MALAYSIAN JOURNAL OF SCIENCE
MJS
ISSN 1394-3065
Regular Issue
Malaysian Journal of Science 45(2): 75-86 (June 2026) https://mjs.um.edu.my

Comparative Performance Analysis of Interpolative Decomposition and Singular Value Decomposition for Image Reduction

Wei Shean Ng[1a]*, How Hui Liew[2a], Huey Voon Chen[3a] and Hao Tian Ng[4a]

Abstract: Matrix decomposition techniques, such as Singular Value Decomposition (SVD), have traditionally been used for image compression, achieving data reduction while preserving image fidelity. Recently, randomised linear algebra algorithms, including Interpolative Decomposition (ID), have gained attention due to their efficiency in computational expense and memory consumption. This study presents a comparative analysis of ID and randomised SVD for image compression using a large collection of images from the USC-SIPI Image Database. Experimental results demonstrate that both methods effectively reduce image sizes, with SVD yielding greater compression ratios, while ID better preserves image quality at lower computational costs. The analysis also demonstrated that both techniques faced limitations when compressing images with dark or low-contrast areas, with ID performing best on images exhibiting repetitive or structured patterns. These findings indicate that ID is more suitable for memory-limited applications, such as reducing large tabular data, which is important in marketing data analysis, whereas SVD is preferable for image compression.

Keywords: *Interpolative decomposition, singular value decomposition, image reduction.*

1. Introduction

In the digital era, data reduction plays a crucial role in managing big data (Khoei & Singh, 2024). Much tabular data can contain repeated columns and rows. A collection of survey data is one such example. Duplicate data is important since it represents the information collected from the general public. Randomised matrix decomposition has emerged as an essential approach for data decomposition. By deconstructing complex matrices into simpler components, these techniques provide a framework for analysing and approximating data, improving computational efficiency and insight (Muravev et al., 2018; Kawamura & Suda, 2021).

Among randomised matrix decomposition methods (Achlioptas & McSherry, 2001), interpolative decomposition (ID) and randomised Singular Value Decomposition (SVD) have been found to be versatile in data analysis. SVD works by decomposing a matrix into singular values and orthogonal vectors, making it highly effective for tasks such as dimensionality reduction, collaborative filtering, and image compression (Su et al., 2018). Recent research by Parameshachari et al. (2023) explored SVD-based image compression using a modified truncated SVD approach, where truncated components were set to zero. Their findings showed that this method effectively reduces storage requirements while maintaining image fidelity. Hybrid models combining SVD with convolutional neural networks have shown potential in balancing fidelity and compression rate (Li et al., 2025).

Further advancements in SVD-based image compression include its application to hexagonal grid images, as demonstrated by Varghese and Saroja (2021). Hexagonal grids require fewer sampling points than conventional square grids, reducing storage and computational demands (Mersereau, 1979). Inspired by the hexagonal patterns in human retinal photoreceptors, this approach converts a square-lattice image into a pseudo-hexagonal structure using methods described by Wüthrich and Stucki (1991). By applying SVD and rank-k approximations, the study demonstrated improved image compression performance on hexagonal grids compared to conventional square grids. A recent advancement by Li et al. (2024) introduced an efficient online randomised ID algorithm with a single-pass error estimator, enabling real-time image compression in streaming and large-scale applications. Moreover, matrix factorisation and low-rank approximations have also demonstrated effectiveness in addressing complex numerical models in thermal analysis and image domains (Aruchunan et al., 2022). These foundational contributions support the motivation to extend interpolative and singular value decompositions to high-dimensional data applications, including image compression.

The Interpolative Decomposition (ID), a lesser-known randomised matrix decomposition method, is found to be a basic randomised algorithm for the low-rank approximation of data with repeated entries. ID selects representative columns (or rows) from the matrix A, known as the skeleton matrix C, and interpolates across these columns to generate low-rank approximations $A \approx CZ$, where Z is called the interpolation matrix. Since the skeleton matrix C is a random selection of the columns from A, there is no requirement to allocate computer memory to store C. This approach greatly reduces computational

Authors information:
[a]Department of Mathematical and Actuarial Sciences, Lee Kong Chian Faculty of Engineering and Science, Universiti Tunku Abdul Rahman, Jalan Sungai Long, Bandar Sungai Long, Cheras 43000, Kajang, Selangor, MALAYSIA. E-mail: ngws@utar.edu.my[1]; liewhh@utar.edu.my[2]; chenhv@utar.edu.my[3]; haotianng@utar.my[4]
*Corresponding Author: ngws@utar.edu.my

Received: December, 2024
Accepted: August, 2025
Published: June, 2026

DOI:https//doi.org/10.22452/mjs.vol45no2.8
Malaysian Journal of Science 45(2): 75-86 (June 2026)

complexity and memory usage while retaining essential data properties. Its potential for tabular data is obvious, but we are interested in experimenting with it for image reduction.

In the domain of data reduction, ID is the foundation for the CUR matrix factorisation, $A \approx CUR$, where C is the subset of the columns of A, R is the subset of the rows of A, and U refines the approximation. Beyond data reduction, ID has found applications in fields such as quantum mechanics. For example, Damle et al. (2017) employed ID in the Kohn-Sham Density Functional Theory (DFT) for solving nonlinear eigenvalue problems. They used the Selected Columns of the Density Matrix (SCDM) technique to efficiently approximate solutions by focusing on sparse and localised columns of the density matrix. Additionally, Lu and Ying (2015) demonstrated ID's effectiveness in compressing low-rank matrices representing electron systems through its use in the electron repulsion integral tensor, combined with Fast Fourier Transforms (FFTs). These studies highlight ID's ability to decrease computational demands in quantum simulations.

In the area of signal processing, Tang et al. (2023) examined ID's role in musical instrument source separation, comparing its performance with Non-negative Matrix Factorisation (NMF) and Convolutive Non-negative Matrix Factorisation (CNMF). Their results showed that ID underperformed relative to NMF and CNMF due to its random column selection, which may introduce redundant or irrelevant data into the factorised matrix. This finding highlights the importance of understanding ID's inherent limitations and adapting its application to specific problems.

