https://doi.org/10.2352/CIC.2025.33.1.35

An Empirical Evaluation of Down- and Up-sampling in Jacobi Retinex

Afsaneh Karami and Graham Finlayson. University of East Anglia, Norwich, England.

Abstract

The McCann99 Retinex is a particular algorithmic implementation of Retinex theory. Retinex algorithms attempt to compress the dynamic range of an input image to make an output which can be displayed and has a pleasing rendition. The McCann99 Retinex algorithm employs a multi-grid structure, processing the image in a pyramidal manner from coarse to fine resolutions. This hierarchical approach is designed to enhance computational efficiency and accelerate convergence. In recent work, the structure of the McCann99 Retinex has - by appealing to classical work in numerical analysis - been modified to be a per pyramid level Jacobi iteration and as such is not only faster, but aligns Retinex with other theories of early visual processing (including centre-surround operators and path-based theories).

In this paper, we conduct a complementary systematic investigation into how images are down- and up-sampled as part of the multi-grid computation. The McCann99 and Jacobi Retinex algorithms use a simple box-filter approach, which means that down-sampled image data must be aliased. We investigate a set of other commonly used sampling methods, including bilinear, bicubic, Lanczos2 and 3, sinc with rectangular, Hamming and Kaiser windows. We are interested in using image quality metrics to determine whether substituting a different sampling approach into the Jacobi Retinex leads to better quality outputs being generated. Our investigation leads us to an interesting conclusion: while more complex down- and up-sampling algorithms can improve image quality, the simple box method performs well. If image quality is the primary concern, sinc is found to be the best method for down- and up-sampling. However, it requires an appropriate window function to taper the truncation and avoid ringing artefacts.

Introduction

A digital image is a two-dimensional array of discrete numerical values representing light intensity or colour at specific spatial locations, which serves as a sampled approximation of a real-world continuous scene. To ensure that the discretised image preserves all the essential information from the original scene, the Nyquist–Shannon sampling theorem provides a critical theoretical foundation. According to the theorem, a continuous image can be perfectly reconstructed from its samples if it is band-limited and sampled at a rate at least twice its highest frequency in both spatial dimensions [7]. If this condition is not satisfied, aliasing occurs, causing high-frequency components to fold into the low-frequency range, leading to signal distortion [7]. This results in visual artefacts such as wavy patterns, jagged edges, checkerboard effects, and a loss of image fidelity. It is a fundamental issue when converting a continuous image to a digital one or when down-sampling images without a proper sampling rate. High-frequency components that cannot be accurately represented at lower resolutions must be removed to prevent aliasing and preserve visual quality.

The Jacobi Retinex [9]- an improvement of the McCann99 algorithm [16] - employs a multi-resolution pyramid of an HDR image, ranging from a coarse representation at the top to fine details at the bottom, as represented in Fig. 2. This hierarchical structure facilitates the separation and correction of both local and global illumination effects and speeds up the calculation process. In the Jacobi Retinex [9], down-sampling with the scale of 0.5 - each level of the pyramid is half the resolution of the level below - is used to construct the multi-resolution pyramid and up-sampling is used to give the initial solution to the Jacobi reintegration part at the top level. However, potentially, the choice of sampling methods—such as nearest-neighbour, bilinear, bicubic, box, Lanczos, or frequency-based filters—plays a critical role in determining the quality of both the pyramid and the final Retinex output, as well as the speed of convergence.

Each method introduces varying degrees of blurring, aliasing, halo artefacts, edge preservation, noise amplification, and information loss, which can significantly influence the algorithm's ability to enhance image detail while preserving naturalness. For illustration, Fig. 1 compares the Jacobi Retinex results when four different sampling strategies are used. The top row displays the outputs of box-nearest (box for down-sampling and nearest-neighbour for up-sampling) with its zoomed-in patch to show artefacts. Notice there are dark halos around the edges of the window and the shadow in front of the window. Of course, halos such as these are intrinsic to Retinex processing. But to what extent are they dependent on the down- and up-sampling used in building the processing pyramid?

In the second row of the figure, we show the output of down- and up-smapling, using sinc with a Hamming window [8]. While halos remain, they are somewhat reduced in conspicuity. Rows 3 and 4 show processing outputs using bilinear [7] and Lanczos2 kernels [4]. In terms of Retinex-related artefacts, such as halos, the sinc with a Hamming window produces noticeably cleaner results compared to others. Bilinear performs better than Lanczos2, possibly due to its smoother interpolation kernel, which may suppress high-frequency components more effectively. Lancoz2 has noticeably more conspicuous halos than Box sampling.

