

Assessment of perivascular space filtering methods using a three-dimensional computational model

Jose Bernal^{1,5,6}, Maria d. C. Valdés-Hernández^{1,3,*}, Javier Escudero², Roberto Duarte¹, Lucia Ballerini¹, Mark E. Bastin^{1,3}, Ian J. Deary³, Michael J. Thrippleton¹, Rhian M. Touyz⁴, and Joanna M. Wardlaw^{1,3}

¹ Centre for Clinical Brain Sciences, The University of Edinburgh, Edinburgh, UK² Institute for Digital Communications, The University of Edinburgh, Edinburgh, UK

³ Lothian Birth Cohorts group, Department of Psychology, The University of Edinburgh, UK

⁴ Research Institute of the McGill University Health Centre, McGill University, Montréal, Canada

⁵ Institute of Cognitive Neurology and Dementia Research, Otto-von-Guericke University Magdeburg, Magdeburg, Germany

⁶ German Centre for Neurodegenerative Diseases (DZNE), Magdeburg, Germany

* Corresponding author

Abstract

Growing interest surrounds the assessment of perivascular spaces (PVS) on magnetic resonance imaging (MRI) and their validation as a clinical biomarker of adverse brain health. Nonetheless, the limits of validity of current state-of-the-art segmentation methods are still unclear. Here, we propose an open-source three-dimensional computational framework comprising 3D digital reference objects and evaluate the performance of three PVS filtering methods under various spatiotemporal imaging considerations (including sampling, motion artefacts, and Rician noise). Specifically, we study the performance of the Frangi, Jerman and RORPO filters in enhancing PVS-like structures to facilitate segmentation. Our findings were three-fold. First, as long as voxels are isotropic, RORPO outperforms the other two filters, regardless of imaging quality. Unlike the Frangi and Jerman filters, RORPO's performance does not deteriorate as PVS volume increases. Second, the performance of all "vesselness" filters is heavily influenced by imaging quality, with sampling and motion artefacts being the most damaging for these types of analyses. Third, none of the filters can distinguish PVS from other hyperintense structures (e.g. white matter hyperintensities, stroke lesions, or lacunes) effectively, the area under precision-recall curve dropped substantially (Frangi: from 94.21 [IQR 91.60, 96.16] to 43.76 [IQR 25.19, 63.38]; Jerman: from 94.51 [IQR 91.90, 95.37] to 58.00 [IQR 35.68, 64.87]; RORPO: from 98.72 [IQR 95.37, 98.96] to 71.87 [IQR 57.21, 76.63] without and with other hyperintense structures, respectively). The use of our computational model enables comparing segmentation methods and identifying their advantages and disadvantages, thereby providing means for testing and optimising pipelines for ongoing and future studies.

Keywords

Digital reference object; Perivascular spaces; Spatio-temporal imaging artefacts; Perivascular space filtering; Cerebral small vessel disease

1. Introduction

Perivascular spaces (PVS) in the brain are small fluid-filled tubular structures that surround cerebral microvessels usually less than 3mm maximum diameter (Wardlaw et al., 2020). As part of the brain fluid exchange, or glymphatic system, these are thought to play an essential role in waste elimination, interstitial fluid exchange, and, more broadly, in maintaining healthy brain function (Wardlaw et al., 2020). If enlarged or dilated, PVS become visible on structural magnetic resonance imaging (MRI) (Wardlaw et al., 2020). Even though a few visible PVS may be normal at any age, their abundance and enlargement are associated with ageing, hypertension, altered BBB permeability (Li et al., 2018), and other neuroimaging features of small vessel disease (SVD) (Francis et al., 2019). Accurate quantification of PVS is therefore crucial to determine their involvements in SVD and in the nervous system in a broader context.

The evolution of neuroradiological imaging technologies has enabled visualising enlarged PVS. Quantification for research purposes has, until recently, relied primarily on clinical visual scores (Heier et al., 1987; Patankar et al., 2005; Potter et al., 2015; Rouhl et al., 2008), which approximate the frequency of PVS in several standard brain regions according to ordinal scales (Ballerini et al., 2018). However, the ordinal visual scales are relatively insensitive, may be prone to variation between observers, especially when co-occurring with other SVD features (González-Castro et al., 2017), and do not capture other aspects of PVS such as length, width, individual or total PVS volumes.

The use of computational techniques could help to overcome the limitations of visual scores if accurate, and reduce operator time and subjectivity. However, computational quantification of subtle neuroimaging features such as small visible PVS (i.e., albeit in a macroscopic scale) is extremely challenging due to the small size, the frequency of artefacts, and substandard processing workflows (Ballerini et al., 2020b, 2018). Additionally, from a computational perspective, two additional issues limit their translation to clinics:

- **Proposal validation and lack of a ground truth:** the development of computational processing techniques requires a ground truth, e.g. clinically validated segmentation of the structures of interest. Nonetheless, building a ground truth for PVS is an extremely time-consuming, laborious, and impractical task to perform, especially on a large scale and when dealing with abundant

structures of small size. Moreover, validation of techniques may require correlating *in vivo* and *ex vivo* observations, making it impractical in real life.

- **Subtlety and imaging artefacts:** imaging artefacts arising during MRI scanning can compromise the quantification of subtle neuroimaging features. Although compensating for these visual distortions retrospectively may seem useful, careful testing and tuning is essential to avoid accidentally removing clinically relevant information.

Computational techniques that rely on explicit geometrical information often consider three major steps: pre-processing, PVS filtering, and segmentation (Ballerini et al., 2020a, 2018; Bernal et al., 2021b; Liu et al., 2020; Ramirez et al., 2015; Sepehrband et al., 2021, 2019; Valdés Hernández et al., 2020). At first, all images undergo quality control, inter- or intra-scan intensity standardisation, spatial normalisation, or noise reduction. Then, researchers employ “vesselness” filtering methods to increase the saliency of tubular structures (including PVS). Lastly, binarisation of “vesselness” response maps, segmentation mask refinement, and connected component analysis take place. In this paper, we specifically evaluate the performance and robustness of “vesselness” filtering methods. We scrutinise their behaviour in the presence of typical imaging distortions caused by sampling, motion artefacts, and noise, as well as other pathological features that may co-occur. For this purpose, we propose an open-source computational model that synthesises three-dimensional T2-weighted-like images with PVS-like structures distributed throughout the brain.

2. Materials and methods

Our computational model consisted of the steps illustrated in Figure 1 and is described in the following sections.

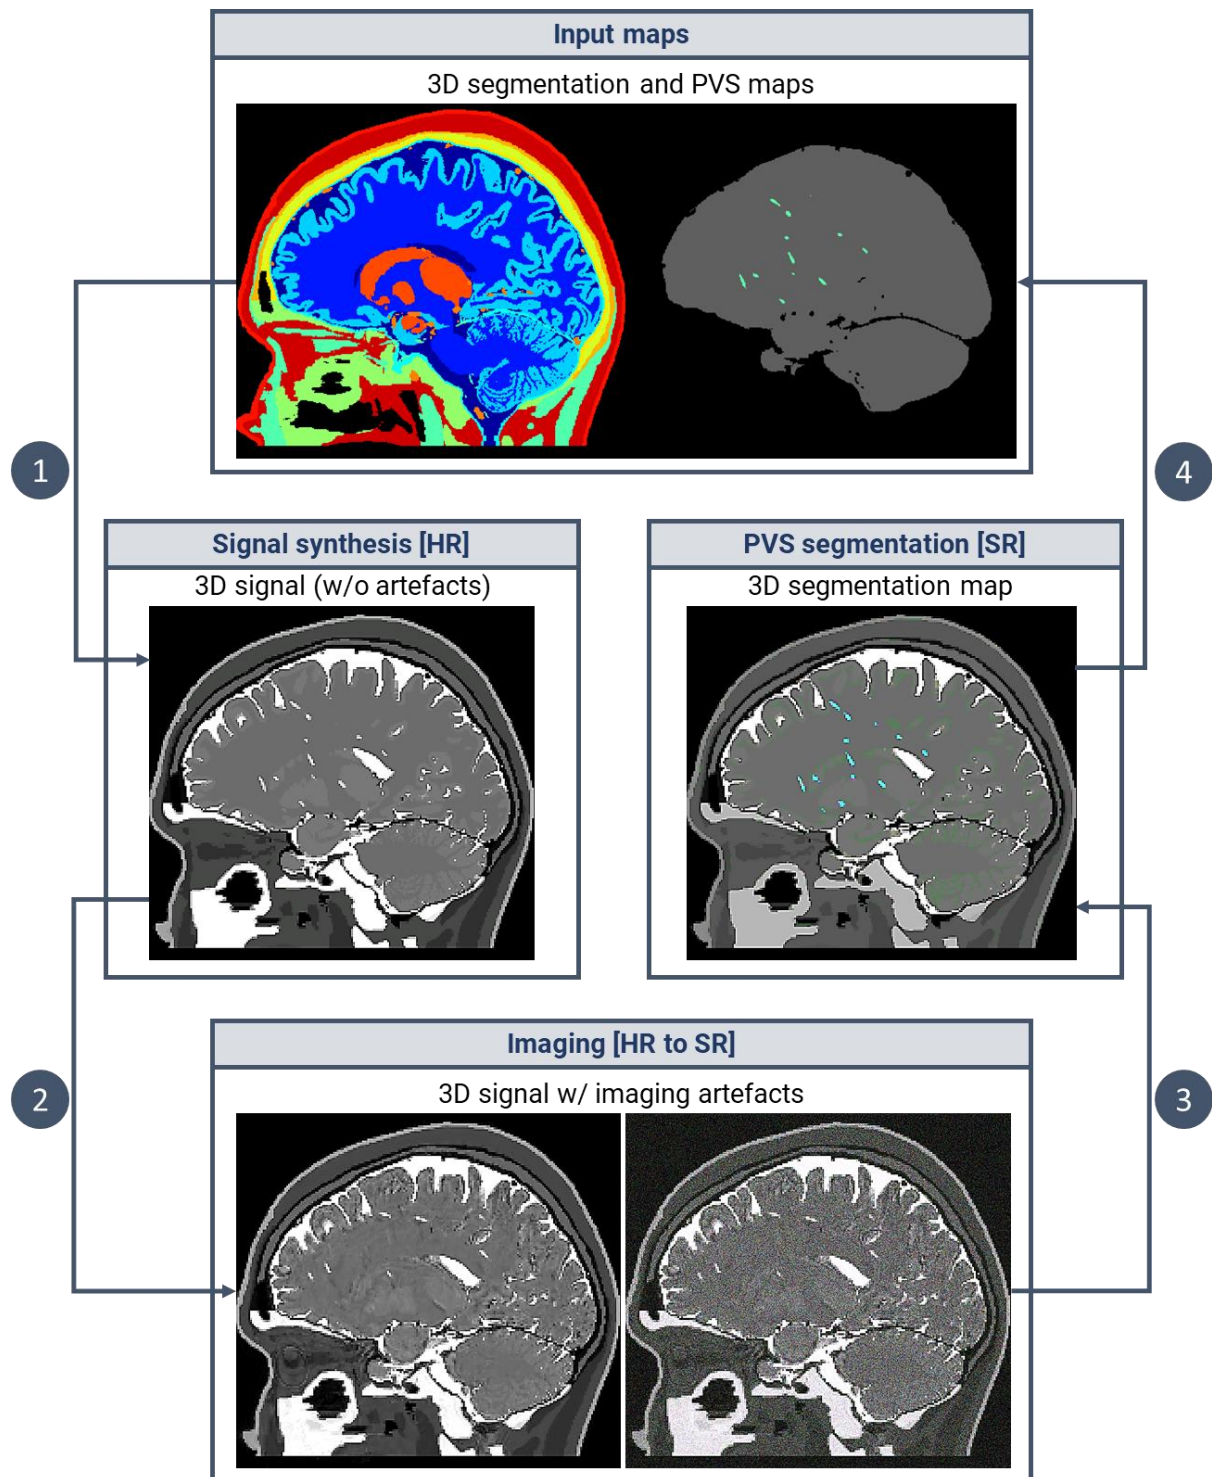


Figure 1. Computational model for assessing PVS quantification. We generated maps with PVS-like structures in different orientations and locations throughout the brain. We used a high-resolution model of the head and the resulting “PVS” maps to synthesise T2-weighted like images. We then sampled these high-resolution digital reference objects (DRO) and incorporated motion and Rician noise to generate the “acquired” T2-weighted image. We finally segmented the “PVS” on this DRO and compared the results against the ground truth. The closer the results to the ground truth, the more

resilient the “vesselness” filtering method against distortion. HR: high-resolution. SR: scanning resolution.

2.1 Computational model

2.1.1 High-resolution reference head model

We leveraged on the MIDA model, a three-dimensional 0.5×0.5×0.5-mm human head and neck atlas of a healthy volunteer (Iacono et al., 2015)¹. This publicly available digital model contained the segmentation of 116 cerebral and non-cerebral structures. We excluded extra-cerebral regions and combined some others to reduce the level of granularity and focus on those of interest: normal-appearing white matter, cortical and subcortical grey matter, and cerebrospinal fluid (Figure 2).

