Optimising the Frangi Filter: SCALR - A Scale Prediction Model to Enhance Perivascular Space Quantification Using Machine Learning

Juan Sebastian Diaz Villota
201731637

Advisor:
Jose Bernal, PhD
Maria Trujillo, PhD

Escuela de Ingeniería de Sistema y Computación
Facultad de Ingeniería
Universidad del Valle
Santiago de Cali, Colombia
November 2024

Abstract

Perivascular Spaces (PVS) are small, fluid-filled tubular structures involved in cerebrospinal and interstitial fluid circulation as well as the clearance of cerebral waste products. Their presence and morphological features have been increasingly linked to neurological diseases, making PVS a significant area of research in understanding brain health and pathology. However, due to their small and thin structure, there is a critical need for precise and effective methods to quantify PVS accurately, enabling further research into their function and implications for neurological conditions.

The Frangi filter has proven to be a valuable tool in medical image analysis for detecting and characterising vascular structures, including PVS. Optimising the filter's scale, however, is essential to achieve accurate PVS quantification. In recent years, multiscale filtering techniques, such as Frangi's filter, have been introduced to improve PVS quantification, facilitating the investigation of PVS function and morphology in both physiological and pathological conditions.

This final degree project introduces a machine learning approach to determine the optimal scale for the Frangi filter in PVS quantification. The SCALR model, a linear regression model, was developed using both synthetic and real MRI data with varying noise levels and voxel sizes to predict the best scale for diverse imaging conditions. Results indicate that this approach significantly improves the precision of PVS detection, particularly in high-contrast settings, underscoring SCALR's capability to enhance PVS detection and provide reliable quantification, even in low signal-to-noise ratio scenarios.

Contents

List of Figures

List of Tables

Glossary

PVS: Perivascular Spaces (PVS) are fluid-filled structures that surround blood vessels in the brain, facilitating cerebrospinal fluid exchange and waste removal.

Hessian Matrix: A square matrix of second-order partial derivatives of a function, used in image processing to analyse the curvature of structures within an image.

Frangi Filter: A Hessian-based multiscale vesselness filter used to enhance tubular structures in images, particularly useful for detecting perivascular spaces in MRI scans.

Gaussian Distribution: A statistical distribution characterised by its bell-shaped curve, commonly used in image processing to model noise and smooth variations in intensity.

Voxel: The three-dimensional equivalent of a pixel, representing a value on a regular grid in three-dimensional space, commonly used in volumetric imaging.

Scale-Space Theory: A framework in image processing that represents images at multiple scales to capture relevant features, allowing for the detection of structures that may vary in radius sizes (measured in points or voxels).

Scale or Gaussian Sigma (σ_S): The sigma scalar value of the Gaussian distribution, representing the scale of the Frangi filter. This parameter corresponds approximately to the radius, measured in points or voxels, of the target structure being enhanced by the Frangi filter.

Rician Noise: A type of noise commonly encountered in MRI data, arising from the combination of Gaussian noise with a non-negative signal. It is characterised by a Rician distribution, which leads to biased estimation of low-intensity signals and can affect the accuracy of image analysis.

Noise Sigma (σ_N): The sigma scalar value representing the standard deviation of intensities, used to quantify the average variation of the Rician noise in the background of an MRI image.

Signal-to-Noise Ratio (SNR): A measure used in imaging that compares the average intensity of the desired signal to the level of background noise (σ_N), indicating the quality of the image. Higher SNR values correspond to clearer images with less noise interference.

Region of Interest (ROI): A specific, manually or automatically selected portion of an image that is of particular interest for analysis. In medical imaging, the ROI typically focuses on regions where pathologies, such as lesions or PVS, are expected to be found.

Tubular Structures: Structures in an image that have a long, thin shape, such as blood vessels or PVS, which can be enhanced using specific filtering techniques.

Image Artifacts: Distortions or errors in an image that can arise from various sources, such as equipment limitations or patient movement, potentially affecting the accuracy of image analysis.

Machine Learning: A subset of artificial intelligence that uses algorithms and statistical models to improve task-specific predictions by identifying patterns in data, enabling the prediction of unknown values.

MRI (Magnetic Resonance Imaging): A non-invasive imaging technique that uses strong magnetic fields and radio waves to generate detailed images of organs and tissues inside the body, particularly useful for brain imaging.

Image Processing: A method of performing operations on an image to enhance it or extract useful information, often involving techniques such as filtering, segmentation, and feature extraction.

Segmentation: The process of partitioning an image into distinct regions or segments to simplify its analysis, often used to isolate specific structures like PVS in medical images.

1 Introduction

Perivascular Spaces (PVS) are thin and small fluid-filled tubular structures (smaller than 3 millimeters in diameter) which surround cerebral arterioles and venules. These structures are located in the brain, mainly in the basal ganglia (BG), centrum semiovale (CSO) and midbrain. PVS are thought to be related to the glymphatic system, working as drainage channels that facilitate protein and toxic metabolite elimination from the brain [1].

Even though every brain contains PVS, they are not typically visible through MRI. However, these spaces may become enlarged (ePVS) and visible in MRI due to aging, disease, or trauma. In recent years, the scientific community has become increasingly aware of the implications that these enlarged structures may have on brain health, labeling them as a "biomarker of clearance failure within the central nervous system" [2].

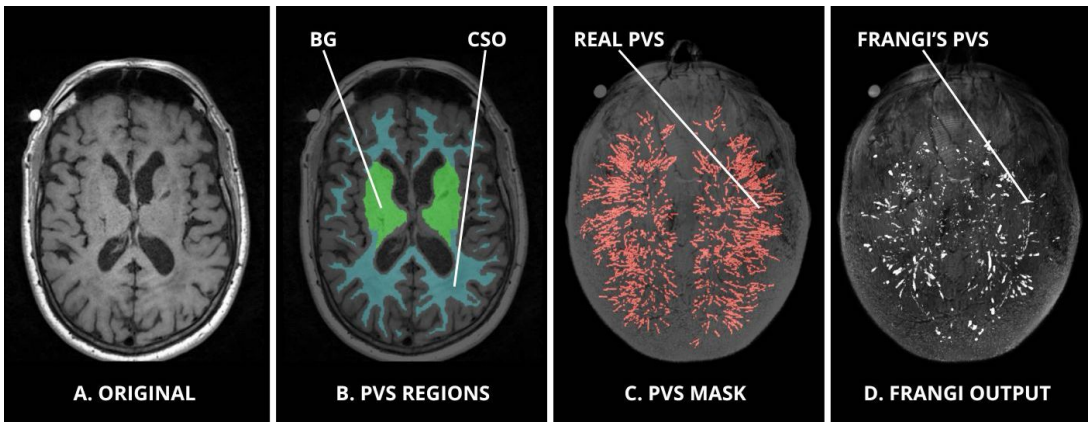

Figure 1: (A) Axial slice of an MRI. (B) Same slice with highlighted regions where PVS are typically observed: Basal Ganglia (green) and Centrum Semiovale (blue). (C) Volume rendering of the PVS, overlaid on the original image, with PVS visualised in red. (D) Frangi filter output overlaid on the original image.

This final degree project focuses on optimising the scale parameter for the Frangi filter, a key tool in PVS quantification. The inherent variability in MRI image quality and resolution makes manual parameter tuning both labor-intensive and prone to inconsistency. To address this, a linear regression model, SCALR, was used to predict an optimal scale parameter across different imaging conditions. The model was calculated with artificial and real MRI images, incorporating diverse noise levels and voxel resolutions, and applying a third-degree polynomial fit for scale prediction.

The results indicate SCALR's efficacy in improving detection accuracy, especially under challenging conditions with low signal-to-noise ratios. By reliably identifying the optimal scale, SCALR minimises the need for trial and error, streamlining the process of PVS quantification. This approach not only saves processing time but also enables further refinement of other filter parameters, establishing SCALR as a promising tool for efficient and precise PVS analysis [36].

2 Research Problem

Detecting and quantifying PVS in MRI can be difficult due to their small size and limited visibility. Numerous approaches have been developed to improve detection, with Frangi's filter recognised as one of the most effective [3]. This Hessian-based multiscale method enhances tubular structures like PVS, but its performance depends on selecting the optimal parameters and MRI quality.

The filter parameters are as follows:

- α: Controls the filter's sensitivity to the roundness of the structure. A high α value emphasises flat shapes, while a low α value favours inflated shapes.
- β: Adjusts sensitivity to the length of the structures. Larger values allow detection of longer tubular structures, while smaller values focus on shorter features.

- c: Acts as a noise reduction factor. A higher c value suppresses noise aggressively but risks removing small PVS, whereas a lower value retains finer details at the expense of increased noise.
- *Scale (σ_S)*: Represents the size of the targeted structures. Larger scales detect larger objects (e.g., enlarged PVS), while smaller scales focus on finer details (e.g., smaller PVS and potential noise).
- *Black Vessels*: A boolean parameter specifying whether to search for dark vessels on a bright background (set to `True`) or bright vessels on a dark background (set to `False`).

Fine-tuning the parameters of Frangi's filter to achieve optimal results can be both complex and time-consuming, often leading to suboptimal outcomes. Poor parameter choices may result in the overestimation of PVS size or the misidentification of noise as genuine PVS structures [6]. A more in-depth explanation of these filter parameters and their role in achieving accurate PVS quantification will be provided in a later section.

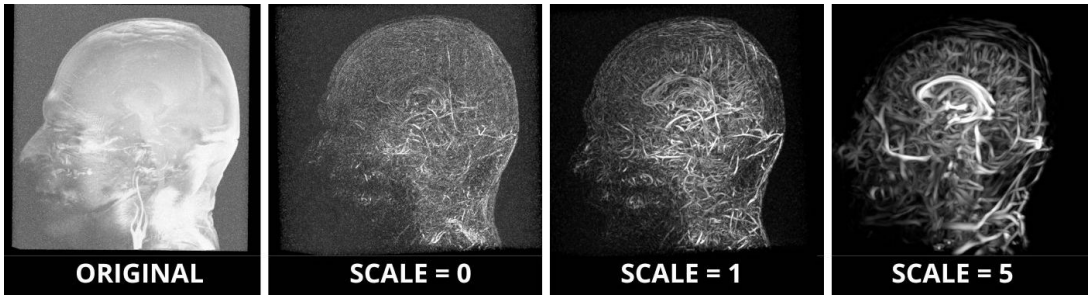

Figure 2: MRI sample (left) and the corresponding outputs following the application of the Frangi filter at various scales in Volume Render.

This tuning problem highlights the need for an efficient method to optimize these parameters and improve the accuracy of PVS quantification.

The scale parameter in Frangi's filter is particularly challenging to define due to its significant impact on the level of detail preserved in the filtered image. Unlike other parameters—α, β, c, and *black vessels*—which have more straightforward definitions based on their geometric roles, the scale parameter controls the filter's sensitivity to structures of different sizes, making its selection crucial for accurate quantification. This effect is illustrated in Figure 2, where varying scale choices lead to noticeable differences in detail retention.

Given these challenges, the central research question arises: Can machine learning techniques be used to automate the selection of the scale parameter in Frangi's filter to enhance PVS quantification?

3 Objectives

3.1 General Objective

Propose a machine learning-based method to optimise the local scale parameter for Frangi's multiscale vesselness filter, with the goal of improving PVS quantification.

3.2 Specific Objectives

- Characterise the behaviour of the Frangi filter on images with varying noise levels, quality, and dimensional characteristics to determine the optimal scale for each image.
- Establish the relation between the measured data and the expected scale to develop a linear regression model that predicts the optimal scale for the Frangi filter, based on the signal-to-noise ratio and the voxel dimensions in MRI images.
- Evaluate the performance of the proposed model in enhancing the detection and quantification of perivascular spaces in medical images.

4 Theoretical Background

4.1 Imaging Techniques on MRI

Magnetic Resonance Imaging (MRI) is a primary modality for visualising PVS due to its high-resolution images of soft tissues. MRI measures the average concentration of hydrogen protons in each voxel, influenced by the tissue's molecular environment.[22]
Three types of MRI sequences have unique characteristics that affect image contrast and the visibility of structures such as PVS:

- **T1-weighted (T1w) MRI**: This sequence produces images where fat appears bright and water appears dark. T1-weighted images are useful for assessing anatomical structures and provide detailed information on tissue morphology. However, in T1w images, PVS appear dark or hypointense, which may limit their visibility due to subtle contrast with surrounding tissues.

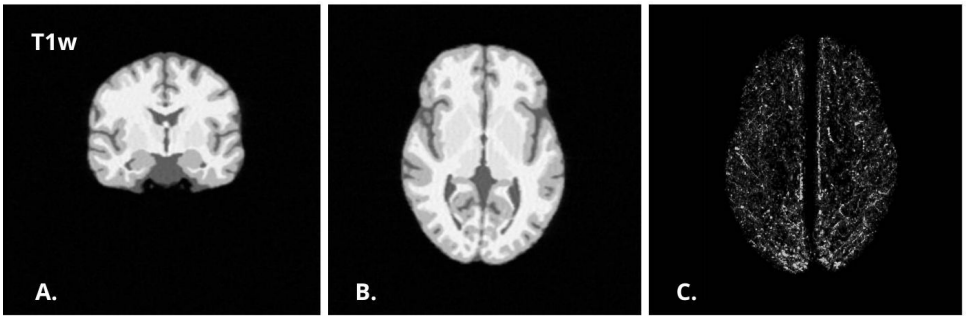

Figure 3: Slice sample of a T1-weighted MRI in coronal view (A), axial view (B), and the corresponding Frangi filter output at a 1mm scale (C).

- **T2-weighted (T2w) MRI**: In T2-weighted images, water appears bright and fat appears dark. This type of imaging is beneficial for identifying pathological changes in brain tissue and enhances PVS visibility.