In this paper, we will investigate the effect of various down- and up-sampling filters on the quality of images produced by Jacobi Retinex with a focus on naturalness [3], structural fidelity [25], and the formation of halo artefacts. Broadly, we find that - despite its inherent aliasing - the box down-sampling method works well. Arguably, Sinc down-sampling may be a touch better in terms of Image quality metrics, but this modest improvement is at the cost of expensive down-sampling.

Figure 1: Jacobi Retinex using different resizing methods with the zoomed-in patches on the left highlighting the halo artefact.

The remainder of this paper is structured as follows: the background section provides an overview of the Jacobi Retinex algorithm and sampling methods. The methodology section outlines our experimental setup and evaluation strategy. The results section presents and analyses our findings. Finally, the paper concludes with a summary of key insights and potential directions for future work.

Background

Retinex is both a theory [13] of aspects of colour perception and an image processing algorithm. In terms of image processing, it can be used to compress the dynamic range [24], solve for colour correction [14] and enhance an image for preference [18]. The Jacobi Retinex - appealing to well-known theories of equation solving in numerical analysis [6] - reformulated the McCann99 Retinex algorithm [16] into an iterative convolution-based centre-surround structure. This new formulation offers several advantages: improved compatibility with receptive field processing in the human visual system, faster convergence, enhanced image quality, and reduced halo artefacts [9].

The Jacobi Retinex structure, detailed in Algorithm 1, consists of two main steps. First, it computes the gradient field of the image at multiple resolutions. Second, it reconstructs the desired image by iteratively reintegrating this gradient field using the Jacobi method, referred to as incomplete integration. This process of calculating derivatives and integration is repeated progressively in the image pyramid in Fig. 2, from top to bottom. This algorithm retains three fundamental components from McCann99: ratio, reset, and averaging. As for most Retinex algorithms, we assume that all image processing is in the log-domain, where the non-log domain is the interval $(0, 1]$.

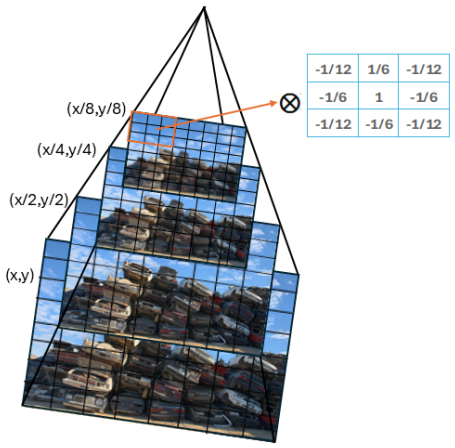

Figure 2: Multi-resolution pyramid of image in Jacobi Retinex

The gradient computation is carried out in line 3 of Algorithm 1, where the image $I(u, v)$ at resolution level L is convolved with the Laplacian operator ($\mathscr{L}$)—this operation is equivalent to the ratio step in McCann99. The down-sampling step is implemented at line 2.

Line 10 in Algorithm 1 represents the Jacobi iteration for an *incomplete* reintegration process. The up-sampling is implemented here for the initial guess at the Jacobi iteration for the next finer resolution of the pyramid. At the coarsest level ($L = 1$), the initial estimate is a zero matrix (line 5), whereas for higher levels ($L > 1$), the initial guess is obtained by up-sampling the solution from the previous level $L - 1$ by a factor of two (line 7). The reset and averaging steps from McCann99 are also implemented here. For details on how this algorithm links to the McCann99 Retinex and its relationship to Jacobi iteration in numerical analysis, see [9].

For down-sampling - the key focus of this paper - in line 2

Algorithm 1 Jacobi Retinex

for L=1:F **do**
  $I^L(u,v) = \downarrow_{box} I^{L+1}(u,v)$
  Calculate ratio $R^L(u,v) = I^L(u,v) \bigotimes \mathscr{L}$
  if L=1 **then**
    Initialise $O^L(u,v) = \mathbf{0}^L$
  else
    Initialise $O^L(u,v) = \uparrow_{nearest} O^{L-1}(u,v)$
  end if
  for k=1:iterations **do**
    $O^L_k(u,v) = \frac{1}{2}(O^L_{k-1}(u,v) + \min(R^L(u,v) + O^L_{k-1}(u,v) \bigotimes \mathscr{C}, 0))$
  end for
end for

Algorithm 1, a box filter is used (as it is in McCann99[16]). A box filter averages uniformly over a 2×2 pixel region, and then the image is decimated. Each output pixel is the average of four neighbouring pixels in the original image, reducing the image resolution by a factor of 2 while smoothing local variations. For up-sampling, nearest-neighbour interpolation is used. Here, we replicate each row and column of pixels to double the image size. This combination of box and nearest-neighbour methods aligns with the approach used in the original McCann99 Retinex algorithm. In this paper, we aim to investigate how alternative down- and up-sampling techniques affect the image quality produced by the Jacobi Retinex algorithm.

