

UDC 004.932:004.62:512.64

# On the Theory and Application of Singular Value Decomposition in Image Processing and Data Analysis

Gayane G. Ghazaryan<sup>1</sup> and Artashes L. Ghazaryan<sup>2</sup>

<sup>1</sup>Institute of Physics, Yerevan State University, Yerevan, Armenia

<sup>2</sup> Technical University of Munich, Munich, Germany

e-mail: gayane.ghazaryan@ysu.am, artashes.ghazaryan17@gmail.com

## Abstract

Singular Value Decomposition (SVD) is a fundamental factorization method central to image processing, dimensionality reduction, and numerical linear algebra. Its practical strength lies in truncating small singular values, thereby reducing storage while preserving the essential structure. The key challenge is choosing the truncation rank, which determines the trade-off between accuracy, efficiency, and stability.

We introduce an adaptive rank selection method that combines cumulative-energy thresholds with automated elbow detection, yielding a principled alternative to ad hoc rules. The framework is validated on image compression, denoising, and principal component analysis (PCA), with benchmarks against eigenvalue decomposition (EVD) and QR factorization. The results show that the adaptive rule improves fidelity and reproducibility, especially for noisy or ill-conditioned data.

All code, figures, and tables are released as open source, ensuring that the experiments can be reproduced in a clean environment. This unifies theory, automation, and reproducibility, reaffirming that SVD is both mathematically optimal and practically versatile for contemporary data analysis.

**Keywords:** Singular Value Decomposition, Adaptive rank selection, Cumulative energy, Elbow detection, Image denoising, Image compression, Principal component analysis.

**Article info:** Received 7 November 2025; sent for review 9 December 2025; accepted 25 May 2026.

## 1. Introduction

Singular Value Decomposition (SVD) is a cornerstone of numerical linear algebra and a central tool in data science, machine learning, and image processing [1-3]. Its optimality for low-rank approximation, established by the Eckart–Young–Mirsky theorem [1, 4], explains its broad use in image compression [5,6], noise reduction [7-10], and dimensionality reduction via principal component analysis (PCA) [11-13].

Despite its elegance, a persistent practical difficulty is the choice of truncation rank  $k$ . This parameter determines the balance between fidelity, compression ratio, and computational cost [14, 15]. Existing strategies rely either on cumulative-energy thresholds, which require task-specific tuning, or elbow rules, which tend to underestimate the rank when singular values decay gradually [16]. Consequently, applications often rely on ad hoc heuristics, which reduce reproducibility.

In this work, we propose an adaptive rank selection strategy that integrates cumulative energy and elbow detection within a unified SVD pipeline. Concretely, we adopt a conservative hybrid rule

$$k^* = \max(k_\tau, k_{\text{elbow}}),$$

which guards against under-selection while preserving interpretability. Defaults such as  $\tau \approx 0.995$  are reliable for clean images.

On standard images (e.g., `astronaut`, `cameraman`), the hybrid rule avoids elbow under-selection and maintains a higher structural similarity at comparable ranks (see Section 4, Fig. 2).

**Positioning vs. alternatives.** Classical fixed-rank tuning lacks reproducibility. Energy thresholds are interpretable but require task-dependent choices of  $\tau$ . Elbow rules adapt automatically but tend to under-select when the singular values decay slowly. Denoising-specific criteria such as optimal hard thresholding [15] assume particular noise models and are not directly comparable across tasks, such as compression or PCA. In contrast, our hybrid rule is task-agnostic, reproducible, and guards against under-selection, without requiring prior calibration.

**Contributions.** The main contribution of this work is methodological and reproducibility-oriented.

- **Unified framework.** A reproducible SVD workflow for image compression, denoising, and PCA with standardized inputs and metrics (PSNR, SSIM, energy retention, and runtime).
- **Adaptive rank selection.** A principled hybrid of cumulative energy and elbow criteria, robust to slow spectral decay.
- **Systematic benchmarking.** Comparisons against EVD and QR factorization under identical conditions, clarifying the trade-offs in accuracy and runtime.
- **Open-source reproducibility.** All Python scripts can regenerate every figure and table from a clean environment using fixed random seeds, explicit data sources, and version-pinned dependencies.

*Scope and limitations.* We focus on standard truncated SVD for moderate-scale images and matrices; randomized or streaming variants are beyond the scope of this study.

## 2. Mathematical Foundations of Singular Value Decomposition

### 2.1 Definition and Existence

For any matrix  $A \in \mathbb{R}^{m \times n}$  with rank  $r \leq \min(m, n)$ , the **singular value decomposition (SVD)** is

$$A = U \Sigma V^\top,$$

where  $U \in \mathbb{R}^{m \times m}$  and  $V \in \mathbb{R}^{n \times n}$  are orthogonal, and  $\Sigma \in \mathbb{R}^{m \times n}$  is diagonal (rectangular) with entries  $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_r > 0$  on the main diagonal and zeros elsewhere. The SVD exists for every real (and complex) matrix, regardless of its shape or rank [1,17]. The singular values are unique; the singular vectors are unique up to orthogonal transformations within degenerate singular subspaces (e.g., sign flips or rotations when some  $\sigma_i$  coincide) [18].<sup>1</sup>

## 2.2 Relation to Eigendecomposition

SVD is tied to the eigendecomposition of symmetric positive semidefinite matrices  $A^\top A$  and  $AA^\top$  [1,18]:

- Columns of  $V$  are eigenvectors of  $A^\top A$  with eigenvalues  $\lambda_i = \sigma_i^2$ ;
- Columns of  $U$  are eigenvectors of  $AA^\top$  with the same nonzero eigenvalues;
- Hence  $\sigma_i = \sqrt{\lambda_i(A^\top A)}$ .

Consequently,  $\{u_i\}$  and  $\{v_i\}$  form orthonormal bases for the column and row spaces of  $A$ , respectively, and describe the directions of maximal variance/energy in the data representations.