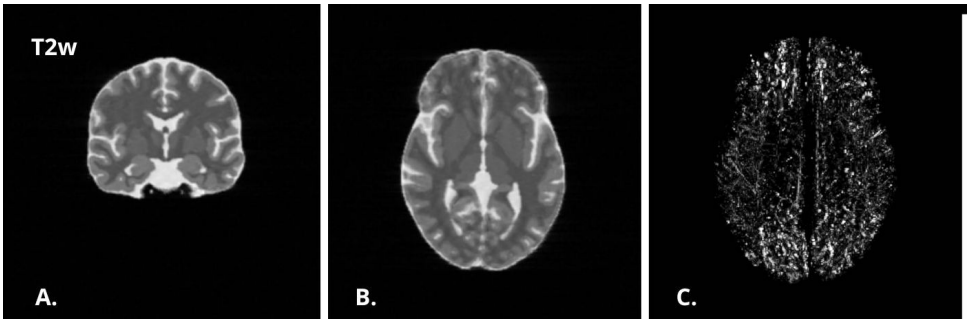

Figure 4: Slice sample of a T2-weighted MRI in coronal view (A), axial view (B), and the corresponding Frangi filter output at a 1mm scale (C).

- **Fluid-Attenuated Inversion Recovery (FLAIR) MRI**: FLAIR is designed to suppress the signal from cerebrospinal fluid (CSF), facilitating the visualization of lesions and small structures. FLAIR images are often used to highlight PVS by reducing background noise from CSF. However, PVS may not always be clearly visible in FLAIR images, as the suppression of CSF signal can lead to a loss of detail, particularly for small PVS, if not properly optimized.

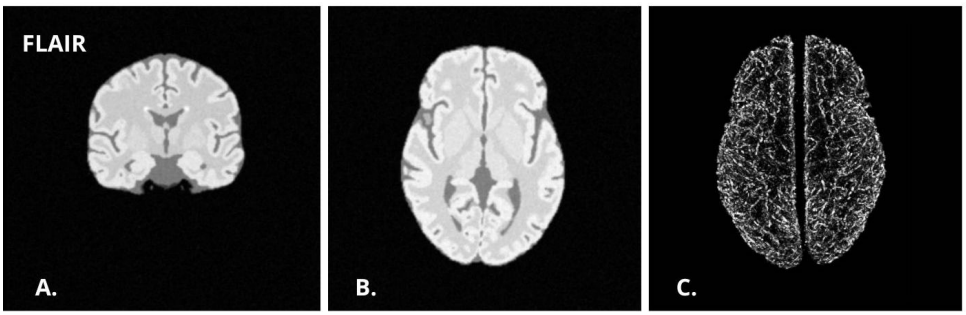

Figure 5: Slice sample of a FLAIR MRI in coronal view (A), axial view (B), and the corresponding Frangi filter output at a 1mm scale (C).

Despite these imaging advantages, detecting PVS remains challenging due to their small size and noise interference. MRI scans often experience low SNR and insufficient spatial resolution, obscuring small structures and affecting quantification accuracy. Noise can arise from thermal fluctuations in the radiofrequency coils, patient motion, and inherent limitations of the imaging technique. One effective way to discern between noise and PVS is by enhancing tubular structures using the Frangi filter, which emphasizes vessel-like structures while suppressing noise. Therefore, discerning between noise and pvs is essential for improving the reliability of PVS detection and quantification.

4.2 Scale-Space Theory

Scale-space theory, initially developed by the computer vision community, emphasises the multiscale nature of real-world objects. It posits that objects become meaningful only when viewed at their appropriate scale [28]. As the perception of objects in images depends heavily on the scale at which they are observed, adjusting the scale is essential for capturing relevant features [28, 29]. The theory suggests representing and processing images at multiple scales, especially when the scale of the target structures is unknown. This approach involves applying filters to progressively remove finer details, *avoiding noise introduction* [28]. The ***Gaussian distribution*** and its derivatives are commonly used for this representation due to their effectiveness in detecting high-contrast regions [30].

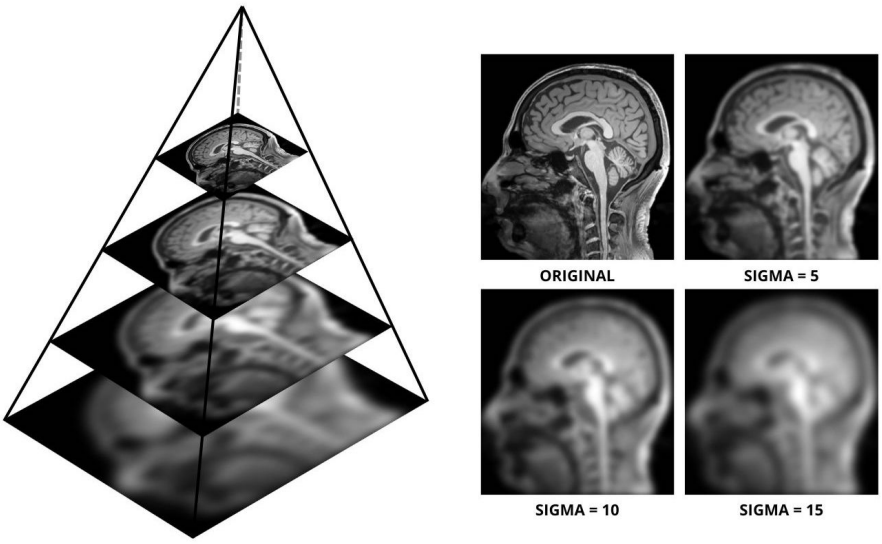

Figure 6: Example of an MRI slice at different scales. This variation in scale is usually visualised as a pyramid of scales as can be seen on the left image.

The scale at which one views an image depends on the value of the Gaussian blur applied to the image. Scale-space theory is fundamental to understanding how structures of different sizes can be detected in images. It involves analysing the image at multiple scales using Gaussian blur to account for features that may only be visible at certain resolutions. In the context of Frangi's filter, scale-space

theory enables the detection of PVS at different diameters by varying the scale parameter. The filter then selects the most prominent response from the multiple scales, ensuring that PVS of varying sizes are captured accurately.

In this case, the smallest possible scale (σ_S) will be used in the filter to identify the required structures in the MRI scan. The first challenge is determining the optimal σ_S for detecting these structures. If σ_S is too small, details such as noise or image artifacts might be misidentified as PVS; conversely, if σ_S is too large, the PVS may be too small to be detected (see Figure 2).

The objective is to identify the optimal scale at which the maximum amount of information is preserved, rather than concentrating on a range. While PVS generally possess diameters of **less than 3 millimeters**, the aim is to propose a single optimal scale for the detection of these tubular structures. This scale will be determined during the model training phase, thereby streamlining the process.

To achieve this, it is necessary to quantify the quality of the image to ensure that the chosen scale is larger than the noise but still capable of detecting small structures.

4.3 The Gaussian Distribution

To gain a deeper understanding of the significance of the Gaussian distribution in image processing, it is essential to examine its behaviour in one dimension before extending this understanding to higher dimensions.

The one-dimensional Gaussian function is defined as follows:

$$g(x) = \frac{1}{\sigma_S\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma_S}\right)^2} \tag{1}$$

In this equation, the mean (μ) is considered equal to zero for simplicity, and the standard deviation is represented as σ_S.

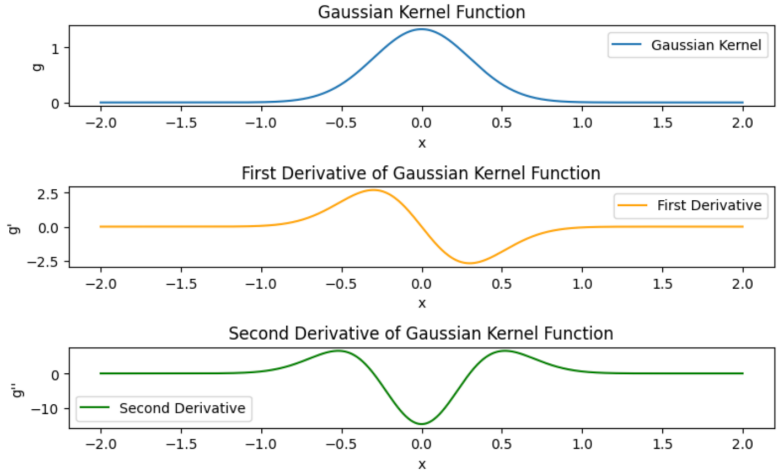

Figure 7: Example of a 1D Gaussian distribution and its derivatives (with $\sigma_S = 0.3$).

The parameter σ_S measures the variability of the data around the mean μ. A larger σ_S results in a wider curve, while a smaller σ_S yields a sharper peak (see Figure 8). The second derivative of the Gaussian function reveals critical information regarding the curvature of the function at specific points.

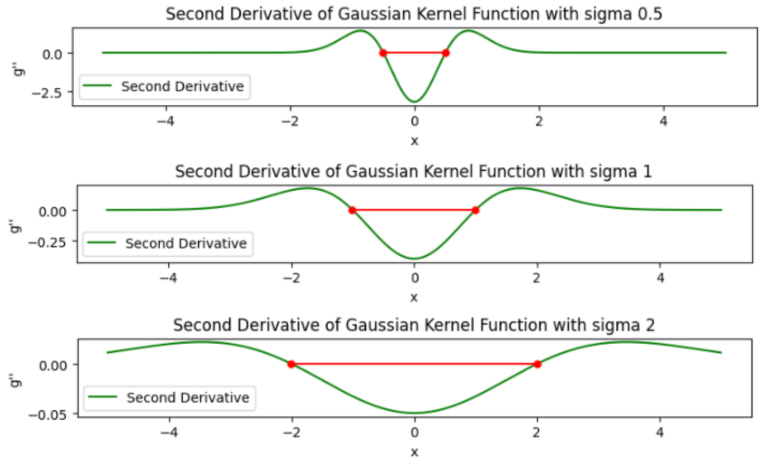

Figure 8: Comparison of the distances between the zeros in the Gaussian function for various σ_S values.

The Gaussian distribution is utilised for detecting structures, such as bumps, along an axis in an image. By smoothing variations in intensity, it identifies changes at pixel values that correspond to these structures.

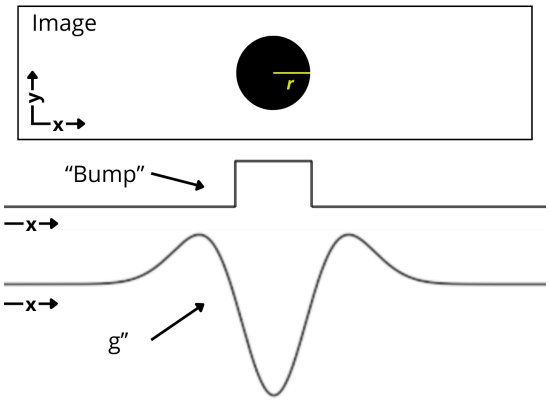

Figure 9: Representation of the second derivative of the Gaussian distribution and a circle of ideal size.

If a bump conforms to a smooth Gaussian profile and falls within the defined radius of σ_S, the shape will be detected. Moreover, when the second derivative of the Gaussian function reaches its maximum values, it signifies that the point of interest lies within the area defined by the level curves of the Gaussian's second derivative. These maxima indicate regions where the curvature is greatest around the central point.

To derive the three-dimensional Gaussian distribution from the one-dimensional case, the process begins with the one-dimensional Gaussian function, which is then extended to three dimensions by assuming equal variance in x, y, and z. The three-dimensional Gaussian function can be expressed as the product of the Gaussian functions in each dimension:

$$g(x,y,z) = g(x) \cdot g(y) \cdot g(z) = \left(\frac{1}{\sigma_S\sqrt{2\pi}} e^{-\frac{x^2}{2\sigma_S^2}}\right) \cdot \left(\frac{1}{\sigma_S\sqrt{2\pi}} e^{-\frac{y^2}{2\sigma_S^2}}\right) \cdot \left(\frac{1}{\sigma_S\sqrt{2\pi}} e^{-\frac{z^2}{2\sigma_S^2}}\right)$$

This simplifies to:

$$g(x,y,z) = \frac{1}{(\sigma_S\sqrt{2\pi})^3} e^{-\frac{1}{2\sigma_S^2}(x^2+y^2+z^2)}$$

Thus, the final form of the three-dimensional Gaussian function is given by:

$$g(x, y, z; \sigma_S^2) = \frac{1}{(2\pi)^{3/2}\sigma_S^3} e^{-\frac{x^2+y^2+z^2}{2\sigma_S^2}} \tag{2}$$

The three-dimensional distribution of the Frangi filter can facilitate the detection of high-contrast regions while minimizing the influence of noise, a principle highlighted by Frangi [6].

To achieve effective detection, the diameter of the target object (or the distance between its edges) should align with the zeros of the Gaussian function (see Figure 8). These zeros represent where the Gaussian function cuts the horizontal axis, indicating the size of structures that are most sensitive to the filter. As shown in Table 1, the sigma value (σ_S) of the Gaussian function corresponds to the radius of the object of interest.

$$\sigma_S \approx \frac{Diameter}{2} = \text{radius} \tag{3}$$

It is important to note that the boundaries of these distances are not sharp. Objects that are slightly larger or smaller than the ideal size can still be detected accurately. Therefore, while the value of σ_S should be close to the desired diameter, it does not need to be exact for effective detection.

σ_S	Zeros Distance
0.5	1.0010010010010006
1	2.0020020020020017
2	4.004004004004004

Table 1: Comparison of the sigma value of the Gaussian function and the distance between its zeros in the second derivative (see Figure 8).

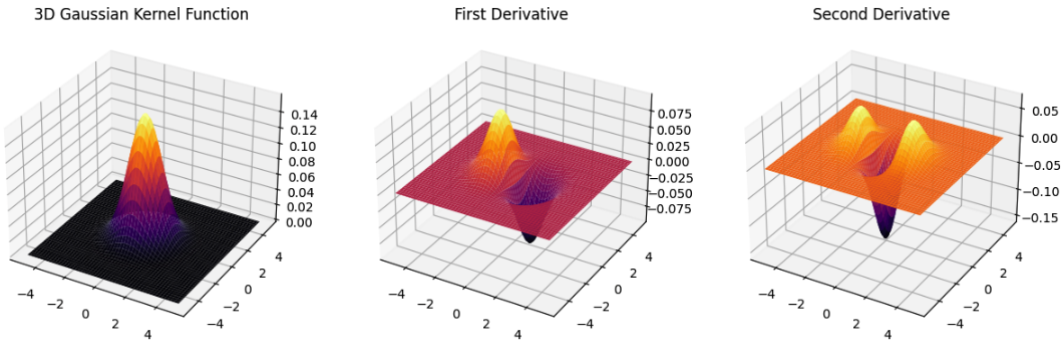

Figure 10: 3D plot of the 2D Gaussian distribution and it's derivatives in x with the intensity value mapped in z ($\sigma_S = 1$).