Method

To evaluate the impact of different down- and up-sampling methods on the performance of the Jacobi Retinex, we considered three aspects: (1) Retinex convergence speed, assessed by the image quality - after 24 iterations - by metrics TMQI [27] and FSITM [28]; (2) image similarity to the McCann99 output, quantified using the colour difference metric ΔE_{76} [19]; and (3) the ability of sampling methods to produce high-quality images regardless of iteration count, evaluated using metrics TMQI and FSITM.

TMQI [27] assesses image quality by combining multiscale structural fidelity TMQI_S and statistical naturalness TMQI_N. Structural fidelity captures how well the tone-mapped image retains details from the original HDR scene, while statistical naturalness evaluates how visually plausible the image appears based on the distribution of intensities found in natural scenes. Higher TMQI values correspond to better perceptual quality. FSITM [28] measures similarity between the enhanced image and the original HDR image by analysing local phase information. It compares locally weighted mean phase angle maps to assess how well structural and perceptual features are preserved. FSITM provides separate scores for the red, green, and blue channels; in this study, we report the average across all three channels as the final score.

Changing the image size typically involves weighted averaging—either to combine multiple pixels into one during down-sampling or to estimate new pixel values when inserting rows or columns during up-sampling. As an example, in Bilinear down-sampling, the image is convolved with a bilinear kernel, and then the image is decimated (selecting every second pixel along both the horizontal and vertical directions). In bilinear up-sampling, we replicate every pixel so have a double-resolution image where each 2x2 block of pixels has the same value. Then this replicated image is blurred with the same bilinear kernel.

This weighted averaging - and the difference between many down- and up-sampling methods - is governed by the interpolation kernel, which determines both the number of neighbouring pixels involved and the weights assigned to them based on their spatial distance from the target location. For instance, the Lanczos3 kernel - down-sampling to half resolution - uses six neighbouring pixels (three on each side), whereas methods like bilinear interpolation use only two. Each interpolation kernel impacts image properties differently, influencing factors such as sharpness, structure, noise, artefacts, and perceived naturalness [5].

To visually assess the performance of down- and up-sampling methods, Fig. 3 shows results of the image that was first down-sampled by a factor of 0.25 and then up-sampled by a factor of 4 and their TMQI value. Aliasing artefacts are visible in all, particularly around high-contrast edges such as electrical cables, except in the sinc-based. While sinc with a rectangular window introduces noticeable ringing artefacts [23], these can be mitigated by using window functions such as Kaiser [2] or Hamming [8], albeit at the cost of increased blurring. Ultimately, the key consideration is how each interpolation method influences image properties when integrated into the Jacobi Retinex framework.

All the interpolation methods in the spatial domain and the sinc-based approach in the frequency domain used in our experiments are summarised in Table 1. In this table, Nearest stands for nearest-neighbour [7], sinc-R is sinc with a rectangular window, sinc-K denotes sinc with a Kaiser window, and sinc-H indicates sinc with a Hamming window. Bilinear interpolation [7] uses two neighbours, while bicubic [10] and Lanczos2 [4] use four. Lanczos3 [4], being more extensive, uses six.