¹ The MIDA human head model can be downloaded from www.itis.ethz.ch/MIDA/

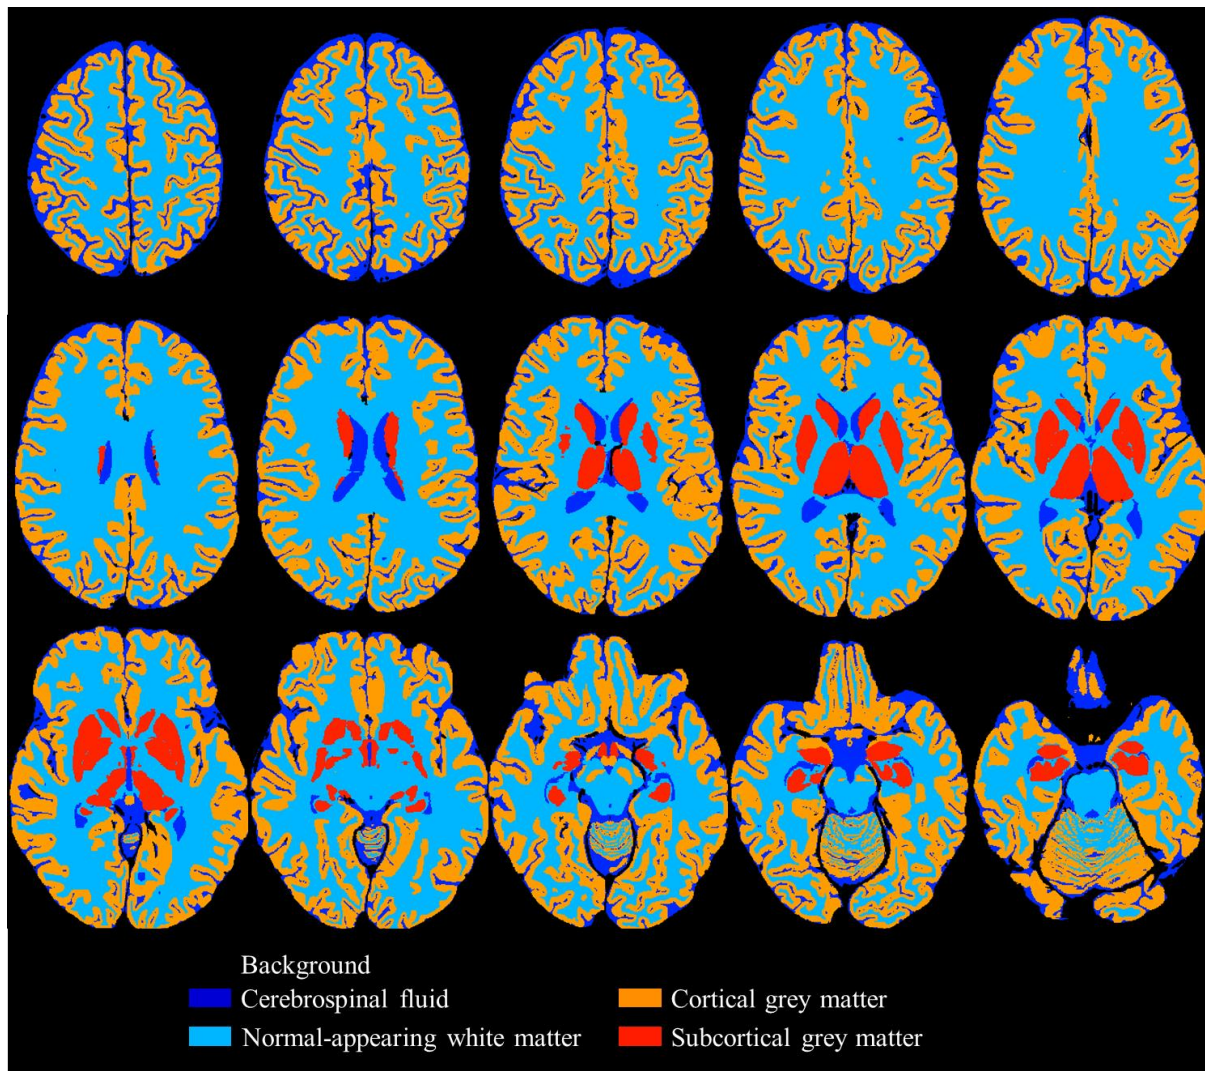


Figure 2. Brain model of reference. In total, we considered five healthy regions of interest: normal-appearing white matter, cortical and subcortical grey matter, and cerebrospinal fluid. We did not synthesise PVS in the cerebellum.

Pathological regions

The presence of white matter hyperintensities (WMH) near or around PVS may influence the ability of vesselness filters to highlight PVS appropriately. Thus, in a few experiments, we synthesised them using spatial occurrence templates extracted from patient data (<https://doi.org/10.7488/ds/2716>).

2.1.2 High-resolution ground truth masks

We created multiple high-resolution binary masks with “PVS” distributed across centrum semiovale, basal ganglia, and pontomesencephalic junction — where PVS are often seen (Wardlaw et al., 2020). To minimise circularity, we refrained from using population-based occurrence and orientation templates.

PVS model

Clinical publications have described PVS as tiny tubular structures in the brain (Wardlaw et al., 2020). We thus modelled PVS as cylinders using the following equations: $x^2 + y^2 \leq r^2$ and $z \leq h$, where $2r$ and h represent the width and height of a PVS, respectively ($r, h \in \mathbb{R}^+$).

PVS orientation

The orientation of a PVS was defined by its location inside the brain. We particularly oriented them towards the geometrical centre of the brain.

PVS location

We generated a set of random locations inside the brain parenchyma and then tested whether we could synthesise PVS at each of them. Given that a uniform sampling strategy would lead to PVS being clustered around the centre of the brain, we relied on stratified sampling instead (jittered sampling; cubic strata of length determined by the length of a PVS). Only regions of the white matter or subcortical grey matter where PVS could fit entirely were considered as valid. We created PVS in such a way they did not overlap. We opted for such an approximation to later investigate the effect of partial volume averaging on adjacent but not overlapping PVS.

2.1.3 High resolution to scanning resolution

We leveraged a recent, open-source, and publicly-available four-dimensional computational model to simulate MRI data acquisition (Bernal et al., 2021a).

k-space sampling

We calculated the k -space representation of the high-resolution digital reference object (DRO) as the three-dimensional inverse Fourier transform. We resampled it to obtain k -space data with the acquired field of view and spatial resolution. We assumed three-dimensional Cartesian k -space sampling. To emulate band filtering and slab-selective excitation, we suppressed signals from outside the field of view in the frequency- and slice-encoding directions. Finite k -space sampling results in information loss, visually manifesting in Gibbs ringing artefacts and partial volume effects.

Motion artefacts

Ghosting or blurring caused by motion affect the visual quality of the MRI scans, potentially compromising the quantification of subtle neuroimaging features. We studied their effect on PVS quantification using a composite k -space model (Bernal et al., 2021a; Shaw et al., 2020): we extracted random proportions of successive k -space lines with the head on different positions and put them together to create a composite k -space (Figure 3). We generated three head positions by rotating the DRO random angles ($\pm 2.5^\circ$ for rotations).

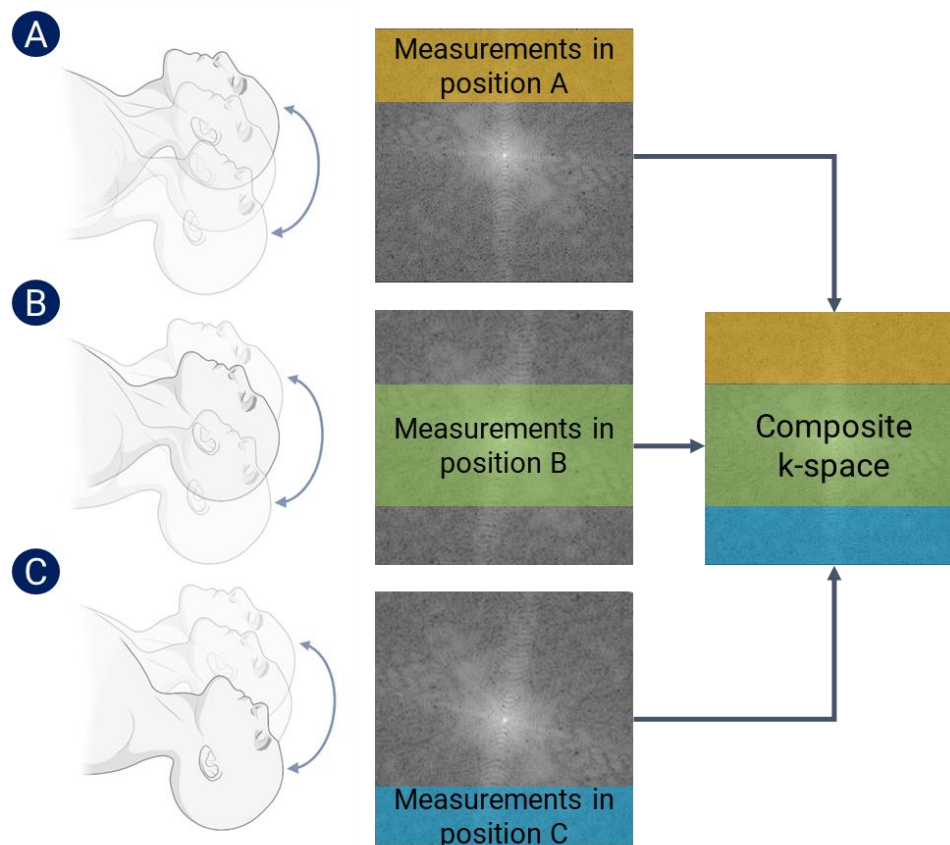


Figure 3. Simulation of motion artefacts. We simulated motion artefacts by generating composite k -spaces: we took random portions of successive k -space lines from position A, B, and C. (Figure created in part with BioRender.com).

Rician noise

The signal-to-noise ratio may reduce the performance of the PVS filtering. We investigated the effect of Rician noise by adding uncorrelated additive white Gaussian noise to the real and imaginary channels of the k -space data. We then used the three-dimensional Fourier transform and computed the magnitude to yield the DRO image, now including k -space sampling artefacts (partial volume effects and Gibbs ringing), motion artefacts (ghosting or blurring), and Rician noise.

Scanning-resolution PVS ground truth

All high-resolution binary PVS masks underwent k -space sampling to provide the PVS ground truth in the scanning resolution. We binarised the resulting maps with a threshold of 0.5, guaranteeing that voxels marked as "PVS" are mostly PVS.

2.2 In-silico experiments

2.2.1 MR protocol parameters and signal intensities for generating DROs

We estimated signal intensities from the T2-weighted (T2-w) brain MRI scans of the Lothian Birth Cohort 1936 (LBC1936) study, a large-scale study from 700 community-dwelling individuals (mean age 72.6 [SD 0.7]; 48.3% female; estimated mean signal intensities in Table 2). Data from this study have already been used to demonstrate meaningful and sensitive relations between computational PVS measurements and WMH burden and risk factors (Ballerini et al., 2020a, 2020b). Written informed consent was obtained from each participant under protocols approved by the Lothian (REC 07/MRE00/58) and Scottish Multicentre (MREC/01/0/56) Research Ethics Committees (<http://www.lothianbirthcohort.ed.ac.uk/>) (Deary et al., 2007).

Structural MRI scans were obtained using a 1.5-Tesla GE Signa Horizon HDx scanner (General Electric, Milwaukee, WI). Table 1 summarises the parameters of the structural MRI sequences used for the segmentation of normal-appearing and pathological brain regions in the LBC1936 study. Further details about the acquisition can be found in (Wardlaw et al., 2011).

Table 1. Parameters of the structural MRI sequences used for the segmentation of normal-appearing and pathological brain regions. Further details can be found in (Wardlaw et al., 2011).

Sequence	Acquisition method	Field of view [mm]	Matrix	Slices	Acquired resolution [mm]	TR/TE/TI [ms]
T2-w	Fast spin echo	256 × 256	256 × 256	80	1 × 1 × 2	11320/105/--
T2*-w	Gradient echo	256 × 256	256 × 192 ^T	80	1 × 1 × 2	940/15/--
FLAIR	Fast spin echo	256 × 256	256 × 192 ^T	40	1 × 1 × 4	9002/147/2200
T1-w	3D IR-Prep fast spoiled gradient echo	256 × 256	256 × 192 ^T	160	1 × 1 × 1.3	10/4/500

Legend ^T: zero-filled to 256 × 256.

