https://doi.org/10.2352/CIC.2025.33.1.10

Three dimensional non-local anisotropic diffusion

Ali Alsam, Hans Jakob Rivertz; Norwegian University of Science and Technology, Trondheim, Norway,
Ali.Alsam@ntnu.no, Hans.J.Rivertz@ntnu.no

Abstract

We present a novel anisotropic diffusion algorithm for noise reduction in Magnetic Resonance Imaging (MRI). The method integrates two key concepts: (1) diffusion is explicitly constrained to avoid increases in local image gradients, thereby preserving edges and fine structural details; and (2) a sequence of filters with exponentially increasing radii is applied, each maintaining a fixed number of non-zero coefficients. These filters allow the algorithm to evaluate whether pixels distant from the target location can contribute to smoothing without degrading local gradients. As a result, the method aims to balance between preserving local details and averaging global similarities.

In contrast to traditional denoising techniques based on local filtering or total variation minimization, the proposed algorithm enables controlled non-local diffusion and naturally extends to three-dimensional voxel arrays, making it well-suited for volumetric MRI data. The framework also permits the integration of additional geometric constraints, such as curvature, further enhancing its ability to preserve anatomical structures and surfaces.

The effectiveness of the proposed method is demonstrated on real MRI data from a macaque monkey. The experimental results indicate PSNR values comparable to those of our previous approach, while providing substantially better suppression of low-frequency noise, absence of visible artifacts, and faithful preservation of critical image features.

Introduction

Magnetic Resonance Imaging (MRI) is a widely used medical imaging technique that generates detailed images of internal body structures through the interaction of strong magnetic fields and radio frequency pulses. In essence, MRI operates by aligning the hydrogen protons in water and fat molecules using a powerful magnetic field. Subsequent radio frequency pulses temporarily disturb this alignment, and as the protons return to their original state, they emit signals that are captured and processed to reconstruct an image. This simplified description underscores a key aspect of MRI: the presence of noise in the resulting images is intrinsic to the physical and technical processes involved.

Image noise, and the challenge of reducing it, is a central concern in signal processing. Textbook denoising algorithms typically rely on some form of filtering, where a pixel's intensity is replaced by a weighted combination of its neighbours. These weights may reflect simple statistical measures such as the arithmetic mean or median, or incorporate the distance between the pixel and its neighbours.

Advanced denoising algorithms aim to suppress noise while preserving essential image features—an objective that is both critical and technically challenging. Several prominent approaches have been developed to achieve this balance, including anisotropic diffusion [1], bilateral filtering [2], total variation minimization [3], non-local means [4], and Block-Matching and 3D Filtering (BM3D) [5], which is regarded as one of the most effective non-deep learning-based methods.

In recent years, the field of image denoising has been increasingly dominated by deep learning approaches, particularly Convolutional Neural Networks (CNNs), which have demonstrated state-of-the-art performance across a range of denoising tasks [6]. There has also been work on a PDE based non-local diffusion method (NLD) [7].

In this paper, we integrate two previously published ideas [8, 9] into a unified algorithm for denoising MRI images without requiring access to ground truth data. The central principle of the proposed method is that any modification to a pixel's value—intended to smooth the image surface—should not result in an increase in the surrounding image gradients. This principle is operationalized by imposing a constraint on diffusion from neighbouring pixels toward the target pixel. We show that when diffusion is restricted to directions that decrease local gradient magnitudes, the result naturally preserves edges and fine structures. Furthermore, the diffusion constraint is generalized to include non-local pixel contributions. This is implemented using a series of filters that sample pixel values at increasing distances from the target location and assess whether their inclusion affects local gradient behaviour. The algorithm is thus capable of allowing diffusion from the entire image plane, provided that doing so does not distort the local image detail.