Table 1: Different methods and their interpolation functions

Method	Interpolation Function
Nearest	$\mathsf{f}(x) = \mathsf{round}(x)$
Bilinear	$f(x) = \begin{cases} 1-\lvert x\rvert, & \text{if } \lvert x\rvert \le 1 \\ 0, & \text{otherwise} \end{cases}$
Bicubic	$f(x) = \begin{cases} 1.5\lvert x\rvert^3 - 2.5\lvert x\rvert^2 + 1, & \lvert x\rvert \le 1 \\ -0.5\lvert x\rvert^3 + 2.5\lvert x\rvert^2 - 4\lvert x\rvert + 2, & 1 < \lvert x\rvert \le 2 \\ 0, & \text{otherwise} \end{cases}$
Box	$f(x) = \begin{cases} 1, & \text{if } \lvert x\rvert < 0.5 \\ 0, & \text{otherwise} \end{cases}$
Lanczos2	$f(x) = \mathsf{sinc}(x) \cdot \mathsf{sinc}\left(\frac{x}{2}\right),\ \lvert x\rvert < 2$
Lanczos3	$f(x) = \mathsf{sinc}(x) \cdot \mathsf{sinc}\left(\frac{x}{3}\right),\ \lvert x\rvert < 3$
Sinc-R	$f(x) = \mathsf{sinc}(x) \cdot w_{\mathsf{Rectangular}}(x)$
Sinc-K	$f(x) = \mathsf{sinc}(x) \cdot w_{\mathsf{Kaiser\text{–}Bessel\text{-}derived}}(x)$
Sinc-H	$f(x) = \mathsf{sinc}(x) \cdot w_{\mathsf{Hamming}}(x)$

For the sinc-based approach in the down-sampling step, we first transform the image into the frequency domain and retain only the central low-frequency components, typically half the frequency content along each dimension. A window with the same

(a) Box-nearest (0.884) (b) Sinc-Rectangular (0.909)

(c) Bicubic (0.892) (d) Lanczos3 (0.893)

(e) Lanczos2 (0.893) (f) Bilinear (0.863)

(g) Sinc-Kaiser (0.814) (h) Sinc-Hamming (0.729)

Figure 3: Different down- and up-sampling methods with TMQI

size as a down-sample, such as a Hamming window, applies to these selected components. For example, if the original image size is $N \times N$, the down-sample size is $\frac{N}{2} \times \frac{N}{2}$, we extract a frequency band (of the down-sample size) centred around the origin and multiply with window. The inverse Fourier transform returns the image to the spatial domain. This process effectively suppresses high-frequency components and helps prevent aliasing in the image, which can be seen in Fig. 3. For up-sampling, we perform zero-padding in the frequency domain to increase the size of the spectrum. After zero-padding, we apply a window function and transform the result back to the spatial domain using the inverse Fourier transform.

From a frequency-domain perspective, preserving most of the power spectral density (PSD) of the original image is critical for maintaining structural and textural fidelity during down-sampling [20]. Natural images typically exhibit $\frac{1}{f^2}$ spectral distribution, meaning most of their information is concentrated in low frequencies [22]. For down-sampling, the ideal low-pass filter is the sinc function, as it has a perfect rectangular response in the frequency domain, which ensures maximal preservation of low-frequency content below the Nyquist limit while suppressing higher frequencies (beyond Nyquist), thereby minimising aliasing distortions [5]. It omits several high-frequency components, leading to a loss of fine details compared to the original image [5]. Consequently, based on PSD analysis, we might expect sinc down-sampling to be theoretically optimal in preserving the most informative components of natural images. But we should consider that an ideal sinc filter has infinite support, while the spatial domain of the image is limited, and even with a considerable amount of harmonics, we still have ripples around the cut-off areas in the frequency domain —known as the Gibbs phenomenon [23] - leading to ringing artefacts around sharp transitions in the spatial domain (Fig. 3 image b). In our work, we also apply Kaiser–Bessel-derived and Hamming windows to the sinc function to preserve more image information while reducing the roll-off (undesirable ringing artefacts). These windows reduce the amplitude of side lobes (which cause ringing artefact) while maintaining a reasonably sharp frequency cutoff: (1) The Kaiser–Bessel-derived window offers adjustable parameters (via the β value), allowing control over the balance between main lobe width and side lobe attenuation, and we chose the $\beta = 40$. (2) The Hamming window provides strong suppression of side lobes and is widely used due to its simplicity and smooth profile. Nevertheless, because the windowing slightly distorts the ideal frequency response, it remains an open question of how these trade-offs affect final image quality, one of the key motivations for our experimental comparison.

How can we determine the most suitable method for image resizing? The optimal choice depends on the specific application and the primary objective—whether it is faithful reconstruction, high perceptual quality, or preservation of structural details. A variety of metrics have been proposed in the literature to evaluate the performance of different interpolation kernels for down- and up-sampling. Common choices include: (1) SSIM [26] and MSSIM [26], which measure structural similarity and are used for perceptual quality assessment; (2) MSE [7] and PSNR [15], which quantify pixel-wise errors and signal fidelity; (3) PEE (Percentage Edge Error) [1], which assesses the preservation of edges after resizing; and (4) frequency-domain criteria such as the error spectrum derived from the power spectral density of the image and the frequency response of the interpolation kernel [20]. In this research, we selected TMQI [27] (as it is widely used and in spite of its limitations [11]) and FSITM [28] as our evaluation metrics. TMQI is especially suitable for HDR tone-mapped images, combining structural fidelity (similar in spirit to SSIM) with a naturalness measure that reflects how visually pleasing the image appears. FSITM focuses on preserving visual features during tone mapping, making it well-aligned with our goal of evaluating

perceptual and structural image quality in the context of Retinex-based enhancement.

