar{\theta}_\ell = 1.5\theta_\ell$.
2: **for** $\ell = 0, \ldots, H - 1$ **do**
3: Sample I, J and threshold $(X - L)_{:,J}$ and $(X - L)_{I,:}$ into S.
4: Form C, R, U and $Z = U_r^\dagger R$ using Eq. (3).
5: Compute column scores with Eq. (5).
6: Compute neighbors, interpolation, and confidence with Eqs. (6)–(7).
7: Apply Eqs. (8)–(9) and replace L by L^+.
8: **end for**
9: **return** Final L, S. Entries of S outside the sampled blocks retain earlier values.

3.4 Structural properties

We characterize how each gate responds to its inputs and establish the norm and rank properties of the coefficient correction. With the default exponent $\eta = 2$, the gate derivatives are

$$\frac{\partial \beta_k}{\partial o_k} = \beta_h - \beta_l \geq 0, \qquad \frac{\partial T_k}{\partial o_k} = -2o_k(1 - c_k)^2 \leq 0, \qquad \frac{\partial T_k}{\partial c_k} = 2o_k^2(1 - c_k) \geq 0. \tag{10}$$

Thus a larger residual score increases neighbor mixing and strengthens attenuation, while higher confidence moderates attenuation. At $o_k = 0$, $T_k = 1$ and mixing remains at the baseline level β_l. As c_k approaches one, T_k approaches one even for a high-score column.

Proposition 1 (Neighborhood norm bound and rank preservation)**.** *Assume Eq. (3) is well defined, $\mathcal{N}_k$ is nonempty, $t > 0$, and $0 \leq \beta_l \leq \beta_h \leq 1$. For $o_k, c_k \in [0, 1]$ and $\eta > 0$, the update satisfies*

$$\left\|z_k^+\right\|_2 \leq T_k \left[(1 - \beta_k)\left\|z_k\right\|_2 + \beta_k \sum_{j \in \mathcal{N}_k} w_{kj}\left\|z_j\right\|_2\right] \leq \max_{j \in \{k\} \cup \mathcal{N}_k} \left\|z_j\right\|_2 . \tag{11}$$

Moreover, $Z^+ = ZB$ for an input-dependent matrix B, and $\operatorname{rank}(L^+) \leq r$.

The norm bound follows from convex mixing and bounded attenuation. The representation $Z^+ = ZB$ keeps the corrected coefficients within the original coefficient span, so the reconstructed matrix satisfies the target rank bound. Appendix A provides the proof.

3.5 Computational cost

Sampling restricts the SVD to the $n_r \times n_c$ intersection matrix. Computing all cosine affinities costs $O(n^2 n_c)$ operations and requires $O(n^2)$ storage. Aggregating K neighbors per column costs $O(nKn_c)$ operations, and explicit reconstruction $L = CZ^+$ costs $O(mn_c n)$. The pairwise affinity calculation therefore depends quadratically on the number of observations, while aggregation depends on the selected neighborhood size. This breakdown separates the costs of coefficient comparison, interpolation, and reconstruction.

4 EXPERIMENTS

4.1 EXPERIMENTAL SETUP

We consider retinal image restoration and video background separation in traffic, shopping-mall, and indoor scenes. These settings involve observations that share persistent structure despite corruption or scene changes. The evaluation is organized around reconstruction quality, separation of the low-rank component L and residual component S, and execution time.[1]

The comparison includes IALM (Lin et al., 2010), AltProj (Netrapalli et al., 2014), ScaledGD (Tong et al., 2021), LRPCA (Cai et al., 2021c), eRPCA (Fang et al., 2023), and BPCA (Chen et al., 2025), covering convex optimization, nonconvex recovery, learned unfolding, sampled CUR, and prior subspace information.

Reconstruction quality is summarized using MSE, Euclidean-normalized NRMSE, PSNR, and SSIM. Metrics are computed per RGB image or frame and averaged separately, within each corruption group where applicable. Appendix B describes the data arrangement, metric implementation, and aggregation procedure.