Remember that in the Frangi filter, the scale at which one views an image is determined by the dispersion (value of σ_S) of the Gaussian filter. The level of detail retained in an image after applying the Gaussian filter is inversely proportional to the value of σ_S (see Figure 6).

In a manner analogous to the one-dimensional scenario, the selection of the appropriate value of σ_S for the proposed method in three-dimensional space is contingent upon the radius of the target structure. However, a complication arises in three dimensions due to the potential variation in radius along each axis.

Given the emphasis on quantifying PVS, it is imperative that the σ_S for the Gaussian blur adequately encompasses the established dimensions (in millimeters) of these spaces in three-dimensional space.

$$V(x) = \frac{x}{V_s} \tag{4}$$

Equation 4 defines the function $V(x)$, which provides the number of voxels for a given distance x in millimeters, where V_s represents the voxel's dimensions in millimeters. This function is used to express distances in the image independently of the actual resolution of the image, ensuring spatial coherence between all input images. This information can be recovered from the NIfTI metadata.

For instance, consider an MRI with a resolution of 1 millimeter^3 per voxel, wherein PVS exhibit a diameter of 3 millimeters or less. Given that the σ_S value delineates details with diameters of $2\sigma_S$ or close, the desired scale range would be established from $\sigma_S = 0.5$ to $\sigma_S = 1.5$.

In the context of the three-dimensional image, the two-dimensional Gaussian function (see Figure 10) will be convolved along the x, y, and z axes during the image processing stage. This *"smoothed image"* is employed to compute the Hessian matrix for each voxel as part of Frangi's vesselness function.

4.4 Hessian Matrix and Filtering Techniques

The scientific community has developed visual rating scales, such as the one proposed by Potter [4], to measure MRI-visible PVS. While these scales provide a standardised approach, they limit measurement precision, are time-consuming, and are prone to human error, given the reliance on subjective visual assessments.

In contrast, more advanced neuroimaging techniques have emerged that minimise human interaction, focusing instead on volumetric methods to identify and quantify PVS automatically. Among these techniques, Hessian filters are particularly effective due to their ability to leverage the tubular structure of PVS. [5] This approach models such structures by computing second-order derivatives of the image, encapsulated in what is known as the Hessian matrix.

The Hessian matrix is a second-order derivative matrix utilised in image processing to detect edges, ridges, and tubular structures. It is particularly relevant in the quantification of PVS, as Hessian-based filters are employed to enhance features with tubular morphology. These specialised image processing techniques emphasise the second-order partial derivatives represented by the Hessian matrix, allowing for the identification of continuous edges within an image, including structures such as blood vessels, wrinkles, and rivers.

By examining the eigenvalues of the Hessian matrix, valuable insights into the local attributes of the image can be obtained, facilitating the detection of areas with pronounced edge-like characteristics. Among the Hessian-based filters, Frangi's filter stands out as especially well-suited for detecting PVS, utilising eigenvalue analysis to identify and highlight structures that resemble blood vessels and tubular spaces. This filter operates across multiple scales, making it effective in detecting PVS of varying sizes.

The Hessian matrix, denoted as H, of a function f at point x is a square matrix where each element H_{ij} represents the second partial derivative of f with respect to the variables x_i and x_j, where n is the number of dimensions of the function.

$$H_{ij} = H_{ij}(x) = \frac{\partial^2 f(x)}{\partial x_i \partial x_j} \qquad (1 \leq i, j \leq n) \tag{5}$$

To better understand the Hessian filter process, it is important to highlight the steps described below prior to the image processing:

1) Calculation of Spatial Derivatives: To identify tubular structures, it is necessary to calculate partial derivatives in the three spatial directions (x, y, and z) of the 3D image. This process involves determining the rate of change of image intensity along each direction using finite differentials.

2) Construction of the 3D Hessian Matrix: The 3D Hessian matrix is an extension of the 2D Hessian matrix into three dimensions. It is constructed by computing the second partial derivatives in the three spatial directions.

$$H_{3D}(x, y, z) = \begin{bmatrix} f_{xx} & f_{xy} & f_{xz} \\ f_{yx} & f_{yy} & f_{yz} \\ f_{zx} & f_{zy} & f_{zz} \end{bmatrix} \tag{6}$$

This matrix provides valuable geometric information about the local curvature of the function, analogous to the role of the second derivative in one-dimensional functions. In the context of image processing, the Hessian matrix offers insights into the curvature of multidimensional functions, particularly in volumetric images such as MRI scans [5].

To compute second-order derivatives effectively, it is advisable to convolve the image with the second derivative of the Gaussian distribution (see Figure 7). This approach is consistent with scale-space theory, which suggests that image analysis benefits from progressively smoothing the image at different scales, as illustrated in Figure 6. Convolution with the second derivative of the Gaussian ensures the continuity and differentiability of the image function, reduces noise, and enhances the detection of tubular structures by leveraging the characteristic shape of the Gaussian's second derivative [9, 5].

3) Calculation of Eigenvalues: After forming the 3D Hessian matrix, its eigenvalues, λ_1, λ_2, and λ_3, are obtained. These eigenvalues represent the characteristic curvature at each point along the three axes and are essential for analysing linear transformations that affect the shape and orientation of geometric figures.

Computing the Hessian matrix and its eigenvalues gives geometric information about the local curvature of the function, which aids in identifying structures in volumetric images, such as MRI scans. The sign and magnitude of the eigenvalues characterise local curvature and provide insights into the intensity and shape of structures within the voxels.

Frangi's vesselness function utilises the geometric relationship between these eigenvalues and an ellipsoid, offering a clear understanding of the local direction of curvature [6].

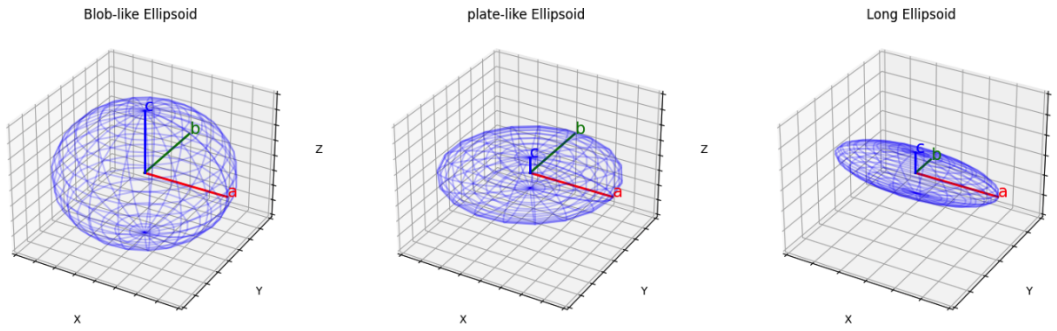

Figure 11: Ellipsoids with variations on a ($\sqrt{|\lambda_1|}$), b ($\sqrt{|\lambda_2|}$), and c ($\sqrt{|\lambda_3|}$) obtained form the 3D Hessian matrix.

The 3D Hessian matrix describes how the partial derivatives of a function change in all directions at a given point. For a function $H(x, y, z)$, the Hessian matrix is a 3x3 matrix where each element represents the second partial derivative with respect to each variable (see Equation 6).

Given the ellipsoid equation:

$$\frac{x^2}{a^2} + \frac{y^2}{b^2} + \frac{z^2}{c^2} = 1 \tag{7}$$

One can obtain different types of ellipsoids based on the values of a, b, and c.

Note that the eigenvalues λ are inversely proportional to the lengths of their corresponding radii. According to Frangi's method for identifying tubular structures, the desired eigenvalues should satisfy the following conditions: $|\lambda_1| \approx 0$, $|\lambda_1| << |\lambda_2|$, and $|\lambda_2| \approx |\lambda_3|$, which indicate thin, elongated ellipsoids.

λ_1	λ_2	λ_3	Description
L	L	H-	Plate-like structure (bright)
L	L	H+	Plate-like structure (dark)
L	H-	H-	Tube-like structure (bright)
L	H+	H+	Tube-like structure (dark)
H-	H-	H-	Blob-like structure (bright)
H+	H+	H+	Blob-like structure (dark)

Table 2: Interpretation of eigenvalues. Each λ represents an eigenvalue, where H+ indicates high positive values, H- indicates high negative values, and L represents low values (close to 0).

5 State of the Art

5.1 Frangi's Formula

This filter belongs to the family of Hessian-based filters, and as previously stated, it works based on the information provided by the eigenvalues of the Hessian matrix [5, 6]. The method is based upon the Taylor series approximation of an image already convolved with the Gaussian distribution around a point a, from which one can extract the Hessian matrix and gradient of the function.

$$f^s(a+x) \approx f^s(a) + x^\top \begin{pmatrix} f^s_x(a) \\ f^s_y(y) \end{pmatrix} + \frac{1}{2} x^\top \begin{pmatrix} f^s_{xx}(a) & f^s_{xy}(a) \\ f^s_{yx}(a) & f^s_{yy}(a) \end{pmatrix} x \tag{8}$$

After computing the Hessian matrix and its eigenvalues, the eigenvalues are sorted as $\lambda_1 < \lambda_2 < \lambda_3$ to account for all possible rotations. This ordering allows for the extraction of three key properties [6]:

$$R_b = |\lambda_1| / \sqrt{|\lambda_2 \lambda_3|} \tag{9}$$

$$R_a = |\lambda_2| / |\lambda_3| \tag{10}$$

$$S = \sqrt{\lambda_1^2 + \lambda_2^2 + \lambda_3^2} \tag{11}$$

Each measure provides information regarding a point in an image; R_b is particularly useful for discriminating blobs. More specifically, considering the ellipsoid, it can be stated that as R_b approaches zero, λ_1 increases, resulting in a more tubular structure. (R_a) distinguishes line and plate structures, in this case the closer that (R_a) gets to one the more tubular the structure gets because this would mean that $(|\lambda_3| \approx |\lambda_2|)$. Finally (S) which filters low contrast structures meaning that this property can differentiate between the background of the image(low contrast) and the vessels which are brighter(high contrast). With this measures Frangi proposed the following vesselness function:

$$V(R_a, R_b, S) = \begin{cases} 0 & if \quad \lambda_2 > 0 \quad or \quad \lambda_3 > 0 \\ (1 - e^{-{R_a}^2/2a^2}) \cdot (e^{-{R_b}^2/2b^2}) \cdot (1 - e^{-S^2/2c^2}) & otherwise \end{cases} \tag{12}$$

Frangi's filter is multiscale, meaning it works with the concept of scale-space theory to detect PVS at different scales (see Figure 6). The vesselness function is evaluated over a defined range of Gaussian blur sigmas (see Figure 8), which correspond to a range of scales. The function outputs a 3D matrix of intensity values for each scale. Once all matrices within the given range are computed, the maximum intensity value at each point (voxel) is selected. A value close to zero indicates the absence of PVS, whereas a value close to one suggests that the structure may potentially be a PVS [6].

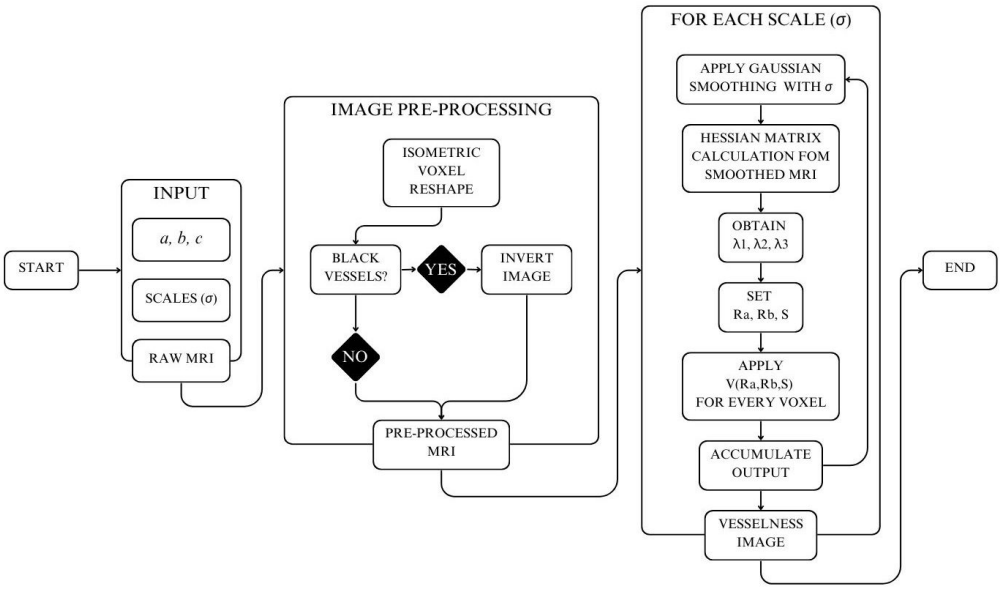

Figure 12: Flowchart illustrating the Frangi vesselness filtering method for MRI images across one or multiple scales (σ_S), using the function V in equation 12.

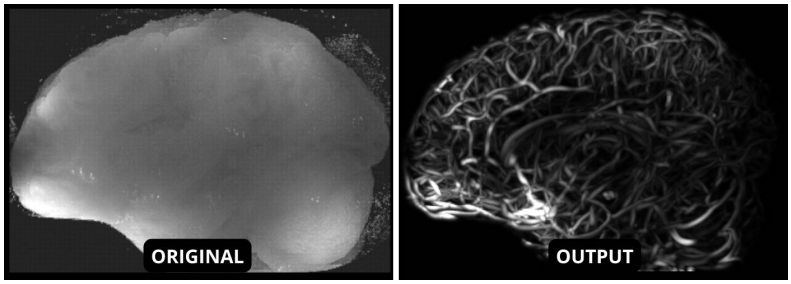

Figure 13: Volume Render of Frangi's vesselness enhancement filter applied to a sample MRI (left), highlighting tubular structures approximately 3 millimeters in diameter (right).