Experiments

Dataset: To evaluate the performance of our Retinex methods, we use the LVZ-HDR dataset [17], which offers a large and varied collection of high dynamic range images. It includes 457 HDR images captured under a wide range of conditions, encompassing different times of day, locations, and environments. The images are grouped into four major categories—Indoors (71), Nature (173), Nighttime (80), and River-side-sunset (133)—providing broad scene diversity for a comprehensive assessment.

IQ at a fixed number of iterations: We evaluated the TMQI and FSITM values of different down- and up-sampling methods after 24 iterations. This number of iterations was selected because the baseline Jacobi Retinex—using box filtering for down-sampling and nearest-neighbour for up-sampling—produces images that are similar to those of McCann99, based on the ΔE metric. Table 2 presents the results for the TMQI and FSITM metrics.

Among all methods, sinc-rectangular and box-nearest produced the highest TMQI scores (0.874). However, they differed in their subcomponents: sinc-rectangular achieved the highest naturalness score (TMQI_N) but had a slightly lower structural score (TMQI_S), whereas box-nearest had the highest structure but slightly lower naturalness. Among interpolation kernels, Lanczos3 has the highest TMQI. Regarding FSITM, box-nearest has the highest value.

Interestingly, combining sinc-rectangular down-sampling with nearest up-sampling (sinc with rectangular-nearest) slightly improved naturalness further (0.486) compared to sinc-rectangular, but this came at the cost of lower structure and FSITM. The interpolation-based kernels Lanczos3, Lanczos2, bicubic, and bilinear showed similar trends: using nearest up-sampling instead of their default method had a negligible effect on TMQI and FSITM, implying that the down-sampling method plays a more critical role in convergence and quality than the up-sampling method. The sinc-Kaiser and sinc-Hamming methods underperformed in both TMQI and FSITM compared to the sinc-rectangular approach. These methods, which are robust to aliasing and ringing artefacts - from a single down- and up-sampling viewpoint - are further from the starting image than other methods. Empirically, there is slower convergence for Retinex, and so, to achieve better quality scores, these sampling methods need to be run for more iterations. In Fig. 4, we plot per method the TMQI against FSITM scores. Notice box-nearest is furthest to the top right, indicating the best combined quality and has the fastest convergence speed.

Delta E: We evaluated the minimum number of iterations required for each method - Jacobi Retinex plus a particular sampling approach - to achieve a perceptual colour difference (CIELAB ΔE [19]) below 3, using McCann99 as the reference (since for complex images a mean CIELAB $\Delta E_{76} < 3$ is not visually significant[21]). The McCann99 Retinex outputs were all run for a fixed number of iterations (32). In most cases, according to the author's preference, this produced pleasing outputs. Of all the down-sampling methods implemented, only box-nearest can meet this 3 Delta E threshold, likely due to its alignment with the original McCann99 method. For all other methods, *DeltaE* remained above 3, decreasing initially but rising again with further iterations. The lowest *DeltaE* values for each method are reported in Table 3 together with the number of iterations needed to meet this lower-bound.

Table 2: Image quality results of the Jacobi Retinex using different sampling methods with a fixed number of 24 iterations

Method (down-up)	TMQI	TMQI_S	TMQI_N	FSITM
sinc_R-sinc_R	**0.874**	0.854	0.478	0.885
box-nearest	**0.874**	**0.862**	0.466	**0.897**
sinc_R-nearest	0.873	0.844	**0.486**	0.862
lanc3-lanc3	0.873	0.850	0.476	0.885
lanc3-nearest	0.873	0.850	0.475	0.885
lanc2-lanc2	0.869	0.853	0.448	0.885
bicubic-bicubic	0.869	0.853	0.447	0.885
bicubic-nearest	0.869	0.853	0.447	0.885
lanc2-nearest	0.869	0.853	0.447	0.885
bilinear-nearest	0.859	0.848	0.407	0.880
bilinear-bilinear	0.859	0.847	0.403	0.879
sinc_K-sinc_K	0.854	0.843	0.385	0.875
sinc_H-sinc_H	0.809	0.808	0.211	0.847

