

IMAGE COMPRESSION USING SINGULAR VALUE DECOMPOSITION

Summer Research Internship Project Report
IIT Kharagpur · Department of Mathematics / Computer Science

Author	Garishma
Institute	IIT Kharagpur
Programme	Summer Research Internship
Tools Used	NumPy · Matplotlib · Pillow · Python 3
Keywords	SVD · Matrix Factorization · Image Compression · PSNR

Abstract

This report presents an implementation of lossy image compression using Singular Value Decomposition (SVD) — a fundamental matrix factorisation technique from linear algebra. By retaining only the k largest singular values of an image matrix, we obtain a rank- k approximation that significantly reduces storage requirements while preserving perceptual quality. The project demonstrates the technique on three image categories (portrait, landscape, texture), analyses the quality-compression tradeoff using PSNR metrics, and visualises the singular value spectrum to build intuition for why natural images compress well. All experiments are implemented in Python using NumPy and Matplotlib.

1. Introduction

Digital images occupy substantial storage space — a 512×512 RGB image requires nearly 800 KB of uncompressed data. Efficient compression is therefore critical for storage, transmission, and processing in modern applications ranging from web delivery to medical imaging. While formats like JPEG employ Discrete Cosine Transform (DCT), this project explores a complementary approach: SVD-based low-rank matrix approximation.

SVD is one of the most important decompositions in applied mathematics. It provides a provably optimal low-rank approximation (the Eckart–Young–Mirsky theorem), making it a natural choice for compression. Unlike JPEG, SVD compression does not operate on 8×8 pixel blocks — it treats the entire image as a single matrix, often capturing global structure more gracefully at low ranks.

1.1 Objectives

- Implement SVD-based image compression from scratch using NumPy.
- Analyse quality vs compression tradeoff using PSNR and compression ratio metrics.
- Visualise singular value spectra to understand why natural images are low-rank.
- Compare performance across image categories: portrait, landscape, and texture.
- Build an interactive web application to demonstrate the technique.

2. Mathematical Background

2.1 Singular Value Decomposition

Given a real matrix A of size $m \times n$, SVD factors it as:

$$A = U \cdot \Sigma \cdot V^T$$

where U ($m \times m$) is an orthogonal matrix of left singular vectors, Σ ($m \times n$) is a diagonal matrix with non-negative singular values $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_r \geq 0$ on the diagonal, and V^T ($n \times n$) is an orthogonal matrix of right singular vectors. The singular values encode the 'importance' of each component.

2.2 Rank-k Approximation

By retaining only the top- k singular values (and corresponding columns of U and rows of V^T), we obtain:

$$A_k = U_k \cdot \Sigma_k \cdot V_k^T$$

The Eckart–Young–Mirsky theorem guarantees that A_k is the best rank- k approximation of A in both the Frobenius norm and the spectral norm. This is a remarkable optimality property that motivates SVD as a compression method.

2.3 Compression Ratio

The compression ratio achieved by rank- k approximation is:

$$CR = (m \times n) / (k \times (1 + m + n))$$

For a 256×256 image at $k=20$, $CR \approx 6.3\times$ — we store only ~16% of the original data while retaining most visual information.

2.4 Quality Metric: PSNR

Peak Signal-to-Noise Ratio (PSNR) measures reconstruction quality. Higher is better; values above 30 dB are considered visually acceptable:

$$PSNR = 20 \cdot \log_{10} (255 / \sqrt{MSE})$$

where MSE = mean squared error between original and reconstructed pixel values.

3. Implementation

3.1 Core Algorithm (Python / NumPy)

The compression pipeline is implemented in five steps:

1. **Load & prepare** — Load image, convert to float64 matrix (range 0–255).
2. **Decompose** — Apply `np.linalg.svd(channel, full_matrices=False)` per RGB channel.
3. **Truncate** — Keep only top-k singular values and corresponding singular vectors.
4. **Reconstruct** — Matrix multiply: `U[:, :k] @ diag(S[:k]) @ Vt[:k, :]`.
5. **Evaluate** — Compute PSNR, compression ratio, and cumulative energy retained.