Each matrix decomposition method offers distinct strengths, weaknesses, and compromises in terms of computational efficiency, accuracy, and suitability for particular fields. This paper aims to provide a comparative analysis of Interpolative Decomposition (ID) and Singular Value Decomposition (SVD) for image reduction and compression, investigating their performance, benefits, and limitations in image processing.

2. Matrix Decompositions

Randomised numerical linear algebra methods are gaining popularity in modern data analysis because they typically use less memory and have lower computational complexity (Muravev et al., 2018). They involve sampling column vectors and row vectors from the original matrix and subsequently performing random embedding. Interpolative Decomposition (ID) and Singular Value Decomposition (SVD) are examined in detail in the following sections.

Interpolative Decomposition (ID)

ID is calculated by modifying the pivoted QR, as proposed by Golub (1965). The computation cost of ID is $O(mnk)$, which is cheaper than SVD ($O(mn \cdot \min(m,n))$), where m, n, k are the number of rows, the number of columns, and the rank of the original matrix, respectively. As $k \leq \min(m,n)$, ID is beneficial when solving rank-deficient least squares problems. Many algorithms have been proposed to solve the interpolative decomposition problem (Liu & He, 2019; Bhaskara et al., 2020; Advani & O'Hagan, 2022). Python includes a library called scipy.linalg.interpolative, which calculates the interpolative decomposition of a given matrix.

Algorithm 1. An outline of the randomised interpolative decomposition (ID) for an $m \times n$ matrix A (Liberty et al., 2007):

Step 1. Randomly sample $p = 1.2k$ nonnegative integers from the range 0 to $n-1$ as indices, idx.

Step 2. Perform pivoted QR decomposition of the submatrix $A[:, \mathrm{idx}]$ to obtain Q, R, P.

Step 3. $C = A\left[:, \mathrm{idx}\left[P[:k]\right]\right]$ where $P[:k]$ is the first k values in P.

Step 4. Let $R_k = R[:k, :k]$ be the $k \times k$ submatrix of R, obtain the matrix Z by solving $\left(R_k^T R_k\right) Z = C^T A$ and return the ID approximation CZ as well as the indices $\mathrm{idx}\left[P[:k]\right]$ and the $k \times n$ interpolation matrix Z.

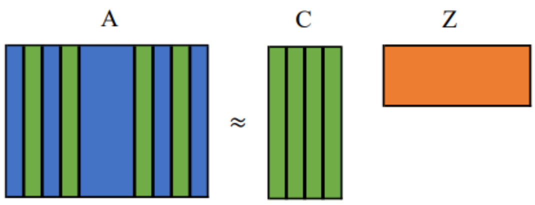

Figure 1. Illustration of ID.

Figure 1 illustrates the Interpolative Decomposition (ID) process, where the original $m \times n$ matrix A is decomposed into a skeleton matrix C of size $m \times k$ and an interpolation matrix Z of size $k \times n$, for any real number k satisfying $k \leq m$ and $k \leq n$.

Singular Value Decomposition (SVD)

Let A be an $m \times n$ matrix. SVD decomposes A into a product of three matrices, $A = U\Sigma V^T$, where U is an $m \times m$ orthogonal matrix, where its columns represent the left singular vectors, Σ is an $m \times n$ diagonal matrix with decreasing singular values, and V^T is an $n \times n$ orthogonal matrix, whose rows represent the right singular vectors.

Both U and V^T are unitary matrices, i.e., $U^T = U^{-1}$ and $V^T = V^{-1}$. The singular values in Σ indicate the importance of the corresponding singular vectors in the decomposition, with higher values showing greater importance (Strang, 2006). With a higher k value, more singular values and singular vectors from the original matrix are retained, which increases the accuracy of the approximation and reduces matrix norms (e.g., Frobenius norm or spectral norm). Consequently, the compression ratio is lower to store more information of the original data, and the compressed representation is less compact. SVD and its variants are commonly applied to perform data compression (Bentbib et al., 2022) and dimensionality reduction (Libal et al., 2020).

The randomised SVD can better deal with the standard SVD because it embeds the original large $m \times n$ matrix A to subspaces for memory-efficient low rank- k approximation $A \approx U\Sigma V$.

DOI:https//doi.org/10.22452/mjs.vol45no2.8
Malaysian Journal of Science 45(2): 75-86 (June 2026)

Algorithm 2. An outline of the randomised SVD for a matrix A (Halko et al, 2010):

Step 1. Draw an $n \times k$ Gaussian random matrix Ω and Ψ.

Step 2. Form the $m \times k$ sample matrix $Y = A\Omega$ and $Z = A\Psi$.

Step 3. Find orthogonal matrices Q and W such that $Y = QQ^*Y$ and $Z = WW^*Z$.

Step 4. Solve for T, the linear system $QY = T(W^*\Omega)$ and $W^*Z = T^*(Q^*\Psi)$.

Step 5. Compute the SVD of the small matrix T: $T = \hat{U}\Sigma\hat{V}^*$.

Step 6. Form $U = Q\hat{U}$ and $V = W\hat{V}$, and return the triple (U, Σ, V).

2. Methodology

To enable image compression via matrix approximation, we represent a grayscale image as an $m \times n$ matrix A, or a colour image as $A = [R, G, B]$, where R, G, B denote the red, green, and blue channels, respectively. Once the image is expressed in matrix form, matrix decomposition techniques, such as Interpolative Decomposition (ID) or Singular Value Decomposition (SVD), are employed to determine A with a lower-rank matrix $\tilde{A}$. The decomposition rank k, which determines the degree of compression, is varied from $\min(m, n)$ down to 1 to determine the smallest value of k that preserves visual quality above a defined threshold. The visual similarity between the original and compressed image is measured using the Structural Similarity Index Measure (SSIM).

Figure 2 presents the procedure for determining the optimal compression parameters using ID. The process begins with image loading and matrix extraction, followed by the definition of candidate lists for the decomposition rank k and a sparsity threshold. An empty result list is initialised to store the compression outcomes. For each combination of k and threshold value, the following steps are executed:

1. Apply ID to the matrix A;
2. Apply thresholding to the resulting coefficient matrix Z to induce sparsity;
3. Reconstruct the compressed matrix;
4. Compute key metrics, including reconstruction error, compression ratio, and SSIM; and
5. Record the results.