2.2.2 Segmentation of regions of interest in the study of reference

Structural MRI scans of each participant were co-registered to the corresponding T2-w image using a 3D rigid-body transformation in FSL-FLIRT (Jenkinson et al., 2002; Jenkinson and Smith, 2001) and subsequently resampled to a 1 mm in-plane resolution.

Cerebrospinal fluid, grey matter, white matter, and white matter hyperintensities were segmented using FreeSurfer (Fischl, 2012), FSL-FAST (Zhang et al., 2001), and a semi-automated multi-spectral fusion method, which capitalised on multimodal information provided by T2-w, T1-w, T2*-w and FLAIR MRI sequences (Valdés Hernández et al., 2010; Wardlaw et al., 2011; Wheeler et al., 2021). The normal-appearing white matter was defined as regions of the white matter where no white matter hyperintensity was identified. A trained image analyst edited segmentation masks manually using the Object Extraction Tool in Analyze (AnalyzeDirect, Mayo Clinic, Rochester, Minnesota).

PVS in the normal-appearing white matter were segmented on T2-w MRI using a thoroughly validated method that uses the Frangi filter to facilitate their identification (Ballerini et al., 2020a, 2018). In addition, after segmentation, an experienced image analyst visually inspected all binary segmentation masks to remove erroneous ones. Images were considered acceptable if the method detected a reasonable amount of visible PVS and did not flag too many artefacts as potential PVS candidates (< 20% of false positives and false negatives). Around 77% (n=540) of segmentation masks passed this test (Ballerini et al., 2020a).

All of the aforementioned image analyses were performed blindly to participants' characteristics.

*Table 2. Absolute signal intensities of the T2-w MRI scans from the Lothian Birth Cohort 1936 study (Ballerini et al., 2020a). The original T2-w images were acquired with the following parameters: 11320 ms repetition time, 104.9 ms echo time, 20.83 kHz bandwidth, 2 mm slice thickness, and 256 × 256 mm field-of-view. We also present signal intensity variations between perivascular spaces and other regions of interest (ROI) as a measure of contrast (i.e. $100 * (\text{Intensity}_{\text{ROI}} / \text{Intensity}_{\text{PVS}} - 1)$). We report mean and standard deviation.*

Region of interest	Signal intensity	$100 * \left(\frac{\text{Intensity}_{\text{ROI}}}{\text{Intensity}_{\text{PVS}}} - 1 \right)$
Cerebrospinal fluid	1152.03 (236.69)	110.41 (59.59)
Normal-appearing white matter	395.54 (69.86)	-27.76 (19.00)
White matter hyperintensity	657.27 (118.05)	20.04 (31.82)
Grey matter	450.02 (82.09)	-17.80 (21.94)
Perivascular spaces	547.52 (106.74)	0.00 (0.00)

2.2.3 Experiments

We generated high-resolution T2-w like DROs with PVS of different sizes distributed all over the brain. The length of each PVS ranged from 0.5 to 10 mm and their width from 0.5 to 3 mm. We considered three main aspects of interest. First, we tested three filters on these ideal high quality high-resolution DROs to determine whether their performance varied with PVS dimensions. Second, we performed

in-silico experiments to investigate the impact of k-space sampling, motion artefacts, and Rician noise on filtering results. Third, we synthesised DROs with “WMH” to study whether the filter evaluated were able to discern between tubular and non-tubular hyperintense structures. We ran each combination (aspect of interest, PVS length, PVS width). We repeated each experiment ten times and reported mean and standard deviation (RStudio v1.2.5019 with R v3.5.1). We ran all experiments on a 189GB RAM computer running Scientific Linux 7.3 (Nitrogen; Arch x86 64; 56 CPUs Intel(R) Xeon(R) CPU E5-2683 v3 @ 2.00GHz).

2.3 “Vesselness” filtering methods

2.3.1 Multiscale Hessian-based filtering methods

PVS filtering relies heavily on the analysis of the Hessian matrix, which carries geometrical information that makes it possible to distinguish between tubular structures and round or flat ones (Frangi et al., 1998; Jerman et al., 2015; Lamy et al., 2020).

The functioning of Hessian-based filters can be divided into four major steps (Lamy et al., 2020). First, the input scan, $f: \Omega^3 \rightarrow \mathbb{R}$, is convolved with a Gaussian kernel, $\mathcal{G}_{\sigma^2}: \Omega^3 \rightarrow \mathbb{R}$ with standard deviation σ , to guarantee the continuity and differentiability required for computing second order derivatives, i.e. $f * \mathcal{G}_{\sigma^2}$. Second, for each voxel, second-order partial derivatives are computed to produce the Hessian matrix H , i.e. $(H(f * \mathcal{G}_{\sigma^2}))_{ij} = \frac{\partial^2 (f * \mathcal{G}_{\sigma^2})}{\partial x_i \partial x_j}$. Third, the eigenvalues of the Hessian matrix are found. The sign and magnitude of the eigenvalues, $\lambda_1, \lambda_2, \lambda_3 \in \mathbb{R}$, with $|\lambda_1| \leq |\lambda_2| \leq |\lambda_3|$, indicate whether a voxel contains a bright structure and it is tubular or not. In T2-w MRI sequences, the search for PVS candidates can be limited to locations where $|\lambda_1| \approx 0$ and $\lambda_2 \approx \lambda_3 \ll 0$, as these constraints limit responses to (the hyperintense in this case) cylindrical structures (Lamy et al., 2020). Fourth, further eigenvalue analysis allows for the selection of possible candidates based on certain qualities such as contrast and eccentricity. The output of this step is a heatmap, often referred to as the “vesselness” response V .

Step 1 of the Hessian-based filtering consists of convolving the input image with a Gaussian kernel. To identify objects at different scales (different sizes), the input images are convolved with a family of Gaussian kernels with different variances. The distance between the zero-crossings of the Laplacian of the Gaussian kernel increases with σ , which permits enhancing larger objects of interest. The maximum response along the considered scales is produced:

$$V_{multiscale}(f; \alpha, \beta, \gamma, \sigma_{min}, \sigma_{max}) = \max_{\sigma \in [\sigma_{min}, \sigma_{max}]} V(f * \mathcal{G}_{\sigma^2}; \alpha, \beta, \gamma), \quad (1)$$

where $\sigma_{min}, \sigma_{max} \in \mathbb{R}^+$ relate to the size of the structures of interest.

2.3.1.1 The Frangi filter

The Frangi filter (Frangi et al., 1998) — a popular Hessian-based filtering technique (Lamy et al., 2020) — examines three additional properties $R_b = |\lambda_1|/\sqrt{|\lambda_2\lambda_3|}$, $R_a = |\lambda_2|/|\lambda_3|$, and $S = \sqrt{\lambda_1^2 + \lambda_2^2 + \lambda_3^2}$ to filter blobs, lines, and low contrast structures, respectively (Frangi et al., 1998; Lamy et al., 2020). These three aspects are jointly evaluated in the following “vesselness” response:

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

and the sensitivity of the filter to each of them is adjusted via a , b , and $c \in \mathbb{R}^+$.

We used the implementation of the Frangi filter available in MathWorks (www.mathworks.com/matlabcentral/fileexchange/24409-hessian-based-frangi-vesselness-filter).

2.3.1.2 The Jerman filter

The Jerman filter (Jerman et al., 2015) is also a Hessian-based filter. This filter takes inspiration from the volume ratio, a scalar index used in diffusion tensor imaging to gauge structural anisotropy (Pierpaoli and Basser, 1996; Vilanova et al., 2006). The “vesselness” response of the Jerman filter is the following:

$$V = \begin{cases} 0 & \text{if } \lambda_2 > 0 \text{ or } \lambda_3 > 0, \\ \lambda_2^2 \cdot (\lambda_3^* - \lambda_2) \cdot \left(\frac{3}{\lambda_2 + \lambda_3^*}\right)^3 & \text{if } \lambda_2 \leq 0.5 \cdot \lambda_3^*, \\ 1 & \text{otherwise,} \end{cases} \quad (3)$$

where λ_3^* is a regularised version of λ_3 which prevents division by zero.

We used the implementation available in MathWorks (www.mathworks.com/matlabcentral/fileexchange/63171-jerman-enhancement-filter).

2.3.2 Ranking Orientation Responses of Path Operators (RORPO)

Unlike the previous methods, RORPO does not rely on differential information, but on mathematical morphology, path opening in particular (Merveille et al., 2018, 2014). Given an input binary image $X: \mathbb{R}^3 \rightarrow \{0, 1\}$ and a connectivity operator Γ , path opening $\alpha_L^\Gamma(X)$ consists of finding all possible paths of length L in X (Heijmans et al., 2005), understanding “path” as a set of successively connected voxels according to Γ . RORPO utilises this notion of path opening to detect tubular structures in an eight-bit image $I: \mathbb{R}^3 \rightarrow \mathbb{R}$. First, it applies path opening in multiple directions using seven connectivity operators $\{\Gamma_i\}_{i \in [1,7]}$. Since this operation applies to binary images, RORPO thresholds I on all possible

intensity levels (e.g. $\tau \in [0, 2^8 - 1]$ for an eight-bit image). This step results in seven response images $\{A_L^{\Gamma_i}(I)\}_{i \in [1,7]}$, where each $A_L^{\Gamma_i}(I) = \max_{\tau \in [0, 2^8 - 1]} \{\tau | (x, y, z) \in \alpha_L^{\Gamma_i}(I \geq \tau)\}$. Second, RORPO computes the difference between the maximum and median response for every voxel in $\{A_L^{\Gamma_i}(I)\}_{i \in [1,7]}$. The higher the difference, the more likely a voxel pertains to a thin object rather than a blob since the response of an isotropic structure will be similar in all directions. This difference image is the output of the filter (Figure 4).




Figure 4. Illustration of the functioning of RORPO on a 2D binary image. First, RORPO applies path opening with connectivity operators $\Gamma_{i \in [1,n]}$ in n directions, leading to n response images. In this example, we assume $n = 4$ and the length of the path $L = 3$. Second, RORPO computes and outputs the difference between the maximum and median response for each voxel. A voxel remaining in all response images receives a value of zero (both the maximum and median response equal to one) whereas those in only one receive a value of one (maximum and median response equal to one and zero, respectively).

The parameter L , which controls the length of the paths, requires tuning depending on the problem at hand. However, since it has a physical meaning related to the length of the object of interest, it is

easily adjustable. We applied the implementation of RORPO available in GitHub (<https://github.com/path-openings/RORPO>).

2.4 Performance metrics

We evaluated filtering results using three metrics based on the confusion matrix, precision, recall, F-score, and the area under the precision-recall curve (AUPRC). We opted for precision-recall curves instead of the popular receiver operating curves due to the dramatic disproportion between positive (i.e. PVS) and negative samples.

2.4.1 Precision

Precision allowed us to examine the proportion of correctly segmented PVS. Given a binary segmentation map S and a binary ground truth GT , we computed precision by dividing the number of voxels classified as PVS in both images by the number of voxels classified as PVS in S , i.e.

$$\text{Precision} = \frac{\sum S \cap GT}{\sum S}. \quad (4)$$

The values for precision lay between zero and one, being one the maximum and best precision score a method can obtain.

2.4.2 Recall

Recall enabled us to examine the proportion of PVS the filter actually identified. We computed recall by dividing the number of voxels classified as PVS in both images by the number of voxels classified as PVS in GT , i.e.

$$\text{Recall} = \frac{\sum S \cap GT}{\sum GT}. \quad (5)$$

The minimum and maximum recall scores are zero and one respectively.

2.4.3 Area under precision-recall curve (AUPRC)

The AUPRC summarised the precision and recall values obtained by the filtering method for different thresholds applied on the filter response maps. We computed the AUPRC by integrating over the precision-recall curve. Consequently, a good PVS segmentation method should obtain AUPRC values close to one.

3. Results

3.1 Baseline performance

We first segmented the PVS-like structures on clean high-resolution images to measure the performance of all filters in the absence of distortion. We noticed PVS dimensions influenced filtering effectiveness. First, RORPO outperformed both Hessian-based filters in most cases (AUPRC - RORPO:

98 [IQR 98, 99]; Frangi: 96 [IQR 95, 97]; Jerman: 96 [IQR 95, 97]), except when dealing with the smallest 1.0×0.5 or 1.5×0.5 mm PVS (AUPRC around 94.5 for RORPO vs ~ 100 for Frangi and Jerman). We noticed that RORPO produced low, partial, or null responses for PVS neighbouring the boundaries of the region of interest (Figure 5). Second, Hessian-based filters performed worse with increasing PVS volume; e.g. their AUPRC values were close to 100 for 3.5×2.0 mm PVS but 95 for 10.0×3.0 mm PVS. Inspection of filter response maps revealed that this outcome was a consequence of the poor localisation of Hessian-based filters – not only voxels comprising PVS but also those around them had a positive response (Figure 5).