The visualisation presented in Figure 13 exemplifies the efficacy of Frangi's method in enhancing tubular structures, wherein the 'thickness' radius is determined by the selected scale.

To characterise the behavior of Frangi's filter on geometric shapes, a synthetic 3D image will be generated, containing long ellipsoids with diameters ranging from 1 to 9 voxels (1 millimeter per voxel).

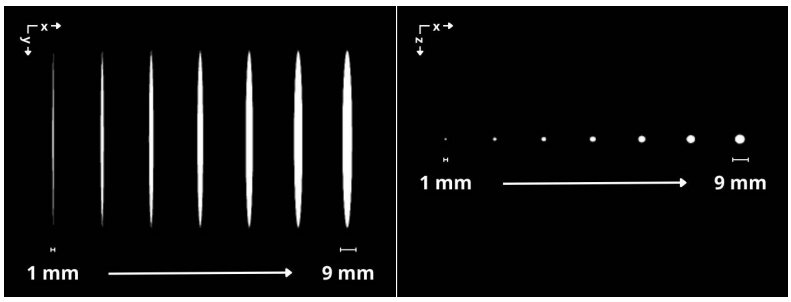

Figure 14: 3D image sample of seven ellipsoids with their diameters going from 1 to 9 millimeters progressively in Coronal view (left), and Axial view (right).

Frangi's filter will then be applied to the image, and the intensity response of the known shapes will be observed to assess the filter's performance.

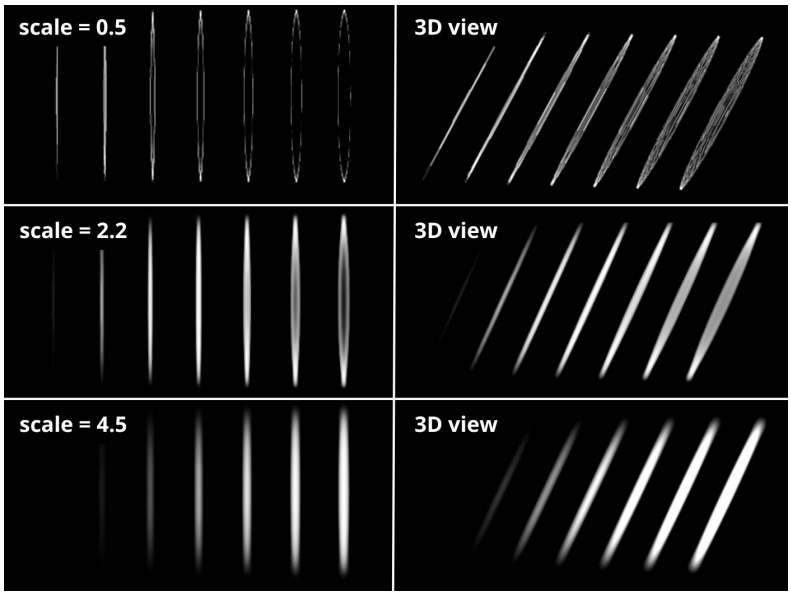

Figure 15: Illustration of Frangi's method applied to the ellipsoids in Figure 14, each row with scales around the radius of each ellipsoid in Axial view (left), and Volume Render (right).

To assess the effect of the filter on the target shapes, the contrast of each ellipsoid will be measured

at various scales, emphasising the scale at which each ellipsoid attains its maximum intensity.

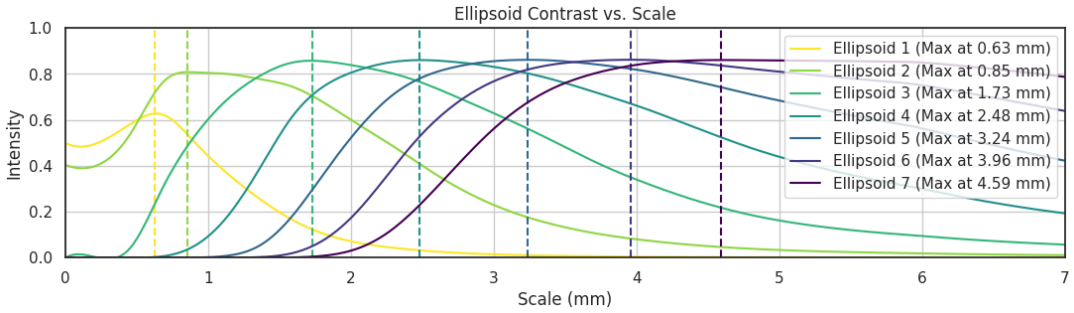

Figure 16: Plot depicting the average contrast of the ellipsoids, indicating the scale at which each ellipsoid reaches its maximum intensity (scales ranging from 0 millimeters to 7 millimeters).

From the measurements shown in Figure 16, it can be observed that the suggested scales, corresponding to the diameters of the ellipsoids (ranging from 0.5 to 4.5 millimeters in radius), exhibit significantly high intensity values, approaching unity, while values for other scales remain comparatively lower. This representation facilitates the identification of the optimal scale (σ_S) ranges that enhance the target structures.

Despite its remarkable performance on synthetic images, such as the one illustrated in Figure 15, the method's effectiveness diminishes when applied to real-world images, such as MRIs, primarily due to the presence of inherent noise and artifacts.

Consequently, it is necessary to fine-tune the filter's variables to enhance the output. These parameters modulate Frangi's formula to either amplify or penalise R_a, R_b, and S, in addition to the σ_S.

α **Parameter**	**Result Description**
Close to 1	The method emphasises more flat structures.
Close to 0	The method emphasises more round structures.
β **Parameter**	**Result Description**
Close to 1	The method emphasises longer, slender structures.
Close to 0	The method emphasises more round, blob-like structures.
c **Parameter**	**Result Description**
Close to 1	The method enhances noise reduction, sacrificing finer details.
Close to 0	The method may retain more noise, but keeps finer details.
Scale **or** σ_S **Parameter**	**Result Description**
Structure radius (in voxels)	The method emphasises structures with a diameter approximately equal to $2 \times \sigma_S$.
Black Vessels **Parameter**	**Result Description**
True	The method finds dark structures in a bright background.
False	The method finds bright structures in a dark background.

Table 3: Description of the α, β, c, σ_S and Black-Vessels parameters.

The selection of these parameters is a key challenge in this project. The parameters α, β, and c control geometric properties: α adjusts for roundness, β for elongation, and c reduces noise. These parameters can be selected with relative ease. Additionally, the identification of black vessels can be automatically adjusted based on the MRI sequence (enabled for T1-weighted and FLAIR, and disabled for T2-weighted images). However, the scale, which determines the size of the structures detected, is more difficult to define. A machine learning approach to automatically determine the optimal scale would significantly improve the filter's accuracy and efficiency.

The methodology outlined in this project focuses exclusively on processing MRIs using Frangi filtering in 3D. The background reduction parameter c is computed as half of the maximum value S in the image, following a commonly used approach in the state of the art and in scientific libraries such as scikit-image and SimpleITK for image processing [7, 8]. The results are influenced by tunable parameters, including

α (ideally close to 0), β (preferably close to 1), and the range of small scales (with a wider range is often considered better).

5.2 About MRI Quality and PVS Quantification

Recent work has introduced frameworks for addressing imaging artifacts, enhancing the reliability of PVS quantification. Bernal et al. (2020) [25] proposed a texture analysis and total variation optimization approach to reduce noise-related artifacts, significantly improving quantification accuracy. Similarly, Bernal et al. (2021) [26] developed a radiomics and k-space reconstruction technique specifically targeting motion artifacts, further refining PVS detection. These methodologies underscore the need for tailored preprocessing techniques to isolate PVS structures more effectively.

Despite these advances, there is still a need for robust, automated techniques that can reliably quantify PVS across diverse patient data, which is where machine learning (ML) models show promising potential. ML approaches have been increasingly applied to enhance artifact reduction, segmentation accuracy, and consistency in PVS quantification, forming the basis for ongoing research in this domain. However, while ML has been valuable for these applications, not all ML techniques are effective for every task, underscoring the importance of selecting or developing models suited to the specific challenges of PVS quantification.

6 Proposed Approach

The characteristics of the data will be analysed to establish a relationship between image properties and the optimal parameters for the Frangi filter. Measurements and machine learning techniques, such as noise quantification and linear regression, will be employed to determine the most appropriate scale for the filter, tailored to images with varying quality levels. This determination will be specifically based on the SNR and the minimum detectable size, corresponding to the length of one pixel.

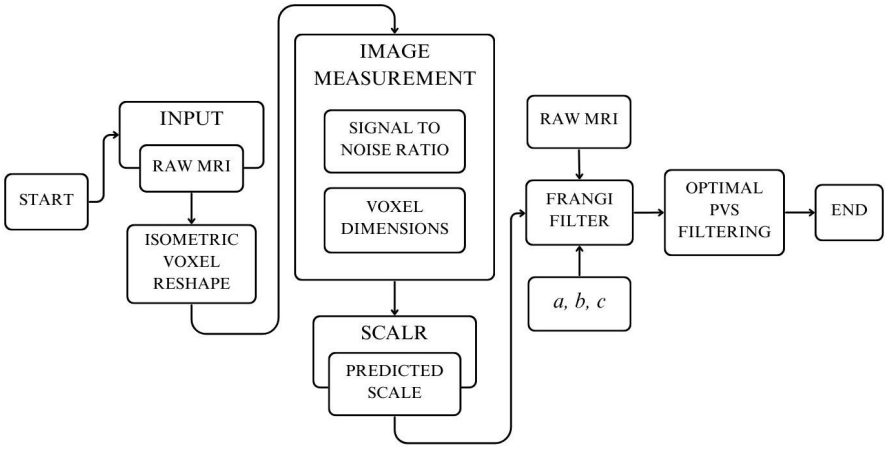

Figure 17: Flowchart illustrating the SCALR workflow for multiscale vessel enhancement in MRI analysis.

6.1 Image Preprocessing

Before applying the model to the MRI data, the image must be transformed and standardised to ensure consistency in the input. The preprocessing steps include:

1) Ensuring isometric voxels: If the voxels are not cubic, a reshape transformation is applied. The dimension information for each voxel in millimeters is obtained from the image metadata. By making the voxels isometric, the problem of considering three best scales—one for each axis—can be minimized. This allows for focusing on a single scale across all rotations, simplifying the parameter selection and ensuring consistency in the analysis.

2) Image Standardisation: While the Signal-to-Noise Ratio (SNR) can provide an initial approximation of the optimal scale parameter, it exhibits exponential behaviour, in contrast to the linear relationship of noise, quantified by its standard deviation σ_N. For improved precision in scale prediction,

histogram matching can be employed to align the input image's intensity distribution with the average histogram of the training dataset. This standardisation ensures consistency in intensity values, thereby enhancing the robustness of the predicted scale. By quantifying noise through σ_N, a more reliable and linear measure of noise level is achieved, facilitating finer calibration of the scale parameter. This aspect will be explored in future work (see Figure 32).

6.2 Finding the Optimal Scale

In Frangi's method, the scale parameter, σ_S, is essential for emphasising fine structures while reducing noise interference. Proper selection of σ_S ensures precise segmentation and visualisation of the structures of interest. For this project, the focus is on identifying the smallest and thinnest structures within the image, which are constrained by the voxel size itself (with a minimum size of 1 voxel). However, due to the SNR, the contrast of the ROI can shift, necessitating careful adjustment of the scale. After standardising the MRI input and determining voxel dimensions, it becomes crucial to assess the SNR of each MRI post-processing. This approach enables fine-tuning of the scale, facilitating a balance between noise reduction and the visibility of the smallest structures, with typical scale values ranging from the voxel size up to 3 millimeters, depending on the measured SNR.

6.2.1 Quantification of Noise in MRI

A common method for quantifying noise in medical imaging is to calculate the standard deviation of the noise within a homogeneous background region. This standard deviation provides a noise σ_N value, which approximates the range of intensity variations introduced by noise in the image. In addition to this, a ROI sample is extracted, a cube centered in the middle of the image, representing 20% of the total image size. This region is chosen because, in most MRI scans, the essential tissue structures are typically centered. By combining the standard deviation from the homogeneous background with the intensity values in the ROI sample, a better assessment to the SNR can be made for each MRI. This measurement aids in adjusting the optimal scale parameter, ensuring that even the smallest structures remain visible despite noise interference.

To evaluate this approach, the same ellipsoids depicted in Figure 14 will be employed. Rician noise will be added at progressively increasing levels, each with a known σ_N value, until the signal-to-noise ratio (SNR) reaches 1. This methodology allows for a rigorous validation of the noise quantification process and its influence on identifying the optimal scale parameter.

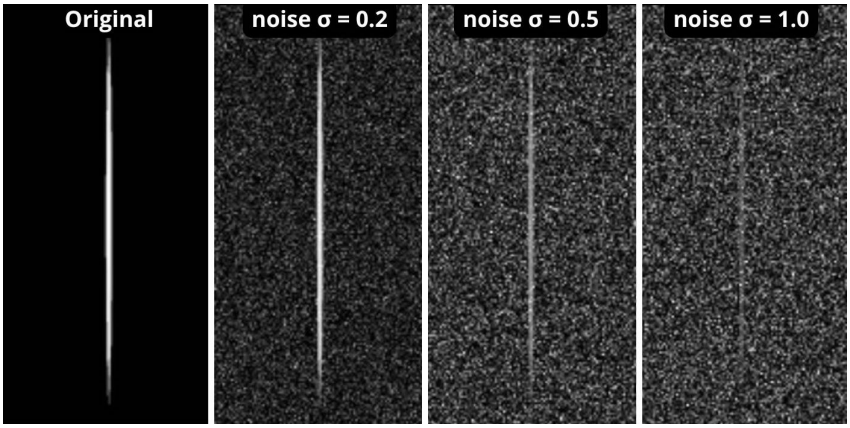

Figure 18: Thin ellipsoid sample from Figure 14 (ROI intensity of 1) with different levels of Rician noise applied.

As the level of noise increases, the degradation in image quality becomes more pronounced, making it progressively harder to distinguish the target structures from the surrounding noise. This is especially noticeable in regions where the SNR is low, as the contrast between the ROI and the background diminishes.