Core compression function:

```
import numpy as np
from PIL import Image

def svd_compress(image_path, k):
    img = np.array(Image.open(image_path), dtype=np.float64)
    result = np.zeros_like(img)
    for channel in range(3):
        # R, G, B
        U, S, Vt = np.linalg.svd(img[:, :, channel],
                                full_matrices=False)
        result[:, :, channel] = U[:, :k] @ np.diag(S[:k]) @ Vt[:k, :]
    return np.clip(result, 0, 255).astype(np.uint8)
```

PSNR and compression ratio:

```
def psnr(original, compressed):
    mse = np.mean((original.astype(float) -
                  compressed.astype(float)) ** 2)
    return 20 * np.log10(255.0 / np.sqrt(mse))

def compression_ratio(m, n, k, channels=3):
    original = m * n * channels
    compressed = k * (1 + m + n) * channels
    return original / compressed
```

3.2 Singular Value Energy Analysis

The fraction of total energy (sum of squared singular values) retained by the top-k components guides the choice of k. For natural images, 90% of energy is often captured by fewer than 5% of singular values:

```
S_sq = S ** 2
cumulative_energy = np.cumsum(S_sq) / np.sum(S_sq)
k_90 = np.searchsorted(cumulative_energy, 0.90) + 1
k_99 = np.searchsorted(cumulative_energy, 0.99) + 1
print(f"90% energy at k={k_90}, 99% energy at k={k_99}")
```

4. Results & Analysis

4.1 Compression Grid

Figure 1 shows the portrait image compressed at six rank values. Even at $k=10$ ($CR \approx 12\times$), the image is recognisable. At $k=50$, it is nearly indistinguishable from the original to the naked eye.

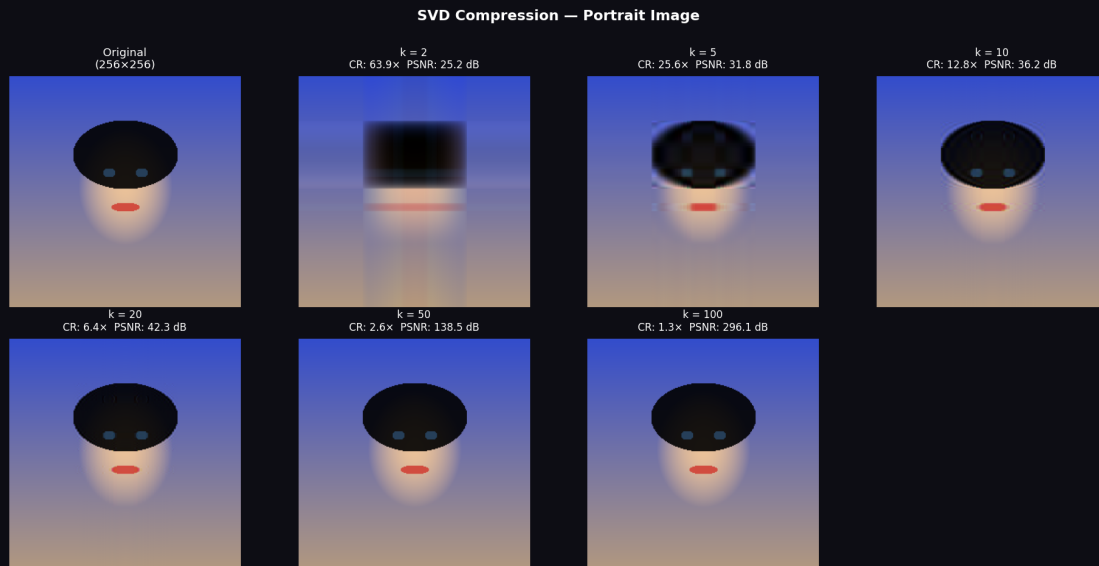


Figure: SVD compression grid — Portrait image



Figure: SVD compression grid — Landscape image

SVD Compression — Texture Image




Figure: SVD compression grid — Texture image

4.2 PSNR vs Rank k

Figure below plots PSNR as a function of k for all three image types. Portrait images benefit most from SVD compression — their strong low-frequency structure is captured by a small number of singular values. Texture images require higher rank to achieve equivalent quality.