Figure 5. Filter response maps on clean $0.5 \times 0.5 \times 0.5$ mm images. The green boundary in “Zoom” delimits the region of interest. Yellow arrows point to actual PVS-like structures (referred to as PVS for simplicity). In maps produced by the Frangi and Jerman filters, voxels around PVS have a non-zero response (white arrows). RORPO produces lower responses for PVS around boundaries of the region of interest (PVS appears red in top row vs blue in bottom row). Outputs cannot be compared between filters as values vary in range and scale.

3.2 Effect of imaging considerations

3.2.1 Slice thickness

We present DROs before and after k -space sampling to represent slice thickness of 0.5, 1, and 2 mm in Figure 6. Note the slice thickness of the real T2-w images used to derive the DRO is 2 mm. Thicker slices increased partial volume averaging effects, which in turn reduced PVS saliency and sharpness.

Moreover, slice thickness prevents visibility in some cases: PVS of length lower than 2 mm or width lower than 1 mm were not visible at $1 \times 1 \times 1$ and $1 \times 1 \times 2$ mm resolutions.




Figure 6. *k*-space sampling leads to partial volume averaging, which limits PVS visibility. This situation is particularly evident for 2-mm length and 1-mm width PVS (yellow arrow). Imaging considerations: *k*-space sampling.

We then segmented PVS on 1 × 1 × 1 and 1 × 1 × 2 mm images (Figure 7). Our findings were four-fold. First, both precision and recall decreased in thicker slices. For PVS of 4 mm length and 2 mm width, the maximum recall at the maximum precision in 1 × 1 × 1 mm resolution images was at least 30%

(Frangi: 30%; Jerman: 30%; RORPO: 58%), whereas it was between 7% and 26% in $1 \times 1 \times 2$ mm (Frangi: 7%; Jerman: 26%; RORPO: 14%) (Figure S1). Second, Gibbs ringing artefacts often exhibited a low but non-zero response (Figure 8a), i.e. they may be incorrectly picked up depending on the threshold. Third, filtering responses depend on how a PVS is oriented with respect to the sampling grid and acquisition plane. Long PVS that run diagonal to the acquisition plane (rather than parallel or perpendicular) are more likely to be broken into multiple segments after binarisation (Figure 8b). Fourth, due to partial volume effects, PVS in close proximity may be identified as a single one (Figure 8c).




Figure 7. Mean AUPRC (colour bar) in relation to PVS dimensions (width: x-axis; height: y-axis) and slice thickness (bottom row: 1 mm; top row: 2 mm). White regions represent cases that were not considered (width \geq length, eccentricity < 0.8 , or lack of PVS visibility). Our findings were four-fold. First, PVS with lengths and widths less than 2 and 1 mm, respectively, could not be properly quantified at either of these resolutions (bottom left corner of each panel). Second, while it was generally true that larger PVS were easier to identify than smaller ones (in each panel, AUPRC scores tended to rise from left to right and bottom to top), we also noted that Hessian-based filters performed worse with larger PVS (in each panel, top right corner). Third, PVS filtering performance decreased with slice thickness. This performance drop was caused by an increase in false positive rates and a reduction in true positive rates (Figure S1). RORPO was in general incapable of dealing with anisotropic voxels (AUPRC values for 1.0 or 1.5 mm wide PVS range between 20 and 60). Fourth, independent of length or filtering

technique, PVS wider than 2 mm can be detected with acceptable AUPRC values (90-100 %) in $1 \times 1 \times 1$ and $1 \times 1 \times 2$ mm scans. Imaging considerations: *k*-space sampling.

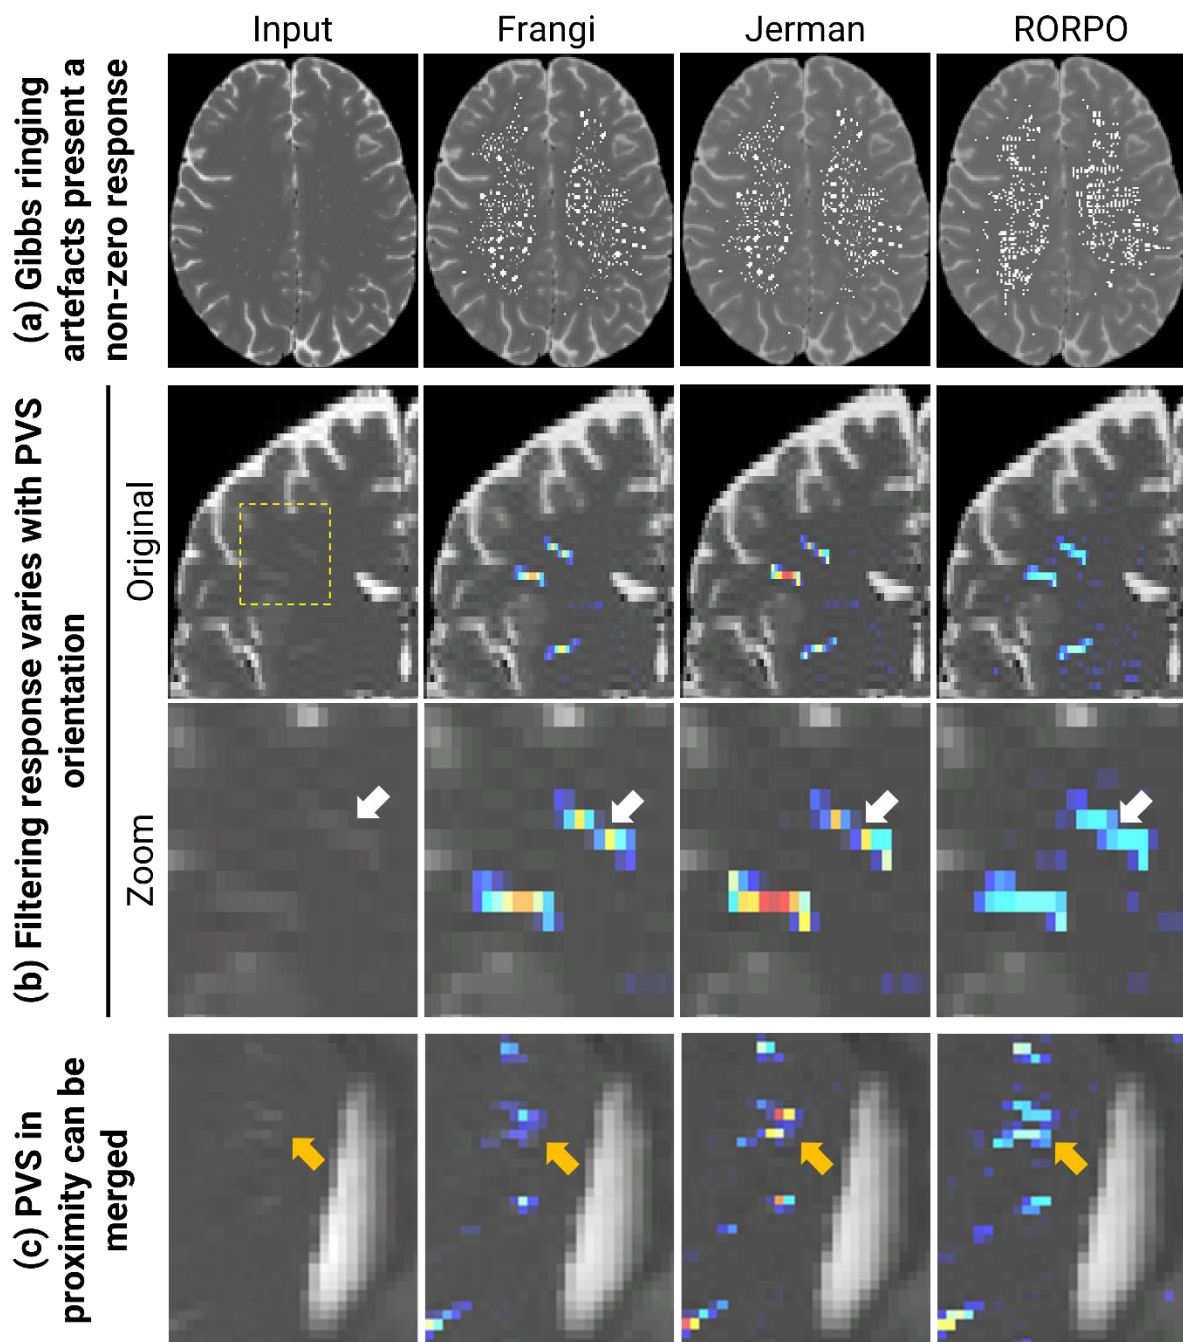


Figure 8. Gibbs ringing and partial volume effects caused by *k*-space sampling affect PVS quantification. (a) PVS filters pick up Gibbs ringing artefacts propagating from T2-w hyperintense regions. For visualisation purposes, we excluded filtering response around lateral ventricles, sulcal cerebrospinal fluid, and cortical grey matter. (b) Filtering response varies depending on the orientation of a PVS with respect to the sampling grid and acquisition plane. Long PVS that run diagonal — rather than parallel or perpendicular — to the acquisition plane are more likely to be split into multiple parts

after hard thresholding (potential breaking point pointed by the white arrow). (c) PVS in close proximity may be identified as a single one (yellow arrow) due to partial volume effects. Imaging considerations: *k*-space sampling.

3.2.2 Motion artefacts

We used the composite *k*-space model to simulate motion artefacts (Figure 9). Our computational model produced ringing and ghosting artefacts, in alignment with the theory. Note PVS remained visible after inducing motion artefacts.




Figure 9. Motion causes ghosting artefacts, which could easily be confused with “tubular” structures. Imaging considerations: *k*-space sampling and motion artefacts.

We then gauged the effect motion artefacts could have on PVS segmentation to determine whether quantification was still feasible on motion corrupted scans (Figure 10). Motion caused variation in quantification. When visual artefacts were subtle, the performance was similar to that obtained after sampling. However, when these artefacts were evident (as in Figure 11), false positives were unavoidable for Hessian-based filters regardless of the considered threshold (max precision: Frangi: 86%; Jerman: 88%; RORPO: 100%) (Figure S2). This situation took place as ghosting artefacts caused by motion were detected as potential tubular structures (Figure 11).




Figure 10. Mean AUPRC (colour bar) in relation to PVS dimensions (width: x-axis; height: y-axis), slice thickness (bottom row: 1 mm; top row: 2 mm), and motion artefacts. White regions represent cases that were not considered (width \geq length, eccentricity < 0.8 , or lack of PVS visibility). The performance of PVS filtering techniques dropped substantially with the presence of motion artefacts. For example, for a 2 mm wide PVS, AUPRC values ranged from 90 to 100% in k-space sampling simulations (Figure 7) whereas from 85 to 100% when motion is incorporated. False positives were unavoidable for Hessian-based filters regardless of the considered threshold (max precision: Frangi: 86%; Jerman: 88%; RORPO: 100%). Imaging considerations: k-space sampling and motion artefacts.