4.2 IMAGE RESTORATION

Table 1 summarizes reconstruction quality at the two corruption levels. At $\rho = 0.5$, the MIAT-RPCA entry has the highest listed PSNR (24.76 dB) and matches LRPCA in MSE, NRMSE, and SSIM at the reported precision. At $\rho = 0.3$, its MSE is 0.0049 and SSIM is 0.93. The two groups summarize the restoration outcomes under each corruption setting.

The visual comparison complements these aggregate metrics. In Figure 2, the recovered images suppress visible block artifacts while retaining the large-scale retinal structure. Figure 3 presents the baseline reconstructions; its PSNR and SSIM annotations describe the displayed images, whereas Table 1 reports separately averaged metrics.

Table 1: Image restoration across corruption levels. Metrics concern the low-rank output L; arrows indicate the preferred direction. Best values within each corruption level are bold, with ties determined at the reported precision. MSE and PSNR are averaged separately.

Method	MSE ↓	NRMSE ↓	PSNR ↑	SSIM ↑
$\rho = 0.3$				
IALM	0.012	0.24	19.55	0.88
AltProj	0.0055	0.15	23.50	0.90
ScaledGD	0.0059	0.14	23.01	0.92
LRPCA	**0.0024**	**0.089**	26.61	**0.96**
eRPCA	0.0062	0.14	**30.02**	**0.96**
BPCA	0.0071	0.16	21.62	0.91
MIAT-RPCA	0.0049	0.13	25.21	0.93
$\rho = 0.5$				
IALM	0.072	0.50	11.70	0.75
AltProj	0.061	0.47	12.17	0.70
ScaledGD	0.049	0.41	13.01	0.81
LRPCA	**0.011**	**0.22**	19.61	**0.91**
eRPCA	0.070	0.51	11.53	0.80
BPCA	0.074	0.51	11.34	0.42
MIAT-RPCA	**0.011**	**0.22**	**24.76**	**0.91**

[1] Figures, metrics, and timings come from prior experimental records. Their correspondence to the prescribed-threshold variant has not yet been verified; no new benchmark runs are reported here.

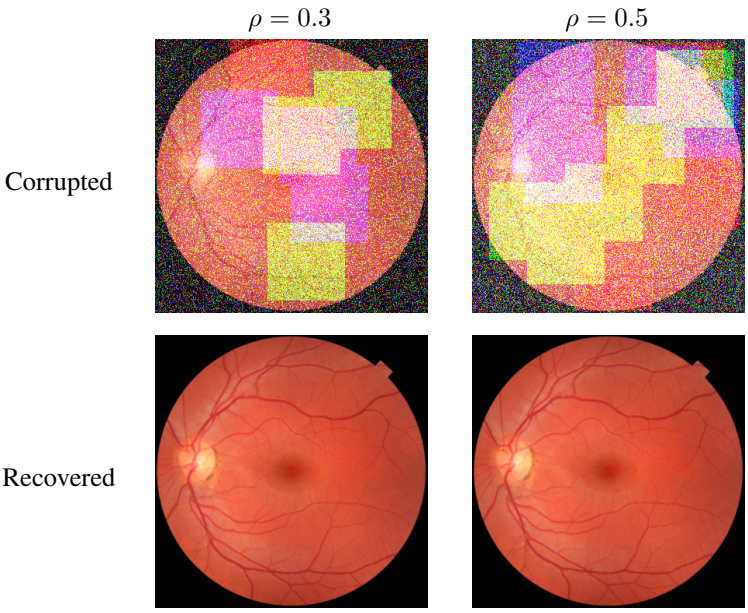

Figure 2: Retinal restoration at $\rho = 0.3$ (left) and $\rho = 0.5$ (right). Rows show corrupted inputs and recovered images, respectively.