After evaluating all possible combinations, the results with SSIM ≥ 0.8 are retained to ensure acceptable image quality. The configuration that yields the maximum compression ratio among the filtered results is selected as the optimal setting. This approach facilitates a comprehensive exploration of the parameter space, while maintaining a balance between compression efficiency and visual fidelity.

Figure 2. Flowchart to obtain the optimal compression parameters of an image using ID.

Image Pre-processing

Images are loaded into Python using libraries like OpenCV or Pillow as pixel grids. Grayscale images form a 2D matrix with intensity values $(0 - 255)$, RGB images have three channels storing intensity data. This matrix representation enables detailed analysis of pixel colour intensities for several image-processing tasks. After loading, the image is converted to grayscale, and its matrix A is extracted. If the number of rows m exceeds the number of columns n, the matrix is transposed. The prepared image is then ready for reduction and compression.

DOI:https//doi.org/10.22452/mjs.vol45no2.8
Malaysian Journal of Science 45(2): 75-86 (June 2026)

Image Reduction and Compression

An iterative process determines optimal parameters k (rank) and a threshold for the Interpolative Decomposition (ID). The matrix A is approximated as $A \approx CZ$, where C is an $m \times k$ matrix of selected columns from A, and Z is the $k \times n$ interpolation matrix. Elements in Z below the threshold are set to their minimum value after normalisation using a min-max scaler. This adjustment reduces less impactful weights in Z to simplify reconstruction without greatly affecting image quality. The compressed matrix $\tilde{A}$ is then reconstructed as $CZ \approx \tilde{A} = CZ'$, where Z' is the modified interpolation matrix.

Performance Metrics

The compression error, denoted as ϵ_c, of the compressed matrix is calculated using the Frobenius norm, $\|A - \tilde{A}\|_F^2 = \epsilon_c$. Mean Squared Error (MSE) is defined as

$$\mathrm{MSE} = \frac{1}{m \times n} \sum_{i=1}^{m} \sum_{j=1}^{n} \left(\tilde{A}_{i,j} - A_{i,j}\right)^2, \tag{1}$$

where $\tilde{A}_{i,j}$ is the (i,j)-entry of the compressed matrix $\tilde{A}$, and $A_{i,j}$ is the (i,j)-entry of the original matrix A, while Peak Signal-to-Noise Ratio (PSNR) is defined as

$$\mathrm{PSNR} = 10 \times \log_{10} \frac{(2^N - 1)^2}{\mathrm{MSE}}, \tag{2}$$

where N is the bit depth per channel of the image. MSE and PSNR are the two metrics applied to evaluate the quality of image reconstruction tasks. MSE quantifies the average squared difference between corresponding pixels in the original image and the reconstructed image following compression. A lower MSE value indicates a smaller overall difference and potentially improved reconstruction quality. In contrast, PSNR offers a measure of the signal-to-noise ratio expressed in decibels (dB). A higher PSNR value signifies a higher ratio of the original signal power (image content) to the noise power (introduced by compression artifacts).

The Structural Similarity Index Measure (SSIM) evaluates image quality, with iterative adjustments applied to refine parameters (Wang et al., 2003). The process is adapted for RGB images with modified error computation. ID's compression performance is then compared to SVD using the same image set.

4. Results and Discussion

For a fixed k, SVD is always the best approximation for an image matrix A, theoretically. ID is a poorer approximation of A. However, for the miscellaneous set (the baboon, peppers, and other favourites) of USC-SIPI Image Database, we obtain Table 1. The table consists of image label, image size (gray-scale images are pairs, while colour images are triplets with the third position 3 indicating red, green, and blue colours), the chosen rank k with the accuracy parameter eps $= 0.05$, and the MSE calculated by using Equation (1).

According to Table 1, we can observe that the majority of k values are over 20% of the image size. However, further compression may be possible by removing visually similar or redundant images. Only images with strong horizontal similarity (e.g., gray21.512.tiff) achieve low MSE. A diverse set of images from the USC-SIPI Image Database was considered for analysis, encompassing grayscale and colour images with varying levels of detail and texture. For experimental evaluations, four representative images, “Baboon”, “Sailboat on Lake”, “Peppers”, and “Airplane F-16” were selected to perform detailed compression analysis. These images were selected for their distinct characteristics, which provide a comprehensive understanding of how Interpolative Decomposition (ID) and Singular Value Decomposition (SVD) handle varying image complexities. However, the generalised parameters derived are intended to extend beyond these four examples, as demonstrated in the “Splash” (image test) and shown across the complete dataset in Table 1.

Table 1. k and Mean Squared Errors for UCS-SIPI misc images.

Image label	Image size	k with eps $= 0.05$	Mean Squared Error
4.1.01.tiff	(256, 256, 3)	128	14.76
4.1.02.tiff	(256, 256, 3)	123	9.32
4.1.03.tiff	(256, 256, 3)	41	24.46
4.1.04.tiff	(256, 256, 3)	32	51.83
4.1.05.tiff	(256, 256, 3)	56	48.93
4.1.06.tiff	(256, 256, 3)	118	42.43
4.1.07.tiff	(256, 256, 3)	30	30.17
4.1.08.tiff	(256, 256, 3)	44	34.56
4.2.01.tiff (Splash (Test Image))	(512, 512, 3)	51	36.98
4.2.03.tiff (Baboon)	(512, 512, 3)	329	52.23
4.2.05.tiff (Airplane (F-16))	(512, 512, 3)	103	45.58
4.2.06.tiff (Sailboat on lake)	(512, 512, 3)	217	47.24
4.2.07.tiff (Peppers)	(512, 512, 3)	156	43.33

DOI:https//doi.org/10.22452/mjs.vol45no2.8
Malaysian Journal of Science 45(2): 75-86 (June 2026)