The proposed method extends seamlessly to three-dimensional voxel arrays, making it well-suited for volumetric MRI data. We evaluate the algorithm on real MRI data of a Macaque monkey, made available through the project of de Castro et al. [10]. The proposed method introduces no visible artifacts, improves the signal-to-noise ratio, and preserves edges and fine image details. Unlike traditional diffusion-based methods such as total variation, our approach permits diffusion from non-local locations and can be readily adapted to enforce constraints not only on gradient magnitude but also on geometric features such as curvature. Compared to deep learning methods, the algorithm requires no training data, is independent of ground truth images, and—perhaps most significantly—is simple to interpret and explain. This last attribute is particularly valuable, as technical advancements in imaging hardware

often benefit from conceptual insights derived from transparent, analytical algorithms. A detailed description of the algorithm is presented in the next section.

Method

The diffusion condition

An arbitrary target voxel I_0 is surrounded by 26 neighbouring voxels $I_1,\dots,I_{26}$. The directional differences are given by $d_i = I_i - I_0$, for $i = 1,2,\dots,26$. Diffusion from a source voxel I to I_0 is denied if a test diffusion from I to I_0 given by $I_0' = (1-s)I_0 + sI$ gives rise to a significant increase in any of the directional differences. The constant s is a test diffusion coefficient. The directional differences after a test diffusion are $d_i' = I_i - I_0' = d_i - sd$, where $d = I - I_0$. More precisely, we have the following condition.

Condition. Diffusion from I to I_0 is not allowed if

$$\frac{{d_i'}^2 - {d_i}^2}{s} > \lambda$$

for any $i = 1,2,\dots,26$, where λ is a threshold parameter.

Based on this, we get the following equivalent condition.

Theorem 1. *Diffusion from I to I_0 is allowed if and only if*

$$sd - \lambda/d \leq 2\min_i\{d_i\} \text{ for } d > 0. \quad (1)$$

$$sd - \lambda/d \geq 2\max_i\{d_i\} \text{ for } d < 0. \quad (2)$$

Corollary. *The diffusion condition of d to I_0 is equivalent to the inequality*

$$\frac{\max_i\{d_i\}}{s} - \sqrt{\frac{\lambda}{s} + \frac{(\max_i\{d_i\})^2}{s^2}} \leq d \leq \frac{\min_i\{d_i\}}{s} + \sqrt{\frac{\lambda}{s} + \frac{(\min_i\{d_i\})^2}{s^2}}. \quad (3)$$

Remark. For $\lambda = 0$, the condition for diffusion is

$$\frac{\max_i\{d_i\} - |\max_i\{d_i\}|}{s} \leq d \leq \frac{\min_i\{d_i\} + |\min_i\{d_i\}|}{s}.$$

When $\max_i\{d_i\} < 0$ or $\min_i\{d_i\} > 0$, diffusion will be allowed for all values of λ. Therefore, the method will effectively remove all salt and pepper noise.

Variable range diffusion

The unrestricted filters G_n, $n = 0,1,2,3,\dots$ are defined by $G_n(i,j,k) = g_n(i)g_n(j)g_n(k)$, where

$$g_n(i) = \begin{cases} \frac{2-|i|}{4} & \text{, for } i = -2^n, 0, 2^n. \\ 0 & \text{, otherwise.} \end{cases} \quad (4)$$

The filter G_n is shown in Equation (5). Only the non-zero coefficients are shown.

$$G_n = \frac{1}{64} \quad (5)$$

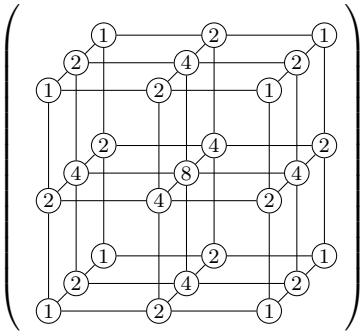

The distances between each non-zero coefficient in G_n are equal to 2^n. The filters $G_0, G_1, G_2, G_3, \dots$ are applied successively by taking the convolution product with the data.