## 2.3 Optimal Low-rank Approximation

Let  $A_k = \sum_{i=1}^k \sigma_i u_i v_i^\top$  be the rank- $k$  truncation. The Eckart–Young–Mirsky theorem [4,19] states that  $A_k$  is the *best* rank- $k$  approximation to  $A$  in all unitarily invariant norms, in particular

$$\|A - A_k\|_F = \min_{\text{rank}(B) \leq k} \|A - B\|_F = \left( \sum_{i=k+1}^r \sigma_i^2 \right)^{1/2}, \quad \|A - A_k\|_2 = \begin{cases} \sigma_{k+1}, & k < r, \\ 0, & k \geq r. \end{cases}$$

These identities justify the use of truncated SVD for compression, denoising, and dimensionality reduction.

## 2.4 Compact and Truncated Forms

Two reduced forms are especially useful [1, 20]:

- **Compact (thin) SVD:**  $A = U_r \Sigma_r V_r^\top$ , with  $U_r \in \mathbb{R}^{m \times r}$ ,  $\Sigma_r \in \mathbb{R}^{r \times r}$  (nonzero singular values),  $V_r \in \mathbb{R}^{n \times r}$ . This is an exact factorization with minimal storage requirement.
- **Truncated SVD:**  $A_k = U_k \Sigma_k V_k^\top$  for  $k < r$ , where  $U_k \in \mathbb{R}^{m \times k}$ ,  $\Sigma_k \in \mathbb{R}^{k \times k}$ ,  $V_k \in \mathbb{R}^{n \times k}$ . This trades accuracy for efficiency by retaining the dominant components.

Two useful projection identities are

$$A_k = U_k U_k^\top A = A V_k V_k^\top,$$

showing that truncation amounts to an orthogonal projection onto the leading left/right singular subspaces.

---

<sup>1</sup>For complex matrices, orthogonal/unitary and transpose/conjugate-transpose are replaced accordingly.

## 2.5 Energy, Variance, and Explained Proportion

SVD provides an orthogonally invariant notion of the total energy as follows:

$$\|A\|_F^2 = \sum_{i=1}^r \sigma_i^2.$$

For a rank- $k$  truncation  $A_k$ , the retained fraction is

$$\eta_k := \frac{\|A_k\|_F^2}{\|A\|_F^2} = \frac{\sum_{i=1}^k \sigma_i^2}{\sum_{i=1}^r \sigma_i^2}, \quad k = 1, \dots, r.$$

This coincides with the explained variance in PCA for mean-centered data [11]. The quantity  $\eta_k$  underlies the cumulative-energy thresholds used later for adaptive rank selection and provides a task-agnostic basis for comparing datasets.

These properties establish SVD as both theoretically optimal and practically versatile. In the next section, we build on them to address the central challenge of choosing the truncation rank  $k$ , and propose an adaptive rule that avoids ad hoc heuristics.

## 3. Adaptive Rank Selection

Choosing the truncation rank  $k$  is the central practical challenge in applying SVD, as it sets the trade-off between the reconstruction fidelity, storage, and runtime. We propose a **unified adaptive rank selection** strategy that replaces ad hoc heuristics with principled and reproducible rules. The method combines two complementary criteria: a cumulative-energy threshold, which is standard in PCA practice [11], and automated elbow detection via the Kneedle algorithm [16]. Together, they yield stable choices for both sharply and slowly decaying singular spectra.

### 3.1 Cumulative-energy Criterion

As introduced in Subsection 2.5, the cumulative-energy ratio  $\eta_k$  measures the fraction of total energy retained by the first  $k$  singular values. The threshold rule selects the smallest  $k_\tau$  such that  $\eta_{k_\tau} \geq \tau$ , with  $\tau \in (0, 1)$ . High thresholds (e.g.,  $\tau \approx 0.995$ ) preserve fine details in compression, whereas lower values may be acceptable when noise dominates the tail. The sensitivity to  $\tau$  is summarized in Appendix B.

### 3.2 Elbow Detection on the Cumulative Curve

To improve robustness, we pair the threshold rule with elbow detection. We apply Kneedle [16] to the *cumulative* spectrum, using the normalized points  $(k/r, \eta_k)$  after light smoothing (Savitzky–Golay) to reduce jitter. Kneedle returns an elbow index  $k_e$  corresponding to the onset of diminishing returns in cumulative energy. Thus, we obtain two candidates:  $k_\tau$  from the threshold rule and  $k_e$  from the elbow detection. We round the continuous elbow location to the nearest integer  $k \in \{1, \dots, r\}$ .




Fig. 1. Cumulative energy  $\eta_k$  with threshold  $k_\tau$  and elbow  $k_e$ . The conservative policy  $k^* = \max(k_\tau, k_e)$  prevents under-selection.

### 3.3 Unified Policy and Scope

Fig. 1 illustrates the complementarity between the two rules. Our *default, conservative* choice is

$$k^* = \max(k_\tau, k_e),$$

which preserves fidelity when the elbow under-selects slowly decaying spectra (Section 2). For practitioners who prefer stronger noise suppression at the cost of detail, we expose an *aggressive* option  $k^\dagger = \min(k_\tau, k_e)$ . Because the procedure depends only on the singular spectrum  $\{\sigma_i\}$ , it applies equally to exact and randomized SVD [12].

**PCA interpretation.** For a mean-centered data matrix  $X \in \mathbb{R}^{m \times n}$  with covariance  $C = \frac{1}{m-1} X^\top X$ , let  $s_i = \sigma_i(X)$  denote the singular values of  $X$ . Then the eigenvalues of  $C$  satisfy  $\lambda_i = s_i^2 / (m-1)$ . The cumulative variance explained by the first  $k$  components is

$$\frac{\sum_{i=1}^k \lambda_i}{\sum_{i=1}^r \lambda_i} = \frac{\sum_{i=1}^k s_i^2}{\sum_{i=1}^r s_i^2} = \eta_k(X).$$

Therefore, the same policy directly selects the number of principal components.