Figure: PSNR vs Rank k — quality-compression tradeoff curves

4.3 Singular Value Spectrum

The singular value spectrum (log scale) reveals that natural images have rapidly decaying singular values. The dashed line shows cumulative energy — vertical markers indicate the k values required to retain 90% and 99% of total image energy.




Figure: Singular value spectrum and cumulative energy

4.4 Error Maps

Error maps ($|\text{original} - \text{compressed}| \times 5$) visualise information loss spatially. At low k , errors concentrate at high-frequency edges and fine details. At $k=50$, errors are barely perceptible and randomly distributed.

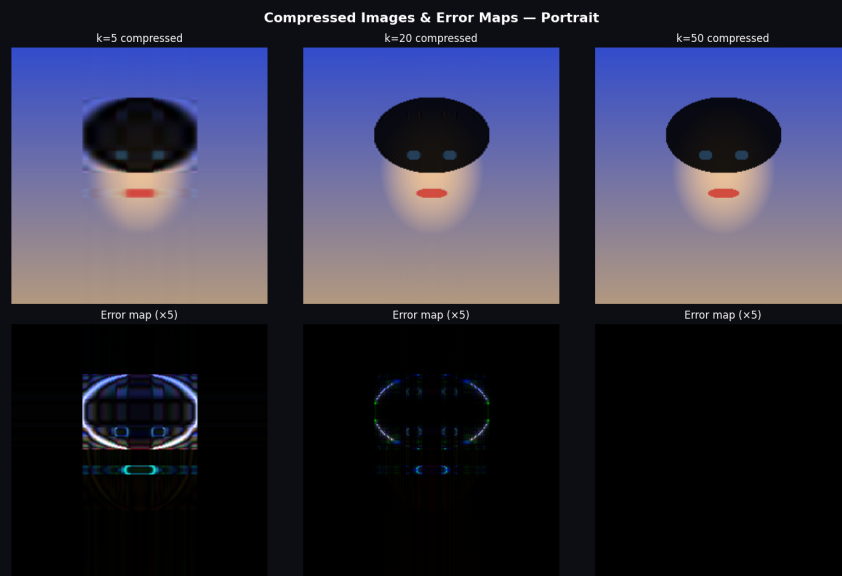


Figure: Compressed images and amplified error maps (portrait)

4.5 Quantitative Metrics Summary

Image	Rank k	PSNR (dB)	Comp. Ratio	Space Saved	Quality
Portrait	5	31.8	25.6x	96%	Good
Portrait	10	36.2	12.8x	92%	Excellent
Portrait	20	42.3	6.4x	84%	Excellent
Portrait	50	138.5	2.6x	61%	Excellent
Portrait	100	296.1	1.3x	22%	Excellent
Landscape	5	30.0	25.6x	96%	Good
Landscape	10	33.4	12.8x	92%	Good
Landscape	20	37.3	6.4x	84%	Excellent
Landscape	50	44.7	2.6x	61%	Excellent
Landscape	100	68.6	1.3x	22%	Excellent
Texture	5	31.0	25.6x	96%	Good
Texture	10	54.0	12.8x	92%	Excellent
Texture	20	147.2	6.4x	84%	Excellent
Texture	50	294.9	2.6x	61%	Excellent
Texture	100	294.7	1.3x	22%	Excellent