5.1.09.tiff	(256, 256)	77	31.66
5.1.10.tiff	(256, 256)	132	20.24
5.1.11.tiff	(256, 256)	46	32.4
5.1.12.tiff	(256, 256)	47	45.71
5.1.13.tiff	(256, 256)	50	44.45
5.1.14.tiff	(256, 256)	120	11.02
5.2.08.tiff	(512, 512)	131	26.11
5.2.09.tiff	(512, 512)	211	31.09
5.2.10.tiff	(512, 512)	258	29.43
5.3.01.tiff	(1024, 1024)	380	19.11
5.3.02.tiff	(1024, 1024)	538	14.17
7.1.01.tiff	(512, 512)	176	13.45
7.1.02.tiff	(512, 512)	52	38.8
7.1.03.tiff	(512, 512)	134	29.13
7.1.04.tiff	(512, 512)	119	21.63
7.1.05.tiff	(512, 512)	271	13.75
7.1.06.tiff	(512, 512)	296	9.67
7.1.07.tiff	(512, 512)	258	13.57
7.1.08.tiff	(512, 512)	99	26.34
7.1.09.tiff	(512, 512)	191	18.71
7.1.10.tiff	(512, 512)	127	20.91
7.2.01.tiff	(1024, 1024)	449	7.83
boat.512.tiff	(512, 512)	156	31.3
gray21.512.tiff	(512, 512)	2	0.19
house.tiff	(512, 512, 3)	160	44.35
ruler.512.tiff	(512, 512)	44	141.35

Optimal Parameters for ID

The optimal compression parameters for the "Baboon" image were determined through 160 iterations, testing 20 k-values (from 1 to the image's rank) with eight threshold values (0.001, 0.005, 0.01, 0.025, 0.05, 0.1, 0.25, and 0.5). Metrics, including MSE (refer to Equation (1)), PSNR (refer to Equation (2)), Structural Similarity Index Measure (SSIM), and compression ratio, were recorded. SSIM and the compression ratio were used to identify the optimal combinations of k and the threshold for effective image compression.

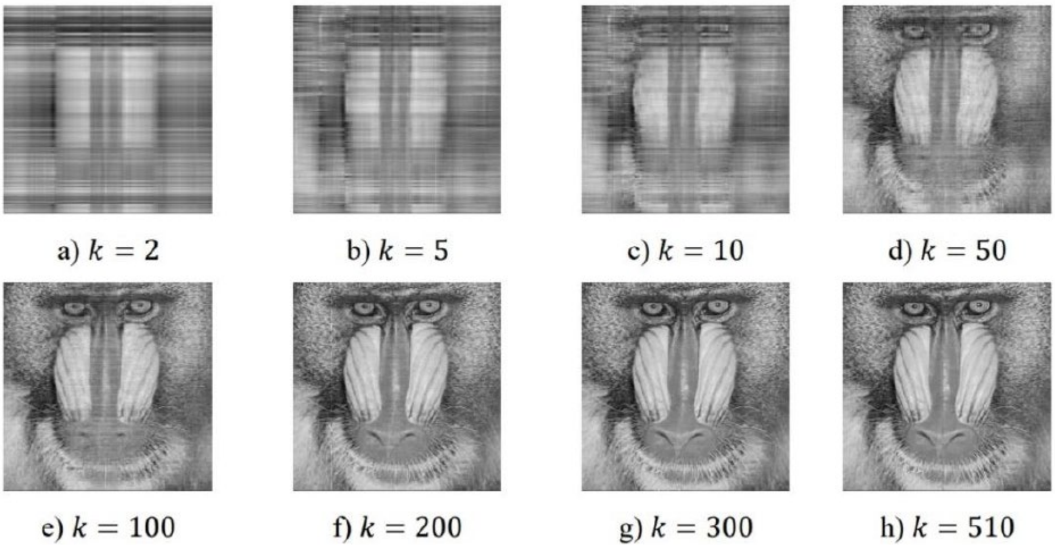

Figure 3. ID compression of different k values with a threshold of 0.01.

DOI:https//doi.org/10.22452/mjs.vol45no2.8
Malaysian Journal of Science 45(2): 75-86 (June 2026)

Table 2. Reconstruction Errors and Compression Ratios for Various k Values in ID Compression of the "Baboon" Image.

k	MSE	PSNR	SSIM	Compression Ratio
2	2010.716	15.0973	0.150149	5.803884
5	1463.848	16.4758	0.176223	4.852780
10	1241.537	17.1912	0.226701	4.352396
50	651.644	19.9907	0.477479	4.183527
100	373.117	22.4001	0.658740	3.584975
200	142.651	26.5842	0.853026	2.589475
300	42.50984	31.8460	0.940425	2.534132
510	25.88132	34.0090	0.989844	2.536124

Further experiments were conducted to investigate the impact of varying rank parameters (k) on image compression and quality. As illustrated in Figure 3 and Table 2, ID decomposition was applied to the "Baboon" image using k values ranging from 2 to 510. The analysis reveals that adjusting the rank parameter has a great effect on reconstruction quality and compression efficiency. As k increases, the MSE consistently decreases, showing improved image fidelity. Simultaneously, PSNR and SSIM values rise, reflecting better preservation of image details. However, this occurs at the cost of a reduced compression ratio, highlighting the trade-off between image quality and storage efficiency. These findings demonstrate that ID can be effectively fine-tuned to balance compression and quality based on specific conditions. While these results exhibit ID's ability to generalize across various parameter settings within the USC-SIPI Image Database, extending this evaluation to other datasets would further confirm its robustness and applicability for diverse image types.

Numerical analysis established that a threshold value of 0.01 is optimal for all four images using ID across most k-values. Figure 4 illustrates a comparison between uncompressed and compressed images using optimal compression parameters in ID for both grayscale and colour images. To generalize a k-value for compressing four different images, the average k-value was calculated and is shown at the bottom of the figure.

DOI:https//doi.org/10.22452/mjs.vol45no2.8
Malaysian Journal of Science 45(2): 75-86 (June 2026)