#### Algorithm.

1. Compute singular values  $\{\sigma_i\}_{i=1}^r$  of  $A$  and  $\eta_k$  for  $k = 1, \dots, r$ .
2. Set  $k_\tau \leftarrow \min\{k : \eta_k \geq \tau\}$ .
3. Smooth  $\eta_k$  (optional), normalize  $k \mapsto k/r$ , and compute  $k_e$  via Kneedle (increasing/concave) on  $(k/r, \eta_k)$ ; round to an integer.
4. Set default  $k^* \leftarrow \max(k_\tau, k_e)$ . (Optional aggressive mode:  $k^\dagger \leftarrow \min(k_\tau, k_e)$ .)

**Output:** Truncation rank  $k \in \{k^*, k^\dagger\}$ .

**Implementation notes.** (i) If  $r = 0$ , return  $k = 0$ ; otherwise clip  $k$  to  $[1, r]$ . (ii) If Kneedle returns no elbow (flat or very steep curves), fall back to  $k_\tau$ . (iii) Use a short Savitzky–Golay window for smoothing and the “increasing/concave” Kneedle setting, then enforce the monotonicity of  $\eta_k$ . (iv) Default thresholds:  $\tau \approx 0.995$  for compression;  $\tau \in [0.95, 0.99]$  for denoising. (v) To avoid brittle ties, always choose the smallest  $k$  that satisfies each of the rules.

**Key insight.** By consolidating cumulative-energy thresholds with elbow detection on the cumulative spectrum, the adaptive rule provides a mathematically grounded and reproducible guideline for rank selection. It removes manual tuning, adapts across tasks, and avoids the elbows under-selection failure mode, improving the fidelity in compression and stability in denoising. On spectra with extremely slow decay, the conservative policy may select a larger  $k^*$ , trading compression for fidelity; in such cases practitioners may opt for the aggressive mode or impose a budget-constrained cap on  $k$ .

## 4. Applications and Experimental Setup

We evaluate SVD on three canonical low-rank tasks: image compression, image denoising, and principal component analysis (PCA). Each task highlights a different facet of low-rank modeling: storage reduction, noise suppression, and dimensionality reduction. Unless stated otherwise, we use the adaptive selection from Section 3 with the *default, conservative* policy  $k^* = \max(k_\tau, k_e)$ ; an optional aggressive mode  $k^\dagger = \min(k_\tau, k_e)$  is exposed for heavy-noise denoising.

**Protocol.** All experiments share one codebase, fixed random seeds, and scripts that regenerate figures and tables from a clean environment. The libraries and exact command lines are listed in Appendix B. The metrics are consistent across tasks: PSNR (dB), SSIM, cumulative energy  $\eta_k$ , runtime, and a storage proxy ( $mk+nk+k$ ) corresponding to the parameters of  $U_k, \Sigma_k, V_k$ . SSIM uses an  $11 \times 11$  Gaussian window (default in `skimage.metrics.ssim`) with a data range  $[0, 1]$ .

**Experimental design principles.** Preprocessing, metrics, and defaults are matched across methods (SVD vs. baselines) so differences arise from rank selection or factorization, not setup. When stochasticity is present (e.g., randomized SVD), the seeds are fixed for comparability. We report both numerical fidelity (PSNR,  $\eta_k$ ) and perceptual quality (SSIM).

### 4.1 Image Compression

Given an image matrix  $A$ , retaining the top  $k$  singular values yields

$$A_k = \sum_{i=1}^k \sigma_i u_i v_i^\top,$$

preserving the dominant structure while reducing the storage from  $mn$  to  $mk+nk+k$  parameters. We evaluate on the `astronaut` image (converted to grayscale,  $512 \times 512$ ), comparing fixed manual ranks to adaptive choices.

As a representative case, we contrast a fixed  $k = 50$  with an adaptive rank  $k^*$  from the  $\tau=0.995$  cumulative energy rule (guarded by the elbow via `max`). Fig. 2 shows the cumulative

energy curve with both choices marked and reconstructions annotated with PSNR/SSIM. In this setting, the elbow typically lies below  $k_\tau$ , so  $k^* = \max(k_\tau, k_e) = k_\tau$ , delivering higher fidelity than small fixed ranks with a modest storage increase. For completeness, we also report a guard against under-selection,  $k = \max(k_\tau, k_e)$ , across the images and thresholds. The summary statistics are presented in Appendix Table 2.




Fig. 2. Grayscale astronaut image: manual  $k = 50$  vs. adaptive  $k^* = 59$ . Adaptive rank improves fidelity (PSNR +1.2 dB, SSIM +0.04).

## 4.2 Image Denoising

For denoising, small singular values often capture high-frequency noise; thus, truncation acts as a low-pass spectral filter. We corrupt `astronaut` with additive Gaussian noise ( $\sigma = 0.10$ ) and apply adaptive selection to balance noise removal and detail preservation. Fig. 3 shows the singular-value spectrum of the noisy astronaut image with threshold ( $k_\tau$ ) and elbow ( $k_e$ ) annotations used for adaptive rank selection.

Benchmark comparisons with fixed ranks are presented in Table 1. The adaptive rule achieves PSNR **on par with or exceeding the best fixed choices**, without per-image tuning, and avoids elbow under-selection on slowly decaying spectra. An example of a full reconstruction is shown in Appendix Fig. 6.

Table 1: Comparison of manual fixed ranks and adaptive rank selection ( $k^* = 58$ ) for denoising the noisy `astronaut` image ( $\sigma = 0.10$ ). Reported are PSNR (dB), SSIM, and retained energy  $\eta_k$ . Adaptive rank achieves quality on par with the best fixed choices while avoiding manual tuning.

Method	Rank $k$	PSNR (dB)	SSIM	Energy $\eta_k$
Manual (fixed- $k$ )	40	23.75	0.507	0.967
Manual (fixed- $k$ )	80	23.70	0.467	0.978
Adaptive	58	24.12	0.497	0.973




Fig. 3. Noisy astronaut image ( $\sigma = 0.10$ ). Threshold ( $k_\tau$ ) and elbow ( $k_e$ ) candidates are compared; the default adaptive rule selects  $k^* = \max(k_\tau, k_e)$ .