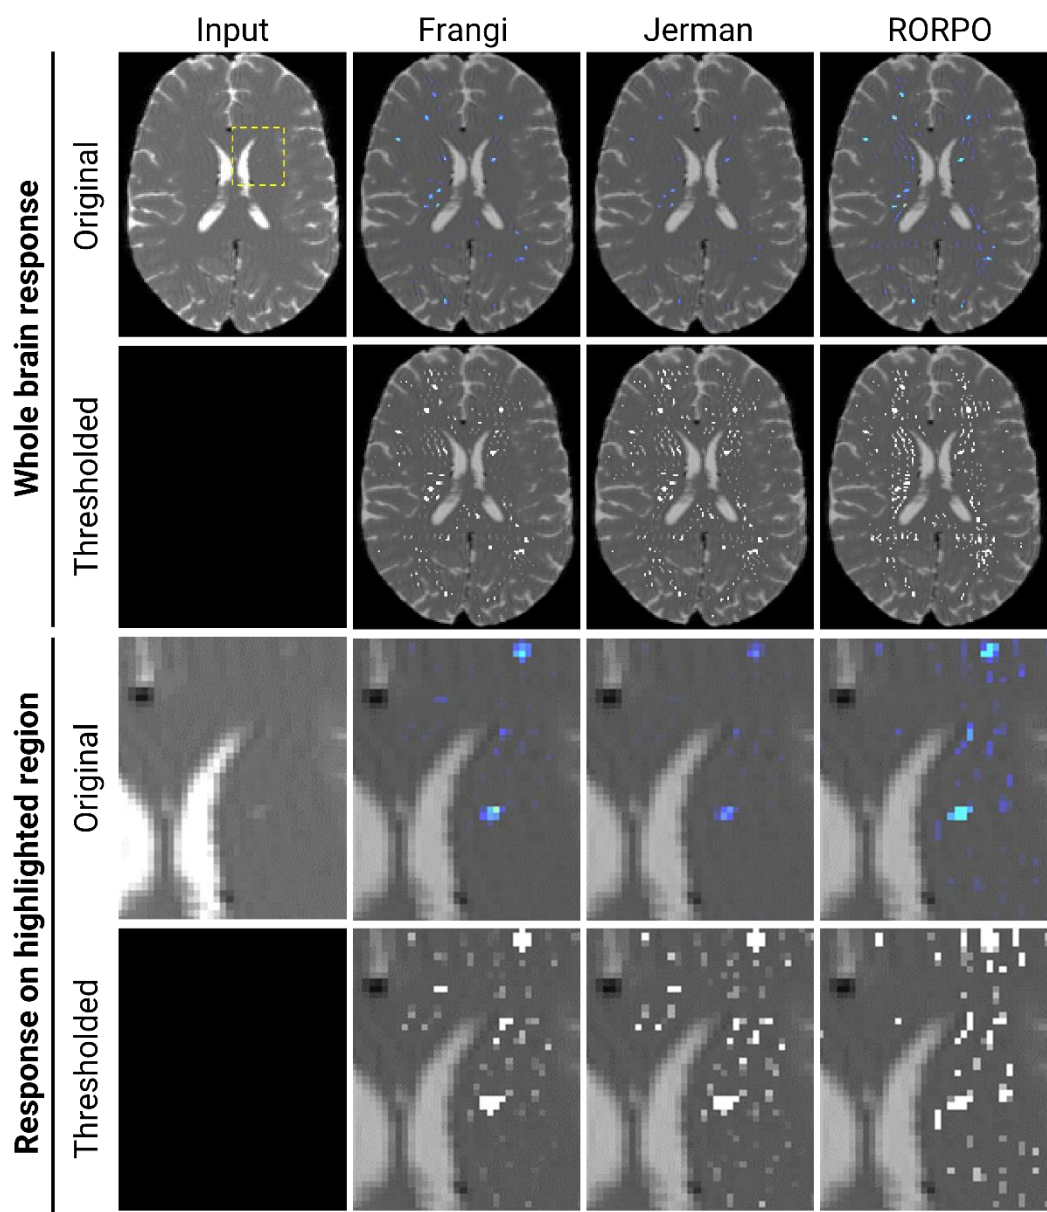


Figure 11. PVS filters detect ghosting artefacts caused by motion. We zoom in on the region indicated by the dashed yellow box to illustrate problems further. For visualisation purposes, we excluded filtering response around lateral ventricles, sulcal cerebrospinal fluid, and cortical grey matter. Imaging considerations: *k*-space sampling and motion artefacts.

3.2.3 Rician noise

We incorporated Rician noise into the simulations to examine whether the level of noise present in the reference study (SNR = 7.14 [IQR 5.93, 8.43]) had any impact on the visibility and quantification of PVS. While PVS remained visible (Figure 12), Rician noise increased false positive rates (Figure 13), being particularly evident in images with anisotropic voxels (Figure S3). Among tested filters, the

German filter exhibited the highest AUPRC drop (German: 3.11 [IQR 6.91, 2.52] vs Frangi: 1.55 [IQR 3.88, 1.18] and RORPO: 1.11 [IQR 2.06, 0.71]).

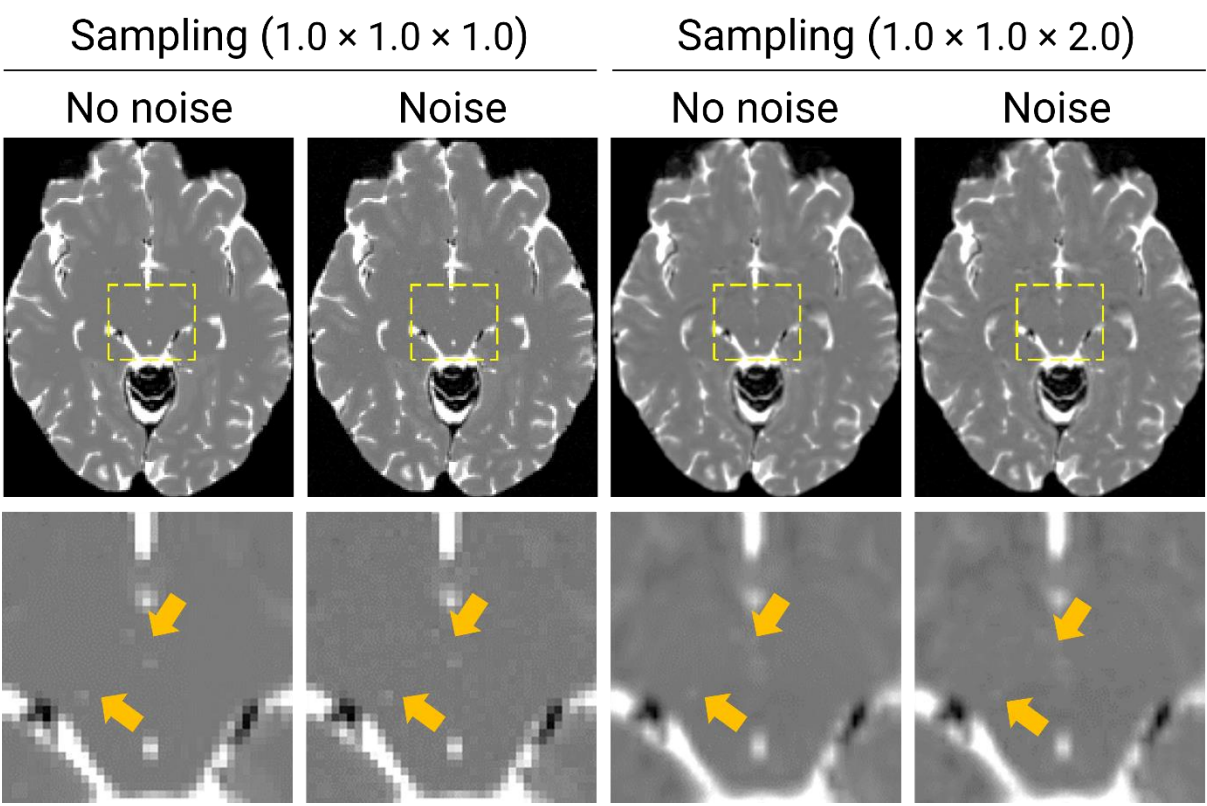


Figure 12. PVS remained visible after incorporating Rician noise into the simulations. Yellow arrows point to a few PVS in the midbrain. We zoom in on the region indicated by the dashed yellow box in the first row to illustrate problems further. Imaging considerations: k -space sampling and Rician noise.




Figure 13. Mean AUPRC (colour bar) in relation to PVS dimensions (width: x-axis; height: y-axis), slice thickness (bottom row: 1 mm; top row: 2 mm), and Rician noise ($SNR = 7.14$ [IQR 5.93, 8.43]). White regions represent cases that were not considered (width \geq length, eccentricity < 0.8 , or lack of PVS visibility). PVS filtering techniques were relatively robust to Rician noise. Among tested filters, the German filter exhibited the highest AUPRC drop (German: 3.11 [IQR 6.91, 2.52] vs Frangi: 1.55 [IQR 3.88, 1.18] and RORPO: 1.11 [IQR 2.06, 0.71]). Imaging considerations: *k*-space sampling and Rician noise.

3.3 Effect of pathological regions

We synthesised WMH to determine whether filters can discern between hyperintense tubular and non-tubular structures in a T2-w contrast (Figure 14). AUPRC values plummeted when we introduced WMH. For example, in $1 \times 1 \times 1$ mm images, median values went from more than 90 to less than 75 (Frangi: from 94.21 [IQR 91.60, 96.16] to 43.76 [IQR 25.19, 63.38]; German: from 94.51 [IQR 91.90, 95.37] to 58.00 [IQR 35.68, 64.87]; RORPO: from 98.72 [IQR 95.37, 98.96] to 71.87 [IQR 57.21, 76.63]). The zigzagging patterns seen in mean AUPRC plots (Figure 14) coincided with increased standard deviation in AUPRC values (Figure S5), suggesting it is only when PVS width surpasses 2 mm, when the response of a PVS and WMH starts to differ.




Figure 14. Mean AUPRC (colour bar) in relation to PVS dimensions (width: x-axis; height: y-axis), slice thickness (bottom row: 1 mm; top row: 2 mm), and white matter hyperintensities. White regions represent cases that were not considered (width \geq length, eccentricity < 0.8 , or lack of PVS visibility). Filters cannot differentiate PVS from other hyperintense structures effectively. We found that the introduction of WMH caused a sharp decline in AUPRC values, and that the response for PVS and WMH was similar until the PVS width crossed 2 mm (AUPRC – less than 2 mm: 0 – 60%; greater than 2 mm: 60 – 85%). Imaging considerations: k-space sampling.

According to response maps, performance decline is a result of two factors (Figure 15). First, segmentation filters detect WMH either completely or partially – depending on their size. Second, contrast change due to overlap between PVS and WMH compromises the detection of PVS. We did not find a single threshold capable of separating WMH from PVS effectively. We also observed that only when PVS width is higher than 2 mm, the response of a PVS and WMH start to differ.

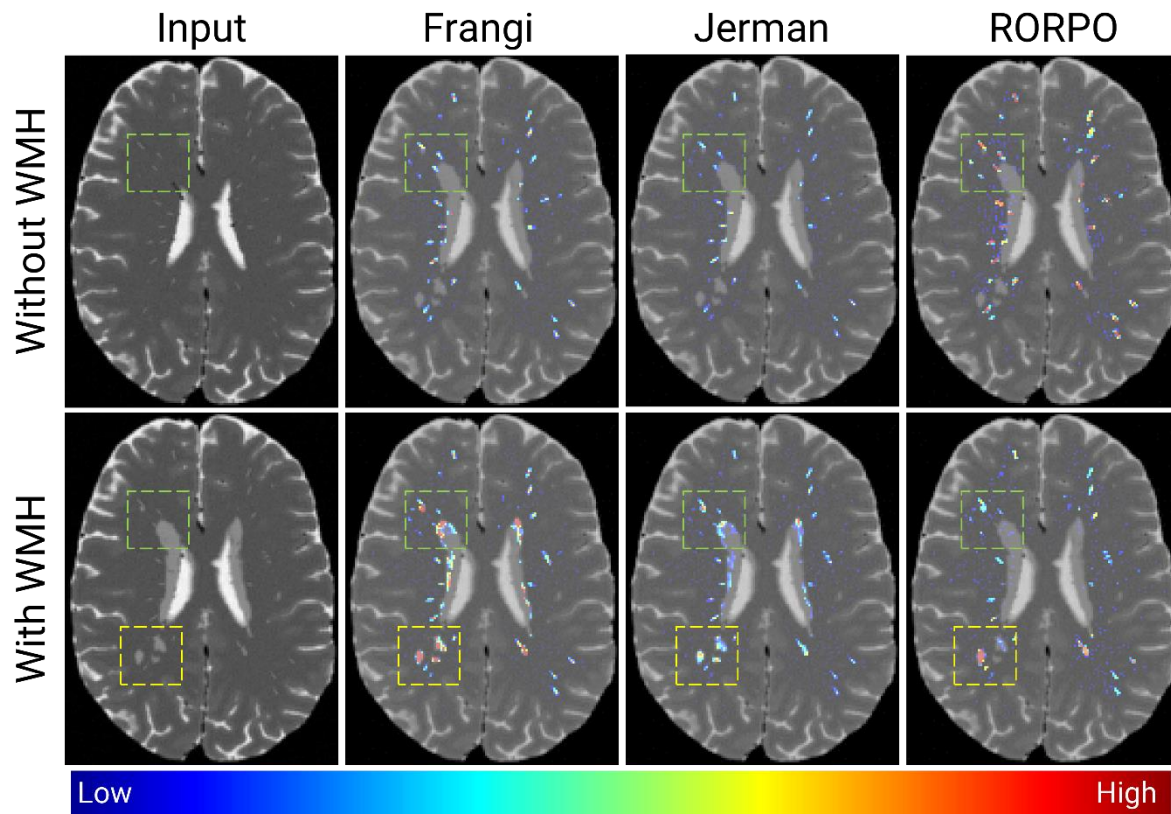


Figure 15. Response maps in the absence and presence of WMH (top and bottom row, respectively). Filters undesirably highlight WMH either partially or fully. Partial detections occur around boundaries of large WMH (green square) whereas full detection takes places when WMH appear as punctate foci (yellow square). Moreover, PVS are either completely missed or fragmented when PVS overlap with WMH (green square). Outputs cannot be compared between filters as values vary in range and scale.

4. Discussion

We presented a computational model for synthesising three-dimensional images with PVS-like structures to quantify the sensitivity and recall of PVS filtering methods – an essential step in PVS segmentation pipelines – particularly when images undergo distortion. With it, we detected key factors that influence PVS filtering – ergo PVS quantification -- and that need to be taken into consideration to better measure PVS characteristics.

4.1 What filtering method to use?

Briefly, the choice of one filter over another depends on whether researchers favour high true positive and false positive rates but low false negative rates (Frangi or Jerman) or the opposite (RORPO).