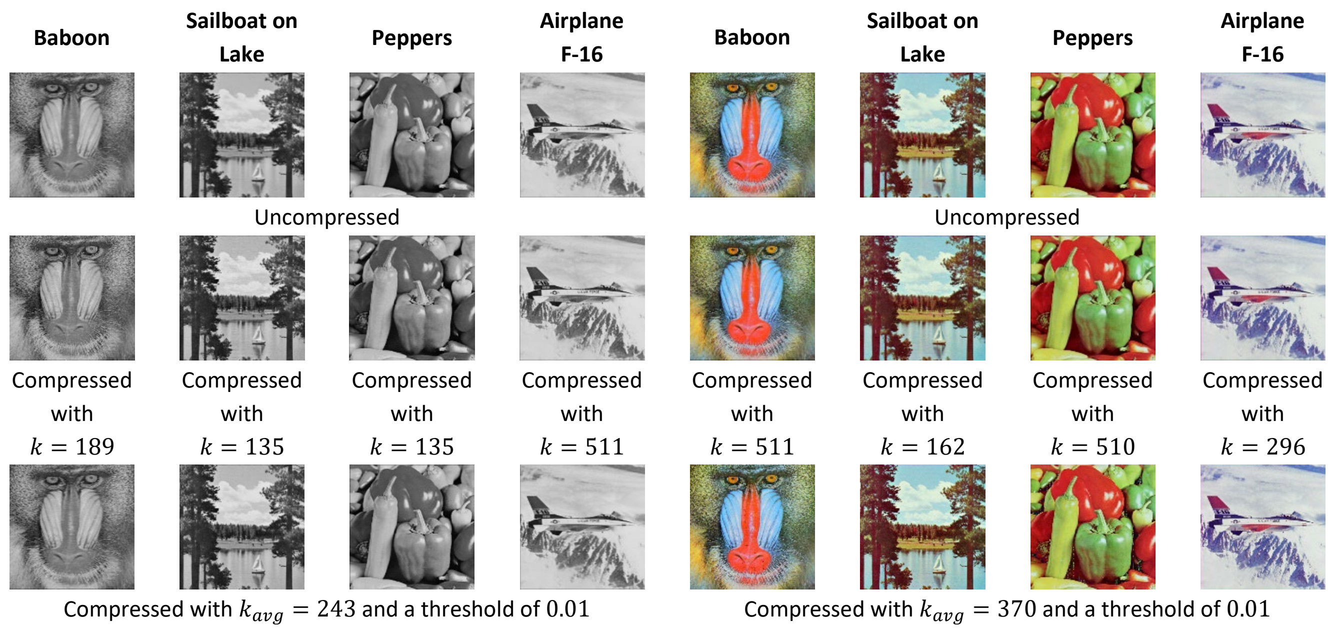

Figure 4. Comparison of uncompressed and compressed images using optimal compression parameters in ID: grayscale images on the left and colour images on the right.

Table 3. Compression results for the four colour images using ID with $k_{avg} = 370$ with a threshold of 0.01.

Image	Baboon	Sailboat on Lake	Peppers	Airplane F-16
MSE	34.57612	7.379738	200.8264	0.687366
PSNR (dB)	33.70509	40.14759	30.47195	49.79675
SSIM	0.966854	0.977783	0.947861	0.992076
File Size (bytes)	123864	90560	82346	63464
Compression Ratio	6.350287	8.685645	9.552037	12.39399

Table 3 shows that all four colour images achieved high SSIM values, indicating good quality. Compression ratios were generally acceptable, except for “Sailboat on Lake” and “Peppers”. For “Sailboat on Lake”, the difference was from the gap between the average k-value and its optimal k-value. For “Peppers”, the larger compression ratio difference was due to lines of dots in darker areas (visible in Figure 4), caused by low-intensity regions being discarded with a small k-value, resulting in a higher MSE and a need for more information to accurately store the image.

To evaluate performance, compression parameters were tested on new images, using the Splash image for both grayscale and colour. The results, compared with their optimal values, are presented in Table 4. Table 4 and Figure 5 demonstrate that grayscale image compression results were acceptable, with improved quality and a tolerable compression ratio decrease despite a larger k-value. However, colour image compression faced issues, similar to those seen with the "Peppers" image, where lines of dots appeared in darker areas. This highlights the challenges of compressing images with dark regions using low-rank approximations. For 512×512 grayscale images, $k = 243$ and a threshold of 0.01 are recommended, while for colour images, $k = 370$ with the same threshold is recommended, although caution is needed for images with dark areas.

Table 4. Comparison results between compressing using optimal and generalised parameters on both grayscale and colour images using ID.

Image Mode	Grayscale		Colour	
Compression Parameters Used	Optimal	Generalised	Optimal	Generalised
k	55	243	484	370
Threshold	0.01	0.01	0.01	0.01
MSE	30.07837	1.971176	0.027419	159.4838
PSNR (dB)	33.34826	45.18355	68.61479	42.46367
SSIM	0.844145	0.978776	0.999614	0.974872
File Size (bytes)	42574	45729	55700	66448
Compression Ratio	3.922676	3.652037	14.12158	11.83741

DOI:https//doi.org/10.22452/mjs.vol45no2.8
Malaysian Journal of Science 45(2): 75-86 (June 2026)

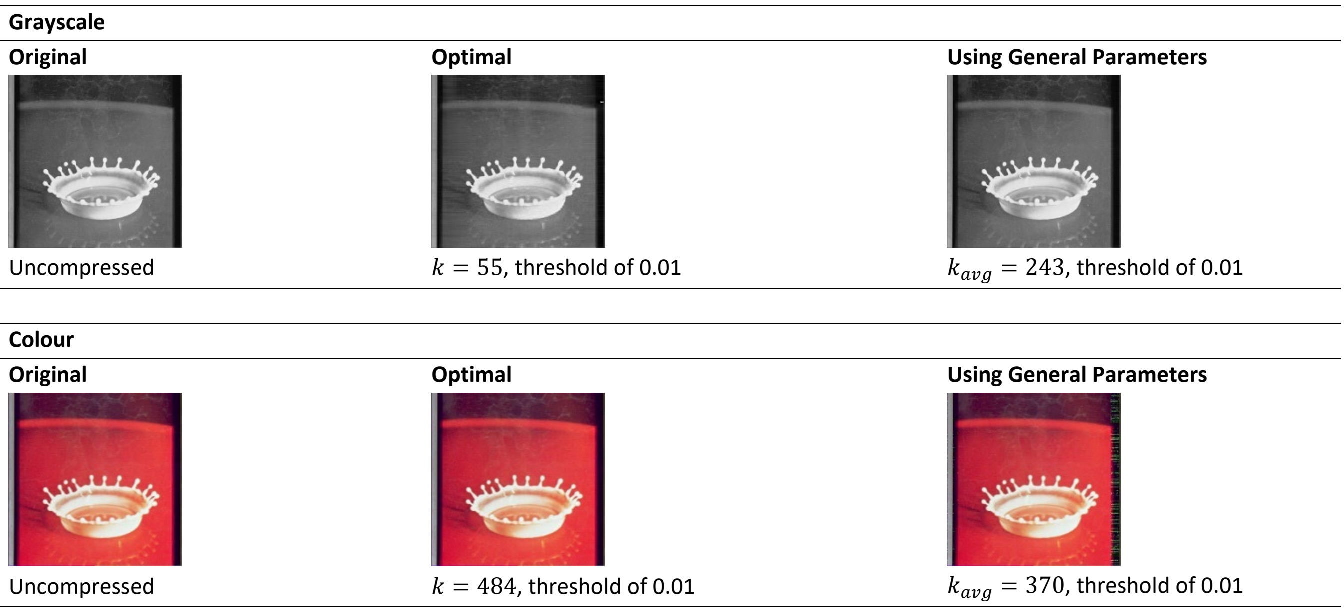

Figure 5. Comparative images of uncompressed, optimally compressed, and generally compressed versions using ID.