Theorem 2. *The convolution product $G_{n-1} * \cdots * G_2 * G_1 * G_0$ is equivalent to the $(2^{n+1}-1)^{\times 3}$ filter $G^{*n}(i,j,k) = g^{*n}(i)g^{*n}(j)g^{*n}(k)$, where*

$$g^{*n}(i) = \begin{cases} \frac{2^n - |i|}{4^n} & \text{, for } -2^n < i,j,k < 2^n \\ 0 & \text{, otherwise.} \end{cases} \quad (6)$$

Since G_n are separable for all n, $g_{n-1} * \cdots * g_2 * g_1 * g_0 = g^{*n}$, which is a triangular filter. To approximate a Gaussian filter, we repeat the filter G^{*n} twice. The following figure shows the shape of g^{*n} and $g^{*n} * g^{*n}$.

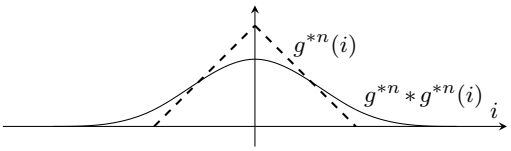

The restricted diffusion process

For each target voxel position $\mathbf{v}$ we define the set

$$A_{\mathbf{v}} = \{\mathbf{w} | \text{Diffusion from } I(\mathbf{w}) \text{ to } I(\mathbf{v}) \text{ is allowed.}\}.$$

We define the restricted diffusion $G_{i,\lambda}$ as

$$G_{i,\lambda}(I)(\mathbf{v}) = I(\mathbf{v}) + \sum_{\mathbf{w} \in A_{\mathbf{v}}} G_i(\mathbf{v} - \mathbf{w})(I(\mathbf{w}) - I(\mathbf{v})).$$

In the iterative diffusion process, $G_{i,\lambda}$ is applied to the entire MRI image for $i = 0,1,2,3,\dots$. The sets $A_{\mathbf{v}}$ need to be recalculated in each iteration.

The new proposed iterative diffusion process is denoted by

$$G_\lambda^{*n} = G_{n-1,\lambda} \circ \cdots \circ G_{1,\lambda} \circ G_{0,\lambda}.$$

In this framework, the old method [9] consists of repeatedly applying $G_{0,\lambda}$. We denote n iterations of the old method by

$$(G_{0,\lambda})^{*n}.$$

Experiment

The sample used in this paper is an MRI image of a macaque monkey from a project by de Castro et al. [10]. We normalize the data by dividing by the largest value outside the top 0.1% of the data. We apply reflection padding at the boundaries.

For all experiments, we used the test diffusion constant $s = 0.1$. In our experiments, we applied the filter G_n for 8 iterations, cycling through the values $n = 0, 1, 2, 3$ in order, two times (i.e. $n = 0, 1, 2, 3, 0, 1, 2, 3$). Effectively, this is the filter $G^{*4} * G^{*4}$, which unrestrictedly smooths the image with a Gaussian filter with $\sigma = 9.22$. The filtering was performed for $\lambda = 1.0 \times 10^{-3}, 5.0 \times 10^{-4}, 2.0 \times 10^{-4}, 1.0 \times 10^{-4}, 5.0 \times 10^{-5}, 2.0 \times 10^{-5}, 1.0 \times 10^{-5}, 5.0 \times 10^{-6}, 2.0 \times 10^{-6}, 1.0 \times 10^{-6}, 5.0 \times 10^{-7}, 2.0 \times 10^{-7}, 1.0 \times 10^{-7}$.

Results

Quantitive results

We begin with a statistical analysis of the difference images produced by subtracting. As was shown in [9], the method effectively removes salt-and-pepper noise, and this holds true for any non-negative diffusion threshold $\lambda \geq 0$. In particular, when $\lambda = 0$, the filter only removes salt-and-pepper artifacts, leaving other types of noise untouched.

This observation gives a natural method to study how the method removes the remaining noise: We compare the filter output at $\lambda = 0$ with its output at a higher diffusion threshold. Let $\mathcal{I}_\lambda$ denote the result image obtained by filtering with threshold λ. The difference image

$$\Delta_\lambda = \mathcal{I}_0 - \mathcal{I}_\lambda \tag{7}$$

captures the change in residual noise as a function of the diffusion threshold.