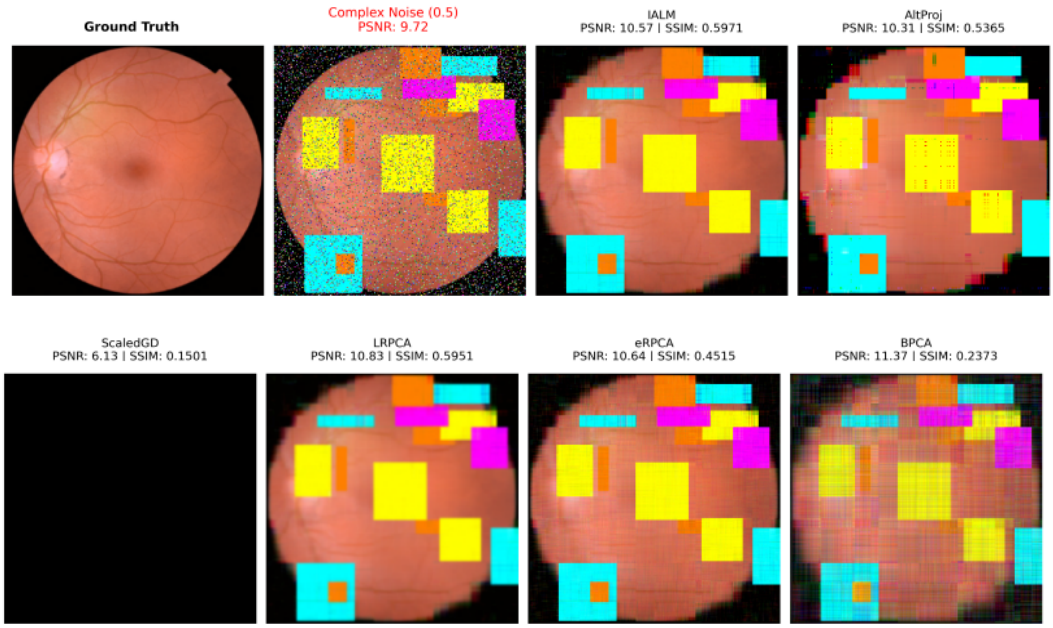

Figure 3: Retinal baseline reconstructions for the displayed example at $\rho = 0.5$. Panels show the reference image, corrupted input, and six comparison methods. PSNR and SSIM annotations refer to the individual images shown.

4.3 Video Background Separation

Traffic and shopping-mall scenes. The traffic and shopping-mall examples in Figure 4 separate persistent scene structure from changing regions. The low-rank panels retain the background, while the residual panels emphasize changes across the displayed frames.



Figure 4: Background separation in traffic (left) and shopping-mall (right) scenes. Within each panel, rows show input frames, low-rank outputs L, and residuals S for the first five frames, followed by the next five frames.

Indoor scene. In Figure 5, the low-rank output retains the static room layout, and the residual contains prominent person-shaped regions. Figure 6 shows the associated baseline reconstructions.



Figure 5: Indoor background separation. Rows show input frames, low-rank backgrounds L, and residuals S at the labeled frame indices.



Figure 6: Indoor baseline reconstructions. The two rows contain the reference image and outputs from six comparison methods.

Table 2 summarizes the indoor reconstruction metrics. The MIAT-RPCA entry has the lowest listed MSE (0.0030) and NRMSE (0.1054), together with the highest PSNR (25.295 dB) and SSIM (0.9538). The visual examples separately illustrate background structure and residual patterns.

Table 2: Recorded reconstruction metrics and execution times for the indoor sequence. Best values are bold; time is reported in seconds.

Method	Time (s)	MSE ↓	NRMSE ↓	PSNR ↑	SSIM ↑
IALM	57.28	0.0041	0.1184	24.683	0.9308
AltProj	83.85	0.0345	0.3595	14.621	0.8156
ScaledGD	98.65	0.2909	1.0428	5.376	0.0002
LRPCA	55.30	0.0047	0.1299	23.695	0.9261
eRPCA	95.88	0.0055	0.1398	23.058	0.9173
BPCA	1.72	0.0079	0.1670	21.581	0.8894
MIAT-RPCA	**0.11**	**0.0030**	**0.1054**	**25.295**	**0.9538**

4.4 Runtime comparison

Table 2 also lists the recorded execution times, including 0.11 seconds for the MIAT-RPCA entry. Section 3.5 complements the timing record with an operation-count analysis, separating the contributions of coefficient comparison, neighborhood aggregation, and reconstruction.