This analysis highlights the relationship between SNR and the ability to discern structures within the ROI. A lower SNR results in a greater degree of image degradation, reducing the visibility of key structures, while a higher SNR allows for a clearer distinction between the signal and noise.

With the true value of the **Rician noise sigma** known (σ_N), the next step involves recovering this value by analysing the image.

6.2.2 Calculating the Rician Noise and SNR

The first step involves calculating the standard deviation of the background noise. A homogeneous region, typically the black background in MRI images, is selected for this purpose. This region contains no anatomical structures and only noise contributes to the intensity values. By sampling approximately 5-10% of the total image area from this region, an accurate estimate of the noise σ_N can be obtained.

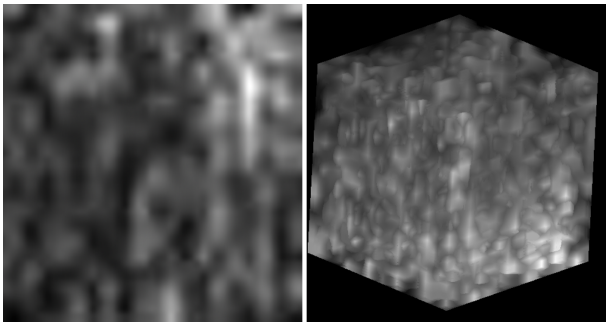

Figure 19: Background noise sample for σ_N calculation taken from the corner of one of the noisy images in Figure 18.

The standard deviation of the noise is calculated using the following equation:

$$\sigma_N = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(x_i - \bar{x})^2} \tag{13}$$

where x_i represents the intensity of each pixel in the sampled region, $\bar{x}$ is the mean intensity of the sample, and N is the total number of pixels in the sample.

By applying this formula, the resulting standard deviation from each image can be compared to both the **known noise** sigma value and the calculated one.

Known Noise σ_N	Calculated Noise σ_N	Percentual Diff (%)
0	0.0000	0.00
0.1	0.11396	13.96
0.2	0.23340	16.70
0.3	0.35408	18.03
0.4	0.45983	14.96
0.5	0.57420	14.84
0.6	0.69753	16.25
0.7	0.82252	17.50
0.8	0.91696	14.62
0.9	1.05118	16.80
1.0	1.13228	13.23
Average		15.90

Table 4: Comparison of the Noise sigma value from the sample image and its calculation via Standard deviation

Although there may be a tendency for the standard deviation method to slightly overestimate the noise σ_N, this overestimation is consistent.

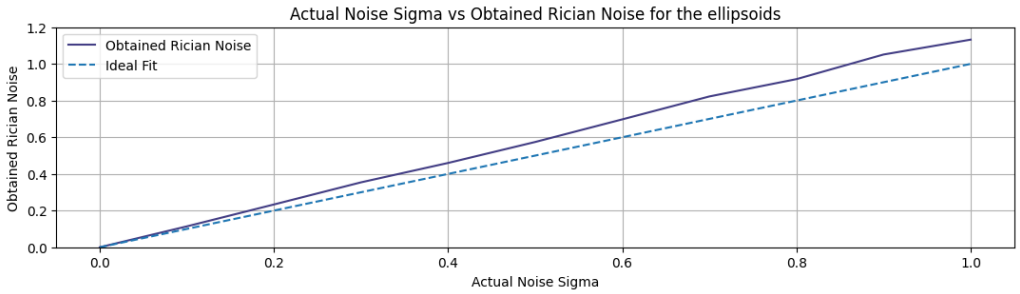

Figure 20: Illustration of actual Noise Sigma vs the calculated one from table 4.

By applying a correction factor based on the average percentage of overestimation, the precision of the noise measurement can be significantly improved.

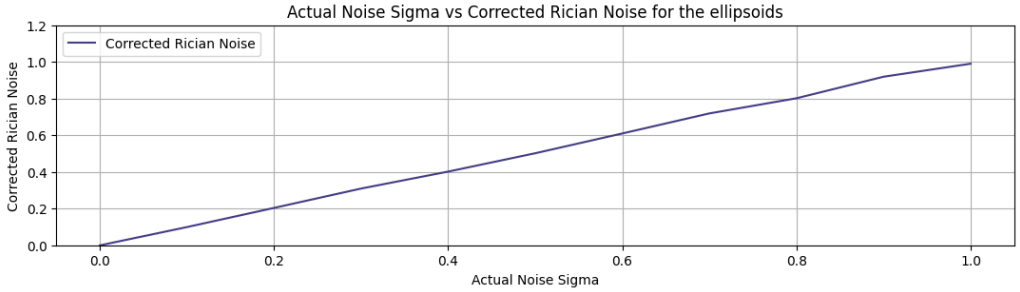

Figure 21: Illustration of the corrected Noise sigma from Figure 19.

Once the background noise σ_N is precisely calculated, the SNR can be calculated for the ROI. For the ellipsoids, the average intensity of the ROI is 1, but in real images The ROI can be defined as a cube located at the center of the image, encompassing approximately 20% of the total image volume.

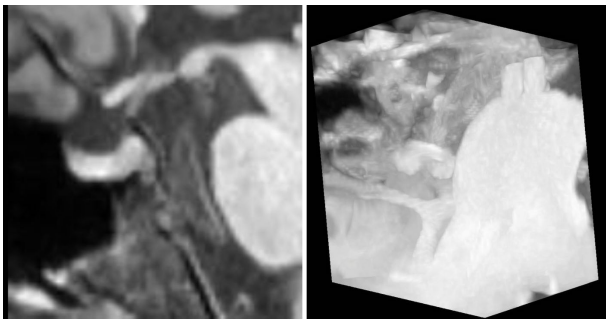

Figure 22: Region Of Interest (ROI) for SNR calculation – cubic sample from a real MRI.

This central region is used for the SNR, computed as follows:

$$\text{SNR} = \frac{\mu_{\text{ROI}}}{\sigma_N} \tag{14}$$

where μ_{ROI} represents the mean intensity of the ROI and σ_N is the standard deviation of the noise obtained from the background sample.

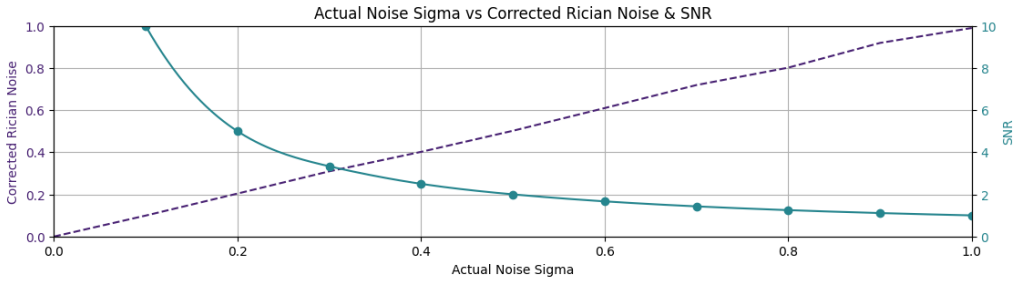

Figure 23: Plot illustrating the change in Noise Calculation and SNR when the actual noise increases for the ellipsoid from Figure 18.

It is important to note that either the SNR or the calculated noise sigma can be mapped directly to the actual noise present in a given image.

When the intensity of the ROI is constant and known (as demonstrated by the ellipsoids in Figure 18), σ_N provides a reliable measure due to its linear increase, making it less sensitive to minor errors. Conversely, while the SNR offers a more robust and general measurement when the ROI intensity is variable, it exhibits an exponential behaviour, causing even small errors to lead to significant inaccuracies. Both noise variation (σ_N) and SNR can therefore be used to estimate the optimal scale within the image, with σ_N offering greater stability in scenarios with minimal error sensitivity.

6.2.3 How Noise Affects the Scale

With the improved ability to quantify image noise more accurately, it becomes possible to analyse how increasing noise levels affect the optimal sigma value for the Frangi filter. For the sample image shown in Figure 14, the plot in Figure 16 displays the best scales for each ellipsoid in the absence of noise. When noise is introduced, these optimal scales are expected to shift due to the presence of additional unwanted information.

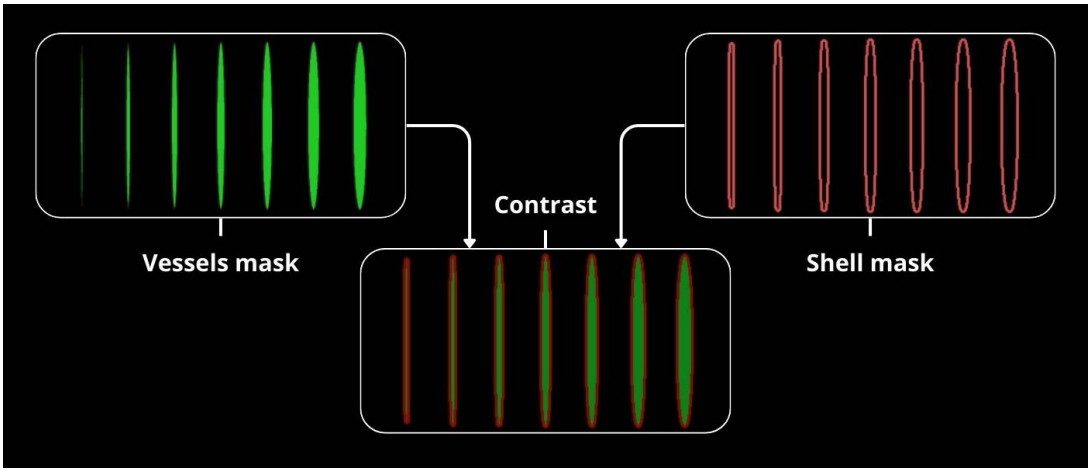

Figure 24: Illustration of the foreground and surrounding background of the shapes from Figure 14, used to measure contrast at different scales.

The new optimal scale is determined by analysing the contrast between the foreground and background regions across a range of scales and identifying the highest value.

Furthermore, examining how this optimal scale shifts as noise increases provides a deeper understanding of the impact of noise on the filter's performance.

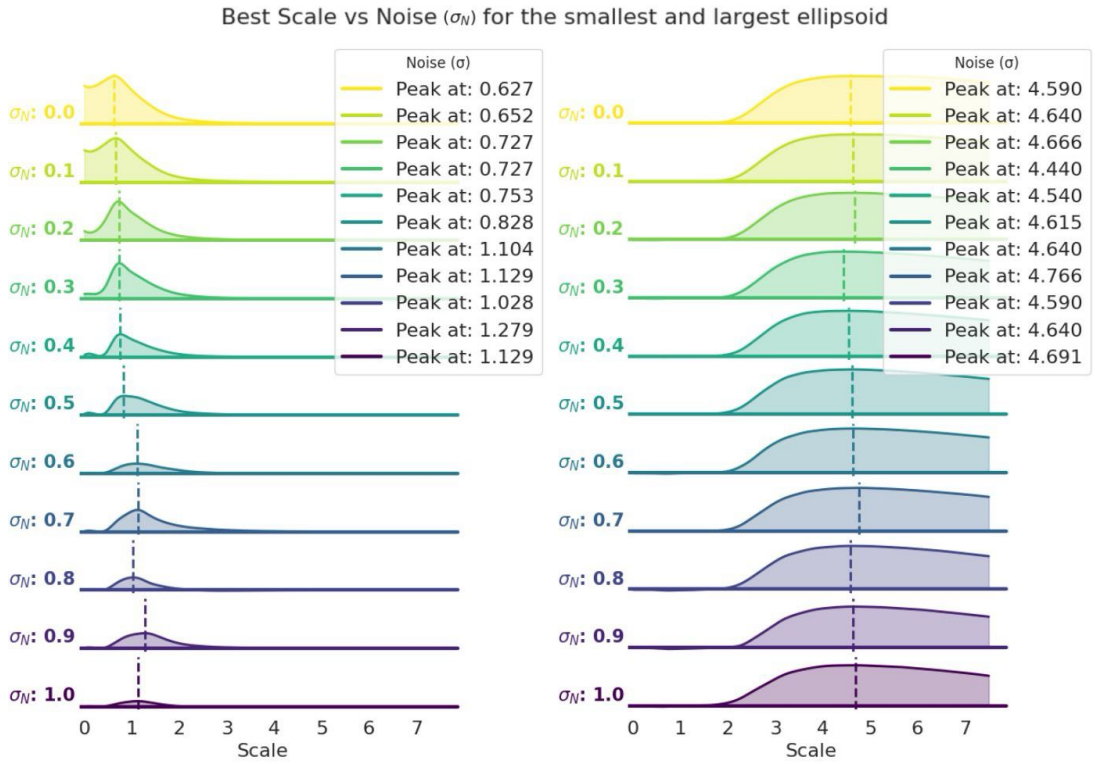

Figure 25: Plot illustrating the shift in the best scale for the smallest ellipsoid (1 millimeters in diameter) and the largest one (9 millimeters in diameter) as σ_N increases.

After analysing the test data, a relationship between scale shifting and shape diameter can be established. The results indicate that when the diameter of the target shape is significantly larger than the volume of the noise particle, the target shape is less affected by noise, resulting in a wider and more stable optimal scale range with similar intensity. In contrast, for smaller diameters, the optimal scale range is narrower, less intense and shifts to the right, indicating a greater sensitivity to noise.

By applying the method outlined in Figure 16 to determine the scale that produces the maximum contrast in the 10 noisy images generated, a table like the one shown below can be derived. 1.0, 2.33, 3.67, 5.0, 6.33, 7.67, 9.0

Noise Level	Ellipsoid Diameters						
σ_N	1.0 mm	2.3 mm	3.6 mm	5.0 mm	6.3 mm	7.6 mm	9.0 mm
0.0	0.6271	0.8528	1.7308	2.4833	3.2358	3.9632	4.5903
0.1	0.6522	0.8779	1.7057	2.4833	3.2358	3.9632	4.6405
0.2	0.7274	0.8278	1.7308	2.5084	3.2358	3.9632	4.6656
0.3	0.7274	1.1538	1.7057	2.4331	3.1355	3.8629	4.4398
0.4	0.7525	1.1789	1.7559	2.5334	3.1605	3.6873	4.5401
0.5	0.8278	1.2291	1.7308	2.5836	3.1605	3.7124	4.6154
0.6	1.1037	1.3545	1.8813	2.5836	2.9599	3.7375	4.6405
0.7	1.1288	1.2542	1.9064	2.5084	3.0351	3.8378	4.7659
0.8	1.0284	1.2793	2.0318	2.4833	3.1355	3.8378	4.5903
0.9	1.2793	1.5050	1.8060	2.4833	3.1355	3.9130	4.6405
1.0	1.1288	1.5050	1.9565	2.4833	3.1856	3.8378	4.6906

Table 5: Collected data of peak values of the ellipsoids as the noise level increases illustrated in figure 25.