Figure 1 displays the standard deviation σ of Δ_λ for different values of λ. We used 4 and 8 iterations in both the present method and the old method [9]. The model $\lambda = \left(\frac{\sigma - \sigma_0}{m}\right)^2$ gives a good least squares approximation of λ as a function of σ. These curves are also plotted in Figure 1. Table 1 shows that the old and new methods perform equally well with respect to PSNR.

λ	old	new
1.0×10^{-3}	29.6	29.4
5.0×10^{-4}	30.6	30.4
2.0×10^{-4}	31.4	31.2
1.0×10^{-4}	31.8	31.6
5.0×10^{-5}	32.0	31.8
2.0×10^{-5}	32.2	31.9
1.0×10^{-5}	32.2	31.9
0	32.3	32.0

Table 1: The table shows the PSNR of the noise removed by the new and old methods. Seven iterations were performed for each method, and the top 0.1% of values in the MRI image were removed.

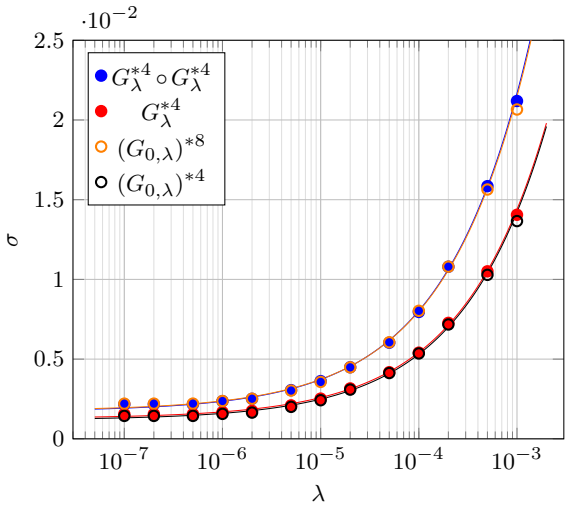

Figure 1: Semi-log plot showing the standard deviation σ of the difference images Δ_λ as a function of the diffusion threshold λ.

Qualitative results

Figure 4 shows the result of applying G_λ^{*4} twice to an MRI image, using two different values of the diffusion threshold λ. In the green-circled region, several bright voxels are noticeably suppressed, indicating the filter's effectiveness in damping outliers. In contrast, the yellow-circled area reveals an important trade-off: for $\lambda = 0.001$, some fine structural details are lost, but for $\lambda = 2.0 \times 10^{-5}$, these features are preserved. Outside these highlighted regions, the overall anatomical structure remains intact for both threshold settings, demonstrating the filter's ability to preserve global image content while varying its sensitivity to local features.

Figure 5 uses gamma adjustment ($\gamma = 0.3$) to highlight low-level noise smoothing. We compare the new and old methods using 4 and 8 iterations and two values of λ. The new method more effectively removes low-frequency noise, due to its long-range diffusion and higher variance. This difference is further illustrated in Figure 3.

Figure 5 (with $\gamma = 0.3$) highlights how low-frequency noise is smoothed. The new method, compared to the old [9] at 4 and 8 iterations and two λ values, achieves better noise reduction thanks to its longer-range diffusion and higher variance—clearly visible in Figure 3. Figure 2 shows two absolute-difference images between the original and the processed MRI image. Noise is effectively removed when $\lambda = 0.001$, whereas for $\lambda = 2.0 \times 10^{-5}$, noise is removed only in areas with very little change.

Conclusion

The proposed method preserves structural detail while effectively removing both high- and low-frequency noise. By adjusting the diffusion threshold λ, the method allows control over the trade-off between smoothing and detail retention. The method provides a flexible tool for denoising, with minimal loss of critical image features. The efficiency of the new method in terms of PSNR is indistinguishable

from that of the old method; however, the new method performs better at removing low-frequency noise, which often occurs in MRI images.

Future work

In addition to basing diffusion on gradient control, many other geometric properties could be explored in future investigations. An MRI image can be modelled as a three-dimensional submanifold of $\mathbb{R}^4$, and therefore, examining the curvature tensor as a metric may be an interesting direction for further investigation.