5 Conclusion

We presented MIAT-RPCA to incorporate neighborhood information retained across corrupted observations into sampled CUR recovery. Its central contribution is an explicit coefficient correction that couples residual-guided interpolation with confidence-modulated attenuation. This design uses reconstruction inconsistency to regulate information drawn from neighboring coefficients and neighborhood agreement to moderate attenuation before reconstruction. The correction preserves the target rank bound and satisfies a neighborhood norm bound. Together with the gate-response and computational analyses, these properties characterize how local information can be integrated into the low-rank representation without training data or parameter learning.

Future work will explore recovery guarantees across iterations and applications to broader data settings.

AI USE STATEMENT

AI use in the original manuscript was limited to improving the language style of selected passages. During revision, an AI assistant additionally supported manuscript restructuring, English editing, literature search and citation integration, reference checking, LaTeX formatting, figure arrangement, code inspection, diagnostic scripting, and formulation and checking of the equations and structural arguments. The numerical benchmark values were retained from the original experimental record.

REPRODUCIBILITY STATEMENT

Section 3 and Algorithm 1 specify the coefficient update. Appendix C collects the threshold schedule, default settings, and neighborhood-selection rule; Appendix A provides the structural proof; and Appendix B describes the data arrangement and metric aggregation.

ETHICS STATEMENT

The retinal image is used as an image-restoration example; the study makes no clinical claims. The visual examples are used to illustrate low-rank decomposition and background separation.

REFERENCES

Antoni Buades, Bartomeu Coll, and Jean-Michel Morel. Non-local means denoising. *Image Processing On Line*, 1:208–212, 2011. doi: 10.5201/ipol.2011.bcm_nlm.

HanQin Cai, Keaton Hamm, Longxiu Huang, Jiaqi Li, and Tao Wang. Rapid robust principal component analysis: CUR accelerated inexact low rank estimation. *IEEE Signal Processing Letters*, 28:116–120, 2021a. doi: 10.1109/LSP.2020.3044130.

HanQin Cai, Keaton Hamm, Longxiu Huang, and Deanna Needell. Robust CUR decomposition: Theory and imaging applications. *SIAM Journal on Imaging Sciences*, 14(4):1472–1503, 2021b. doi: 10.1137/20M1388322.

HanQin Cai, Jialin Liu, and Wotao Yin. Learned robust PCA: A scalable deep unfolding approach for high-dimensional outlier detection. In *Advances in Neural Information Processing Systems*, volume 34, pp. 16977–16989, 2021c.

HanQin Cai, Zehan Chao, Longxiu Huang, and Deanna Needell. Robust tensor CUR decompositions: Rapid low-Tucker-rank tensor recovery with sparse corruptions. *SIAM Journal on Imaging Sciences*, 17(1):225–247, 2024. doi: 10.1137/23M1574282.

HanQin Cai, Chandra Kundu, Jialin Liu, and Wotao Yin. Deeply learned robust matrix completion for large-scale low-rank data recovery. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 48(6):6541–6556, 2026. doi: 10.1109/TPAMI.2026.3659041.

Emmanuel J. Candès, Xiaodong Li, Yi Ma, and John Wright. Robust principal component analysis? *Journal of the ACM*, 58(3):11:1–11:37, 2011. doi: 10.1145/1970392.1970395.

Liangyu Chen, Xiaojie Chu, Xiangyu Zhang, and Jian Sun. Simple baselines for image restoration. In *European Conference on Computer Vision*, volume 13667 of *Lecture Notes in Computer Science*, pp. 17–33, 2022. doi: 10.1007/978-3-031-20071-7_2.

Lin Chen, Li Ge, Xue Jiang, and Hongbin Li. Boosting RPCA by prior subspace. *IEEE Transactions on Signal Processing*, 73:2170–2186, 2025. doi: 10.1109/TSP.2025.3569861.

Shun Fang, Zhengqin Xu, Shiqian Wu, and Shoulie Xie. Efficient robust principal component analysis via block Krylov iteration and CUR decomposition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1348–1357, 2023. doi: 10.1109/CVPR52729.2023.00136.

Yunlong Gao, Xinjing Wang, Jiaxin Xie, Jinyan Pan, Peng Yan, and Feiping Nie. Robust principal component analysis based on fuzzy local information reservation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(12):9321–9337, 2024. doi: 10.1109/TPAMI.2024.3418983.