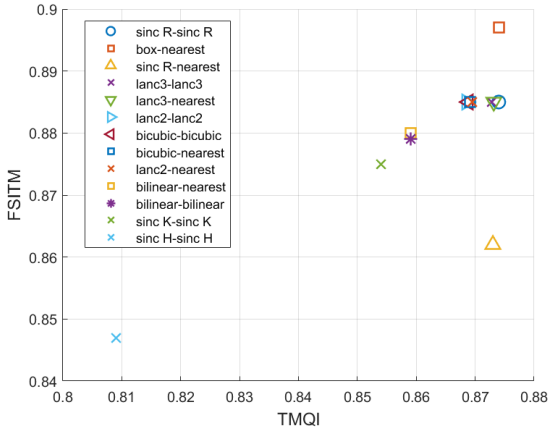

Figure 4: FSITM-TMQI for different methods at 24 iterations

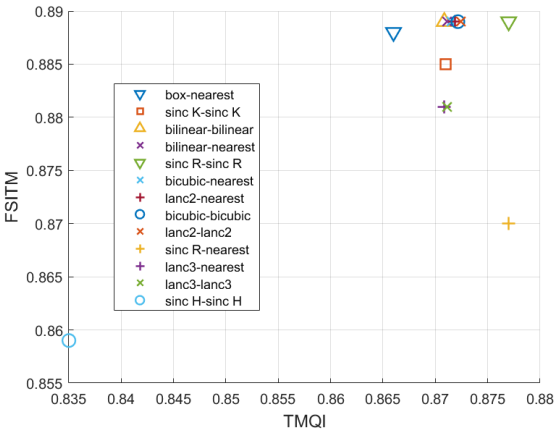

Figure 5: FSITM-TMQI for different methods at lowest delta E

Allowing different numbers of iterations for each method does change how well the methods perform in terms of our IQ measures. Now **every** method (other than sinc with Hamming

Table 3: Result of delta E (CICED76) for different methods

Method (down-up)	#its	TMQI_Q	FSITM	Min ΔE
box-nearest	16	0.866	0.888	**2.6**
sinc_K-sinc_K	40	0.871	0.885	3.5
bilinear-bilinear	40	0.871	**0.889**	3.9
bilinear-nearest	38	0.871	**0.889**	4.0
sinc_R-sinc_R	30	**0.877**	**0.889**	4.4
bicubic-nearest	30	0.872	**0.889**	4.9
lanc2-nearest	30	0.872	**0.889**	4.9
bicubic-bicubic	30	0.872	**0.889**	4.9
lanc2-lanc2	30	0.872	**0.889**	4.9
sinc_R-nearest	30	**0.877**	0.870	5.5
lanc3-nearest	20	0.871	0.881	6.1
lanc3-lanc3	20	0.871	0.881	6.1
sinc_H-sinc_H	40	0.835	0.859	9.1

window) has a higher TMQI score than the box. Though typically a few more iterations are needed. In Fig. 5, we plot per method the TMQI against FSITM scores. Notice that most methods have slightly higher TMQI and FSITM scores than the box. Thus, if the iteration count is not an issue (we can run the Retinex for a few more iterations), then changing the sampling method generally improves over the box. The standard sinc with rectangular window works best overall.

IQ as a function of the number of iterations: Fig. 6 and 7 illustrate, respectively, the TMQI and FSITM scores for various resizing methods for different numbers of iterations. The sinc method with a rectangular window achieved the highest TMQI, while the box-nearest method yielded the highest FSITM score.