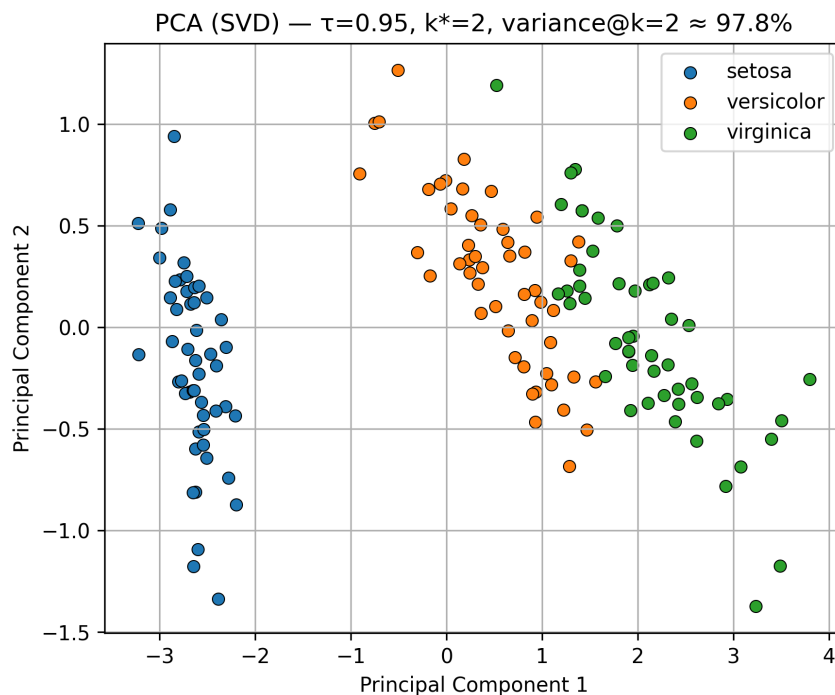


Fig. 4. PCA on the Iris dataset. With  $\tau = 0.95$ , the adaptive rule selects  $k = 2$ , retaining 97.8% variance.

### 4.3 Principal Component Analysis (PCA)

PCA is equivalent to applying SVD to a mean-centered data matrix [11]. For  $X \in \mathbb{R}^{m \times n}$  with  $X = U\Sigma V^\top$  and covariance  $C = \frac{1}{m-1}X^\top X$ , the eigenvalues satisfy  $\lambda_i = \sigma_i^2 / (m-1)$ ; the cumulative variance explained equals  $\eta_k$  (Section 3).

We illustrate this on the `iris` dataset. Using  $\tau = 0.95$ , the adaptive rule selects  $k = 2$ , capturing  $\approx 97.8\%$  of the variance. Projecting  $Z = XV_k$  yields a two-dimensional embedding that preserves the class structure (Setosa is clearly separated; Versicolor/Virginica partially overlap), as shown in Fig. 4.

The variance-retention curve (Appendix Fig. 7) confirms the stability of the adaptive choice across seeds and preprocessing variants.

**Summary.** Across all tasks, adaptive rank selection provides a reproducible alternative to manual tuning: (i) In **compression**, it matches or exceeds fixed- $k$  fidelity at modest storage. (ii) In **denoising**, it balances detail and noise suppression while avoiding elbow under-selection; (iii) In **PCA**, component selection is automated to meet variance targets without per-dataset tuning.

## 5. Comparative Analysis of Matrix Decomposition Methods

To contextualize SVD, we benchmark against eigenvalue decomposition (EVD) and QR factorization on the  $512 \times 512$  *grayscale astronaut* image. We study both the *fixed-rank* and *adaptive-rank* settings, evaluating the reconstruction fidelity (PSNR and SSIM) and runtime.

### 5.1 Benchmark Setup

All methods use NumPy/SciPy and identical hardware; BLAS/LAPACK backends are held fixed (Appendix B). Fixed-rank runs use  $k \in \{10, 30, 50, 100\}$ ; adaptive selection follows Section 3 with the default policy  $k^* = \max(k_\tau, k_e)$  and  $\tau = 0.995$  for high-fidelity compression. The elbow-only choices  $k_e$  are reported in the appendix. Metrics are averaged over repeated runs (error bars show  $\pm 1$  std). The runtime is measured on noisy inputs ( $\sigma = 0.10$ , images normalized to  $[0, 1]$ ); PSNR/SSIM are always computed against a clean reference.

**Reproducibility metadata.** All figures and tables were generated using a unified Python-based experimental pipeline with fixed random seeds and standardized preprocessing. Although absolute runtimes may vary depending on hardware and numerical backends, the relative trends remained stable across different environments.

#### Rank- $k$ formulations.

- **SVD:**  $A_k = U_k \Sigma_k V_k^\top$ , the optimal rank- $k$  approximation in Frobenius and spectral norms [4].
- **EVD (Gram-based):** Form  $A^\top A = V \Lambda V^\top$  with  $\lambda_i = \sigma_i^2$ . Taking  $V_k = [v_1, \dots, v_k]$  and  $\Sigma_k = \text{diag}(\sigma_1, \dots, \sigma_k)$  yields

$$A_k = AV_k V_k^\top = U_k \Sigma_k V_k^\top, \quad U_k = AV_k \Sigma_k^{-1} (\sigma_i > 0).$$

This is algebraically equivalent to the truncated SVD, although explicitly forming  $A^\top A$  may amplify conditioning errors.

- **QR (projection form):** With column-pivoted thin QR,  $A\Pi = QR$ . Let  $Q_k$  be the first  $k$  columns of  $Q$ . Then

$$A_k \approx Q_k Q_k^\top A = Q_k R_{k,:} \Pi^\top,$$

i.e., an orthogonal projection onto  $\text{span}(Q_k)$  (since  $Q_k^\top A = R_{k,:} \Pi^\top$ ). This is generally *non-optimal* for a fixed  $k$ .

## 5.2 Benchmark Summary

Fig. 5 plots PSNR versus rank. SVD attains the highest fidelity at each  $k$ , consistent with optimality, and EVD reconstructions match SVD to numerical precision (when computed via  $AV_k V_k^\top$ ). See Fig. 8 for the same curve with the adaptive  $k^*$  indicated. QR is faster but yields lower fidelity at the same  $k$ . The exact timings depend on the backend and cache effects; see Appendix Table 5.