Samuel Hurault, Arthur Leclaire, and Nicolas Papadakis. Gradient step denoiser for convergent plug-and-play. In *International Conference on Learning Representations*, 2022a. URL https://openreview.net/forum?id=fPhKeld3Okz.

Samuel Hurault, Arthur Leclaire, and Nicolas Papadakis. Proximal denoiser for convergent plug-and-play optimization with nonconvex regularization. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 9483–9505. PMLR, 2022b.

Zhouchen Lin, Minming Chen, and Yi Ma. The augmented Lagrange multiplier method for exact recovery of corrupted low-rank matrices. *arXiv preprint arXiv:1009.5055*, 2010. URL https://arxiv.org/abs/1009.5055.

He Lyu, Ningyu Sha, Shuyang Qin, Ming Yan, Yuying Xie, and Rongrong Wang. Manifold denoising by nonlinear robust principal component analysis. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

Youssef Mansour and Reinhard Heckel. Zero-shot Noise2Noise: Efficient image denoising without any data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14018–14027, 2023.

Praneeth Netrapalli, U. N. Niranjan, Sujay Sanghavi, Animashree Anandkumar, and Prateek Jain. Non-convex robust PCA. In *Advances in Neural Information Processing Systems*, volume 27, 2014.

Vaishnav Potlapalli, Syed Waqas Zamir, Salman H. Khan, and Fahad Shahbaz Khan. PromptIR: Prompting for all-in-one image restoration. In *Advances in Neural Information Processing Systems*, volume 36, pp. 71275–71293, 2023. doi: 10.52202/075280-3121.

Rachel S. Y. Teo and Tan M. Nguyen. Unveiling the hidden structure of self-attention via kernel principal component analysis. In *Advances in Neural Information Processing Systems*, volume 37, pp. 101393–101427, 2024. doi: 10.52202/079017-3215.

Tian Tong, Cong Ma, and Yuejie Chi. Accelerating ill-conditioned low-rank matrix estimation via scaled gradient descent. *Journal of Machine Learning Research*, 22(150):1–63, 2021.

Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Deep image prior. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 9446–9454, 2018.

Stéfan van der Walt, Johannes L. Schönberger, Juan Nunez-Iglesias, François Boulogne, Joshua D. Warner, Neil Yager, Emmanuelle Gouillart, Tony Yu, and the scikit-image contributors. scikit-image: image processing in Python. *PeerJ*, 2:e453, 2014. doi: 10.7717/peerj.453.

Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4): 600–612, 2004. doi: 10.1109/TIP.2003.819861.

Huan Xu, Constantine Caramanis, and Sujay Sanghavi. Robust PCA via outlier pursuit. *IEEE Transactions on Information Theory*, 58(5):3047–3064, 2012. doi: 10.1109/TIT.2011.2173156.

Zhengqin Xu, Rui He, Shoulie Xie, and Shiqian Wu. Adaptive rank estimate in robust principal component analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6577–6586, 2021.

Ming Yin, Junbin Gao, and Zhouchen Lin. Laplacian regularized low-rank representation and its applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(3):504–517, 2016. doi: 10.1109/TPAMI.2015.2462360.

Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Restormer: Efficient transformer for high-resolution image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5728–5739, 2022.

Kai Zhang, Yawei Li, Wangmeng Zuo, Lei Zhang, Luc Van Gool, and Radu Timofte. Plug-and-play image restoration with deep denoiser prior. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(10):6360–6376, 2022. doi: 10.1109/TPAMI.2021.3088914.

Rui Zhang and Hanghang Tong. Robust principal component analysis with adaptive neighbors. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

A Proof of the Structural Properties

For Proposition 1, $o_k, c_k \in [0,1]$ and $\eta > 0$ imply $T_k \in [0,1]$. The interpolation gate satisfies $\beta_k \in [\beta_l, \beta_h] \subseteq [0,1]$. For $t > 0$ and nonempty $\mathcal{N}_k$, the softmax weights satisfy $w_{kj} \geq 0$ and $\sum_{j\in\mathcal{N}_k} w_{kj} = 1$. Applying the triangle inequality to Eq. (9) gives the first inequality in Eq. (11). Replacing each norm by $M_k = \max_{j\in\{k\}\cup\mathcal{N}_k} \|z_j\|_2$ gives

$$\left\|z_k^+\right\|_2 \leq T_k\big[(1-\beta_k)+\beta_k\big]M_k \leq M_k. \tag{12}$$