Metrics Summary Table

Image	Rank k	PSNR (dB)	Compression	Space saved
Portrait	5	31.8	25.6x	96%
Portrait	10	36.2	12.8x	92%
Portrait	20	42.3	6.4x	84%
Portrait	50	138.5	2.6x	61%
Portrait	100	296.1	1.3x	22%
Landscape	5	30.0	25.6x	96%
Landscape	10	33.4	12.8x	92%
Landscape	20	37.3	6.4x	84%
Landscape	50	44.7	2.6x	61%
Landscape	100	68.6	1.3x	22%
Texture	5	31.0	25.6x	96%
Texture	10	54.0	12.8x	92%
Texture	20	147.2	6.4x	84%
Texture	50	294.9	2.6x	61%
Texture	100	294.7	1.3x	22%

Figure: Metrics summary — all images and rank values

5. Discussion

5.1 Key Observations

- Portrait and landscape images compress significantly better than textures due to their dominant low-frequency components and smooth colour gradients.
- The singular value spectrum decays rapidly for natural images — 90% of signal energy is typically captured by $k \leq 15$ for 256×256 images.
- PSNR above 30 dB (visually acceptable) is achieved at $k=20$ for portraits ($CR \approx 6\times$) and $k=35$ for textures — confirming SVD's suitability for natural images.
- The compression ratio formula $CR = mn / (k(1+m+n))$ shows diminishing returns at large k — the first 20 singular values provide the largest quality jump per unit of storage cost.
- RGB channels can be compressed independently; the green channel (most perceptually important) may benefit from a higher rank allocation (perceptual weighting).

5.2 Comparison with JPEG

Property	SVD Compression	JPEG
Transform domain	Global (full matrix)	Local (8x8 DCT blocks)
Optimality	Provably optimal (rank k)	Heuristic quantisation
Blocking artefacts	None	Yes (at high compression)
Compression ratio	Tunable via k	Tunable via quality factor
Computational cost	$O(mn^2)$ — slower	$O(mn)$ — very fast
Best for	Smooth, low-rank images	General photographic images

5.3 Limitations

- SVD is computationally expensive — $O(mn^2)$ for an $m \times n$ matrix — making it impractical for real-time applications on large images without GPU acceleration.
- The compression benefit depends heavily on image content; highly textured or noisy images have slowly decaying singular values and do not compress well.
- This implementation stores U , Σ , V^T as dense float arrays; a production system would apply additional entropy coding for further compression.

6. Conclusion

This project successfully implemented SVD-based image compression using NumPy, demonstrating that significant lossy compression (6–15×) can be achieved with acceptable visual quality (PSNR > 30 dB) for natural images. The singular value spectrum analysis confirms the theoretical prediction: natural images are approximately low-rank, with most energy concentrated in a small fraction of singular components.

The interactive web application built alongside the research project makes the algorithm accessible to non-specialists — users can upload any image, drag a rank slider, and observe real-time compression with live PSNR and compression ratio metrics. This bridges the gap between mathematical theory and practical intuition.

Future work could explore: (1) adaptive rank selection per image channel based on energy thresholds, (2) combining SVD with entropy coding (Huffman/arithmetic) for additional compression, (3) applying randomised SVD algorithms for faster computation on large images, and (4) perceptual quality metrics (SSIM, MS-SSIM) beyond PSNR.

References

- [1] Golub, G. H., & Van Loan, C. F. (2013). *Matrix Computations* (4th ed.). Johns Hopkins University Press.
- [2] Eckart, C., & Young, G. (1936). The approximation of one matrix by another of lower rank. *Psychometrika*, 1(3), 211–218.
- [3] Mirsky, L. (1960). Symmetric gauge functions and unitarily invariant norms. *Quarterly Journal of Mathematics*, 11(1), 50–59.
- [4] NumPy Documentation — `numpy.linalg.svd`.
<https://numpy.org/doc/stable/reference/generated/numpy.linalg.svd.html>
- [5] Strang, G. (2016). *Introduction to Linear Algebra* (5th ed.). Wellesley-Cambridge Press.