Complete numerical results: Table 3 (fixed- $k$ ), Table 4 (adaptive with  $\tau = 0.995$ ), and Table 5 (runtime).

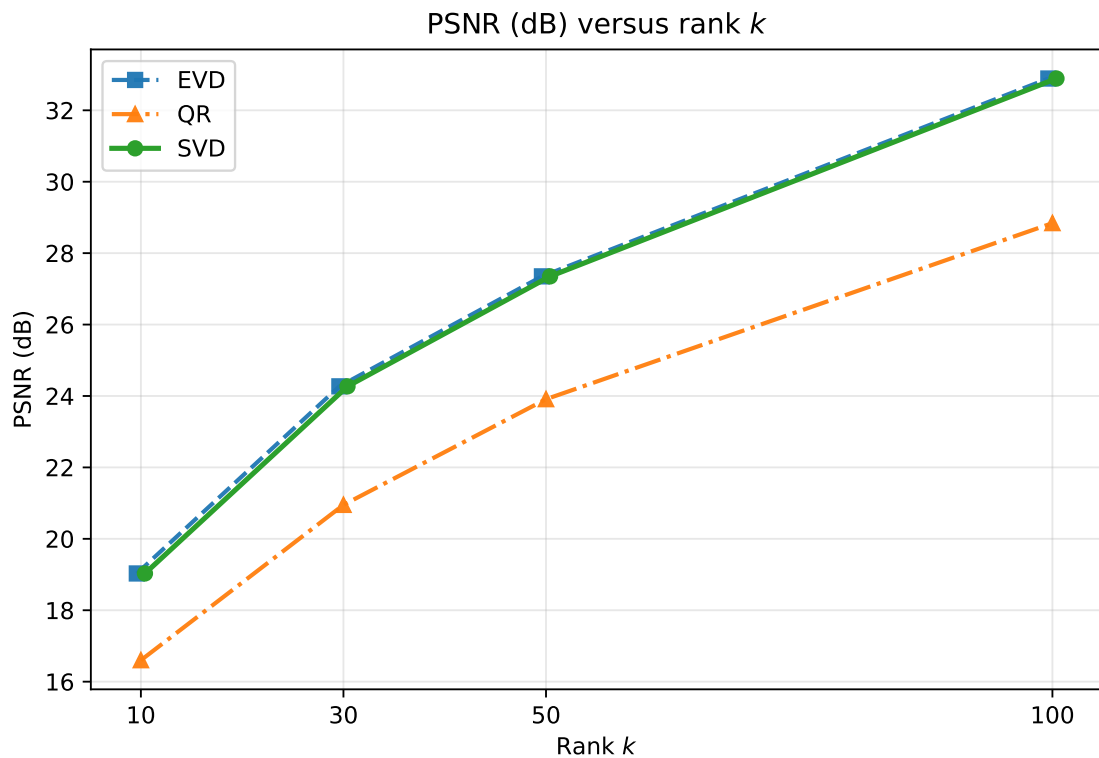


Fig. 5. PSNR vs. rank  $k$  for SVD, Gram-based EVD, and pivoted QR. SVD/EVD attain the highest fidelity; QR is faster but less accurate.

## 5.3 Key observations

- **SVD optimality.** At every tested  $k$ , SVD achieved the highest PSNR/SSIM values. For example, at  $k = 50$ , SVD reached 27.35 dB / 0.747 vs. QRs 23.91 dB / 0.624; at  $k = 100$ , 32.89 dB / 0.895 vs. 28.84 dB / 0.777 (Table B2).
- **EVD equivalence.** Reconstructions based on  $A^\top A$  (via  $AV_k V_k^\top$ ) matched SVD within rounding error, as expected; directly forming  $A^\top A$  can, however, increase sensitivity to conditioning.

- **Adaptive rank** ( $\tau = 0.995$ ). The rule selected  $k^* = 59$  in our run, improving fidelity over  $k = 50$  by  $\approx 1.2$  dB PSNR (28.54 vs. 27.35) and  $\approx 0.04$  SSIM (0.787 vs. 0.747). The QR at  $k^*$  remained lower in quality (24.72 dB / 0.648).
- **Speed–quality trade-off**. QR was  $\sim 2\times$  faster than SVD at moderate  $k$  (e.g.,  $k \in \{10, 50, 100\}$ ), at the cost of noticeable accuracy loss.

## 5.4 Conclusion

These benchmark experiments confirm the trade-offs. SVD (and Gram-based EVD) achieves the best fidelity at a given rank. QR offers faster runtime at reduced accuracy, and adaptive selection with  $\tau = 0.995$  provides a practical balance. These observations motivate the unified framework summarized in the overall conclusion section.

## 6. Conclusion and Future Work

We presented a unified and reproducible framework for applying SVD to image compression, denoising, and PCA, with systematic comparisons to EVD and QR. Our central contribution is an **adaptive rank selection** strategy that combines cumulative energy thresholds with elbow detection on the cumulative spectrum, eliminating manual tuning while preserving interpretability.

**Key findings.** (i) SVD consistently achieves the highest reconstruction fidelity, in line with Eckart–Young [4]; (ii) Gram-based EVD reconstructions are numerically indistinguishable from SVD. (iii) QR offers speed advantages at reduced fidelity; (iv) Adaptive selection improves quality over fixed settings while standardizing the decisions across tasks.

Our contribution is to **consolidate** commonly used heuristics into a principled, data-driven, and fully reproducible framework for adaptive rank selection. For clean images, a 99.5% cumulative-energy threshold emerges as a reliable default that balances perceptual fidelity and compression efficiency.

The proposed framework is primarily methodological and does not introduce a new matrix decomposition or new theoretical optimality guarantees. In addition, the elbow-detection component remains partially empirical and may depend on the spectral characteristics of the data. Nevertheless, the approach provides a unified and reproducible strategy for adaptive rank selection across multiple SVD-based applications.