For the span argument, define $W \in \mathbb{R}^{n\times n}$ by $W_{jk} = w_{kj}$ when $j \in \mathcal{N}_k$, and zero otherwise. Let $D_\beta = \mathrm{diag}(\beta_1, \ldots, \beta_n)$ and $D_T = \mathrm{diag}(T_1, \ldots, T_n)$. Then

$$Z^+ = Z\big[(I - D_\beta) + WD_\beta\big]D_T = ZB. \tag{13}$$

Since $Z = U_r^\dagger R$ has rank at most r, the product $L^+ = CZB$ also has rank at most r. This completes the proof. The argument concerns a single update and uses the explicit conditions in Proposition 1.

B Data Representation and Evaluation

Matrix representation. For an RGB sequence, each frame contributes three consecutive channel columns. With frame height h, width w, and N observations, the matrix has size $(hw)\times(3N)$. In the repeated-image restoration construction, the clean matrix repeats the same three channel vectors and therefore has rank at most three. Candidate neighbors in the interpolation step include all coefficient columns.

Corruption generation. The accompanying synthetic generator min–max scales a clean image, adds sparse positive uniform noise and positive random blocks, and clips the observation to $[0,1]$. Its parameter ρ controls the noise density and number of blocks; clipping and overlapping blocks determine the final corrupted support.

Metrics and aggregation. For each RGB image or frame i, the metric utility uses scikit-image (van der Walt et al., 2014) to compute MSE, Euclidean-normalized NRMSE, PSNR with data range one, and SSIM (Wang et al., 2004) with `channel_axis=-1`. Each metric is then averaged separately, by corruption group where applicable:

$$\overline{\mathrm{MSE}} = \frac{1}{N}\sum_i \mathrm{MSE}_i, \qquad \overline{\mathrm{PSNR}} = \frac{1}{N}\sum_i [-10\log_{10}\mathrm{MSE}_i], \tag{14}$$

Because the logarithm is applied before averaging PSNR, $\overline{\mathrm{PSNR}}$ generally differs from $-10\log_{10}\overline{\mathrm{MSE}}$. Image-specific annotations in the visual comparisons and averages in the tables therefore describe different levels of aggregation.

C Implementation Details

Threshold function and schedule. The sampled sparse update uses

$$h_{\theta,\bar{\theta}}(x) = \operatorname{sign}(x)\left[b\ \mathrm{sp}_\kappa(|x| - \theta) - (b-1)\,\mathrm{sp}_\kappa(|x| - \bar{\theta})\right], \quad b = \frac{\bar{\theta}}{\bar{\theta} - \theta}, \tag{15}$$

where $\mathrm{sp}_\kappa(u) = \log(1 + \exp(\kappa u))/\kappa$ and $\kappa = 40$. Thresholds follow the prescribed schedule $\theta_\ell = \theta_0(0.7)^\ell$ and $\bar{\theta}_\ell = 1.5\theta_\ell$, with $\theta_0 > 0$ chosen before decomposition (default: $\theta_0 = 1$). This ensures $\bar{\theta}_\ell > \theta_\ell > 0$ throughout the iteration.

Residual normalization and neighbor selection. The empirical quantile Q_q in Eq. (5) uses linear interpolation. For $n \geq 2$ and $K \geq 1$, the implemented neighborhood selection retains the largest $\min(K+1, n)$ affinities and removes the first ranked entry to obtain $\mathcal{N}_k$. Weights are normalized within this set, and candidate neighbors span all coefficient columns.

Default settings. The neighborhood settings are $K = 10$, $t = 0.1$, and $\tau_d = 0.5$. The interpolation and attenuation settings are $(\beta_l, \beta_h) = (0.1, 0.5)$ and $\eta = 2$. The threshold schedule and these hyperparameters are specified before recovery. Each iteration computes coefficients, residual scores, similarities, and gates from the current observation, without backpropagation, an optimizer, or pretrained weights.

Output updates. At the end of each iteration, L is explicitly reconstructed from CZ^+. Entries of S outside the sampled blocks retain their preceding values.