From the data collected in Table 5, the effect of the shifting can be visualised for each ellipsoid from the synthetic image.

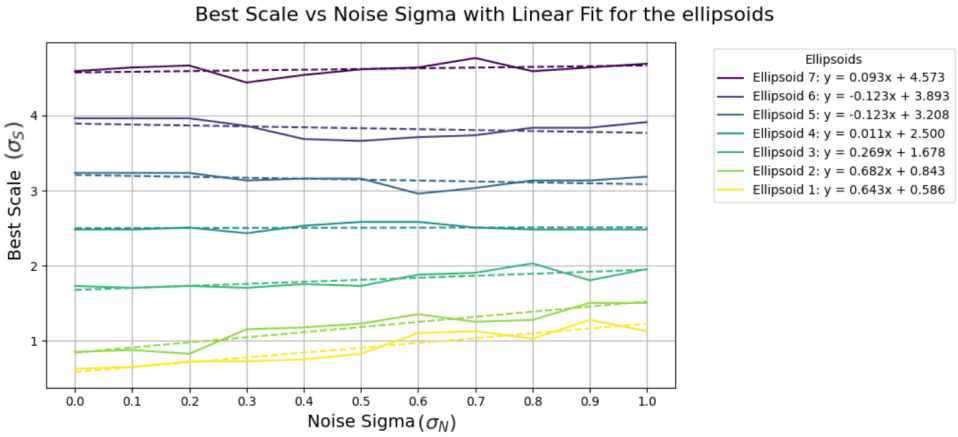

Figure 26: Plot illustrating the shift in the best scale for the 7 ellipsoids (see Figure 14) as noise increases recovered from Table 5, with a linear fit applied to each shape.

Note that after 5 millimeters, the noise effects impact less on the peak scale shifting, the focus will now shift to understanding the impact of noise on the smallest structures, which are of primary interest for this project. Therefore, a second synthetic image will be generated, similar to that in Figure 14, but with thinner ellipsoids ranging from 1 millimeters to 5 millimeters in diameter.

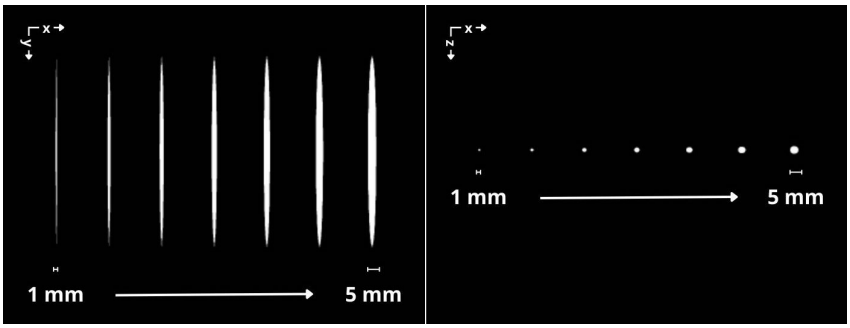

Figure 27: 3D image sample of seven ellipsoids with their diameters going from 1 to 5 millimeters progressively in Coronal view (left), and Axial view (right) considering 1 millimeters per voxel.

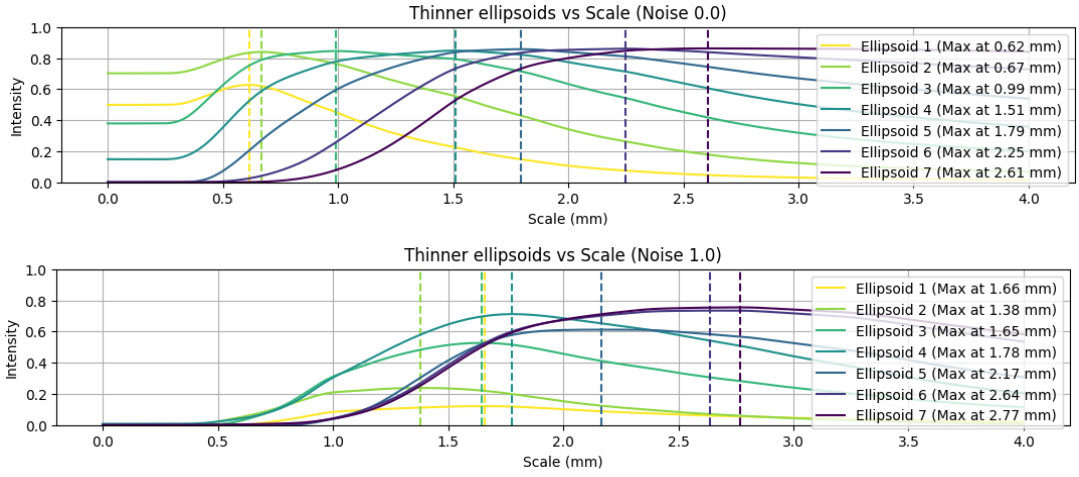

Figure 28: Plot depicting the average contrast of the thinner ellipsoids from Figure 27, before (top) and after (bottom) adding the rician noise.

After repeating the same process applied to the previous ellipsoids (see Figure 16) to the new thinner

ones, the mapping from Figure 26 can be completed.

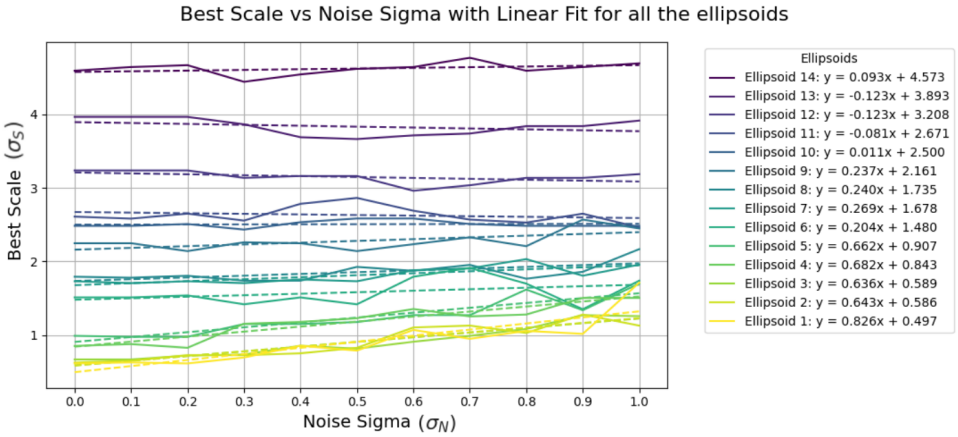

Figure 29: Completed Plot from Figure 26 including the measurements for the new thinner ellipsoids.

Based on the approximations obtained from the data analysed, it is possible to extrapolate the results to fit any noise sigma and shape diameter of intensity 1, allowing for an estimation of the best scale as a function of image noise (σ_N) and a target diameter size (potentially voxel size). This can be visualised as:

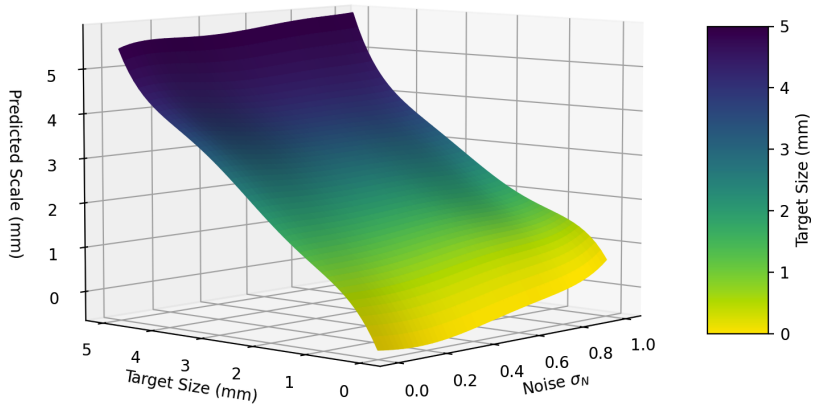

Figure 30: Illustration of the interpolated noise-based model for scale prediction in ellipsoids (based on σ_N).

The *predicted Scale* refers to the suggested scale based on an estimated noise level, σ_N, and a specified target size (where yellow corresponds to smaller shapes and blue to larger shapes). It can be observed that the yellow region of the noise-based model is more significantly impacted as the noise level increases, while the blue region remains comparatively less affected, as illustrated in Figure 29. Furthermore, every scale under the surface of the plot represents a scale where nothing can be detected, only noise. This is crucial for understanding how much information can be accessed from an image under the same conditions (resolution and quality, i.e., noise level), highlighting the limitations of detecting PVS in noisy or low-resolution images.

It is important to emphasise that this analysis aims to identify a trend from images containing random Rician noise. In real-world scenarios, the values produced by the model provide an estimate, rather than fixed or deterministic outcomes. Consequently, any replication of this experiment may yield results that differ slightly; however, these variations are expected to remain close to the original findings.

Synthetic ellipsoids of intensity 1 were used to visualise how the noise sigma influences the optimal scale, but this approach applies specifically to images where the average ROI intensity is consistently 1. In these cases, only the noise sigma is considered, but the true cause of the scale shift is not solely the noise itself. Instead, it is the **proportion** or contrast between the σ_N and the average intensity of the ROI.

In order to develop a more accurate and general model, the training process incorporates both synthetic and real MRI images, alongside their corresponding masks and actual σ_N values. The σ_N for each image is computed as illustrated in Figure 19, and the ROI is extracted as depicted in Figure 22, enabling the calculation of the SNR. By establishing this proportional relationship, the SNR is mapped to the corresponding σ_N (see Figure 23). This adjustment ensures that the model accounts for variations in the contrast within the ROI, rather than solely focusing on noise intensity.

6.3 Training Data Set

Due to privacy constraints associated with patient data, access to a substantial number of real MRI scans is limited. To address this limitation and fulfil the need for an extensive dataset, two distinct datasets are utilised.

The first dataset comprises 5 real, denoised T1-weighted (T1w) MRI images with PVS manually annotated by experts. These real images are used to generate multiple Rician noise levels for training by artificially adding noise to simulate various conditions.

Similar to the noise-based model for ellipsoids, which was trained on various target shapes representing different ellipsoid diameters, this SNR-based model emphasises voxel sizes. Here, the focus is on detecting the smallest structures—particularly based on voxel size—and thus includes images with voxel sizes ranging from 0.7 millimeters^3 to 1.8 millimeters^3.

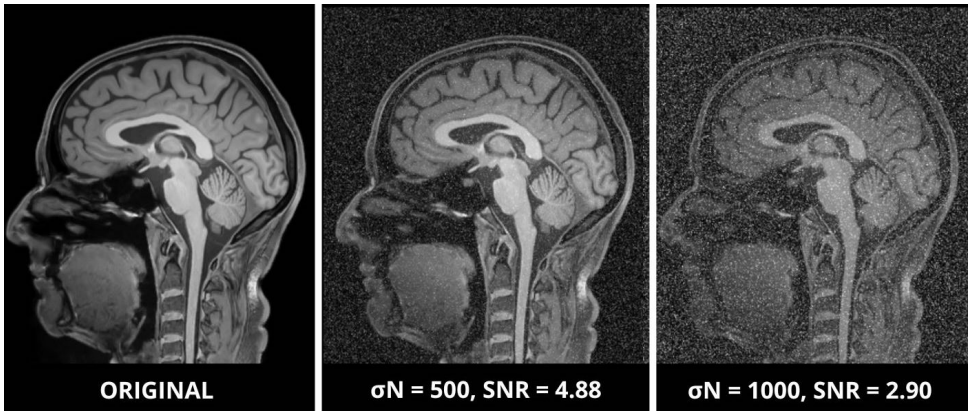

Figure 31: MRI sample from the real image dataset with varying levels of applied noise.

The filter demonstrates similar behaviour on real images as observed with the ellipsoid simulations, particularly in noise measurement.

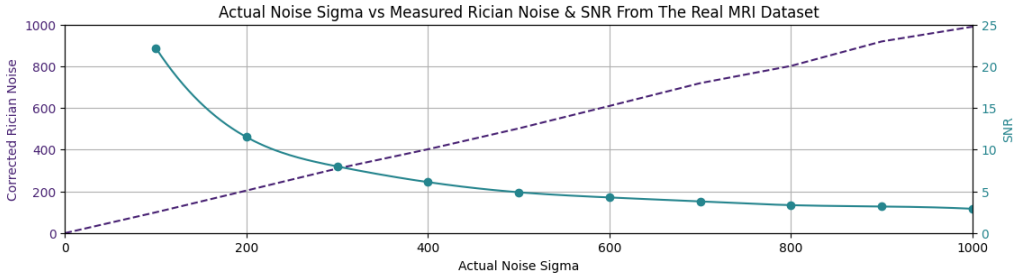

Figure 32: Plot illustrating the change in Noise Calculation and SNR from the Figure 31.

This consistency supports the reliability of the SNR calculations as a robust input for the SNR-based model.

The second dataset is derived from a synthetic MRI generator that produces realistic representations of PVS within each generated image. This synthetic dataset includes multiple MRI sequences (FLAIR, T1w, T2w) and corresponding high-resolution PVS masks. The synthetic images already contain varying levels of noise, which makes them suitable for training.

The synthetic MRI generator is a reliable tool and has been previously used successfully in training other models, making it a valuable option for this project where extensive data is required for robust model development.