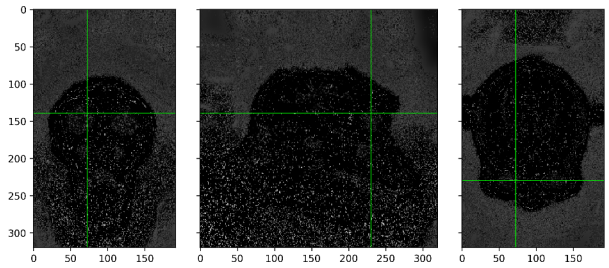

(a) $\lambda = 2.0 \times 10^{-5}$.

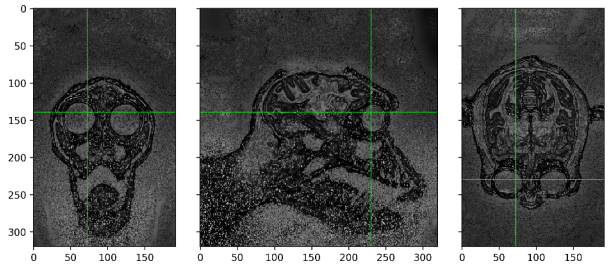

(b) $\lambda = 0.001$.

Figure 2: The figures (a) and (b) show the absolute difference between the original image and the results after applying the filter $G_\lambda^{*4} \circ G_\lambda^{*4}$ with different values for λ. We have used a gamma adjustment with $\gamma = 0.3$ to visualize the differences.

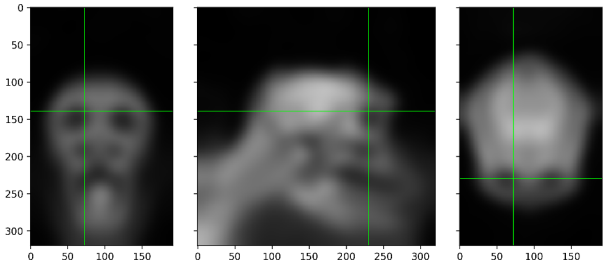

(a) New method

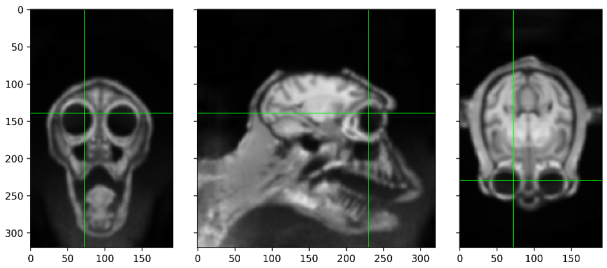

(b) Old method

Figure 3: The result of the (a) unrestricted filter $G^{*4} * G^{*4}$ and (b) the unrestricted filter G_0 applied 8 times.

References

[1] P. Perona and J. Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(7):629–639, 1990.

[2] Carlo Tomasi and Roberto Manduchi. Bilateral filtering for gray and color images. In *Sixth international conference on computer vision (IEEE Cat. No. 98CH36271)*, pages 839–846. IEEE, 1998.

[3] Leonid I Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: nonlinear phenomena*, 60(1-4):259–268, 1992.

[4] Antoni Buades, Bartomeu Coll, and Jean-Michel Morel. A review of image denoising algorithms, with a new one. *Multiscale modeling & simulation*, 4(2):490–530, 2005.

[5] Kostadin Dabov, Alessandro Foi, Vladimir Katkovnik, and Karen Egiazarian. Image denoising by sparse 3-d transform-domain collaborative filtering. *IEEE Transactions on image processing*, 16(8):2080–2095, 2007.

[6] Jaakko Lehtinen, Jacob Munkberg, Jon Hasselgren, Samuli Laine, Tero Karras, Miika Aittala, and Timo Aila. Noise2noise: Learning image restoration without clean data. *arXiv preprint arXiv:1803.04189*, 2018.