The entire framework, including the code and figure-generation scripts, is openly available.<sup>2</sup>

**Implementation updates.** All scripts were refactored into a unified modular structure with standardized outputs and harmonized argument parsing. These updates ensure that every figure and table can be reproduced from a clean environment with a single command execution.

**Future work.** Extensions include color images and video, streaming and distributed settings, and hybrids that pair SVD with learned priors. An especially promising direction is the integration of *randomized and streaming SVD algorithms* [12], enabling scalable low-rank approximation on large datasets while preserving theoretical guarantees. Beyond imaging,

---

<sup>2</sup>All code and reproducibility materials are publicly accessible at <https://github.com/gayaneghazaryan-dot/svd-image-analysis> and archived on Zenodo with DOI: <https://doi.org/10.5281/zenodo.20418974> .

applications in NLP, recommender systems, and bioinformatics are the natural next steps. More broadly, SVD remains foundational for modern machine learning embeddings and representation learning, suggesting further opportunities for cross-disciplinary impact.

## Appendix

### A. Compression Quality Metrics

For completeness, we summarize the metrics used in the main text as follows.

- **Compression ratio (CR).** For an  $m \times n$  image stored as a dense array, a rank- $k$  SVD uses  $mk + nk + k$  parameters (for  $U_k, \Sigma_k, V_k$ ) versus  $mn$  for the original.<sup>3</sup> We therefore use the proxy

$$\text{CR}_{\text{param}} = \frac{mn}{mk + nk + k}.$$

- **Peak signal-to-noise ratio (PSNR).**

$$\text{PSNR} = 10 \log_{10} \left( \frac{\text{MAX}^2}{\text{MSE}} \right),$$

where  $\text{MAX} = 1$  for our normalized images, and  $\text{MSE}$  is the mean squared error vs. the clean reference.

- **Structural similarity (SSIM).** A perceptual score in  $[0, 1]$  combining luminance, contrast, and structure; higher values indicate better fidelity.
- **Energy retention.**

$$\eta_k = \frac{\sum_{i=1}^k \sigma_i^2}{\sum_{i=1}^r \sigma_i^2} = \frac{\|A_k\|_F^2}{\|A\|_F^2},$$

coinciding with the cumulative variance explained in PCA for mean-centered data.

These metrics are used consistently in Subsections 4.1 and 4.2.

### B. Detailed Benchmarking and Supplementary Figures

#### B1. Image Compression (Extended Results)

The extended results illustrate how rank selection influences reconstruction quality and storage; see Table 2.

These results support the main-text default:  $\tau = 0.995$  balances compression and perceptual fidelity, whereas  $\max(k_\tau, k_e)$  guards against elbow under-selection.

#### B.2 Image Denoising (Extended Results)

To complement the main-text results, we illustrate adaptive rank selection on a noisy *astronaut* image. Fig. 6 shows the cumulative energy curve with the selected  $k^*$ , along with the noisy input and the corresponding adaptive reconstruction.

---

<sup>3</sup>Byte-level ratios depend on datatypes and entropy coding; we report a parameter-count proxy.

Table 2: Comparison of manual and adaptive rank selection on the  $512 \times 512$  grayscale astronaut. Manual fixed- $k$  uses  $k = 50$ . Adaptive results are grouped by threshold ( $\tau = 0.95$  and  $\tau = 0.995$ ), with elbow-only and elbow+guard rules shown separately.

Method	Rank $k$	PSNR (dB)	SSIM	Energy $\eta_k$
<b>Manual</b>				
Fixed- $k$ (50)	50	27.35	0.747	0.993
<b>Adaptive (<math>\tau = 0.95</math>)</b>				
Energy threshold	9	18.58	0.459	0.951
<b>Adaptive (<math>\tau = 0.995</math>)</b>				
Energy threshold	59	28.54	0.787	0.995
Elbow only	35	25.13	0.671	0.989
Elbow + guard ( $\max(k_\tau, k_e)$ )	59	28.54	0.787	0.995

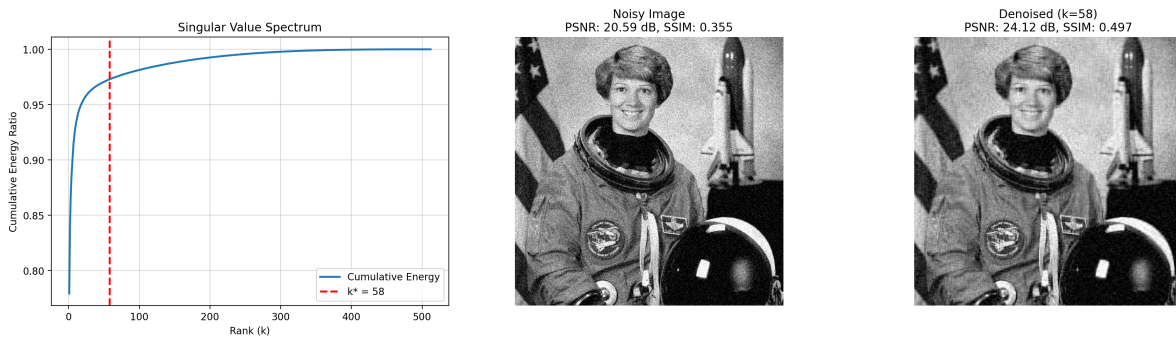


Fig. 6. Noisy astronaut with adaptive rank  $k^*$ . Left: cumulative energy curve; middle: noisy input; right: reconstruction with PSNR/SSIM.

### B.3 Adaptive Rank Selection in PCA




Fig. 7. Cumulative variance explained for the *iris* dataset. With  $\tau = 0.95$ , the adaptive rule selects  $k = 2$ , consistent with Section 4.3.




Fig. 8. PSNR vs. rank  $k$  for fixed and adaptive policies. Adaptive  $k^* = 59$  improves fidelity by  $\approx 1.2$  dB over fixed  $k = 50$ .

## B.4 Benchmarking Tables and Figures

Table 3: Fixed-rank reconstructions on the  $512 \times 512$  grayscale `astronaut` (clean). SVD and Gram-based EVD coincide; QR achieves lower fidelity.