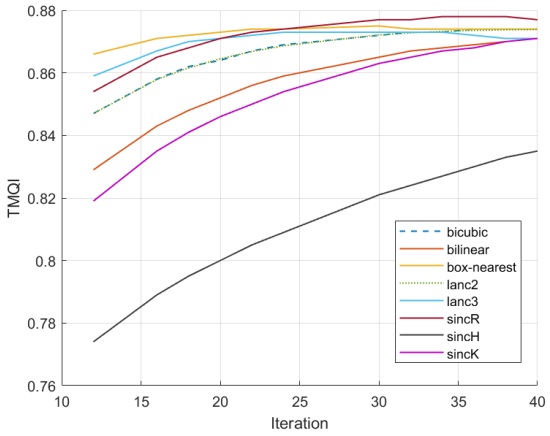

Figure 6: TMQI result for different methods at different iterations

IQ of Down- and Up-sampling: Finally, we investigated the effect of different down- and up-sampling methods on image quality using the widely used Kodak dataset [12] of 24 photographic images, evaluating performance with the PSNR, TMQI_Q, and FSITM metrics. Images were down-sampled by a factor of 0.25 and then up-sampled by a factor of 4, and their quality was reported in Table 4. The results show that sinc with a rectangular window outperforms the other methods in terms of both PSNR and TMQI. The lowest PSNR is observed with the box-nearest method, while sinc with a Hamming window yields the lowest TMQI_Q score. For FSITM, the differences between methods are minimal; however, Lanczos2, bicubic, and bilinear perform slightly better than the others. Contrasting this result with the

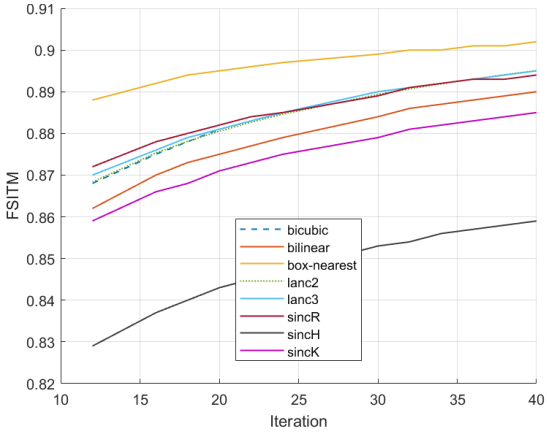

Figure 7: FSITM result for different methods at different iterations

others reported in this paper, we can conclude that methods designed only for effective down-sampling are not always optimal, see Fig. 4, when deployed in the context of the Retinex algorithm (fixed number of iterations), though there appears to be a weak signal that better down-sampling can help Retinex, see Fig. 5 (variable number of iterations).

Table 4: Image quality for different methods on Kodak dataset

Method(down-up)	PSNR	TMQI_Q	FSITM
sinc_R-sinc_R	**25.88**	**0.890**	0.52
box-nearest	22.56	0.856	0.53
lanc2-lanc2	24.34	0.877	**0.54**
lanc3-lanc3	23.64	0.867	0.52
bicubic-bicubic	24.35	0.877	**0.54**
sinc_K-sinc_K	23.81	0.859	0.49
sinc_H-sinc_H	22.93	0.850	0.49
bilinear-bilinear	24.80	0.874	**0.54**

Conclusion

We experimented with various down- and up-sampling methods deployed within the Jacobi Retinex and then compared the output images that were generated using metric-based image quality measures. Our results indicate that the box-nearest (the conventional down-sampling method used in Retinex) and sinc with rectangular window worked similarly when Retinex runs for a fixed number of iterations. Though, significantly, other well-known methods for down-sampling performed less well (surprisingly, so). However, when a per image variable number of iterations was allowed (so Retinex converges to a putative best quality image), there is a small but significant IQ advantage to using sinc and most other down-sampling methods.

From our investigations, we conclude that it is not easy to predict how down-sampling affects Retinex. Whereas down-sampling, using conventional techniques, simply for image rescaling is well understood (and the pros and cons of different approaches are well known), their relative performances do not always predict the performance of Retinex.

References

[1] Al-Fohoum, A. and Reza, A. M. [2001], 'Combined edge crispiness and statistical differencing for deblocking JPEG compressed im-

ages', *IEEE Transactions on Image Processing* **10**(9), 1288–1298.

[2] Bosi, M. and Goldberg, R. E. [2002], *Introduction to digital audio coding and standards*, Vol. 721, Springer Science & Business Media.

[3] Cadik, M. and Slavik, P. [2005], The naturalness of reproduced high dynamic range images, *in* 'Ninth International Conference on Information Visualisation (IV'05)', IEEE, pp. 920–925.

[4] Duchon, C. E. [1979], 'Lanczos filtering in one and two dimensions', *Journal of Applied Meteorology* pp. 1016–1022.

[5] Dumitrescu, D. and Boiangiu, C.-A. [2019], 'A study of image upsampling and downsampling filters', *Computers* **8**(2), 30.

[6] Golub, G. H. and Van Loan, C. F. [2013], *Matrix computations*, Johns Hopkins University Press.

[7] Gonzalez, R. C. and Woods, R. E. [2009], *Digital Image Processing*, 3rd edn, Pearson Education.