Two major drawbacks can also tip the balance one way or another. On the one hand, Hessian-based filters present poorer PVS localisation than RORPO, a consequence of the blurring that takes place before the computation of the Hessian matrix. This situation implies the maximum precision and recall

decrease as PVS size increases; PVS in close proximity will tend to be segmented as a single one; and segmentation errors around hyperintense regions other than PVS (e.g. lateral ventricles or WMH) are to be expected, unless these are perfectly excluded from the analyses. On the other hand, unlike Hessian-based filters, RORPO is unable to handle anisotropic voxels appropriately as it assumes distances between neighbouring voxels are constant in all directions.

Although a mixed approach – RORPO for isotropic voxels and Frangi for anisotropic voxels – would seem reasonable, we do not expect estimates yielded by these filters to be comparable: precision and recall vary depending on PVS dimensions and imaging conditions, and distributions of response map values vary in shape and scale.

4.2 Imaging artefacts lower precision rates

All visual distortions studied in this work reduced precision rates, irrespective of the filtering technique. Artefacts due to k -space sampling or motion were the ones augmenting false positive rates the most, as these present a “tubular” shape. The RORPO filter was better at discerning between Gibbs ringing and ghosting artefacts and true PVS than the Frangi or Jerman filters, as response map values are associated with intensities – higher thresholds would thus separate very hyperintense objects from tubular objects.

PVS of certain dimensions stop being visible after k -space sampling. In $1 \times 1 \times 1$ and $1 \times 1 \times 2$ mm resolution images, PVS of length lower than 2 mm and width lower than 1 mm were simply not observable nor quantifiable. If potential PVS candidates with such dimensions were to be detected, it is reasonable to think they correspond to anything but a PVS. Connected component analysis may be useful to remove such noise.

4.3 The presence of other hyperintense lesions compromises PVS segmentation

The presence of WMH was particularly damaging for PVS quantification among all effects investigated in this study, as filters flagged them as potential PVS candidates. WMH thus boost false positive rates because their filter response values can be similar or greater than those of a PVS. The degree to which WMH are detected is determined on the size of the lesion and the filter. Tiny and focal WMH will be segmented completely, regardless of filtering method. Hessian-based filters will detect the boundaries of large WMH whereas RORPO will disregard them and all PVS within them.

While our observations reflect the damaging effect of WMH on PVS segmentation, these problems can also occur with other hyperintense neuro-radiological features – e.g. ischaemic stroke lesions or lacunes, with wedged shapes or shapes similar to small WMH. Therefore, by synthesising WMH, and

evaluating the filters' ability to differentiate them from PVS, the confound that these other pathological features may pose for the accurate segmentation of PVS could be also estimated.

4.4 Implications, recommendations, and open challenges

Our work has implications for the quantification of PVS on MRI.

4.4.1 Imaging artefacts

Imaging artefacts, caused particularly by k-space sampling and motion artefacts, have a negative impact on the sensitivity and precision of current segmentation methods and, consequently, on any subsequent analysis. The development of imaging protocols reducing their presence and of techniques to recognise them and compensate for them is thus critical for better PVS quantification.

MRI quality control relies mostly on visual inspection: a group of experts assesses the quality of all input scans, categorises them into high vs low quality, and excludes the latter set (Ballerini et al., 2020a, 2020b; Bernal et al., 2021b, 2020; Liu et al., 2020; Sepehrband et al., 2021). This approximation is laborious, impractical in large-scale studies, subjective, and reduces the study's sample size (Alfaro-Almagro et al., 2018; Esteban et al., 2017; Lutti et al., 2021).

The ideal situation is preventing motion from ever occurring. Acclimating users to the acquisition conditions or restricting their motion with head padding is common in clinical settings (Zaitsev et al., 2015). However, the implementation of such strategies does not imply acquisitions will be devoid of artefacts. For example, movement disorders, causing occasional yet involuntary motion, challenge quantification in children and neuropsychiatric patients (Greene et al., 2018). Evidence from ongoing and large-scale studies of the ageing population suggests motion artefacts impeding automatic PVS assessments can affect up to one out of every five MRI images (Ballerini et al., 2020b; Bernal et al., 2021b, 2020).

Retrospective motion correction has been explored to compensate for the detrimental effect of motion artefacts on PVS assessments (Bernal et al., 2021b, 2020; Zong et al., 2021). These studies have demonstrated that retrospective motion correction successfully reduces dependencies between image quality and PVS measurements and can improve the agreement between computational and manual estimation (e.g. visual ratings and manual counts). Nevertheless, motion correction can come at the expense of real PVS being removed (Bernal et al., 2021b) and, to the eyes of the experts, does not necessarily lead to better PVS visualisation (Zong et al., 2021). We therefore recommend careful utilisation of retrospective compensation strategies.

4.4.2 Hyperintense vascular lesions

The striking effect of hyperintense vascular lesions on PVS segmentation suggests that extremely careful – possibly manual – delineation of these other lesions is needed to obtain actual estimates of PVS burden, e.g. mask the WMH first, then assess the PVS, but this will require visual inspection to differentiate PVS from WMH where the two coincide. It should be noted that excluding WMH might have significant statistical implications that require investigation: higher loads of WMH may correlate with lower PVS burden, an implicit correlation caused by pure methodological limitations. Without efforts for addressing these issues, the final goal of understanding the involvement of PVS on brain health function and small vessel diseases on their own is simply not achievable as estimates will incorporate mixed effects.

4.4.3 Regional summaries

We are often interested in reporting regional PVS quantities. Two items need to be kept in mind. First, PVS location and orientation in relation to the voxel grid impacts its visual appearance and, thus, our capacity to correctly segment it: if PVS run orthogonal or parallel to the acquisition planes, they will be better segmented than those diagonal to them (Figure 8). If the subject being imaged lays down in the standard anatomical position and the acquisition plane is axial, we can thus expect PVS coming in from the temporal cortex (parallel to the axial slice) or the posterior frontal or anterior parietal cortex near the vertex on slices superior to the lateral ventricles (perpendicular to the axial slice) to be clearly visible, whereas those on the arc of the brain in between these points to be more difficult to see and segment as they run diagonal to the slice. It should be noted that, while the use of 3D isotropic voxels helps reduce this issue, it does not eliminate it since partial volume averaging is inherent to MRI. Second, obtaining such regional summaries requires parcellating the brain, merging a few small regions to construct those of interest (e.g. centrum semiovale), and using the resulting masks to determine what PVS fall within or outside of the region of interest. Differences in the definition and delimitation of regions of interest across studies, as well as the performance of the selected parcellation algorithms, can influence regional evaluations and hinder study comparability. We thus advocate inspecting parcellations, especially in patients with high burdens of SVD features, given that popular segmentation algorithms are known to struggle with their presence (Dadar et al., 2020). Because the brain is split into non-overlapping regions, PVS fragmentation owing to masking is also anticipated. Analysis of connected components may enable counting each individual PVS only once.

4.4.4 Choosing optimal binarisation thresholds

Correct binarisation remains a challenge as precision and recall rates, therefore “optimal” thresholds, vary with the quality of each individual scan. “Vesselness” response map binarisation influences PVS

measurements. For example, a supra-optimal threshold increases the risk of PVS fragmentation, decreases the chance of merging neighbouring PVS, artificially lowers PVS dimensions (Figure 5 and Figure 8; different thresholds would alter PVS length, width, height, and volume) and alters PVS orientation (Figure 8; different thresholds will cause orientations to artificially vary); the opposite is true with sub-optimal thresholds. Variability in estimates from different research studies, which would naturally limit interpretation and impede direct comparison between works, is therefore expected.

On a more positive note, regardless of how difficult it is to extract measures of PVS, as long as the segments being segmented are part of PVS, it is reasonable to presume these measurements are representative of the condition and should not be completely dismissed. Due to methodological constraints, cautious interpretation is nonetheless advised.

4.5 Can machine learning help to improve PVS quantification?

Machine learning, convolutional neural networks specifically, has been utilised in recent years as an alternative to classical techniques leveraging explicit geometrical information (Boutinaud et al., 2021; Dubost et al., 2020, 2019b, 2019a; Huang et al., 2021; Jung et al., 2019; Lian et al., 2018; Yang et al., 2021; Zong et al., 2021). Its use could certainly be advantageous as these techniques can make use of a plethora of visual cues, including intensity, shape, and size, that would enable them to separate PVS from other lesions (Valdes Hernandez et al., 2013). Nonetheless, while these strategies are certainly powerful, two key issues currently limit their applicability. First, they require vast and heterogeneous training datasets to avoid overfitting (Kushibar et al., 2019). Since they are primarily trained on small and well curated datasets, they may not be able to process medical images in the wild, particularly if scans have been acquired with imaging protocols other than those considered during training. Second, unless unsupervised, training such powerful models requires annotations, which are generally unavailable, especially when it comes to subtle, small, and sometimes abundantly present neuroimaging features.

The computational model developed in this paper has the potential to overcome the restrictions outlined above. From a deep learning point of view, our developments fall into the category of “physics-based data augmentation strategies” because, with them, we are able to generate MRI-like images using biophysical models. Because imaging considerations and distortion levels in our models are easily tuneable, we can train deep learning models with sufficient, heterogeneous, and plausible sets of images, allowing them to be robust against imaging artefacts and generalise to potentially unseen datasets without requiring additional adjustments. Fairly recent work in brain image analysis indicates this topic, referred to in the literature as contrast-agnostic and distortion-robust training schemes (Billot et al., 2021; Shaw et al., 2020) is promising and definitely worth pursuing.

The use of deep learning techniques does not necessarily imply classical methodologies need to be put aside. On the one hand, traditional model-based signal processing approaches offer interpretability and thorough and extensive validation, both of which are critical in healthcare. On the other hand, recent methods leveraging artificial intelligence offer the possibility of training models to be robust against imaging distortions (Shaw et al., 2020) and acquisition settings (Billot et al., 2021). Incorporating prior knowledge from physics, imaging, and signal processing into machine learning algorithms, also referred to as learning known operators, has been proven feasible and can lead to improved quantification, decrease the amount of samples on which a network must be trained as well as the number of parameters that must be manually adjusted (Maier et al., 2019).

4.6 Related projects assessing PVS mapping methods

Members of the Medical Image Computing and Computer-Assisted Interventions (MICCAI) association actively stimulate research in medical image analysis through grand challenges. For example, the VALDO challenge² was established in 2021 with the aim of encouraging the development, assessment, and comparison of vascular lesion segmentation algorithms, including PVS. Because proposals are assessed under the same guidelines, with the same training and test datasets and evaluation metrics, direct comparison across works is feasible, allowing for an understanding of the methodological benefits and disadvantages of the various solutions. Funding bodies are also supporting large consortia dedicated to the study of perivascular spaces, for example, the Leducq Foundation, which sponsors the Transatlantic Network of Excellence for the Study of Perivascular Spaces in Small Vessel Disease (<https://www.small-vessel-disease.org/>), and the Weston Family Foundation through the Weston Brain Institute (<https://westonfoundation.ca/weston-brain-institute/>) among others.

4.7 Limitations and future work

Our computational model is limited by its assumptions and simplifications.

4.7.1 Imaging protocols

Our computational model does not mimic all acquisition processes. However, from a methodological standpoint, our research demonstrates spatiotemporal and imaging considerations influence PVS quantification and lays groundwork for future research. Our computational model might be used in conjunction with MRI simulators, such as MRilab (Liu et al., 2017), to enable testing and optimising imaging protocols for computational PVS assessments. Also, we limited out evaluations to some typical scenarios of imaging acquired under research protocols (e.g. slice thickness only of 0.5, 1, 2,

² <https://valdo.grand-challenge.org/>

and 5 mm). More testing is required to establish the limits of validity of these methods (e.g., minimum length and width of PVS and spatial resolution requirements).

4.7.2 PVS model

For the sake of simplicity, we made certain assumptions about PVS. We assumed PVS could be represented as regular cylinders. Euclidean shapes nonetheless are unlikely representative of the actual curvilinear geometry of PVS observed in real life (see Fig. 2 in (Wardlaw et al., 2013)). We also considered that PVS were oriented toward the brain's geometrical centre. While this assumption may hold in a few regions of the brain, a more realistic version in which PVS appear perpendicular to the surface of the brain and follow the course of perforating brain vessels (e.g. guided by the work of (Huck et al., 2019)) could be implemented in the future.

4.7.3 Focus

As highlighted in the introduction, we focused on the robustness of PVS filtering methods rather than on the pre-processing and segmentation stages. Nonetheless, our computational model can be used to synthesise images “acquired” with multiple imaging protocols, head models, lesions and determine whether common processing techniques are suitable for these types of analyses.