Method	$k = 10$		$k = 30$		$k = 50$		$k = 100$	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
SVD	19.03	0.473	24.27	0.643	27.35	0.747	32.89	0.895
EVD	19.03	0.473	24.27	0.643	27.35	0.747	32.89	0.895
QR	16.60	0.396	20.95	0.525	23.91	0.624	28.84	0.777

Table 4: Adaptive selection with  $\tau = 0.995$  on the grayscale `astronaut`.  $\Delta k$  is relative to fixed  $k = 50$ .

Method	Adaptive $k$	PSNR (dB)	SSIM	$\Delta k$ vs. 50
SVD	59	28.54	0.787	+9
EVD	59	28.54	0.787	+9
QR	59	24.72	0.648	+9

Table 5: Runtime comparison for rank- $k$  approximations on noisy **astronaut** ( $\sigma = 0.10$ ). Values are averages; absolute times depend on hardware/backends, but relative trends are consistent. EVD is excluded from this table as its runtime is dominated by forming  $A^T A$  and is not directly comparable; its reconstruction quality matches SVD to numerical precision (Tables 3,4).

Method	$k = 10$		$k = 50$		$k = 100$		Time (ms)
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	
SVD	19.03	0.473	27.35	0.747	32.89	0.895	102–105
QR	16.60	0.396	23.91	0.624	28.84	0.777	49–52




Fig. 9. Visual comparison at  $k = 50$ . SVD preserves sharper edges; QR exhibits blur and structural loss.

## C. Supplementary Materials

To ensure reproducibility, the supplementary materials accompany this submission. They include:

- Python scripts implementing SVD-based compression, denoising, and PCA;
- Example images and configuration files for complete end-to-end replication of the experiments.

## Data and Code Availability

All Python scripts, figures, and tables used in this study are publicly available at: <https://github.com/gayaneghazaryan-dot/svd-image-analysis>.

Each experiment covering image compression, denoising, benchmarking, and PCA can be reproduced in full by running one of the four Python scripts provided in the `code/` directory. All figures and tables are generated automatically from these scripts, ensuring full reproducibility without the need for Jupyter notebooks or external data files.

## References

- [1] G.H. Golub and C.F. Van Loan, *Matrix Computations*, Johns Hopkins University Press, 4th edition edition, 2013.

- [2] P.C. Hansen, “The truncated svd as a method for regularization”, *BIT Numerical Mathematics*, vol. 27, no. 4, pp. 534–553, 1987. doi: 10.1007/BF02579243.
- [3] M.E. Wall, A. Rechtsteiner, and L.M. Rocha, “Singular value decomposition and principal component analysis”, *A Practical Approach to Microarray Data Analysis*, pp. 91–109, 2003.
- [4] C. Eckart and G. Young, “The approximation of one matrix by another of lower rank”, *Psychometrika*, vol. 1, no. 3, pp.211–218, 1936. doi: 10.1007/BF02288367.
- [5] H. C. Andrews and C. L. Patterson, “Singular value decomposition (svd) image coding”, *IEEE Transactions on Communications*, vol. 24, no. 4, pp. 425–432, 1976. doi: 10.1109/TCOM.1976.1093311.
- [6] H.R. Swathi and G. Sulochana, “Image compression using svd”, *IOP Conference Series: Materials Science and Engineering*, vol. 263, 042082, 2017. doi: 10.1088/1757-899X/263/4/042082.
- [7] P. C. Hansen, *Discrete Inverse Problems: Insight and Algorithms*, SIAM, 2010.
- [8] Q. Guo, C. Zhang, Y. Zhang and H. Liu, “An efficient svd-based method for image denoising”, *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 26, no. 5, pp. 868–880, 2016, doi: 10.1109/TCSVT.2015.2416631.
- [9] S. Gu, L. Zhang, W. Zuo and X. Feng, “Weighted nuclear norm minimization with application to image denoising”, *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2862–2869, 2014, doi: 10.1109/CVPR.2014.366.
- [10] E.J. Cands, X. Li, Y. Ma and J. Wright, “Robust principal component analysis?”, *Journal of the ACM*, vol. 58, no. 3, pp. 1–37, 2011. doi: 10.1145/1970392.1970395.
- [11] I.T. Jolliffe and J. Cadima, *Principal Component Analysis*, Springer, 2nd edition edition, 2016.
- [12] N. Halko, P.G. Martinsson and J.A. Tropp, “Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions”, *SIAM Review*, vol. 53, no. 2, pp. 217–288, 2011. doi: 10.1137/090771806.
- [13] Z. Lin, “A review on low-rank models in data analysis”, *Big Data and Information Analytics*, vol. 1, no. 2-3, pp. 139–161, 2016. doi: 10.3934/bdia.2016001.
- [14] P.C. Hansen, *Rank-deficient and Discrete Ill-posed Problems: Numerical Aspects of Linear Inversion*, SIAM, 1998.
- [15] M. Gavish and D.L. Donoho, “The optimal hard threshold for singular values is  $4/\sqrt{3}$ ”, *Biometrika*, vol. 101, no.3, pp. 587–602, 2014. doi: 10.1093/biomet/asu056.
- [16] V. Satopaa, J. Albrecht, D. Irwin and B. Raghavan, “Finding a “kneedle” in a haystack: Detecting knee points in system behavior”, *31st International Conference on Distributed Computing Systems Workshops (ICDCSW)*, pp. 166–171. IEEE, 2011. doi: 10.1109/ICDCSW.2011.20.
- [17] L.N. Trefethen and D. Bau, *Numerical Linear Algebra*, SIAM, 1997.
- [18] R.A. Horn and C.R. Johnson, *Matrix Analysis*, Cambridge University Press, 2nd edition edition, 2012.
- [19] L. Mirsky, “Symmetric gauge functions and unitarily invariant norms”, *Quarterly Journal of Mathematics*, vol. 11, no. 1, pp. 50–59, 1960. doi: 10.1093/qmath/11.1.50.