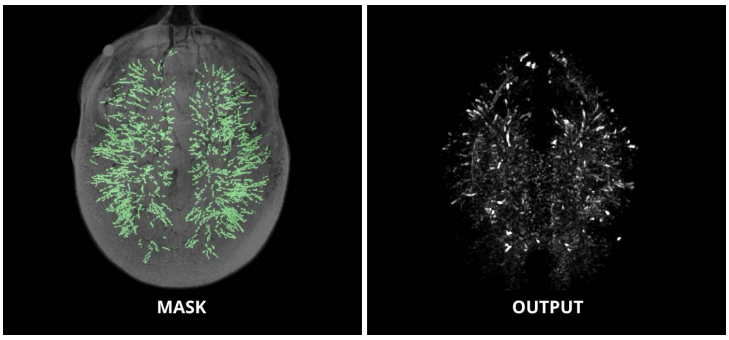

Figure 33: Illustration of a dataset sample: the PVS mask (green) overlaid on the MRI (left) and the optimal CSO-masked output for that image (right) shown in a volume render.

With a substantial collection of MRI scans in FLAIR, T1-weighted, and T2-weighted formats, spanning varying image qualities, the data necessary for model training is now available.

A "shell" of 1.5 millimetres in thickness is generated from each mask to measure the intensity of surrounding voxels.

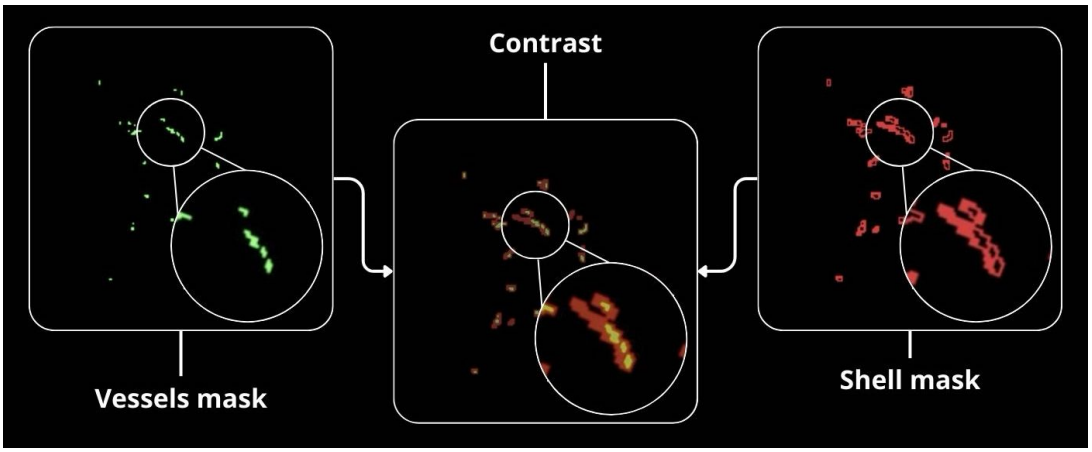

Figure 34: Slice of the Vessels mask (green) and the Shell mask (red), used to measure contrast in PVS across different scales.

Each image in both datasets is labelled according to the method illustrated in Figure 34 with the following attributes: name, calculated noise standard deviation (σ_N), SNR, and optimal scale, measured manually in the same manner as for ellipsoids (see Figure 28).

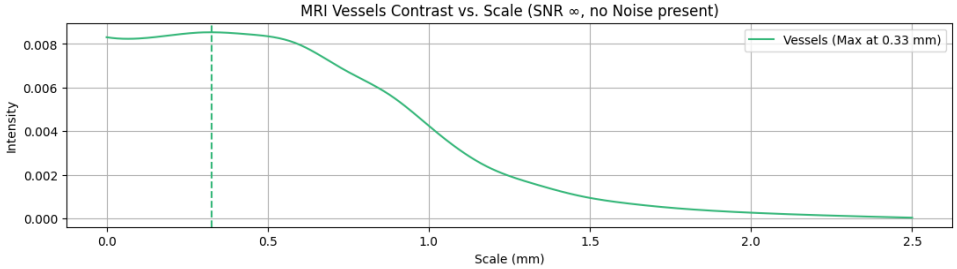

Figure 35: Plot depicting the average contrast of the vessels in the image from Figure 33, highlighting the scale at which maximum intensity is achieved (scales ranging from 0 to 2.5 millimetres).

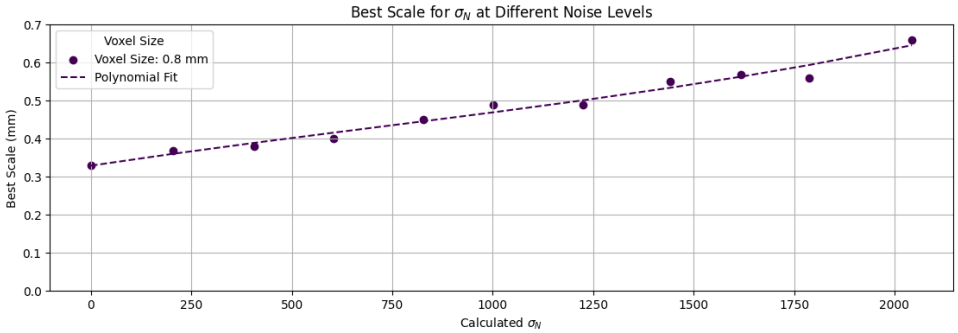

Figure 36: Plot depicting the best Scale for the vessels from Figure 35 at different noise levels.

Consequently, the SNR-based proposed model will be developed based exclusively on measurements derived from MRI images with varying noise levels (σ_N values).

6.4 SCALR Model Implementation

The SCALR model utilises a polynomial regression approach to predict the optimal scale parameter for the Frangi filter. The following steps outline the model workflow:

- **Dataset Processing:** The model loads data from `results.csv`, which contains image measurements such as file name, measured σ_N, SNR, voxel size, and the optimal scale for each MRI image in the dataset. Infinite values in the SNR field are managed by substituting them with a maximum set value of 1000 to avoid processing issues. To manage skewness, a log transformation $\log(1+\text{SNR})$ is applied to the SNR values.

ID	Image Name	STD sigma	SNR	Voxel Size (mm)	Best Scale (mm)
94	PCB_003_noise...	1406.748062	2.354789	0.80	0.621118
196	PCB_005_noise...	1923.001519	2.092076	0.80	0.744147
298	KUU220_62_hm...	2.837888	42.471551	1.68	0.803577
112	PCB_005_noise...	599.456508	4.294113	1.00	0.602007
275	GPI693_62_hm...	1.844610	61.093512	0.84	0.581940
154	PCB_004_noise...	204.210843	12.734197	1.00	0.342545
75	PCB_002_noise...	2009.348494	2.271266	1.00	0.819751
171	PCB_001_noise...	1406.111453	2.497276	1.00	0.674470
280	GPI693_62_hm...	4.301806	26.646066	1.01	0.470497
105	PCB_003_noise...	1406.748062	2.354789	0.80	0.622074

Table 6: Summary of MRI image data with corresponding standard deviation of noise, SNR, voxel size, and optimal Frangi filter scale.

- **Feature Selection:** Two key features are selected as inputs: **SNR** and **voxel size**, forming the input matrix $X_\alpha = [\text{SNR}, \text{Voxel Size}]$. The target variable, y_α, represents the optimal Frangi filter scale. These features encapsulate essential information regarding image quality and spatial resolution, directly influencing the necessary scale for accurate vessel detection.
- **Polynomial Transformation and Model Training:** A train-test split is implemented, partitioning the data into training and testing sets, with 80% allocated for training and 20% for testing. To capture potential non-linear relationships between SNR, voxel size, and optimal scale, a polynomial transformation of degree 3 is applied. The transformed feature set enables the linear regression model to fit a polynomial function of the form:

$$y_\alpha = a_0 + a_1 \cdot \text{SNR} + a_2 \cdot \text{Voxel Size} + a_3 \cdot (\text{SNR})^2 + a_4 \cdot (\text{Voxel Size})^2 + \ldots + a_n \cdot (\text{Voxel Size} \cdot \text{SNR})^3$$

where $a_0, a_1, \ldots, a_n$ are coefficients learned during training. This polynomial fit allows the model to capture interactions and higher-order terms, enabling accurate scale predictions tailored to each MRI's quality.

- **Model Evaluation:** To evaluate predictive accuracy, metrics such as **Mean Squared Error (MSE)** and **R^2 Score** are computed. These metrics offer insights into the model's accuracy and its ability to generalise to unseen data.
- **SCALR Prediction Function:** The SCALR model provides real-time scale predictions for new MRI data. For any new MRI image, the function extracts the SNR and voxel size, applies the polynomial transformation, and returns the predicted scale based on the learned relationship between MRI quality and optimal scale.

6.5 Results

The SCALR model was assessed on a test set comprising 300 MRI images, each varying in SNR and voxel size to ensure robustness across diverse imaging conditions. These images were selected to provide a representative sample of the structural and noise variations typically encountered in clinical MRI datasets.

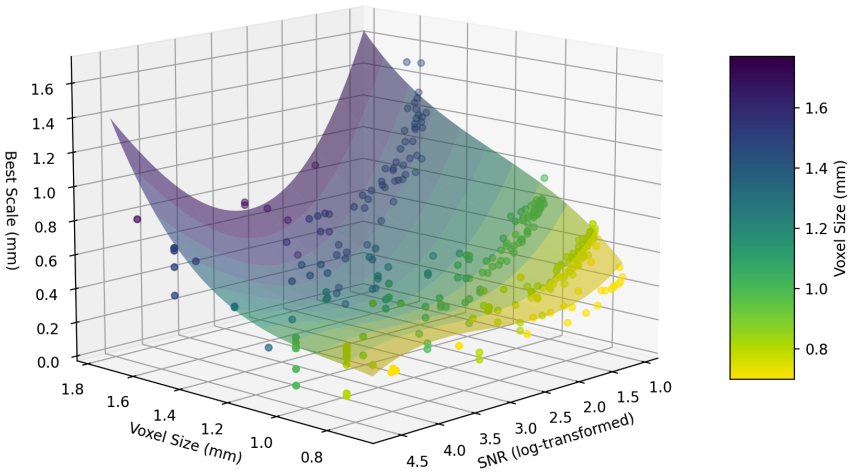

Figure 37: 3D Surface plot of the SCALR model. This visualisation illustrates the relationship between SNR, voxel size, and the predicted optimal scale, showing how the SCALR model dynamically adjusts scale predictions in response to image quality parameters.

Figure 6.5 illustrates a 3D visualisation of the SCALR model's polynomial surface, depicting the relationship between SNR, voxel size, and predicted optimal scale. This plot offers a detailed view of the model's response surface, clarifying how the SCALR model dynamically adjusts scale predictions in accordance with image quality parameters. The model's performance metrics are summarised below:

Metric	Value
Mean Squared Error (MSE)	0.0128
R^2 Score	0.8035

Table 7: SCALR Model Evaluation

- **Mean Squared Error (MSE):** The model achieved a Mean Squared Error (MSE) of 0.012, indicating a low average squared discrepancy between the predicted and actual optimal scale values. This metric reflects the precision with which the model estimates the scale needed for accurate PVS segmentation, minimising scale selection errors.
- **R^2 Score:** An R^2 score of 0.803 was obtained, signifying that approximately 80.3% of the variance in the optimal Frangi filter scale is explained by the model's input features—namely, the SNR and voxel size. This result highlights the model's effectiveness in capturing the relationship between

imaging quality and optimal scale, even when faced with challenges such as non-uniform vessel thickness and stochastic noise.

Subject	SNR	Voxel Size (mm)	Best Scale (mm)	Predicted Scale (mm)
1	1.601631	1.00	0.759231	0.764707
2	11.018970	1.01	0.470497	0.333068
3	3.741848	1.50	0.475920	0.666342
4	17.045681	0.84	0.394646	0.412803
5	14.627818	0.80	0.626873	0.591838

Table 8: SCALR Model Predictions

Tables 6.5 and 8 present quantitative metrics and individual predictions, respectively, demonstrating the SCALR model's performance across varied SNR and voxel sizes. These tables highlight the model's capacity to generalise across data with diverse imaging parameters, affirming its potential utility in automatic scale selection for PVS quantification.

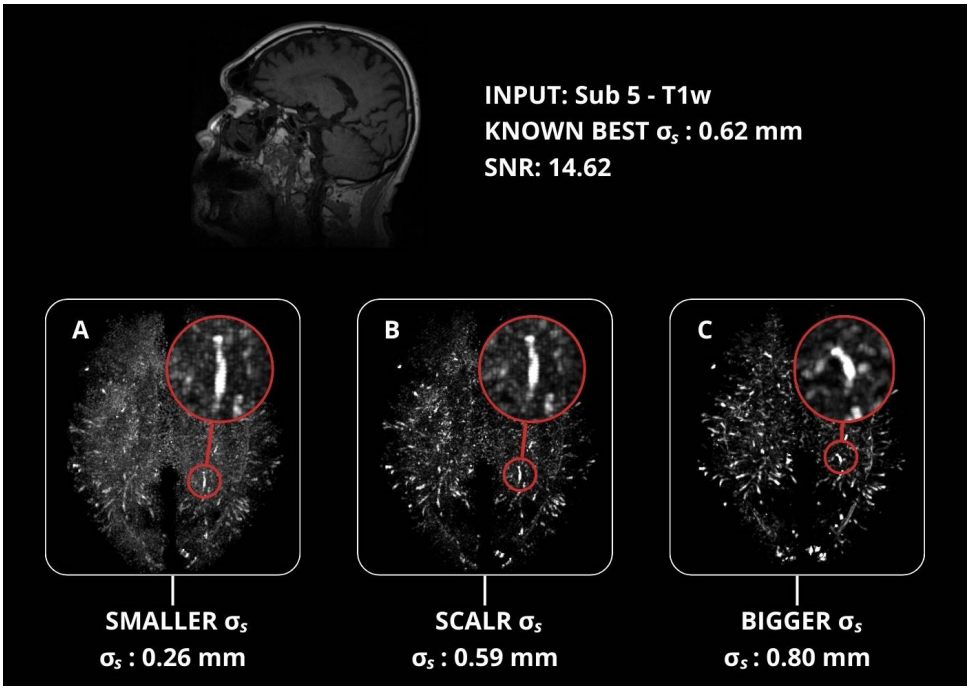

Figure 38: Volume rendering of the segmented PVS for Subject 5 from Table 8, shown at the optimal scale recommended by SCALR (B), with additional renderings at adjacent scales (A and C) for contrast comparison.