[7] Xi-Nian Zuo and Xiu-Xia Xing. Correction: Effects of non-local diffusion on structural mri preprocessing and default network mapping: Statistical comparisons with isotropic/anisotropic diffusion. *PLOS ONE*, 6(11):null, 11 2011.

[8] Ali Alsam and Hans Jakob Rivertz. Long range diffusion with control of the directional differences. In *CIC29: Proceedings of the 29th Color and Imaging Conference*, pages 323–327, Springfield, VA, 2021. IS&T.

[9] Ali Alsam and Hans Jakob Rivertz. Three dimensional surface preserving smoothing. In *CIC30: Proceedings of the 30th Color and Imaging Conference*, pages 135–140, Springfield, VA, 2022. IS&T.

[10] V. De Castro, A. T. Smith, A. L. Beer, C. Leugen, N. Vayssière, Y. Héjja-Brichard, P. Audurier, B. R. Cottereau, and J. B. Dubrand. Connectivity of the cingulate sulcus visual area (csv) in macaque monkeys. *Cerebral Cortex*, 31:1347–1364, 2020.

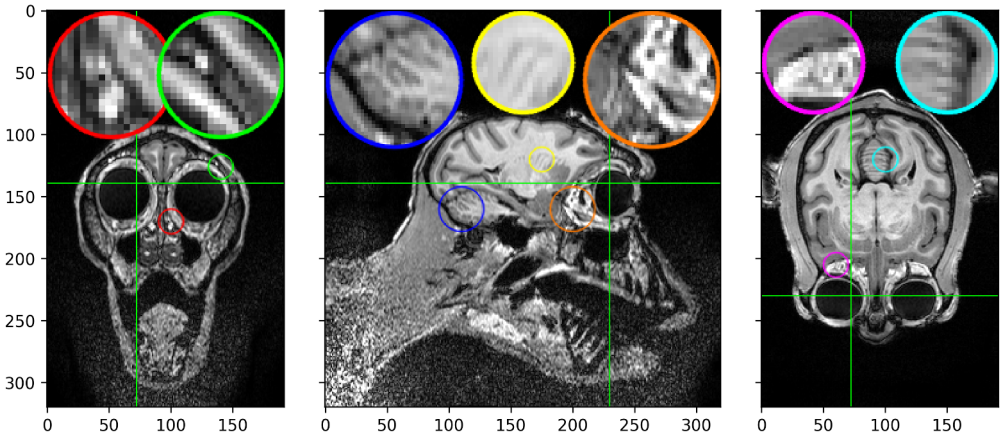

(a) Original

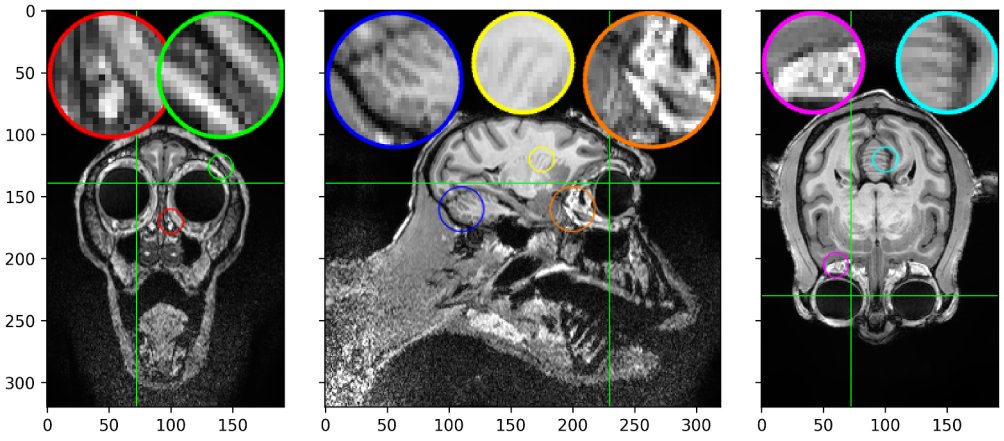

(b) $\lambda = 2.0 \times 10^{-5}$.