- [20] V. Klema and A. Laub, “The singular value decomposition: Its computation and some applications”, *IEEE Transactions on Automatic Control*, vol. 25, no. 2, pp. 164–176, 1980. doi: 10.1109/TAC.1980.1102314.

## Մինգուլյար արժեքների վերլուծության տեսությունն ու կիրառությունները պատկերների մշակման և տվյալների վերլուծության մեջ

Գայանե Գ. Ղազարյան<sup>1</sup> և Արտաշես Լ. Ղազարյան<sup>2</sup>

<sup>1</sup>Ֆիզիկայի ինստիտուտ, Երևանի պետական համալսարան, Երևան, Հայաստան  
Սյունիսենի տեխնիկական համալսարան, Սյունիսեն, Գերմանիա  
e-mail: gayane.ghazaryan@ysu.am, artashes.ghazaryan17@gmail.com

### Ամփոփում

Մինգուլյար արժեքների վերլուծությունը (Singular Value Decomposition, SVD) մատրիցների վերլուծության հիմնարար մեթոդներից է, որը լայնորեն կիրառվում է պատկերների մշակման, չափողականության նվազեցման և հաշվարկային գծային հանրահաշվի խնդիրներում: Մեթոդի գործնական արդյունավետությունը պայմանավորված է փոքր սինգուլյար արժեքների կտրմամբ, ինչը հնարավորություն է տալիս նվազեցնել տվյալների ներկայացման ծավալը՝ միաժամանակ պահպանելով դրանց հիմնական կառուցվածքային առանձնահատկությունները: SVD-ի կիրառման կարևոր խնդիրներից է կտրման ռանգի ընտրությունը, որը որոշում է հավասարակշռությունը վերականգնման ճշգրտության, հաշվարկային արդյունավետության և արդյունքների կայունության միջև:

Աշխատանքում առաջարկվում է ռանգի ադապտիվ ընտրության մեթոդ, որը համատեղում է կուտակային էներգիայի վրա հիմնված չափանիշը և սինգուլյար արժեքների սպեկտրի «ծնկի» (elbow) կետի ավտոմատ հայտնաբերումը: Առաջարկվող մոտեցումը հիմնավորված այլընտրանք է ռանգի ընտրության լայնորեն կիրառվող մոտեցումներին: Մեթոդն ուսումնասիրվել է պատկերների սեղմման, աղմուկի նվազեցման և գլխավոր բաղադրիչների վերլուծության (PCA) խնդիրներում, ինչպես նաև համեմատվել է սեփական արժեքների դեկոմպոզիցիայի (EVD) և QR գործոնավորման հետ: Թվային փորձարկումների արդյունքները ցույց են տալիս, որ ռանգի ընտրության ադապտիվ կանոնն ապահովում է վերականգնման բարձր որակ և նպաստում է հաշվարկային փորձերի վերարտադրելիությանը և հատկապես աղմկոտ կամ վատ պայմանավորված տվյալների դեպքում:

Աշխատանքում օգտագործված ծրագրային կոդը, նկարները և աղյուսակները հրապարակված են բաց հասանելիությամբ, ինչը հնարավորություն է տալիս ամբողջությամբ վերարտադրել ներկայացված փորձերը: Առաջարկվող մոտեցումը միավորում է տեսական հիմքերը, հաշվարկային գործընթացների ավտոմատացումը և վերարտադրելի հետազոտությունների սկզբունքները՝ վերահաստատելով, որ սինգուլյար արժեքների վերլուծությունը մնում է ինչպես մաթեմատիկորեն հիմնավորված, այնպես էլ գործնականում բազմակողմանի գործիք ժամանակակից տվյալների վերլուծության համար:

**Բանալի բառեր՝** սինգուլյար արժեքների վերլուծություն, ռանգի ադապտիվ ընտրություն, կուտակային էներգիա, «ձնկի» մեթոդ (elbow), պատկերների աղմուկի նվազեցում, պատկերների սեղմում, գլխավոր բաղադրիչների վերլուծություն:

## О теории и приложениях сингулярного разложения в обработке изображений и анализе данных

Гаянэ Г. Казарян<sup>1</sup> и Арташес Л. Казарян<sup>2</sup>

<sup>2</sup>Мюнхенский технический университет, Мюнхен, Германия  
e-mail: gayane.ghazaryan@ysu.am, artashes.ghazaryan17@gmail.com

### Аннотация

Сингулярное разложение (Singular Value Decomposition, SVD) является одним из фундаментальных методов матричной факторизации, широко применяемых в обработке изображений, снижении размерности и вычислительной линейной алгебре. Практическая эффективность метода связана с возможностью отсечения малых сингулярных значений, что позволяет уменьшать объём хранимых данных при сохранении наиболее существенной информации. Одной из ключевых задач при использовании SVD является выбор ранга усечения, определяющего компромисс между точностью аппроксимации, вычислительной эффективностью и устойчивостью результатов.

В работе предлагается метод адаптивного выбора ранга, объединяющий энергетический критерий, основанный на накопленной доли энергии, и автоматическое обнаружение точки перегиба спектра сингулярных значений (elbow). Такой подход представляет собой обоснованную альтернативу широко используемым эвристическим правилам выбора ранга.

Предложенный метод исследован на задачах сжатия изображений, подавления шума и анализа главных компонент (PCA), а также сопоставлен с разложением по собственным значениям (EVD) и QR-факторизацией. Результаты численных экспериментов показывают, что адаптивное правило выбора ранга обеспечивает высокое качество восстановления и повышает воспроизводимость вычислительных исследований, особенно в условиях шумных и плохо обусловленных данных.

Исходный код, рисунки и таблицы опубликованы в открытом доступе, что обеспечивает возможность полного воспроизведения всех представленных экспериментов. Предлагаемый подход объединяет теоретические результаты, автоматизацию вычислительных процедур и принципы воспроизводимых исследований, подтверждая, что сингулярное разложение остаётся как математически обоснованным, так и практически универсальным инструментом современного анализа данных.

**Ключевые слова:** сингулярное разложение, адаптивный выбор ранга, накопленная энергия, метод локтя (elbow), подавление шума изображений, сжатие изображений, анализ главных компонент.