Figure 38 provides a comparative volume rendering of the segmented PVS from an MRI image processed at different scales. The middle image (labelled B) shows the optimal scale as predicted by the SCALR model, where the PVS structures are visible with high detail and a reduced level of noise, enabling clear distinction from surrounding tissues. For comparison, renderings at adjacent scales are included: in subfigure A, a smaller scale reveals the PVS but also introduces substantial noise, which complicates the segmentation. In subfigure C, a larger scale reduces the noise to a greater extent, yet the morphology of the PVS appears altered, with finer details less discernible.

7 Conclusions and Future Work

The SCALR model facilitates the application of the Frangi filter in MRI analysis, aiding clinicians and healthcare professionals by efficiently identifying the optimal scale for vessel quantification. This eliminates the need for time-intensive trial-and-error or brute-force methods, allowing users to confidently extract maximal information from MRI images. Consequently, SCALR holds the potential to streamline diagnostic workflows, enabling more time for analysis and interpretation.

The difference between the model trained on synthetic images and that trained on real images lies in the scaling approach: synthetic ellipsoids are scaled linearly (σ_N), whereas a logarithmic scale is applied to real images to account for the exponential behaviour of the SNR. Furthermore, real PVS structures deviate from the idealised cylindrical shape. Despite these variations, both models exhibit similar behaviour in optimal scale selection, suggesting that the ideal scale for segmenting an MRI using the Frangi filter is indeed influenced by the image's resolution and noise level.

As SCALR relies on image quality, certain preprocessing methods can be applied to the images before processing them with SCALR, such as image denoising, intensity rescaling, or histogram matching with high-quality reference MRIs. These steps help SCALR identify the smallest possible optimal scale for vessel detection.

This research demonstrates the role of machine learning in extracting patterns that can reduce resource expenditure in terms of time and computational costs. The model's high R^2 value supports its predictive validity, aligning well with real-world data. Additionally, SCALR is computationally lightweight, ensuring efficient performance across different devices, with potential applications in advancing PVS as a biomarker for neurological conditions.

Some limitations of the model have been identified. Firstly, scale prediction becomes less stable with very low-resolution images (e.g., voxel dimensions over 2 mm per voxel), though this is of limited concern, as MRI scans intended for PVS detection typically fall within SCALR's trained voxel range. If low-resolution scans are used, they inherently restrict the accuracy of PVS quantification due to the need for smaller voxels in high-quality assessments. Secondly, the lack of large, annotated MRI datasets constrains the model's precision compared to larger models trained on extensive data. Nonetheless, SCALR's focused training approach allows for reliable results even with limited data.

Future improvements to SCALR may involve broadening the training dataset to include a more diverse range of MRI images, enhancing the model's predictive accuracy and robustness. Advanced noise quantification methods could further improve SNR measurements, contributing to prediction precision. Additionally, retraining the model with images standardised to consistent intensity ranges, such as 0–1 or 0–1000, could improve stability by minimising SNR discrepancies and focusing on σ_N, a more consistent metric. Preprocessing enhancements may also refine scale determination accuracy, boosting the model's efficacy in vessel detection.

In this research, a reliable method for quantifying image quality (noise or SNR) has been established, along with a strong correlation between image data and optimal scale selection. The success of SCALR, even with a small dataset, indicates potential for further enhancement through extended training. By employing a linear regression model instead of a CNN or other black-box approach, SCALR remains interpretable, allowing for visualisation and understanding of its functional relationships. This transparency facilitates characterisation of optimal and suboptimal cases, making its predictions more dependable than purely statistical models. Consequently, this work aims to contribute to the state-of-the-art in vessel detection, enhancing the accessibility and usability of the Frangi filter as a tool for medical image analysis.

8 References

References

[1] Moses J, Sinclair B, Law M, O'Brien TJ, & Vivash L. Automated Methods for Detecting and Quantitation of Enlarged Perivascular spaces on MRI. J Magn Reson Imaging, 2023. https://doi.org/10.1002/jmri.28369.

[2] J. Ramirez, C. Berezuk, A. McNeely, G. F., J. McLaurin, & S. Black, "Imaging the perivascular space as a potential biomarker of neurovascular and neurodegenerative diseases," 2018. https://www.researchgate.net/publication/298908350_Imaging_the_Perivascular_Space_as_a_Potential_Biomarker_of_Neurovascular_and_Neurodegenerative_Diseases.

[3] Jennifer M.J. Waymont, Maria C. Valdés Hernández, Jose Bernal, Roberto Duarte Coello, Rosalind Brown, Francesca M. Chappell, Lucia Ballerini, Joanna M. Wardlaw "A Systematic Review and Meta-Analysis of Automated Methods for Quantifying Enlarged Perivascular Spaces in the Brain" medRxiv 2024. doi:https://doi.org/10.1101/2024.03.04.24303705.

[4] Chappell Potter, M. Z, & Wardlaw, "Cerebral perivascular spaces visible on magnetic resonance imaging: Development of a qualitative rating scale and its observer reliability,". 2015. doi:10.1159/000375153.

[5] Foruzan, A.H. & Zoroofi, Reza & Sato, Yoshinobu & Hori, Masatoshi. (2011). A Hessian-based filter for vascular segmentation of noisy hepatic CT scans. *International journal of computer assisted radiology and surgery. 7. 199-205.* https://doi.org/10.1007/s11548-011-0640-y.

[6] A. F. Frangi, W. J. Niessen, K. L. Vincken, & M. A. Viergever, "Multiscale vessel enhancement filtering," In *Medical Image Computing and Computer-Assisted Intervention – MICCAI'98*, pp. 130-137, Springer, Berlin, Heidelberg, 1998. https://doi.org/10.1007/BFb0056195

[7] M. W. B. van der Walt, S. C. Colbert, and G. Varoquaux, 2011. "The NumPy Array: A Structure for Efficient Numerical Computation," *Computing in Science and Engineering, vol. 13, no. 2, pp. 22-30.* https://doi:10.1109/MCSE.2011.37.

[8] D. C. J. R. C. C. O. S. S. R. M. B. F. S. P. H. R. A. S. D. A. J. O. W. F. P. M. J. O. W. 2014. "SimpleITK: A simplified layer built on top of the Insight Segmentation and Registration Toolkit," https://simpleitk.readthedocs.io/en/master/.

[9] Osadebey, M.E., Pedersen, M., Arnold, D.L. et al. (2018). Blind blur assessment of MRI images using parallel multiscale difference of Gaussian filters. *BioMed Eng OnLine 17, 76* https://doi.org/10.1186/s12938-018-0514-4

[10] Debnath, Arunabha & Rai, Hari & Yadav, Chahat & Agarwal, Anjali & Bhatia, Ankit. (2013). Deblurring and Denoising of Magnetic Resonance Images using Blind Deconvolution Method. *International Journal of Computer Applications. 81. 7-12.* https://doi.org/10.5120/14046-2209.

[11] Jose Bernal, Maria D.C. Valdés-Hernández, Javier Escudero, Roberto Duarte, Lucia Ballerini, Mark E. Bastin, Ian J. Deary, Michael J. Thrippleton, Rhian M. Touyz, Joanna M. Wardlaw, 2022. Assessment of perivascular space filtering methods using a three-dimensional computational model. *Magnetic Resonance Imaging, Volume 93, 33-51* https://doi.org/10.1016/j.mri.2022.07.016.

[12] H. Adams, G. E. Brancati, D. Bos, & M. A. Ikram, "Machine learning in cerebral small vessel disease imaging: A systematic review," *NeuroImage: Clinical, 30*, 102667, 2021. https://doi.org/10.1016/j.nicl.2021.102667

[13] F. Doshi-Velez, & B. Kim, "Towards a rigorous science of interpretable machine learning," *arXiv preprint arXiv:1702.08608*, 2017.

[14] H. Chen, Q. Dou, L. Yu, J. Qin, & P. A. Heng, "Vessel enhancing techniques for MR angiography using deep learning," *Medical Image Analysis, 51*, 20-34, 2019. https://doi.org/10.1016/j.media.2018.09.008

[15] C. Rudin, "Stop explaining black box machine learning models for high-stakes decisions and use interpretable models instead," *Nature Machine Intelligence, 1*(5), 206-215, 2019. https://doi.org/10.1038/s42256-019-0048-6

[16] S. H. Park, & K. Han, "Methodologic guide for evaluating clinical performance and effect of artificial intelligence technology for medical diagnosis and prediction," *Radiology, 286*(3), 800-809, 2018. https://doi.org/10.1148/radiol.2017171920

[17] G. Chartrand, P. M. Cheng, E. Vorontsov, M. Drozdzal, S. Turcotte, C. J. Pal, & S. Kadoury, "Deep learning: A primer for radiologists," *Radiographics, 37*(7), 2113-2131, 2017. https://doi.org/10.1148/rg.2017170077

[18] A. Esteva, K. Chou, S. Yeung, N. Naik, A. Madani, A. Mottaghi, ... & J. Dean, "Deep learning-enabled medical computer vision," *npj Digital Medicine, 4*(1), 1-9, 2021. https://doi.org/10.1038/s41746-020-00376-2

[19] G. Litjens, T. Kooi, B. E. Bejnordi, A. A. A. Setio, F. Ciompi, M. Ghafoorian, ... & B. van Ginneken, "A survey on deep learning in medical image analysis," *Medical Image Analysis, 42*, 60-88, 2017. https://doi.org/10.1016/j.media.2017.07.005

[20] B. Ebdrup, H. Nørbak, S. Borgwardt, & B. Glenthøj, "Volumetric changes in the basal ganglia after antipsychotic monotherapy: A systematic review," *Current medicinal chemistry*, vol. 20, no. 3, pp. 438-447, 2012. doi:10.2174/0929867311320030015.

[21] V. Priyanka, Kashyap, S. Bhatt, S. J. Bhat, P. Kashyap, & V. Prasad, "MR Parkinsonian index in PSP," Jan. 2017. https://www.researchgate.net/publication/350098663_MR_Parkinsonian_index_in_PSP.

[22] S. Achint, "MRI atlas of normal myelination," 2017. https://www.myelinationmriatlas.com/3-months.html.

[23] C.-Y. Hsu, B. Schneller, M. Ghaffari, A. Alaraj, & A. Linninger, "Medical image processing for fully integrated subject-specific whole brain mesh generation," *Technologies*, vol. 3, pp. 126–141, 2015. doi:10.3390/technologies3020126.

[24] M. Lamy, M. Merveille, A. Kerautret, O. Passat, & G. Vacavant, "Vesselness filters: A survey with benchmarks applied to liver imaging," 2020. https://hal.archives-ouvertes.fr/hal-02544493/document.

[25] Bernal, J. et al. (2020). "A Framework for Jointly Assessing and Reducing Imaging Artefacts Automatically Using Texture Analysis and Total Variation Optimisation for Improving Perivascular Spaces Quantification in Brain Magnetic Resonance Imaging". In: Papież, B., Namburete, A., Yaqub, M., Noble, J. (eds) Medical Image Understanding and Analysis. *MIUA 2020. Communications in Computer and Information Science, vol 1248. Springer, Cham.* https://doi.org/10.1007/978-3-030-52791-4_14

[26] Bernal, J. et al. (2021). "Selective Motion Artefact Reduction via Radiomics and k-space Reconstruction for Improving Perivascular Space Quantification in Brain Magnetic Resonance Imaging". *In: Papież, B.W., Yaqub, M., Jiao, J., Namburete, A.I.L., Noble, J.A. (eds) Medical Image Understanding and Analysis. MIUA 2021. Lecture Notes in Computer Science, vol 12722. Springer, Cham.* https://doi.org/10.1007/978-3-030-80432-9_12

[27] J. Bernal, M. Valdés-Hernández, J. Escudero, R. Duarte, L. Ballerini, M. E. Bastin, I. J. Deary, M. J. Thrippleton, R. M. Touyz, & J. M. Wardlaw, "Assessment of Perivascular Space Filtering Methods Using a Three-Dimensional Computational Model," *Magnetic Resonance Imaging*, 2022. doi:10.1016/j.mri.2022.07.016.

[28] T. Lindeberg, "Scale-space," in *Scale-Space Theory in Computer Vision*, Springer Berlin Heidelberg, 2009, pp. 1–44. doi:10.1007/978-3-642-02465-4_1.

[29] D. Collins, "Lecture 10: Pyramids and scale space," 2021. https://www.cs.cmu.edu/~./enron/docs/ML/pyramids_and_scale_space.pdf.

[30] E. W. Cheney, & D. R. Kincaid, *Numerical Mathematics and Computing*, 7th ed., Cengage Learning, 2007.

[31] S. Pintea, M. L. Ionescu, D. Enescu, & M. Bogdan, "Perivascular spaces: Imaging modalities and their relationship with cognition," *Frontiers in Human Neuroscience*, vol. 15, 2021. 10.3389/fnhum.2021.754891.

[32] Y. Tominari, & K. Shimizu, "Automatic segmentation of perivascular spaces using deep learning in MR images," *Medical Imaging and Health Informatics*, 2023. 10.1016/j.mihi.2023.101894.

[33] K. Ridgeway, M. R. K. Baig, K. Wang, T. K. Sun, H. Wang, & N. D. Le, "Analyzing perivascular spaces and white matter hyperintensities in cognitively healthy older adults using machine learning techniques," *Nature Scientific Reports*, vol. 13, no. 1, 2023. 10.1038/s41598-023-45294-y.

[34] H. Khader, A. K. Bsharat, & S. J. Shaar, "Fuzzy c-means clustering for segmentation of perivascular spaces," *Journal of Medical Systems*, 44, 2020. 10.1007/s10916-020-01725-5.

[35] M. Ismail, A. D. Brown, S. D. O'Connell, & L. P. Collins, "Vesselness filtering method based on deep learning and traditional algorithms for MRA images," *International Journal of Computer Assisted Radiology and Surgery*, 2021. 10.1007/s11548-021-02470-w.

[36] GitHub Documentation: Repository with all the code implemented in this project (open source). [Online]. Available: https://github.com/Riddimental/Frangi-GUI.