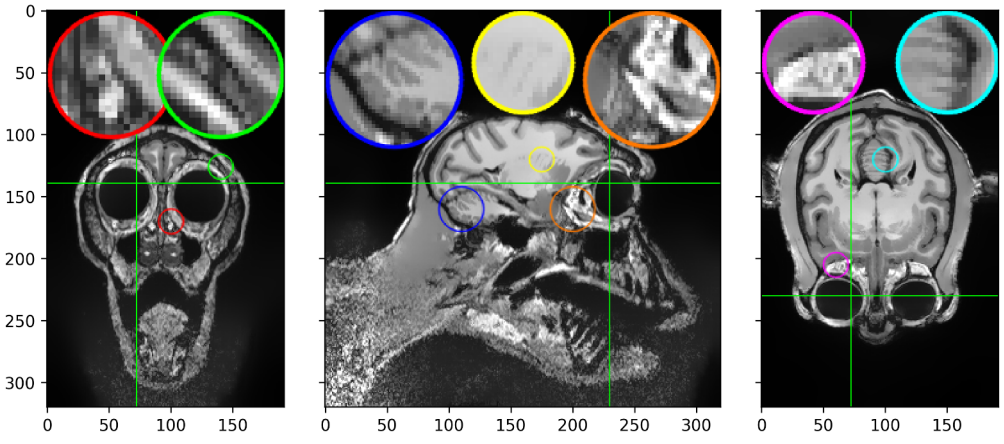

(c) $\lambda = 0.001$.

Figure 4: Figure (a) shows the original MRI image. The figures (b) and (c) show the result after applying the filter $G_\lambda^{*4} \circ G_\lambda^{*4}$ with different values for λ.