Optimal Parameters for SVD

The optimal compression parameters for SVD were determined similarly to ID, but without a threshold parameter. In total, 20 equidistant k-values were tested, and configurations with SSIM values below 0.8 were excluded. The optimal k-value was selected based on the highest compression ratio from the remaining results. Figure 6 presents a comparison of uncompressed and compressed images using optimal SVD compression parameters, with grayscale images shown on the left and colour images on the right.

The compression parameters were tested on the same image to assess their performance. Splash image was employed to evaluate the parameter for both grayscale and colour images. The results were compared with their optimal values and were listed in Table 5. The table shows that using a higher k-value improves image quality for both grayscale and colour images, despite a considerable decrease in the compression ratio. For colour images, the anomaly lines became less noticeable, resulting in significant visual improvements at the cost of larger file sizes.

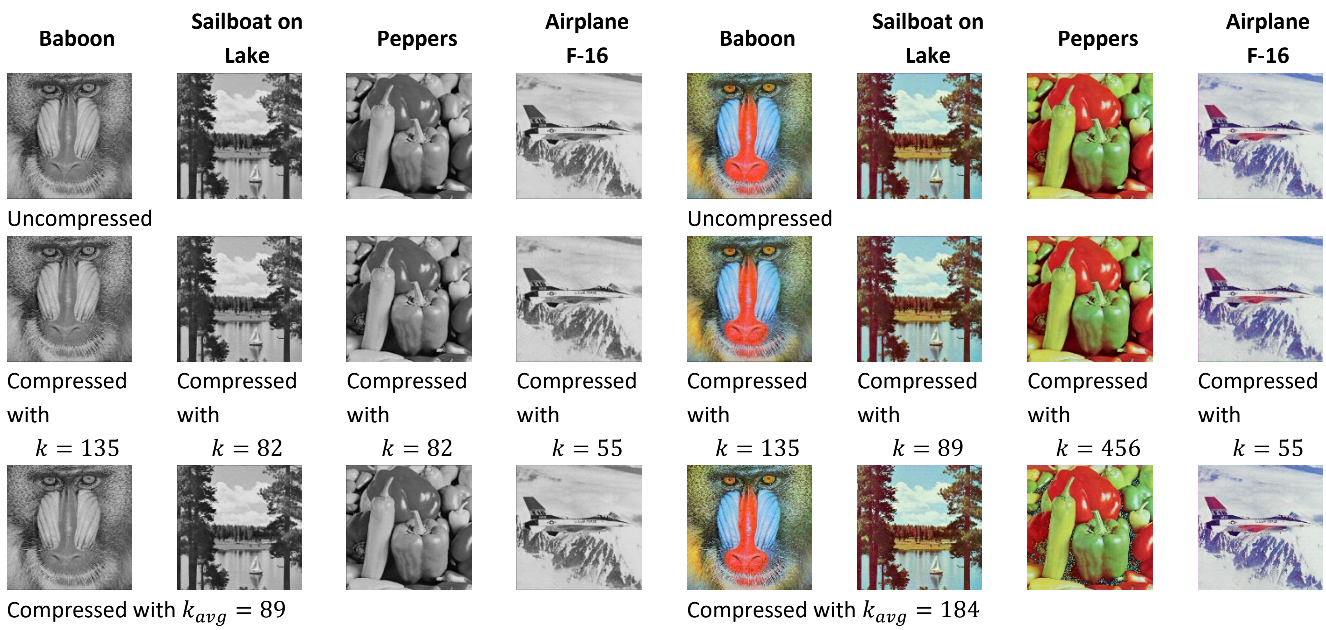

Figure 6. Comparison of uncompressed and compressed images using optimal SVD compression parameters: grayscale images on the left and colour images on the right.

DOI:https//doi.org/10.22452/mjs.vol45no2.8
Malaysian Journal of Science 45(2): 75-86 (June 2026)

Table 5. Comparison results between compressing using optimal and the general parameters for both grayscale and colour images using SVD.

Image Mode	Grayscale		Colour	
Compression Parameters Used	Optimal	Generalised	Optimal	Generalised
k	28	89	28	184
MSE	31.92671	6.080673	451.7643	24.06033
PSNR (dB)	33.08926	40.29129	29.71018	40.92835
SSIM	0.836819	0.937818	0.835039	0.971149
File Size (bytes)	33648	41180	51226	58242
Compression Ratio	4.963267	4.055464	15.35494	13.50524

Comparative Analysis of ID-method and SVD-method

The performance of ID and SVD in image compression highlights a trade-off between image quality (SSIM) and file size (compression ratio), as they are inversely related. Figure 7 shows that ID achieves higher SSIM values than SVD for the tested images when the rank k exceeds the rank of SVD, as determined by our methodology. Grayscale images exhibit lower SSIM values compared to colour images, as colour images retain more information across their three channels.