4.7.4 Bias in scanning-resolution ground truth

By construction, our scanning-resolution ground truth maps, and hence the accompanying performance scores, are impacted by partial volume averaging, a systematic source of error that affects all images of the same resolution similarly. In light of this limitation, we resorted to performance metrics that evaluate all possible segmentation thresholds (AUPRC) – as long as voxels containing PVS partially or fully receive higher “vesselness” scores than those without, there should be a threshold which maximises the coincidence between ground truth and segmentation masks. We also provided additional qualitative and quantitative evidence of filtering performance on each trial (including on clean high-resolution scans) to back up our claims and findings.

5. Conclusions

Our work reveals appropriate processing of MRI signals is necessary to maximise PVS measurement reliability. While filters to enhance and facilitate PVS detection can correctly achieve these goals, they are sensitive to imaging artefacts, such as ringing and motion, and ineffective at distinguishing between PVS from other lesions with similar contrast, shapes and overlapping size. These issues ultimately highlight the importance of masking out other neuroradiological features and of prospective or retrospective image enhancement for better PVS quantification. We have made our

development publicly available for other researchers to build upon it and test their segmentation schemes in advance (<https://github.com/joseabernal/PVSDRO.git>).

6. Data and code availability statement

We implemented our computational model in MATLAB R2018 Update 4 (9.5.0.1067069). We have made our source code publicly available in GitHub: <https://github.com/joseabernal/PVSDRO.git>. The original MIDA model used in this work can be downloaded from the following URL: www.itis.ethz.ch/MIDA.

7. CRediT authorship contribution statement

Jose Bernal: conceptualisation, methodology, software, formal analysis, investigation, writing – original draft preparation, writing – review and editing, visualisation. **Maria d. C. Valdés-Hernández:** patient study data collection, patient study data processing, methodology, resources, patient study data curation, writing – review and editing, supervision, funding acquisition, visualisation. **Javier Escudero:** resources, writing – review and editing, supervision, funding acquisition, visualisation. **Roberto Duarte:** methodology, software, writing – review and editing. **Lucia Ballerini:** patient study data collection, patient study data processing, methodology, software, writing – review and editing. **Mark E. Bastin:** patient study data collection, patient study data processing, study data curation, writing – review and editing. **Ian J. Deary:** funding acquisition, chief investigator of the patient study. **Michael J. Thrippleton:** software, methodology, writing – review and editing. **Rhian M. Touyz:** resources, writing – review and editing, supervision, funding acquisition. **Joanna M. Wardlaw:** resources, study data acquisition and curation, writing – review and editing, supervision, funding acquisition, project administration.

8. Conflict of interest

The authors declare no conflict of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript, or in the decision to publish the results.

9. Acknowledgements

This work is supported by: MRC Doctoral Training Programme in Precision Medicine (JB - Award Reference No. 2096671); The Galen and Hilary Weston Foundation under the Novel Biomarkers 2019 scheme (ref UB190097) administered by the Weston Brain Institute; the UK Dementia Research Institute which receives its funding from DRI Ltd, funded by the UK MRC, Alzheimer's Society and

Alzheimer's Research UK; the Fondation Leducq Network for the Study of Perivascular Spaces in Small Vessel Disease (16 CVD 05); Stroke Association 'Small Vessel Disease-Spotlight on Symptoms (SVD-SOS)' (SAPG 19\100068); The Row Fogo Charitable Trust Centre for Research into Aging and the Brain (MVH) (BRO-D.FID3668413); British Heart Foundation Edinburgh Centre for Research Excellence (RE/18/5/34216); NHS Lothian Research and Development Office (MJT).

The LBC1936 study was funded by Age UK and the UK Medical Research Council (<http://www.disconnectedmind.ed.ac.uk/>) (including the Sidney De Haan Award for Vascular Dementia). LBC1936 MRI brain imaging was supported by Medical Research Council (MRC) grants G0701120, G1001245, MR/M013111/1 and MR/R024065/1. Funds were also received from The University of Edinburgh Centre for Cognitive Ageing and Cognitive Epidemiology, part of the cross council Lifelong Health and Wellbeing Initiative (MR/K026992/1), and the Biotechnology and Biological Sciences Research Council (BBSRC).

We thank the LBC1936 study participants, their families and radiographers at Edinburgh Imaging Facilities. We thank the LBC1936 study team members who recruited the participants of the study.

10. References

- Alfaro-Almagro, F., Jenkinson, M., Bangerter, N.K., Andersson, J.L.R., Griffanti, L., Douaud, G., Sotiropoulos, S.N., Jbabdi, S., Hernandez-Fernandez, M., Vallee, E., Vidaurre, D., Webster, M., McCarthy, P., Rorden, C., Daducci, A., Alexander, D.C., Zhang, H., Dragonu, I., Matthews, P.M., Miller, K.L., Smith, S.M., 2018. Image processing and Quality Control for the first 10,000 brain imaging datasets from UK Biobank. *Neuroimage* 166, 400–424. <https://doi.org/10.1016/j.neuroimage.2017.10.034>
- Ballerini, L., Booth, T., Valdés Hernández, M. del C., Wiseman, S., Lovreglio, R., Muñoz Maniega, S., Morris, Z., Pattie, A., Corley, J., Gow, A., Bastin, M.E., Deary, I.J., Wardlaw, J., 2020a. Computational quantification of brain perivascular space morphologies: Associations with vascular risk factors and white matter hyperintensities. A study in the Lothian Birth Cohort 1936. *NeuroImage Clin.* 25, 102120. <https://doi.org/10.1016/j.nicl.2019.102120>
- Ballerini, L., Lovreglio, R., Valdés Hernández, M.D.C., Ramirez, J., MacIntosh, B.J., Black, S.E., Wardlaw, J.M., 2018. Perivascular Spaces Segmentation in Brain MRI Using Optimal 3D Filtering. *Sci. Rep.* 8, 1–11. <https://doi.org/10.1038/s41598-018-19781-5>
- Ballerini, L., McGrory, S., Valdés Hernández, M. del C., Lovreglio, R., Pellegrini, E., MacGillivray, T., Muñoz Maniega, S., Henderson, R., Taylor, A., Bastin, M.E., Doubal, F., Trucco, E., Deary, I.J., Wardlaw, J., 2020b. Quantitative measurements of enlarged perivascular spaces in the brain are associated with retinal microvascular parameters in older community-dwelling subjects. *Cereb. Circ. - Cogn. Behav.* 1, 100002. <https://doi.org/10.1016/j.cccb.2020.100002>
- Bernal, J., Valdés-Hernández, M., Ballerini, L., Escudero, J., Jochems, A.C.C., Clancy, U., Doubal, F.N., Stringer, M.S., Thrippleton, M.J., Touyz, R.M., Wardlaw, J.M., 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. *Commun. Comput. Inf. Sci.* 1248 CCIS, 171–183.

https://doi.org/10.1007/978-3-030-52791-4_14

- Bernal, J., Valdés-Hernández, M. d. C., Escudero, J., Heye, A.K., Sakka, E., Armitage, P.A., Makin, S., Touyz, R.M., Wardlaw, J.M., Thrippleton, M.J., 2021a. A four-dimensional computational model of dynamic contrast-enhanced magnetic resonance imaging measurement of subtle blood-brain barrier leakage. *Neuroimage* 230, 117786. <https://doi.org/10.1016/j.neuroimage.2021.117786>
- Bernal, J., Xu, W., Valdés-Hernández, M. d. C., Escudero, J., Jochems, A.C.C., Clancy, U., Doubal, F.N., Stringer, M.S., Thrippleton, M.J., Touyz, R.M., Wardlaw, J.M., 2021b. Selective Motion Artefact Reduction via Radiomics and k-space Reconstruction for Improving Perivascular Space Quantification in Brain Magnetic Resonance Imaging, *Lecture Notes in Computer Science* (including subseries *Lecture Notes in Artificial Intelligence* and *Lecture Notes in Bioinformatics*). Springer International Publishing. https://doi.org/10.1007/978-3-030-80432-9_12
- Billot, B., Greve, D.N., Puonti, O., Thielscher, A., Van Leemput, K., Fischl, B., Dalca, A. V., Iglesias, J.E., 2021. SynthSeg: Domain Randomisation for Segmentation of Brain MRI Scans of any Contrast and Resolution.
- Boutinaud, P., Tsuchida, A., Laurent, A., Adonias, F., Hanifelhoul, Z., Nozais, V., Verrecchia, V., Lampe, L., Zhang, J., Zhu, Y.C., Tzourio, C., Mazoyer, B., Joliot, M., 2021. 3D Segmentation of Perivascular Spaces on T1-Weighted 3 Tesla MR Images With a Convolutional Autoencoder and a U-Shaped Neural Network. *Front. Neuroinform.* 15, 1–21. <https://doi.org/10.3389/fninf.2021.641600>
- Dadar, M., Potvin, O., Camicioli, R., Duchesne, S., 2020. Beware of white matter hyperintensities causing systematic errors in grey matter segmentations! *bioRxiv* 1–12. <https://doi.org/10.1101/2020.07.07.191809>
- Deary, I.J., Gow, A.J., Taylor, M.D., Corley, J., Brett, C., Wilson, V., Campbell, H., Whalley, L.J., Visscher, P.M., Porteous, D.J., Starr, J.M., 2007. The Lothian Birth Cohort 1936: A study to examine influences on cognitive ageing from age 11 to age 70 and beyond. *BMC Geriatr.* 7, 1–12. <https://doi.org/10.1186/1471-2318-7-28>
- Dubost, F., Adams, H., Bortsova, G., Ikram, M.A., Niessen, W., Vernooij, M., de Bruijne, M., 2019a. 3D regression neural network for the quantification of enlarged perivascular spaces in brain MRI. *Med. Image Anal.* 51, 89–100. <https://doi.org/10.1016/j.media.2018.10.008>
- Dubost, F., Adams, H., Yilmaz, P., Bortsova, G., Tulder, G. van, Ikram, M.A., Niessen, W., Vernooij, M.W., Bruijne, M. de, 2020. Weakly supervised object detection with 2D and 3D regression neural networks. *Med. Image Anal.* 65, 101767. <https://doi.org/10.1016/j.media.2020.101767>
- Dubost, F., Yilmaz, P., Adams, H., Bortsova, G., Ikram, M.A., Niessen, W., Vernooij, M., de Bruijne, M., 2019b. Enlarged perivascular spaces in brain MRI: Automated quantification in four regions. *Neuroimage* 185, 534–544. <https://doi.org/10.1016/j.neuroimage.2018.10.026>
- Esteban, O., Birman, D., Schaer, M., Koyejo, O.O., Poldrack, R.A., Gorgolewski, K.J., 2017. MRIQC: Advancing the automatic prediction of image quality in MRI from unseen sites. *PLoS One* 12, 1–21. <https://doi.org/10.1371/journal.pone.0184661>
- Fischl, B., 2012. FreeSurfer. *Neuroimage* 62, 774–781. <https://doi.org/10.1016/j.neuroimage.2012.01.021>
- Francis, F., Ballerini, L., Wardlaw, J.M., 2019. Perivascular spaces and their associations with risk factors, clinical disorders and neuroimaging features: A systematic review and meta-analysis. *Int. J. Stroke* 14, 359–371. <https://doi.org/10.1177/1747493019830321>
- Frangi, A.F., Niessen, W.J., Vincken, K.L., Viergever, M.A., 1998. Multiscale vessel enhancement filtering 130–137. <https://doi.org/10.1007/BFb0056195>