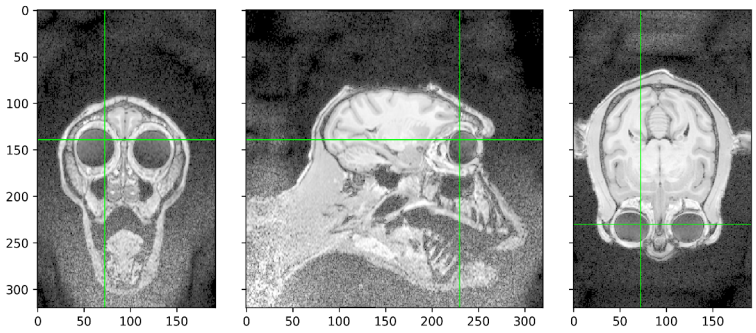

(a) Original

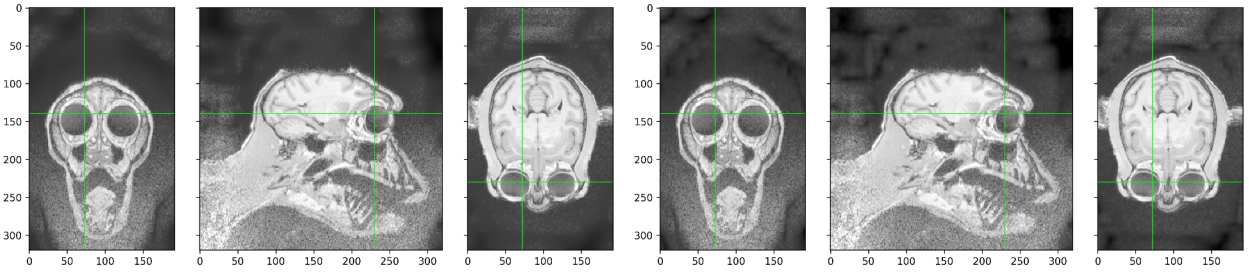

(b) $\lambda = 2.0 \times 10^{-5}$, new method, 4 iterations

(c) $\lambda = 2.0 \times 10^{-5}$, old method, 4 iterations

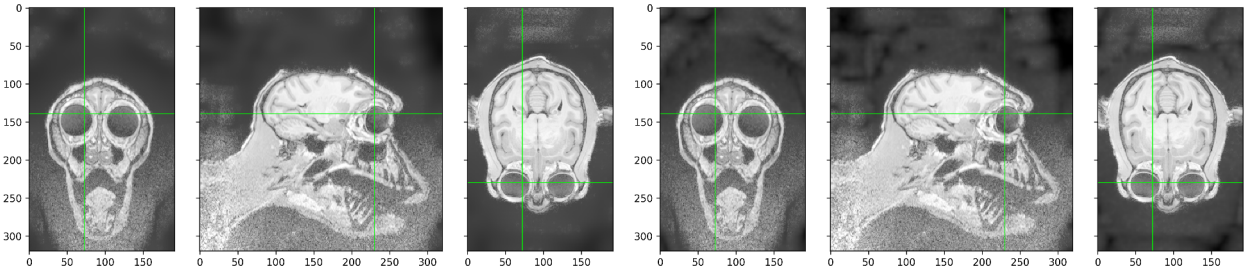

(d) $\lambda = 2.0 \times 10^{-5}$, new method, 8 iterations

(e) $\lambda = 2.0 \times 10^{-5}$, old method, 8 iterations

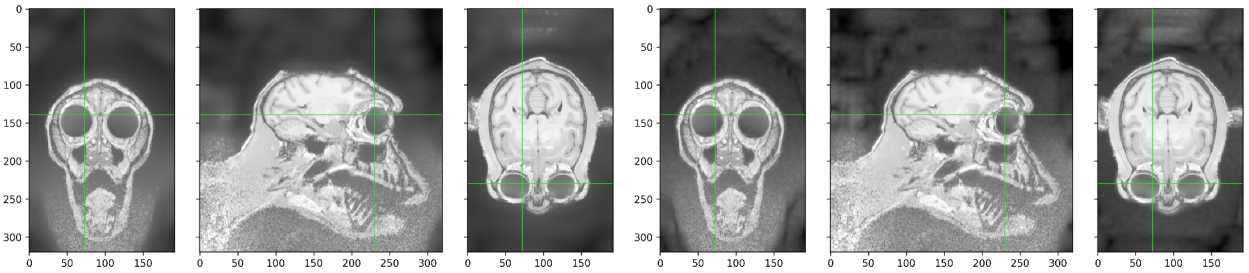

(f) $\lambda = 0.001$, new method, 4 iterations

(g) $\lambda = 0.001$, old method, 4 iterations

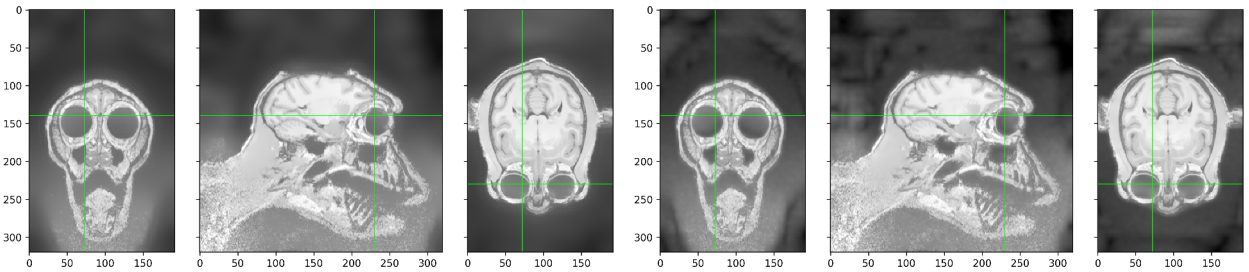

(h) $\lambda = 0.001$, new method, 8 iterations

(i) $\lambda = 0.001$, old method, 8 iterations

Figure 5: We have used a gamma adjustment with $\gamma = 0.3$ in order to visualise the noise. Figure (a) shows the original MRI image. The figures (b) to (f) show the result after applying the filter $G_\lambda^{*4} \circ G_\lambda^{*4}$, G_λ^{*4} or the filter $G_{0,\lambda}$ repeated 4 or 8 times with different values for λ.