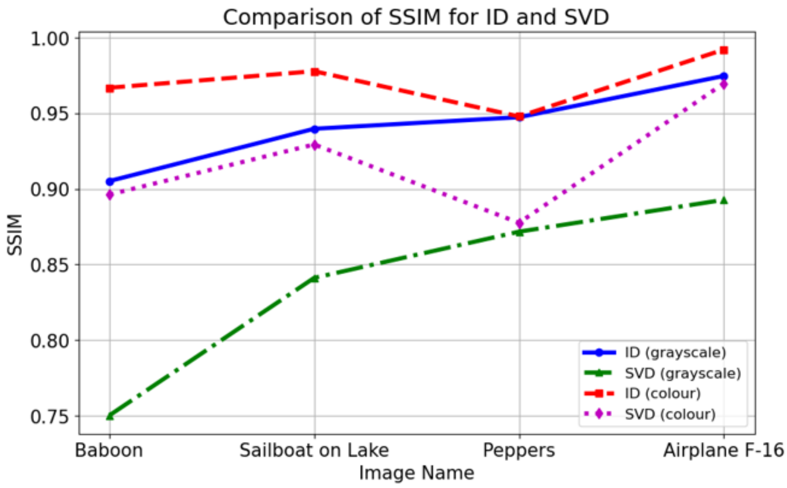

Figure 7. Comparison of SSIM using the generalised parameters for ID and SVD.

5. Computational Performance Comparison

The efficiency of matrix decomposition techniques is typically evaluated not only by their ability to reconstruct data accurately but also by their computational efficiency. This section offers a comparative analysis of the runtime and memory consumption of Interpolative Decomposition (ID), emphasizing its practical implications in image compression tasks.

Experimental Setup

The experiments were conducted on a machine equipped with an Intel Core i7-11700 CPU, 16 GB of RAM, and running Python 3.10 on Windows 11. ID computations were performed using the scipy.linalg.interpolative module. Runtime measurements were recorded with the time module, and peak memory consumption was monitored using the memory profiler package. ID was evaluated on a set of 512×512 colour images from the USC-SIPI Image Database, ensuring consistent matrix rank (k) values based on predefined compression parameters.

Results and Analysis

Table 6 presents the runtime results for ID on selected 512×512 colour images with a rank of $k = 370$. The experimental data demonstrate that Interpolative Decomposition (ID) consistently achieved efficient runtimes for these image decompositions. ID completed matrix decompositions faster for images with detailed features such as textures and edges, as seen with images like Baboon and Peppers.

It is vital to highlight that on standard CPUs, computing the full SVD of a 512×512 matrix typically takes between 1.5 to 3 seconds. This is consistent with theoretical expectations, as the computational complexity of SVD is $O(mn \cdot \min(m, n))$, which is generally more demanding compared to the $O(mk)$ complexity of ID.

DOI:https//doi.org/10.22452/mjs.vol45no2.8
Malaysian Journal of Science 45(2): 75-86 (June 2026)

Table 6. Runtime of ID on the selected 512×512 colour images for $k = 370$.

Image	Runtime (seconds)
Baboon	1.303
Sailboat on Lake	1.088
Peppers	0.958
Airplane F-16	0.841
Splash (Test Image)	0.721

These results are consistent with theoretical expectations, as the computational complexity of ID is $O(mk)$, where m is the number of rows and k is the rank of the approximation. The results suggest that ID not only achieves faster runtimes but also maintains stable performance across different types of image content, even at high resolutions (512×512).

Practical Implications

The computational benefits of ID have significant real-world implications. It is beneficial in applications where speed and memory efficiency are essential, such as real-time image processing, mobile computing, or large-scale data management. ID performs well because it uses less computing power and memory, while still delivering good image quality. This makes it a practical choice for devices with limited resources.

Our study demonstrates that ID achieves faster runtimes compared to typical SVD performance on the same image resolutions while maintaining consistent performance across different types of image content. This highlights its suitability for applications requiring rapid decompositions with minimal resource consumption, especially in real-time processing and edge computing environments.

Interpolative Decomposition (ID) and Singular Value Decomposition (SVD) have demonstrated their effectiveness in various real-world applications. ID is applied in quantum mechanics for solving nonlinear eigenvalue problems, particularly within Kohn-Sham density functional theory (DFT), where it efficiently approximates the density matrix through selected columns (Damle et al., 2017). Additionally, ID is examined in signal processing for musical instrument source separation, although its randomised column selection occasionally introduces redundancy (Tang et al., 2023). In contrast, SVD is extensively employed in medical image compression, video streaming optimization, and AI-based image classification. Its ability to retain key image features while reducing dimensionality makes it suitable for applications requiring both storage efficiency and image fidelity (Parameshachari et al., 2023; Libal et al., 2020). Furthermore, SVD has been adapted for hexagonal grid image compression, inspired by the structure of the human retina, demonstrating enhanced storage efficiency (Varghese & Saroja, 2021). These examples demonstrate the versatile applications of ID and SVD across various domains, underscoring their significance in data reduction tasks.

6. Conclusion

This study examined the effectiveness of Interpolative Decomposition (ID) and Singular Value Decomposition (SVD) for image data reduction through low-rank approximation. ID was shown to maintain matrix sparsity and reusing original columns, resulting in reduced storage and computational cost. However, it requires a higher rank k to achieve acceptable quality for complex or dense images.

Experimental results demonstrated that both ID and SVD yield comparable performance in image compression. Images with high variability or dark regions, such as “Baboon” and “Peppers,” required larger values of k to retain visual quality. SVD consistently achieved better compression ratios, making it suitable where storage efficiency is a priority. In contrast, ID better preserved image fidelity, especially under constrained computational resources. Generalised compression parameters for grayscale and colour images were proposed and tested across various images. For grayscale images, ID with a threshold of 0.01 and SVD with moderate k values produced effective results. For colour images, a similar thresholding was applied, although caution is needed when compressing images with dark regions. For reproducibility, generalised compression parameters were proposed: for 512×512 grayscale images, $k = 243$ with a threshold of 0.01 for ID and $k = 89$ for SVD; for colour images, $k = 370$ with a threshold of 0.01 for ID and $k = 184$ for SVD.

The findings obtained in this study indicate that method selection should be guided by specific application goals, whether prioritizing speed and memory usage (ID) or optimal storage and accuracy (SVD). Future work may explore the application of these methods on real-world health tabular data, such as the CDC Diabetes Health Indicators. By using ID on health data, we can analyse redundant information and highlight important health indicators.

DOI:https//doi.org/10.22452/mjs.vol45no2.8

Malaysian Journal of Science 45(2): 75-86 (June 2026)

7. References

Achlioptas, D., McSherry, F. (2001). Fast computation of low rank matrix approximations. *In Proceedings of the thirty-third annual ACM symposium on Theory of computing (STOC '01). Association for Computing Machinery*, New York, NY, USA, 611-618. https://doi.org/10.1145/380752.380858

Advani, R., O'Hagan, S. (2022). Efficient algorithms for constructing an interpolative decomposition. arXiv.