- González-Castro, V., Valdés Hernández, M. del C., Chappell, F.M., Armitage, P.A., Makin, S., Wardlaw, J.M., 2017. Reliability of an automatic classifier for brain enlarged perivascular spaces burden and comparison with human performance. *Clin. Sci.* 131, 1465–1481. <https://doi.org/10.1042/CS20170051>
- Greene, D.J., Koller, J.M., Hampton, J.M., Wesevich, V., Van, A.N., Nguyen, A.L., Hoyt, C.R., McIntyre, L., Earl, E.A., Klein, R.L., Shimony, J.S., Petersen, S.E., Schlaggar, B.L., Fair, D.A., Dosenbach, N.U.F., 2018. Behavioral interventions for reducing head motion during MRI scans in children. *Neuroimage* 171, 234–245. <https://doi.org/10.1016/j.neuroimage.2018.01.023>
- Heier, L.A., Bauer, C.J., Schwartz, L., Zimmerman, R.D., Morgello, S., Deck, M.D.F., 1987. Large Virchow-Robin Spaces : MR-Clinical Correlation 20, 929–936.
- Heijmans, H., Buckley, M., Talbot, H., 2005. Path openings and closings. *J. Math. Imaging Vis.* 22, 107–119. <https://doi.org/10.1007/s10851-005-4885-3>
- Huang, P., Zhu, Z., Zhang, R., Wu, X., Jiaerken, Y., Wang, S., Yu, W., Hong, H., Lian, C., Li, K., Zeng, Q., Luo, X., Xu, X., Yu, X., Yang, Y., Zhang, M., 2021. Factors Associated With the Dilation of Perivascular Space in Healthy Elderly Subjects. *Front. Aging Neurosci.* 13, 1–9. <https://doi.org/10.3389/fnagi.2021.624732>
- Huck, J., Wanner, Y., Fan, A.P., Jäger, A.T., Grahl, S., Schneider, U., Villringer, A., Steele, C.J., Tardif, C.L., Bazin, P.L., Gauthier, C.J., 2019. High resolution atlas of the venous brain vasculature from 7 T quantitative susceptibility maps. *Brain Struct. Funct.* 224, 2467–2485. <https://doi.org/10.1007/s00429-019-01919-4>
- Iacono, M.I., Neufeld, E., Akinagbe, E., Bower, K., Wolf, J., Oikonomidis, I.V., Sharma, D., Lloyd, B., Wilm, B.J., Wyss, M., Pruessmann, K.P., Jakab, A., Makris, N., Cohen, E.D., Kuster, N., Kainz, W., Angelone, L.M., 2015. MIDA: A multimodal imaging-based detailed anatomical model of the human head and neck. *PLoS One* 10. <https://doi.org/10.1371/journal.pone.0124126>
- Jenkinson, M., Bannister, P., Brady, M., Smith, S., 2002. Improved optimization for the robust and accurate linear registration and motion correction of brain images. *Neuroimage* 17, 825–841.
- Jenkinson, M., Smith, S., 2001. A global optimisation method for robust affine registration of brain images. *Med. Image Anal.* 5, 143–156.
- Jerman, T., Pernuš, F., Likar, B., Špiclin, Ž., 2015. Beyond Frangi: an improved multiscale vesselness filter. *Med. Imaging 2015 Image Process.* 9413, 94132A. <https://doi.org/10.1117/12.2081147>
- Jung, E., Chikontwe, P., Zong, X., Lin, W., Shen, D., Park, S.H., 2019. Enhancement of Perivascular Spaces Using Densely Connected Deep Convolutional Neural Network. *IEEE Access* 7, 18382–18391. <https://doi.org/10.1109/ACCESS.2019.2896911>
- Kushibar, K., Valverde, S., González-Villà, S., Bernal, J., Cabezas, M., Oliver, A., Lladó, X., 2019. Supervised Domain Adaptation for Automatic Sub-cortical Brain Structure Segmentation with Minimal User Interaction. *Sci. Rep.* 9, 1–15. <https://doi.org/10.1038/s41598-019-43299-z>
- Lamy, J., Merveille, O., Kerautret, B., Passat, N., Lamy, J., Merveille, O., Kerautret, B., Passat, N., Vesselness, A.V., Lamy, J., 2020. Vesselness filters : A survey with benchmarks applied to liver imaging, in: *International Conference on Pattern Recognition*.
- Li, Y., Li, M., Zuo, L., Shi, Q., Qin, W., Yang, L., Jiang, T., Hu, W., 2018. Compromised blood-brain barrier integrity is associated with total magnetic resonance imaging burden of cerebral small vessel disease. *Front. Neurol.* 9. <https://doi.org/10.3389/fneur.2018.00221>
- Lian, C., Zhang, J., Liu, M., Zong, X., Hung, S.-C., Lin, W., Shen, D., 2018. Multi-channel multi-scale fully

- convolutional network for 3D perivascular spaces segmentation in 7T MR images Chunfeng. *Med. Image Anal.* 46, 106–117. <https://doi.org/10.1016/j.media.2018.02.009>. Lian
- Liu, C., Habib, T., Salimeen, M., Pradhan, A., Singh, M., Wang, M., Wu, F., Zhang, Y., Gao, L., Yang, G., Li, X., Yang, J., 2020. Quantification of visible Virchow–Robin spaces for detecting the functional status of the glymphatic system in children with newly diagnosed idiopathic generalized epilepsy. *Seizure* 78, 12–17. <https://doi.org/10.1016/j.seizure.2020.02.015>
- Liu, F., Velikina, J. V., Block, W.F., Kijowski, R., Samsonov, A.A., 2017. Fast Realistic MRI Simulations Based on Generalized Multi-Pool Exchange Tissue Model. *IEEE Trans. Med. Imaging* 36, 527–537. <https://doi.org/10.1109/TMI.2016.2620961>
- Lutti, A., Corbin, N., Ashburner, J., Ziegler, G., Phillips, C., Kherif, F., Callaghan, M.F., Di, G., 2021. Restoring statistical validity in group analyses of motion- corrupted MRI data. *bioRxiv* 1–17.
- Maier, A.K., Syben, C., Stimpel, B., Würfl, T., Hoffmann, M., Schebesch, F., Fu, W., Mill, L., Kling, L., Christiansen, S., 2019. Learning with known operators reduces maximum error bounds. *Nat. Mach. Intell.* 1, 373–380. <https://doi.org/10.1038/s42256-019-0077-5>
- Merveille, O., Talbot, H., Najman, L., Passat, N., 2018. Curvilinear Structure Analysis by Ranking the Orientation Responses of Path Operators. *IEEE Trans. Pattern Anal. Mach. Intell.* 40, 304–317. <https://doi.org/10.1109/TPAMI.2017.2672972>
- Merveille, O., Talbot, H., Najman, L., Passat, N., 2014. Tubular structure filtering by ranking orientation responses of path operators. *Lect. Notes Comput. Sci. (including Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinformatics)* 8690 LNCS, 203–218. https://doi.org/10.1007/978-3-319-10605-2_14
- Patankar, T.F., Mitra, D., Varma, A., Snowden, J., Neary, D., Jackson, A., 2005. Dilatation of the Virchow-Robin space is a sensitive indicator of cerebral microvascular disease: Study in elderly patients with dementia. *Am. J. Neuroradiol.* 26, 1512–1520.
- Pierpaoli, C., Basser, P.J., 1996. Toward a quantitative assessment of diffusion anisotropy. *Magn. Reson. Med.* 36, 893–906. <https://doi.org/10.1002/mrm.1910360612>
- Potter, G.M., Chappell, F.M., Morris, Z., Wardlaw, J.M., 2015. Cerebral perivascular spaces visible on magnetic resonance imaging: Development of a qualitative rating scale and its observer reliability. *Cerebrovasc. Dis.* 39, 224–231. <https://doi.org/10.1159/000375153>
- Ramirez, J., Berezuk, C., McNeely, A.A., Scott, C.J.M., Gao, F., Black, S.E., 2015. Visible Virchow-Robin spaces on magnetic resonance imaging of Alzheimer’s disease patients and normal elderly from the Sunnybrook dementia study. *J. Alzheimer’s Dis.* 43, 415–424. <https://doi.org/10.3233/JAD-132528>
- Rouhl, R.P.W., Van Oostenbrugge, R.J., Knottnerus, I.L.H., Staals, J.E.A., Lodder, J., 2008. Virchow-Robin spaces relate to cerebral small vessel disease severity. *J. Neurol.* 255, 692–696. <https://doi.org/10.1007/s00415-008-0777-y>
- Sepehrband, F., Barisano, G., Sheikh-Bahaei, N., Cabeen, R.P., Choupan, J., Law, M., Toga, A.W., 2019. Image processing approaches to enhance perivascular space visibility and quantification using MRI. *Sci. Rep.* 9, 1–12. <https://doi.org/10.1038/s41598-019-48910-x>
- Sepehrband, F., Barisano, G., Sheikh-bahaei, N., Cabeen, R.P., Lynch, K.M., Crawford, M.S., Lan, H., 2021. Volumetric Distribution of Perivascular Space in Relation to Mild Cognitive Impairment. *Neurobiol. Aging* 99, 28–43. <https://doi.org/10.1016/j.neurobiolaging.2020.12.010>. Volumetric
- Shaw, R., Sudre, C.H., Varsavsky, T., Ourselin, S., Cardoso, M.J., 2020. A k-Space Model of Movement

- Artefacts: Application to Segmentation Augmentation and Artefact Removal. *IEEE Trans. Med. Imaging* 0062, 1–1. <https://doi.org/10.1109/tmi.2020.2972547>
- Valdes Hernandez, M. d. C., Piper, R.J., Wang, X., Deary, I.J., Wardlaw, J.M., 2013. Towards the automatic computational assessment of enlarged perivascular spaces on brain magnetic resonance images: A systematic review. *J. Magn. Reson. Imaging* 38, 774–785. <https://doi.org/10.1002/jmri.24047>
- Valdés Hernández, M. del C., Ballerini, L., Glatz, A., Muñoz Maniega, S., Gow, A.J., Bastin, M.E., Starr, J.M., Deary, I.J., Wardlaw, J.M., 2020. Perivascular spaces in the centrum semiovale at the beginning of the 8th decade of life: effect on cognition and associations with mineral deposition. *Brain Imaging Behav.* 14, 1865–1875. <https://doi.org/10.1007/s11682-019-00128-1>
- Valdés Hernández, M.D.C., Ferguson, K.J., Chappell, F.M., Wardlaw, J.M., 2010. New multispectral MRI data fusion technique for white matter lesion segmentation: Method and comparison with thresholding in FLAIR images. *Eur. Radiol.* 20, 1684–1691. <https://doi.org/10.1007/s00330-010-1718-6>
- Vilanova, A., Zhang, S., Kindlmann, G., Laidlaw, D., 2006. An introduction to visualization of diffusion tensor imaging and its applications. *Math. Vis.* 0, 121–153. https://doi.org/10.1007/3-540-31272-2_7
- Wardlaw, J.M., Bastin, M.E., Valdés Hernández, M.C., Maniega, S.M., Royle, N.A., Morris, Z., Clayden, J.D., Sandeman, E.M., Eadie, E., Murray, C., Starr, J.M., Deary, I.J., 2011. Brain aging, cognition in youth and old age and vascular disease in the Lothian Birth Cohort 1936: Rationale, design and methodology of the imaging protocol. *Int. J. Stroke* 6, 547–559. <https://doi.org/10.1111/j.1747-4949.2011.00683.x>
- Wardlaw, J.M., Benveniste, H., Nedergaard, M., Zlokovic, B. V., Mestre, H., Lee, H., Doubal, F.N., Brown, R., Ramirez, J., MacIntosh, B.J., Tannenbaum, A., Ballerini, L., Rungta, R.L., Boido, D., Sweeney, M., Montagne, A., Charpak, S., Joutel, A., Smith, K.J., Black, S.E., 2020. Perivascular spaces in the brain: anatomy, physiology and pathology. *Nat. Rev. Neurol.* 16, 137–153. <https://doi.org/10.1038/s41582-020-0312-z>
- Wardlaw, J.M., Smith, E.E., Biessels, G.J., Cordonnier, C., Fazekas, F., Frayne, R., Lindley, R.I., O'Brien, J.T., Barkhof, F., Benavente, O.R., Black, S.E., Brayne, C., Breteler, M., Chabriat, H., DeCarli, C., de Leeuw, F.E., Doubal, F., Duering, M., Fox, N.C., Greenberg, S., Hachinski, V., Kilimann, I., Mok, V., Oostenbrugge, R. van, Pantoni, L., Speck, O., Stephan, B.C.M., Teipel, S., Viswanathan, A., Werring, D., Chen, C., Smith, C., van Buchem, M., Norrving, B., Gorelick, P.B., Dichgans, M., 2013. Neuroimaging standards for research into small vessel disease and its contribution to ageing and neurodegeneration. *Lancet Neurol.* 12, 822–838. [https://doi.org/10.1016/S1474-4422\(13\)70124-8](https://doi.org/10.1016/S1474-4422(13)70124-8)
- Wheater, E., Shenkin, S.D., Muñoz Maniega, S., Valdés Hernández, M., Wardlaw, J.M., Deary, I.J., Bastin, M.E., Boardman, J.P., Cox, S.R., 2021. Birth weight is associated with brain tissue volumes seven decades later but not with MRI markers of brain ageing. *NeuroImage Clin.* 31. <https://doi.org/10.1016/j.nicl.2021.102776>
- Yang, E., Gonuguntla, V., Moon, W.J., Moon, Y., Kim, H.J., Park, M., Kim, J.H., 2021. Direct rating estimation of enlarged perivascular spaces (Epvs) in brain MRI using deep neural network. *Appl. Sci.* 11. <https://doi.org/10.3390/app11209398>
- Zaitsev, M., Maclaren, J., Herbst, M., 2015. Motion artifacts in MRI: A complex problem with many partial solutions. *J. Magn. Reson. Imaging* 42, 887–901. <https://doi.org/10.1002/jmri.24850>
- Zhang, Y., Brady, M., Smith, S., 2001. Segmentation of brain MR images through a hidden Markov

random field model and the expectation-maximization algorithm. IEEE Trans. Med. Imaging 20, 45–57. <https://doi.org/10.1109/42.906424>

Zong, X., Nanavati, S., Hung, S.C., Li, T., Lin, W., 2021. Effects of motion and retrospective motion correction on the visualization and quantification of perivascular spaces in ultrahigh resolution T2-weighted images at 7T. Magn. Reson. Med. 86, 1944–1955. <https://doi.org/10.1002/mrm.28847>