[8] Harris, F. J. [2005], 'On the use of windows for harmonic analysis with the discrete fourier transform', *Proceedings of the IEEE* **66**(1), 51–83.

[9] Karami, A. and Finlayson, G. [2024], Unifying path and center-surround Retinex algorithms, *in* 'London Imaging Meeting', Society for Imaging Science and Technology, pp. 17–21.

[10] Keys, R. [2003], 'Cubic convolution interpolation for digital image processing', *IEEE Transactions on Acoustics, Speech, and Signal Processing* **29**(6), 1153–1160.

[11] Khan, I. R., Alotaibi, T. A., Siddiq, A. and Bourennani, F. [2022], 'Evaluating quantitative metrics of tone-mapped images', *IEEE Transactions on Image Processing* **31**, 1751–1760.

[12] *Kodak Lossless True Color Image Suite* [n.d.], http://r0k.us/graphics/kodak/. Accessed: 2025-06-06.

[13] Land, E. H. and McCann, J. J. [1971], 'Lightness and retinex theory', *Journal of the Optical Society of America* **61**(1), 1–11.

[14] Li, C., Tang, S., Kwan, H. K., Yan, J. and Zhou, T. [2020], 'Color correction based on CFA and enhancement based on retinex with dense pixels for underwater images', *IEEE Access* **8**, 155732–155741.

[15] Lin, W. and Dong, L. [2006], 'Adaptive downsampling to improve image compression at low bit rates', *IEEE Transactions on Image Processing* **15**(9), 2513–2521.

[16] McCann, J. J. [1999], Lessons learned from Mondrians applied to real images and colour gamuts, *in* 'Color and Imaging Conference', Society for Imaging Science and Technology, pp. 1–8.

[17] Panetta, K., Kezebou, L., Oludare, V., Agaian, S. and Xia, Z. [2021], 'TMO-Net: A parameter-free tone mapping operator using generative adversarial network, and performance benchmarking on large scale HDR dataset', *IEEE Access* **9**, 39500–39517.

[18] Rahman, Z., Jobson, D. J. and Woodell, G. A. [2004], 'Retinex processing for automatic image enhancement', *Journal of Electronic Imaging* **13**(1), 100–110.

[19] Robertson, A. R. [1977], 'CIE recommendations on uniform color spaces, color-difference equations, and metric color terms', *Color Research & Application* **2**(1), 7–11.

[20] Schaum, A. [1993], 'Theory and design of local interpolators', *CVGIP: Graphical Models and Image Processing* **55**(6), 464–481.

[21] Stokes, M., Fairchild, M. D. and Berns, R. S. [1992], 'Precision requirements for digital color reproduction', *ACM Transactions on Graphics* **11**(4), 406—422.

[22] Torralba, A. and Oliva, A. [2003], 'Statistics of natural image categories', *Network: Computation in Neural Systems* **14**(3), 391.

[23] Unser, M. [2002], 'Splines: A perfect fit for signal and image processing', *IEEE Signal Processing Magazine* **16**(6), 22–38.

[24] Wang, L., Horiuchi, T. and Kotera, H. [2007], 'High dynamic range image compression by fast integrated surround retinex model', *Journal of Imaging Science and Technology* **51**(1), 34–43.

[25] Wang, Z. and Bovik, A. C. [2009], 'Mean squared error: Love it or leave it? a new look at signal fidelity measures', *IEEE Signal Processing Magazine* **26**(1), 98–117.

[26] Wang, Z., Bovik, A. C., Sheikh, H. R. and Simoncelli, E. P. [2004], 'Image quality assessment: from error visibility to structural similarity', *IEEE Transactions on Image Processing* **13**(4), 600–612.

[27] Yeganeh, H. and Wang, Z. [2012], 'Objective quality assessment of tone-mapped images', *IEEE Transactions on Image Processing* **22**(2), 657–667.

[28] Ziaei Nafchi, H., Shahkolaei, A., Farrahi Moghaddam, R. and Cheriet, M. [2014], 'FSITM: a feature similarity index for tone-mapped images', *IEEE Signal Processing Letters* **22**(8), 1026–1029.

Author Biography

Afsaneh Karami is a PhD student at East Anglia University under the supervision of Professor Graham Finlayson. A member of the Colour & Imaging Lab. Also, a member of AgriFoRwArdS CDT.

Professor Graham Finlayson is the Director of the Colour & Imaging Lab. He has published 300 conference papers, 100 journal papers and is the inventor of 30+ patents. His interests span perception, colour image processing and physics-based computer vision.