Aruchunan, E., Siri, Z., Noor Aziz, M. H. B., Ab Wahab, M. H. B., Muthuvalu, M. S., Sulaiman, J. (2022). Solution of peak junction temperature with Crank-Nicolson and SOR approach. *In Intelligent Systems Modeling and Simulation II: Machine Learning, Neural Networks, Efficient Numerical Algorithm and Statistical Models* (pp. 225-234). Springer. https://doi.org/10.1007/978-3-031-04028-3_15

Bentbib, A. H., Kreit, K., Labaali, I. (2022). Randomized tensor singular value decomposition for multidimensional data compression. *In 2022 11th International Symposium on Signal, Image, Video and Communications (ISIVC)* (pp. 1-6). El Jadida, Morocco: IEEE.

Bhaskara, A., Lattanzi, S., Vassilvitskii, S., Zadimoghaddam, M. (2020). Residual based sampling for online low rank approximation. In *2020 Information Theory and Applications Workshop (ITA)*. San Diego, CA, USA: IEEE. 1-19. https://doi.org/10.1109/ITA50056.2020.9244974.

Damle, A., Lin, L., Ying, L. (2017). SCDM-k: Localized orbitals for solids via selected columns of the density matrix. *Journal of Computational Physics*, 334, 1-15. https://doi.org/10.1016/j.jcp.2016.12.053.

Golub, G. (1965). Numerical methods for solving linear least squares problems. *Numerische Mathematik*, 7(3), 206-216.

Halko, N., Martinsson, P. G., Tropp, J. A. (2011). Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review,* 53(2), 217-288. https://doi.org/10.1137/090771806.

Kawamura, H., Suda, R. (2021). Least upper bound of truncation error of low-rank matrix approximation algorithm using QR decomposition with pivoting. *Japan Journal of Industrial and Applied Mathematics*, 38(3), 757-779. https://doi.org/10.1007/s13160-021-00459-x.

Khoei, T. T., Singh, A. (2024). Data reduction in big data: a survey of methods, challenges and future directions, *International Journal of Data Science and Analytics*, 2364-4168.

Libal, U., Baras, J. S., Johansson, K. H. (2020). Dimensionality reduction of volterra kernels by tensor decomposition using higher-order SVD. In *2020 59th IEEE Conference on Decision and Control (CDC)*. Jeju, Korea (South): IEEE. 5935-5941. https://doi.org/10.1109/CDC42340.2020.9303951.

Li, S., Lu, J., Hu, Y., Mattos, L. S., Li, Z. (2025). Towards scalable medical image compression using hybrid model analysis. *Journal of Big Data*, 12(45). https://doi.org/10.1186/s40537-025-01073-1

Liberty, E., Woolfe, F., Martinsson, P. G., Rokhlin, V., Tygert, M. (2007). Randomized algorithms for the low-rank approximation of matrices. *Proceedings of the National Academy of Sciences* 104(51), 20167-20172. https://doi.org/10.1073/pnas.0709640104(2007).

Liu, W., He, M. (2019). Accelerating solution of volume-surface integral equations with multiple right-hand sides by improved skeletonization techniques. *IEEE Antennas and Wireless Propagation Letters* 18(10), 2006-2010.

Lu, J., Ying, L. (2015). Compression of the electron repulsion integral tensor in tensor hypercontraction format with cubic scaling cost. *Journal of Computational Physics* 302, 329-335. https://doi.org/10.1016/j.jcp.2015.09.014.

Mersereau, R. M. (1979). The processing of hexagonally sampled two-dimensional signals. *Proceedings of the IEEE* 67(6), 930-949. https://doi.org/10.1109/PROC.1979.11356.

Muravev, A., Tran, D. T., Iosifidis, A., Kiranyaz, S., Gabbouj, M. (2018). Acceleration approaches for big data analysis. In *2018 25th IEEE International Conference on Image Processing (ICIP)*. Athens: IEEE, 311-315. https://doi.org/10.1109/ICIP.2018.8451082.

Parameshachari, B. D., Kumar, D. S., Prafulla, P. S., Yashwanth, J. (2023). Singular Value Decomposition (SVD) based optimal image compression technique. *In 2023 International Conference on Evolutionary Algorithms and Soft Computing Techniques (EASCT)*. Bengaluru, India: IEEE, 1-6.

Strang, G., 2006. Linear algebra and its applications. 4th ed. Boston: Cengage Learning.

Su, Q., Wang, G., Zhang, X., Lv, G., Chen, B. (2018). A new algorithm of blind color image watermarking based on LU decomposition. *Multidimensional Systems and Signal Processing*, 29(3), 1055-1074. https://doi.org/10.1007/s11045-017-0487-7.

Tang, W. K. A., Ng, W. S., Liew, H. H. (2023). Separation of two musical instruments using matrix factorisation techniques. *International Journal of Applied Mathematics* 36(3), 425. http://dx.doi.org/10.12732/ijam.v36i3.8.

DOI:https//doi.org/10.22452/mjs.vol45no2.8
Malaysian Journal of Science 45(2): 75-86 (June 2026)

Varghese, P., Saroja, G. A. S. (2021). Hexagonal image compression using singular value decomposition in Python. In *2021 2nd International Conference on Advances in Computing, Communication, Embedded and Secure Systems (ACCESS)* IEEE. 211-215. https://doi.org/10.1109/ACCESS51619.2021.9563312.

Wang, Z., Simoncelli, E. P., Bovik, A. C. (2003). Multiscale structural similarity for image quality assessment. *In The Thrity-Seventh Asilomar Conference on Signals, Systems & Computers 2003*, 2, 1398-1402. IEEE. https://doi.org/10.1109/ACSSC.2003.1292216.

Wüthrich, C. A., Stucki, P. (1991). An algorithmic comparison between square-and hexagonal-based grids. *CVGIP: Graphical Models and Image Processing*, 53(4), 324-339. https://doi.org/10.1016/1049-9652(91)90036-J.

DOI:https//doi.org/10.22452/mjs.vol45no2.8
Malaysian Journal of Science 45(2): 75-86 (June 2026)